UNIVERSITY OF PAVIA

Cyber risk, operational risk and digital currency: An econometric analysis.

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Electronics, Computer Science and Electrical Engineering

By

Thomas Edward Leach

Supervisor: Paolo Giudici

September 2021

ACKNOWLEDGEMENTS

I would first and foremost like to thank my supervisor Paolo Giudici for his continued guidance, support and encouragement throughout the PhD. Not only with my research but also in improving my 10k times as a runner.

I wish to thank Leonardo Gambacorta and Iñaki Aldasoro from the Bank of International Settlements. They kindly provided data used in my projects and have also been a great help and source of guidance throughout the PhD. I would like to thank other co-authors that it has been a pleasure to collaborate with on other projects throughout my PhD that have not featured in this thesis, namely, Raphael Auer, Rodney Garratt and Simona Ramos. I would like to thank Rob Patalano for giving me the chance to spend a period in his team at the OECD which gave me exposure to overarching policy matters on the topics of digital currency and BigTech.

I would like to thank all of my friends and colleagues at the various places I have spent time during this experience, Pavia, Paris and Basel. Finally, I would like to thank my family and friends back home.

DECLARATION OF AUTHORSHIP

I declare that this thesis was written and composed by myself and is the result of my own work unless clearly stated and referenced. This thesis has not been submitted for any other degrees or professional qualifications.

The work presented in Chapters 1 and 2 is based on my work with my supervisor, Paolo Giudici and two co-authors from the Bank of International Settlements, Leonardo Gambacorta and Iñaki Aldasoro who agreed that the essay can appear within this thesis and that it represents a substantial contribution on my part. Specifically, I undertook out all of the empirical work, coding, analysis of the results and writing of the paper. All co-authors provided guidance, direction were also responsible for edits to some of my original text.

The work presented in Chapters 3 is based on my work with my supervisor, Paolo Giudici and fellow PhD colleague from the University of Pavia Paolo Pagnottoni, who agreed that the essay can appear within this thesis and that it represents a substantial contribution on my part. I carried out major parts of the empirical analyses and contributed to writing up the results. In particular, I obtained the data and solved the optimisation problem, to which the results are contained in 3.4.2.

Thomas Leach

ABSTRACT

The chapters of this thesis comprise three separate studies on topics in cyber risk, operational risk and digital currencies. The first chapter discusses the impact cyber risk to firms and factors that play a role in mitigating or exacerbating costs. The second chapter focuses on the wider operational risks that face firms in the financial sector with additional attention paid to cyber risk. In the third chapter, I look at the design of basket-based digital currencies, their statistical properties and some of the policy implications.

Chapter 1: The drivers of cyber risk Cyber incidents are becoming more sophisticated and their costs difficult to quantify. Using a unique database of cyber events across sectors in the US, we document the characteristics and drivers of cyber incidents. Cyber costs are higher for larger firms and for incidents that impact several organisations simultaneously. Events with malicious intent (i.e. cyber attacks) tend to be less costly, unless they are on the upper tail of the loss distribution. The financial sector is exposed to a larger number of cyber attacks but suffers lower costs, on average. The use of cloud services is associated with lower costs, especially when cyber incidents are relatively small. As cloud providers become systemically important, cloud dependence is likely to increase tail risks. Finally, we document that higher expenditure on IT is associated with future reduced costs from cyber incidents.

Chapter 2: Operational and cyber risk in the financial sector This paper uses a unique cross-country dataset at the loss event level to document the evolution and characteristics of banks' operational risk. Operational risk capital varies substantially – from 2% to 12% of total gross income – depending on the method used, and shows a growing cyber risk component. It takes, on average, more than a year for operational losses to be discovered and recognised in the books. We

show that operational losses depend on macroeconomic conditions and the regulatory environment. Periods of excessively accommodative monetary policy are followed by larger operational losses. Stronger supervision is associated with lower operational losses.

Chapter 3: Libra or Librae: Digital currency baskets. In this part of the thesis, with my coauthors, I attempt to analyse, from an empirical viewpoint, the advantages of a stablecoin whose value is derived from a basket of underlying currencies, against a stablecoin which is pegged to the value of one major currency, such as the dollar. To this aim, we first study the optimal weights of the currencies that comprise the basket. We then employ volatility spillover decomposition methods to understand which foreign currency mostly drives the others. Our empirical findings show that our basket based stablecoin is less volatile than all single currencies. This result is fundamental for policy making, and especially for emerging markets with a high level of remittances: a librae (basket based stable coin) can preserve their value during turbulent times better than a libra (single currency based stable coin).

LAY SUMMARY

Technological developments are continuously presenting new risks to firms and consumers. As a society, we are growing increasingly dependent on digital technologies. Troves of information is stored in the cloud, artificially intelligent machines may replace human workers and payments can be made with the touch of a fingerprint on a smartphone. The papers of which this thesis is comprised look at new, and old, risks that face firms and individuals as we move to an increasingly digital world. Cyber threats are evolving in their sophistication and frequency, making it a vital topic of research. Identifying factors that help firms to better protect themselves from cyber risk can assist policy makers in encouraging firms to adopt practices that are strengthen their resilience against such threats. Sectors of the economy that provide critical infrastructure, like the financial sector, are core to any well-functioning economy. Finance firms have long been a target for criminals and more recently cyber-criminals, as Willie Sutton was quoted, 'because that's where the money is'. A significant cyber attack could lead to significant losses and disruption to the economy. Incumbent banks and investment firms are also confronted with the recent rise of FinTech and rapid financial innovations that are sparking them to undertake more investment into new technologies to keep pace with smaller agile and innovative FinTechs. A particular aspect of finance that is undergoing considerable change is payments. The use of cash is rapidly declining across developing countries as new digital payment technologies make it more convenient for consumers to pay with their smartphones. Not only are FinTechs are trying to amass market share in this space but also larger tech firms are creeping into this space, commonly referred to as BigTech. Google and Apple already offer payment apps for mobile phone users. Facebook, has recently put forth its own digital money scheme that leverages on the idea of a multi-currency backed asset.

The chapters that follow focus on three themes. Specifically, cyber risk, operational risk and

digital currencies. The first chapter discusses the impact cyber risk to firms across different sectors of the economy and the factors that are associated with mitigating or exacerbating costs. The second chapter focuses on the wider operational risks that face banks and investment firms in the financial sector with additional attention paid to cyber risks. The third chapter looks at the design of basket-based digital currencies, their statistical properties and some of the policy implications.

Chapter 1 focuses on the ominous threat of cyber risk that confronts firms globally today. With my co-authors we use a unique database of cyber incidents across all sectors in the US to document trends in cyber risk and identify the drivers of such risks. The frequency of cyber incidents rose strongly in the decade leading up to 2016, but has since seen a slow down. This reduction could reflect an increased investment in cyber security, but may also be the product of delays in discovery of events. Certain economic sectors display a greater resilience to cyber incidents: for example, the financial sector has experienced a higher frequency of cyber incidents but appear, on average, to be less costly. Data breaches emerge as one of the costlier types of incidents that firms face, compounding that concern is the fact data breaches also appear to occur relatively frequently.

Using a linear regression framework we identify the key drivers contributing to the costs of cyber-related events. Firm size – measured in terms of total revenues – is positively correlated with the average cost of an event, implying that larger firms tend to incur larger costs, although they are marginally decreasing. Cyber events can impact multiple firms simultaneously, creating a contagion effect. The results of the regression suggest that events that events associated with multiple entities are also associated with higher costs. Cyber-related incidents may occur unintentionally – e.g. a bug in some internally developed software – or be caused by an actor with malicious intent. We find, malicious cyber attacks have, on average, lower costs. However, a quantile analysis reveals that at the tail of the sample distribution this result is reversed and in fact malicious incidents are associated with higher costs. This finding indicates that, while most at-

tackers are stopped before they can do considerable harm, a successful attacker can go on to cause extensive damage.

We then study the effects of reliance on cloud services and digital technologies more broadly. Cloud technologies have become synonymous with cyber risk as policy institutions grapple with the consequences of having centralised IT storage infrastructures. In principle, reliance on the same service provider by multiple organisations can yield positive externalities by fostering economies of scale and information sharing (Rowe, 2007). Cloud technology can thus reduce IT costs, improve resilience and enable firms to scale better (Financial Stability Board, 2019). However, it also strengthens interdependence, not least given the high concentration of the market for cloud service providers. Our results suggest that, at present, the former effect dominates as firms that could have higher exposure to cloud technology experience lower costs. Whilst this is a promising result for firms adopting the technology, we urge caution as the data may not capture the real 'tail' of these incidents.

This paper also uses data on the level of IT spending across sectors to assess the relationship between investment in IT and the cost of cyber incidents. This analysis can act as a helpful indicator to policymakers as to which sectors may be exposed due to underinvestment in their IT systems. We find that higher expenditure in IT is associated with lower costs at the mean and at the tail of the distribution. Sectors that appear to benefit from this higher level of spending include the manufacturing and the finance and insurance sector. The dividend of such investments is evident through our additional analyses, whereby an annual increase in IT investment is associated with a reduction of costs in the subsequent year.

In Chapter 2 the focus turns to financial sector specifically and the wider set of risks under the umbrella of operational risk. Measuring and understanding operational risk is critical for both banks and public authorities. Operational risk currently represents a significant portion of banks'

risk-weighted assets, second only to credit risk.[1] Regulators, central banks and international organisations, in turn, place the understanding and mitigation of operational risk – and subcomponents such as cyber risk – high in their agendas. While banks use internal data to determine their regulatory capital, there is limited work to identify the relationship between operational risk and the macroeconomic and supervisory environments – not least in an international context. Accordingly, policy discussions on the topic at the wider macroeconomic level tend to lack substantial empirical grounding.

This paper, contributes to filling this gap by analysing a unique cross-country dataset of operational losses. We use data at the loss event level from ORX, a consortium of financial institutions, that facilitates the sharing operational loss risk data in an anonymised fashion in order to benchmark operational risk models. We document that, after a notable increase post-Great Financial Crisis (GFC), operational risk losses in banks have been declining strongly since 2015. Digging deeper in to the type of event behind this aggregate trend shows that one category in particular is responsible for the pattern in cost, namely "Clients, Products & Business Practices". This category includes improper business practices like fiduciary breaches, aggressive sales, breaches of privacy, account churning and misuse of confidential information. These are the type of operational risks that characterise periods of financial excess, with mis-selling of mortgage-backed securities in the mid-2000s being a prime example. Towards the peak of the GFC there is a significant increase in the occurrence of this type of event (especially in North America), which were then recognised in the books of banks a few years later. Importantly, this pattern is observed only in terms of loss amounts and not in terms of frequency of occurrence.

Operational losses are characterised by a fat-tailed distribution.[2] Accordingly, operational risk

[1] Up to 40% of risk-weighted assets can be attributed to operational risk in some jurisdictions (Sands et al. (2018)).

[2] In other words, there are a large number of inconsequential events from a cost perspective and a limited number of very large cost events. The latter group in particular complicates the quantification of operational risks, as such low frequency/high severity events are often cited as being "one-in-a-hundred years" events.

capital estimates can lead to quite different results depending on the method used. Capturing the distribution the extreme values of the distribution of operational losses is a challenge. Indeed, our estimates for operational risk capital using methodologies from the Advanced Measurement Approach (AMA) range from 6 2% to 7% of gross income, against the 12% benchmark of the Basic Indicator Approach. This finding provides some support for the new regulatory framework that proposes the adoption of the Standardised Measurement Approach (SMA) for all banks.

The stylised facts we present point to the existence of a link between operational losses and macroeconomic conditions. Abdymomunov et al. (2017) use data for US banks to document a contemporaneous correlation between macroeconomic conditions and operational risk losses, e.g. operational losses rise during economic downturns. We build on this idea and use a cross-country panel analysis to argue that the ultimate cause of the rising losses during economic downturns lies in the excesses characterising the run-up to the downturn. In other words, favourable conditions during periods of macroeconomic expansion and financial exuberance lead to the occurrence of events that are only discovered when the economic tide turns, and recognised in the books of banks even later.

Using deviations of policy rates from Taylor-rule implied benchmarks, we show that periods of accommodative monetary policy are followed by an increase in operational losses. This appears to be driven by the frequency rather than the severity of events. Periods of excessively accommodative monetary policy can lead to increased risk-taking by banks, which can boost the type of improper business practices that account for the lion's share of operational losses. Finally, in line with the work of De Nicolò and Lucchetta (2013), who find that banks in a higher competition environment increase monitoring efforts and reduce risks, and with Kim (2018) who finds that banks with lower market power take less liquidity risk, we find that periods of intense bank competition are also associated with lower operational losses.

Regulation can also play a role in moderating operational losses. The time pattern of losses stemming from internal fraud and improper business practices suggests that the quality of regulation and supervision can also be related to operational losses in the cross-section of countries. Indeed, we find that better regulation and supervision – as captured by the financial reform index of Abiad et al. (2010) and Denk and Gomes (2017) – is associated with lower operational losses.

The fallout of the financial crisis attracted attention to operational losses caused by people. However, as society moves to a digital age, retail banks are moving from the high street to the world wide web, intensifying interconnectedness through technology. This has led to a growing focus and concerns regarding cyber and IT-related risks. We use the data to construct a proxy range of cyber losses (considered as a subset of operational losses). We document that cyber losses, to date, represent a relatively small share of operational losses. In recent years, however, losses from cyber events saw a spike which aligns with the growing attention cyber risk has been receiving. Despite representing a relatively small share of operational losses, cyber value-at-risk can account for up to a third of total operational value-at-risk.

The third chapter is dedicated to the topic of digital currency and in particular digital currency baskets. Central banks are continuing to grapple with the concept of central bank digital currencies and continue to experiment with their implemention. Meanwhile, the private sector has already begun offering their own solutions to digital money to the public. These are often collected under the term *stablecoin*. In this part of the thesis, with my co-authros, I attempt to analyse, from an empirical viewpoint, the advantages of a stablecoin whose value is derived from a basket of underlying currencies, against a stablecoin which is pegged to the value of one major currency, such as the dollar.

First, we consider the optimal weights of the basket of underlying reference currencies, such as those included in the International Monetary fund Special Drawings Rights (SDR). After com-

puting the optimal weights we construct the historical values of the designed stablecoin (SAC) and compare its volatility against a set of major currencies. For the optimal allocation of weights in the currency basket we follow Hovanov et al. (2004) to compute a currency invariant index. A particular advantage of this approach is that given a fixed set of currencies, the index of a currency will have the same value, regardless of the base currency choice. The index is determined by minimising the variance of a portfolio of currencies, expressed in Reduced Normalised Values (RNVALs). We construct a reference basket that contains the Dollar (USD), the Euro (EUR), the Yen (JPY), the Renminbi (CNY) and the Pound Sterling (GBP), the same currencies that are employed for the determination of the IMF's Special Drawing Rights (SDR) basket. Our empirical findings show that, overall, the stablecoin maintains the lowest volatility, thus could act as a hedge and store of value for overseas workers savings. We also make comparison to the IMF's SDRs, which performs almost as well.

To gain insight into the composition of the basket, we study the currencies which drive the volatility spillovers among exchange rates, using the framework of Diebold and Yılmaz (2014). Specifically, we build a spillover network decomposition analysis of the currencies up to December 2020, thus including the period of the Covid-19 outbreak. Our spillover network decomposition shows that the USD is the currency whose dynamics has the largest impact on the others, especially in terms of exporting contagion. As a consequence, a shock to USD, causes a shock on all currencies that leads to a new lower equilibrium.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

THE DRIVERS OF CYBER RISK

1.1 Introduction

Information technology (IT) has become a critical component of well-functioning economies, underpinning economic growth over the past decades. Organisations of all sizes in both the public and private sector are becoming ever more interconnected and reliant on IT products and services, such as cloud-based systems and artificial intelligence. Accordingly, there is a growing exposure to cyber risks, and public awareness of these threat has been on the rise (see Figure 1.1). Cyber risk commonly refers to the risk of financial loss, disruption or reputational damage to an organisation resulting from the failure of its IT systems.[1] The increasing reliance on cloud technologies exacerbates these risks, as it increases interdependence across firms that have shared exposures to similar (or even the same) cloud service providers.

Firms actively manage cyber risk and invest in cyber security. However, cyber costs are difficult to quantify.[2] In the financial sector, cyber risks are a key "known unknown" tail risk to the system and a potential major threat to financial stability.[3] More broadly, cyber risk in sectors that play a critical role in the economic infrastructure could have systemic implications and can be viewed as

[1] These episodes include malicious cyber incidents (cyber attacks) where the threat actor intends to do harm (e.g. ransomware attacks, hacking incidents or data theft by employees). High-profile attacks such as the WannaCry incident in May 2017 contributed to the growing concern around cyber risk.

[2] The high degree of uncertainty and variability surrounding cost estimates for cyber security incidents has consequences for policy-makers. For example, it is difficult to foster robust insurance markets, as well as to make decisions about the appropriate level of investment in security controls and defensive interventions (Biener et al., 2015; Wolff and Lehr, 2017).

[3] In March 2017, the G20 Finance Ministers and Central Bank Governors noted that "the malicious use of information and communication technologies could disrupt financial services crucial to both national and international financial systems, undermine security and confidence, and endanger financial stability".

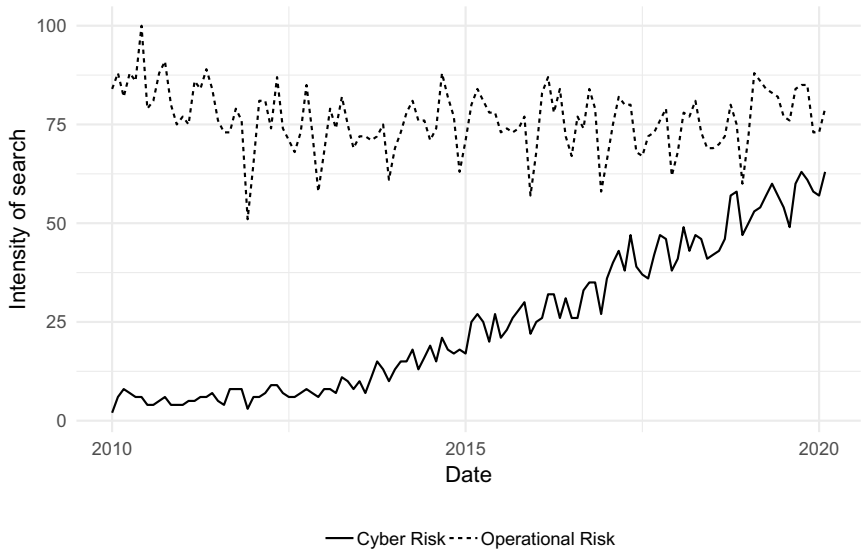

Notes: Number of online searches for "cyber risk" and "operational risk" over the last decade. Worldwide search interest is relative to the highest point (=100). Data accessed on 7 Feb 2020.
Source: Google Trends.

Figure 1.1: Interest on cyber risk is on a par with operational risk.

a matter of national security (Brenner, 2017). Despite such considerations, information concerning the costs, drivers and potential mitigating factors of cyber incidents is relatively scarce.

This paper seeks to help fill this gap by using a sample of 3,705 cyber events across all economic sectors in the US, sourced from the Advisen cyber loss database. We document a series of stylised facts. The frequency of cyber incidents rose strongly in the decade leading up to 2016, but has since moderated somewhat. This reduction could reflect increased investment in cyber security, but also delays in discovery or reporting.[4] We find that certain economic sectors display a greater resilience to cyber incidents: for example, the financial sector has experienced a higher frequency of cyber incidents but these have been on average relatively less costly. Regarding the

[4]This phenomena is widely recognised in the operational risk literature (see Aldasoro et al. (2020); Carrivick and Cope (2013). The dataset used here does not allow us to accurately estimate such "end-of-sample" bias.

type of incident, data breaches, phishing or skimming and security incidents, appear to be most costly. Of particular concern is that data breaches are not only costly, but also relatively frequent.

The paper then documents cyber risk drivers. We first identify the key drivers contributing to the costs of cyber-related events. Firm size – measured in terms of total revenues – is positively correlated with the average cost of an event, implying that larger firms tend to incur larger costs. However, the elasticity is quite low: a 1% increase in total revenues is associated with a 0.2% increase in cyber costs. Cyber events impacting multiple firms at the same time (i.e. "connected" events) are also associated with higher costs. Cyber-related incidents can occur unintentionally – e.g. a bug in some internally developed software – or can also be caused by an actor with malicious intent.[5] Malicious cyber attacks have, *on average*, lower costs. However, a quantile analysis reveals that at the tail of the sample distribution this relationship is reversed and in fact malicious incidents are associated with higher costs. This finding indicates that, while most attackers are stopped before they can do considerable harm, a successful attacker can go on to cause extensive damage.

We then study the effects of reliance on cloud services and digital technologies more broadly. In principle, reliance on the same service provider by multiple organisations can yield positive externalities by fostering economies of scale and information sharing (Rowe, 2007). Cloud technology can thus reduce IT costs, improve resilience and enable firms to scale better (Financial Stability Board, 2019). However, it also strengthens interdependence, not least given the high concentration of the market for cloud service providers. By analysing the cost-benefit trade-off, we

[5]The best known types of cyber attack are: man-in-the-middle attacks, cross-site scripting, denial-of-service attacks, password attacks, phishing, malware and zero-day exploits. Man-in-the-middle attacks occur when attackers insert themselves into a two-party transaction. Cross-site scripting is a web security vulnerability that allows attackers to compromise the interactions a victim has with a vulnerable application. Denial-of-service attacks flood servers with traffic to exhaust bandwidth or consume finite resources. Phishing is the practice of stealing sensitive data by sending fraudulent emails that appear to be from a trustworthy source. Malware (i.e. "malicious software") is a software designed to cause damage to IT devices and/or steal data (examples include so-called Trojans, spyware and ransomware). A zero-day exploit is an attack against a software or hardware vulnerability that has been discovered but not publicly disclosed.

find that the use of cloud services is associated with lower costs of cyber events. While this speaks to the resilience of cloud technology, it should be interpreted with caution. As firms' exposure to cloud services continues to increase and cloud providers become systemically important, cloud dependence is likely to increase tail risks (Danielsson and Macrae, 2019).

Finally, we use data on the level of IT spending across sectors to assess the relationship between investment in IT and the cost of cyber incidents. This analysis can act as a helpful indicator to policymakers as to which sectors may be exposed due to underinvestment in their IT systems. We find that higher expenditure in IT is associated with lower costs at the mean and at the tail of the distribution. Sectors that appear to benefit from this higher level of spending include the manufacturing and the finance and insurance sector. For the latter this could be due to the effects of regulation and a years of experience being a prime target for cyber criminals. The dividend of such investments is evident through our additional analyses, whereby an annual increase in IT investment is associated with a reduction of costs in the subsequent year.

The rest of the paper is organised as follows. Section 1.2 discusses related literature. Section 1.3 contains a description of the data. Section 1.4 discusses our baseline results. 1.5 explores whether exposure to cloud services affects the cost of cyber events. Section 1.6 analyses the optimal amount of IT spending across sectors. Finally, Section 1.7 concludes.

1.2 Related literature

Most of the few empirical studies on cyber risk rely on collected publicly available data sources. Goldstein et al. (2011) study how the exposure to IT operational risk, or the risk of failures of operational IT systems, could translate into significant losses in firms' market value. Biener et al. (2015) emphasise the distinct characteristics of cyber risks compared to other operational risks. The presence of highly interrelated cyber losses, lack of data, and severe information asymmetries,

hinder the development of a sustainable cyber insurance market, an essential element to encourage improvements in cyber resilience. Romanosky (2016) and Chande and Yanchus (2019) use the Advisen dataset to study losses from cyber events across sectors and provide an initial estimate of firm risk by sector. Our paper builds on their work by looking at how characteristics of sectors' management of IT resources can mitigate costs.

An important part of the cost of cyber events is arguably given by the reputational component, which is notably hard to assess. Makridis (2020) find that for the subset of the largest data breaches, brand power (a survey-based measure of reputation) and familiarity decrease by 5-9% after the event (whereas they increase by 26-29% for an average data breach). Further Kamiya et al. (2021) find that a successful data breach can decrease shareholder wealth by 1.09% in the three-day window around the cyber attack. Their findings suggest economically large reputation costs, in that the shareholder wealth loss far exceeds the out-of-pocket costs from the attack.

The literature highlights that the observed heterogeneity in cyber costs across sectors heavily depends on the environment in which each firm operates as well as IT security investments. Kamiya et al. (2018) find that cyber attacks are more likely in industries that face less intense product market competition and in industries with higher growth opportunities. Moreover, controlling for firm characteristics, they find that, among the major industries, cyber attacks are more likely in service industries, wholesale/retail trade, and transportation and communications. Makridis and Dean (2018) find heterogeneity in cyber attack episodes amongst sectors when it comes to data breaches. In particular, companies in the finance, insurance, retail and merchant sectors are the biggest targets. Makridis and Liu (2021) also suggest that higher productivity firms have fewer cyber security vulnerabilities and are able to gain access to more human capital that is better capable of mitigating cyber security vulnerabilities.

Regulation can also play a key role in firms' motives for security investments. Based on a

survey of more than 700 firms, Rowe and Gallaher (2006) find that the vast majority believe that regulation has increased the overall level of security. However, some firms reject this view, because excessive cyber security costs imposed by regulation could stifle firms' ability to innovate (Etzioni, 2011). While our paper does not enter into the debate on who should bear the cost of cyber security, we find that sectors with a more robust policy framework toward cyber risk tend to reap benefits by reducing the costs of cyber incidents.

Some sectors provide critical infrastructure for the functioning of the economy. Cyber attacks on the financial sector could create cascade failures that are not completely understood nor adequately quantified by sector-specific simulations (Brenner, 2017). Kopp et al. (2017)) note that the financial sector is frequently targeted due to its high exposure to IT and its credit intermediation role. Kashyap and Wetherilt (2019) outline some principles for regulators to consider when regulating cyber risk in the financial sector. The Basel Committee has also published guidelines for banks regarding best practice regarding cyber risk.[6] Given that financial institutions tend to maintain better data collection practices due to regulatory reporting, empirical studies focusing on this sector are more developed.

Using a large cross-country panel, Aldasoro et al. (2020) find that cyber losses represent a relatively small share of operational losses for banks. In recent years, however, losses from cyber events saw a spike, with a corresponding increase in risk. The value-at-risk (VaR) associated with cyber events can range from 0.2% to 4.2% of banks' income.[7] This amounts to around a third of operational VaR, despite representing a minor share of the latter in terms of frequencies and loss amounts. The extent of operational and cyber losses depends on the supervisory environment. A higher quality of supervision – as measured by a financial and supervisory quality index – is

[6] See Basel Committee on Banking Supervision (2018a)

[7] Estimates by Bouveret (2018) – based on data collected from media and newspaper articles across countries – point to sizeable potential losses in the financial sector. His estimate of value-at-risk ranges between 14% to 19% of net income.

associated with lower losses, in terms of both frequency and amount. Credit booms and periods of accommodative monetary policy are associated with higher operational losses in the future, but have no effect on cyber losses.

Duffie and Younger (2019) analyse a sample of twelve systemically important U.S. financial institutions and suggest these firms have sufficient stocks of high-quality liquid assets to cover wholesale funding runoffs in a relatively extreme cyber event.[8] From the literature on operational risk, the size of financial institutions is positively linked with the size of operational losses (Shih et al., 2000; Curti et al., 2019b). A large share of banks' operational losses can be traced to a breakdown of internal controls (Chernobai et al., 2011). We devote particular attention to the drivers of cyber risks in the financial sector and how these could differ from other economic sectors.

1.3 Data

The data are obtained from Advisen, a for-profit organisation which collects information from reliable and publicly verifiable sources such as websites, newsfeeds, specialised legal information services, multiple online data breach clearinghouses and federal and state governments in the United States.[9] The entire Advisen database contains a total of 137,164 cyber incidents. Each cyber incident is linked to an ultimate parent company and includes, amongst others, the following characteristics: i) case type (e.g. data breach, phishing); ii) affected count (e.g. in the event of a data breach, how many details were stolen); iii) accident date; iv) source of the loss; v) type of loss; vi) actor (e.g. state-sponsored, terrorist, etc); vii) loss amount; viii) company size (proxied by total revenues); ix) company type (e.g. government, private); x) number of employees; xi) North Amer-

[8] Using a broader network of US banks, Eisenbach et al. (2021) find that the impairment of any of the five most active US banks will result in significant spillovers to other banks, with 38 percent of the network affected on average.

[9] Most cyber incidents go unreported. Typically, only the larger and the more relevant ones become public and are included in the Advisen database.

ican Industry Classification System (NAICS) code identifying the sector of the firm that suffered the cyber incident; and xii) geography (i.e. the area where the incident occurred).

The majority of events reported in the database occur in the Americas region (North, Central and South America). In particular, 86 per cent of the episodes took place in the United States. This is largely due to the fact that information regarding cyber losses is easier to collect there as a result of a higher degree of freedom of information. To remove unobserved country heterogeneity from our analysis we focus on the incidents in the database that occurred in the United States, which leaves a sample of 116,387 incidents. This is the dataset we use when reporting stylised facts that do not require information on loss amounts. However, due to the nature of how the data are collected, it is not possible to obtain all information desirable for each event.

Data on actual loss amounts per event represent only a subset of the larger database. The cost of cyber events can be categorised into three components (Anderson et al., 2019). The *direct cost* is the value of loss, damage and other suffering incurred by the victim of the cyber incident. The *indirect costs* are the losses and opportunity costs borne by society as a consequence of a cyber incident.[10] Firms also bear *mitigation costs*, which include inter alia investment in IT personnel or in security products such as antivirus or cyber threat awareness training for staff. The data from Advisen can best be interpreted as a measure of *direct costs* to a firm as a result of a cyber incident. Individual components of each loss (e.g., fines or penalties from regulators, payments made to a plaintiff in the event of a claim and financial damages) are provided in the data, but are rarely populated in sufficient detail to allow for a meaningful analysis. For our regression analysis, we remove observations missing such critical data, which leaves us with a sample of

[10]Examples of direct costs are those related to the time and effort of repairing IT systems damaged as a result of an incident, the ransom paid to attackers in a successful cyber attack or regulatory fines and penalties. Indirect costs could in turn include reduced uptake by citizens of electronic services whether from companies or governments due to the perceived threat of a cyber incident or the losses incurred by an individual after having their personal data stolen.

3,705 observations for our baseline empirical analysis.[11]

The frequency and costs of cyber events differ across sectors (see Table A.1.2 in the appendix for summary statistics).[12] By frequency, "Financial and insurance activities" (FI) is the most affected sector. However, it shows some resilience, as despite being subject to many attacks, the average cost of a cyber incident is not as high as for other sectors. The sector with the highest average costs is "Wholesale Trade", followed by "Transportation and storage" and "Professional, Scientific and Technical" (PST). The standard deviation of costs across sectors is quite large, implying that most likely the distribution of losses has a heavy tail.

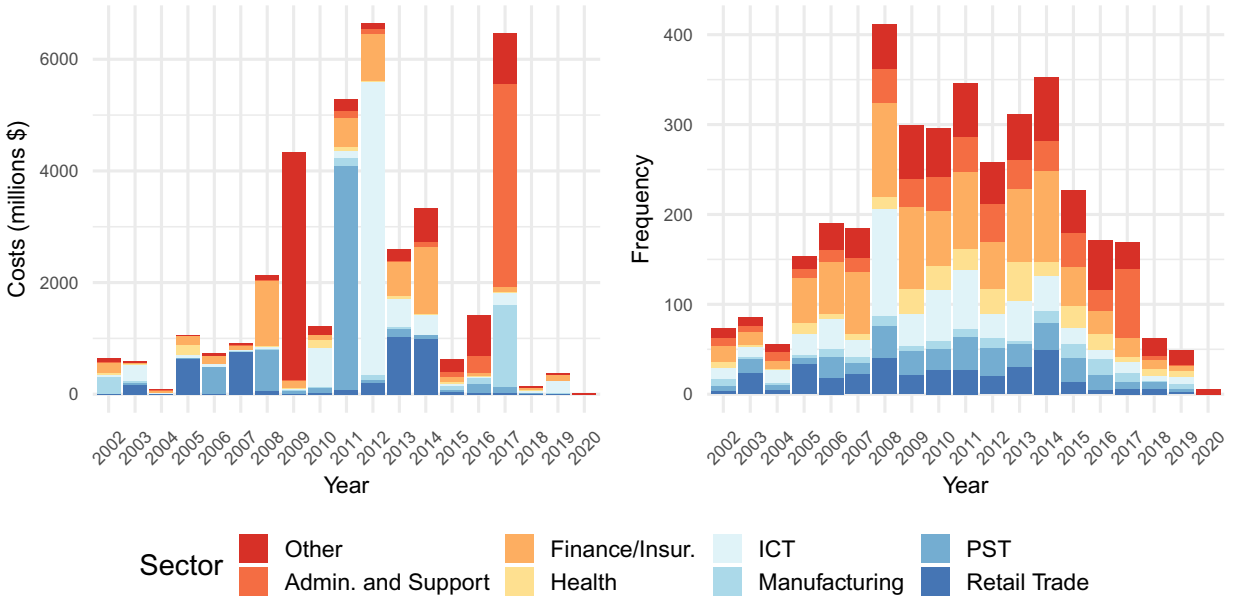

Notes: ICT stands for "Information and Communication Technology", while PST for "Professional Scientific and Technical". The plot shows costs and frequency over time with partitioning by the sectors that suffered the most incidents based on Table A.1.2 in the appendix, remaining sectors are subsumed into the "Other" category.

Figure 1.2: Frequency and cost of cyber incidents across sectors

[11] As we note below, the sample used for the regression analysis (i.e., which includes loss data) is not biased to any particular type of cyber event or sector. This can be seen by comparing the summary statistics from the regression sample (Table 1.1) with those from the full sample (Table A.1.1 in the annex).

[12] The sectors are based on NAICS. For details, see `https://www.census.gov/eos/www/naics/`.

Figure 1.2 shows how events are distributed by sector over time. The overall distribution across sectors in terms of frequency remains relatively stable. Much of the growing frequency of events can be attributed to the FI sector as well as the "Administrative and Support Service" sector. Increases in the frequency of cyber incidents in the FI sector following the great financial crisis may be partly driven by targeted attacks on banks. The peak in costs in 2012 was shared largely amongst the FI and "Information and communication technology" (ICT) sectors.

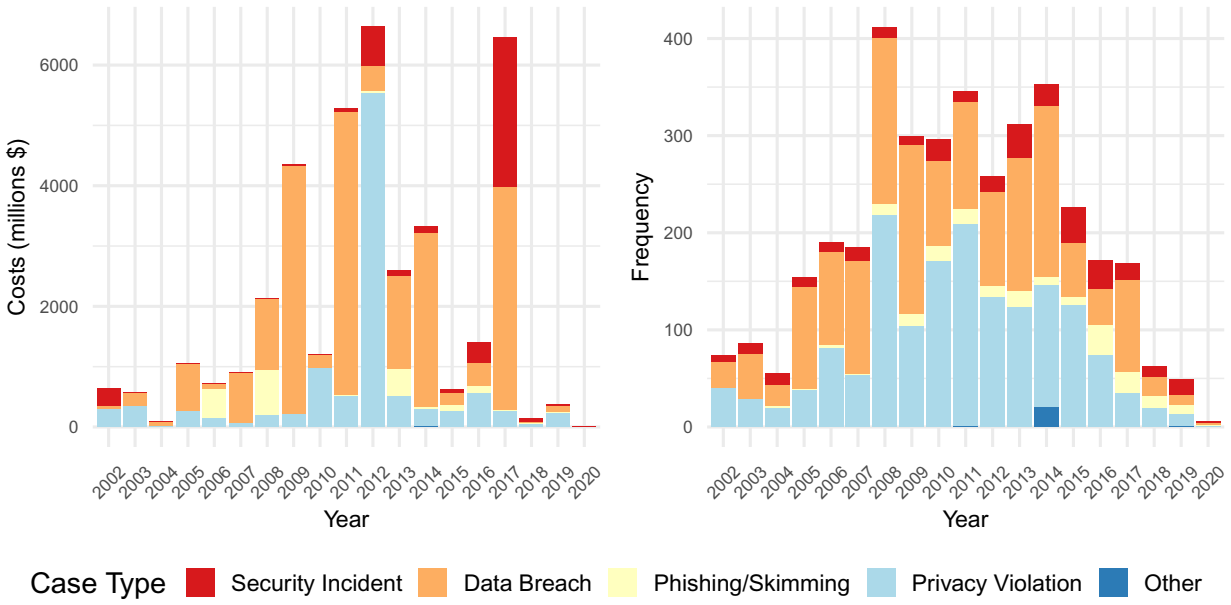

Notes: The graphs show costs and frequency over time by case type. The graphs are based on the 3,705 observations described in Table A.1.2 in the appendix.

Figure 1.3: Frequency and cost of cyber incidents by case type

Figure 1.3 depicts the distribution by case type through time.[13] Privacy violations are the frequent type (44 per cent of the cases). This is likely due to the fact that reporting requirements have been in place for a longer period for such incidents, as well as the relative ease of assigning

[13]Case types are based on the definitions of Romanosky (2016) and are used as fixed effects in our regression (see Section 1.4 for details). Further categorisations are possible, though they do not provide enough variation for econometric analysis.

conclusive responsibility when they occur (Chande and Yanchus, 2019). Data breaches have been responsible for a significant portion of costs over time. The total cost of a data breach grows with the amount of records stolen. Therefore, if hackers are able to obtain large volumes of records, the costs can soar as millions of individuals can be affected.

In terms of frequency, the overall trend has been positive, in line with the growing concern over cyber risks.[14] This is likely driven by a few factors. First, several frameworks and legislation have come into place that encourage the reporting of cyber incidents. Second, the barrier to carrying out cyber attacks has become lower as competent computing skills are no longer required to carry out attacks. The reduction in more recent years could represent the effects of increased investment in cyber security, but should be taken with caution due to the above mentioned reporting bias. Regarding the distribution by frequency over time, the increase has largely been attributed to privacy violations. Costs, on the other hand, peaked in 2011, largely due to spikes in privacy violations and data breaches.

1.4 Identifying the drivers of cyber costs

1.4.1 Empirical approach

Our analysis aims to explain the costs of cyber events by a series of event and firm/sector characteristics.[15] We model the direct costs of a cyber incident through the following regression:

[14] As noted earlier, the most recent end of the data is probably subject to an under-reporting bias, as it takes time for incidents to be discovered and acknowledged. Therefore, we expect the numbers in the most recent years of our sample to increase as more information becomes available in the future.

[15] These include observable direct costs from cyber events. Costs are likely to be a lower bound for a number of reasons. For one, there are many costs associated to any event, which may either be not easily quantifiable nor publicly reported, even if the event has some cost reported. Furthermore, as discussed earlier the costs of cyber events can also be indirect: these refer to the losses and opportunity costs borne by society as a consequence of a cyber incident, and may include hard-to-quantify damage to the reputation of a firm. Finally, companies may also incur mitigating costs.

$$C_{i,f,g} = \beta Z_{i,f,g} + \lambda W_{f,g} + \theta X_g + \eta_k + \alpha_t + u_{i,f,g} \tag{1.1}$$

where, i denotes the individual incident ($N_i = 3705$), f denotes the firm at which the incident took place ($N_f = 2445$) and g the sector of the firm ($N_g = 19$), based on the NAICS sector categorisation. $C_{i,f,g}$ denotes the cost of the incident; X_g denotes sector-level controls; $W_{f,g}$ denotes firm-level variables and $Z_{i,f,g}$ stands for variables that vary at the individual incident level. We control for year fixed effects (α_t) and for fixed effect for incident types (η_k). Finally, $u_{i,f,g}$ denotes the random error term. For clustering of standard errors we take a conservative approach in our baseline estimation by clustering at the sector level.[16]

Firm size is proxied by the revenues of the firm that suffered a cyber incident.[17] Shih et al. (2000) hypothesise a relationship between firm size and costs stemming from operational risks of the form: $C = R^{\alpha}F(\theta)$, where C denotes costs, R stands for the revenues of the firm, θ for a vector of unobserved risk factors that explain the variation in costs not attributed to revenues, and α for the degree of returns to scale in terms of costs.[18] The authors estimate (in log form) an $\hat{\alpha}$ of 0.15, with a low R^2 (0.05). Replicating this equation, we estimate $\hat{\alpha}$ to be 0.23 and similarly a low R^2 of 0.09. The $\alpha < 1$ indicates a decreasing marginal cost with respect to increases in revenues. Inspection of the residuals of this equation does not indicate obvious signs of heterogeneity across firm size or non-normality of the residuals (see the plot of the residuals in Figure A.1.1 in the

[16]The inclusion of $W_{f,g}$ and X_g implies perfect correlation within firm and sector level. Consequently, the error term will be perfectly correlated within clusters, which could lead to bias. This type of clustering has a nested structure, i.e., firm within a sector. The conventional wisdom suggests clustering at the highest level of aggregation, in this case the sector (Cameron and Miller, 2015). We present robustness to alternative clustering choices in the appendix.

[17]We test the robustness of our results by performing regressions using number of employees as an alternative proxy for firm size. Results are unaffected by this choice.

[18]With equation (1.1) we aim to capture some of the unexplained variance (θ) with the inclusion of control variables discussed below. Shih et al. (2000) posit that the unexplained part of this regression could be attributed to variation in firms' attributes regarding risk management, e.g. nature of the business, quality of internal controls, etc. Firm size may implicitly capture the difference in corporate structures and variation in management.

appendix). Contrary to some of the literature (Biener et al., 2015), we do not find evidence in favor of the existence of a U-shape relationship between firm size and the average costs from cyber incidents (details can be found in the annex).

Cyber incidents are likely to exhibit features of contagion: a failure in a firms' IT systems could have spillover effects on other firms (Baldwin et al., 2017; Eisenbach et al., 2021; Crosignani et al., 2020). Incidents that impact multiple firms could contribute to greater costs in the aggregate through other means as well. Affected firms could for instance seek damages and respond by pursuing litigation against the firm at which the incident originated, increasing the costs for the firm that originally suffered the incident. On the other hand, costs could be distributed across firms, thus lowering the average cost across affected firms. We include a variable, *connected events*, that captures how many firms were linked to one specific cyber incident to investigate this effect. To illustrate, if a hacker infiltrated one firm and subsequently managed to penetrate the system of another firm, and both firms recognise they have been affected by the same hacking incident, the *connected events* variable would be 2.[19]

We collect sector level data to estimate the impact of differences in the adoption of information technologies across firms from different sectors. We obtain two variables from the Digital Module of the 2018 Annual Business Survey undertaken by the Bureau of Economic Analysis (BEA). The first variable proxies the *digital share of business activity*. The survey asks: *In 2017, how much of each type of information was kept in digital format at this business?* We collect the percentage of firms that responded that more than 50% of their information is kept in digital format. Therefore, a higher value in this variable indicates a sector with a stronger dependence on digital technologies for its storage of information. Firms with a higher dependence on IT and digital technologies may

[19]The variable does not provide information on the relationship between the root cause and the affected parties. We note that this variable likely acts as a lower bound on the number of related incidents, as some are unable to be traced to a root cause or may have gone unnoticed or unreported by some firms.

expose themselves to more cyber risk (Florakis et al., 2020).[20]

Cloud technologies have become synonymous with cyber risk as policy institutions grapple with the consequences of having centralised IT storage infrastructures. Incidents that involve cloud technology could lead to significant "spillover costs". The survey from BEA also includes data gathered on the penetration of cloud services across sectors. The survey asks: *Considering the amount spent on each of these [IT functions] how much was spent on cloud services services [provided by a third party on–demand via the internet]?* We take an average across all IT functions and collect data on firms that indicated 50-100% was spent on cloud services. The variable proxies an indicator of sectors with a higher exposure to cloud technologies.

Cyber incidents include a broad set of *malicious* and non-malicious events. We test whether cyber attacks (malicious) cause more damage or whether inadvertent incidents are equally damaging. We divide the categorical variable of case types (e.g. DDoS attack, accidental data leak, IT processing error) into two broad categories, malicious and non-malicious, based on whether the incident was done with intent to cause damage or occurred as a result of an accident. Based on this categorisation, we construct a dummy variable labelled *Malicious*, which is equal to one if the event resulted from malicious intent. Around 44% of the incidents recorded fall within this category.

We include in equation (1.1) a set of dummy variables, η_k, for different *types of incidents*, based on the classifications in Romanosky (2016). *Security incident* relates to an incident that compromises or disrupts corporate IT systems (computers or networks) or their intellectual property – examples include hacking and extorting corporate information or a denial of service (DoS) attack. *Data breach* includes unintended disclosure of information (e.g. accidental public dis-

[20]The effect of this variable likely manifests in two ways. First, a higher "digital presence" widens the surface of attack to cyber-criminals, which may increase the likelihood of being attacked. Moreover, it suggests which sectors maintain more of their assets in digital format and thus stand to lose more given a cyber incident.

closure of customer data, improper disposal of information) and/or theft of computers containing personal information of employees or customers of a firm. *Phishing/skimming* are the sending of emails purporting to be from reputable sources in order to convince individuals to reveal personal information to subsequently commit identity theft and the illegal copying of information from the magnetic strips found on credit and debit cards (usually via hardware devices on ATM machines). *Privacy violation* refers to unauthorised collection, use or sharing of personal information – examples include unauthorised collection from cell phones, GPS devices, cookies, web tracking or physical surveillance. This is distinguished from data breaches as an act committed *by* the firm as opposed to *against* the firm. *Other* denote cyber-related losses that were not attributed to one of the above categories.

Table 1.1 contains summary statistics of the variables used in our regressions.[21] The mean cost incurred by a firm is $10.4 million, with a median of $117,000 and standard deviation of $122 million. This implies a high coefficient of variation and is indicative of the heavy-tailed nature of cyber risks. Cyber risk can be considered a subset of firms' operational risk (Aldasoro et al., 2020). The severity of operational losses is typically characterised by a set of long-tailed distributions, including the log-normal, such that $\ln(C_{i,f,g}) \sim \mathcal{N}(\mu, \sigma^2)$. Figure A.1.2 in the appendix shows the density of the costs after the log transformation has been applied. There appears to be bi-modality around the mean, but the data are approximately normally distributed.

The mean of the digital share implies that 14.8% of firms across all sectors' maintain more than 50% of their information in digital format. The value ranges from 9-24%, with sectors at the lower end of the spectrum including Construction and Transportation and Warehousing, and at the top end Manufacturing and Management of Companies and Enterprises. In turn, the average of the cloud variable is 18.6%: roughly a fifth of firms across all sectors spend upwards of 50% of their IT

[21] For comparison, in Table A.1.1 of the appendix, we provide summary statistics based on the full sample where data are available (i.e., including data without loss amounts).

budgets on cloud services. The value ranges from 6 to 26%. Sectors with a lower exposure include Agriculture, Forestry, Fishing and Hunting and Mining, Quarrying, and Oil and Gas Extraction. Those at the other end of the spectrum include Health Care and Social Assistance; Finance and Insurance; Professional, Scientific, and Technical Services; and Information.

	Mean	**Median**	**Std. dev.**	**Minimum**	**Maximum**
Variables varying at individual event level					
Costs ($ mil)	10.4	0.117	122	0[a]	5000
Connected events	4.90	2.00	9.75	0	79.0
Variables varying at firm level					
Firm size (Revenues $ mil)	12,800	27.0	41,700	0[a]	521,000
Variables varying at sector level					
Digital share of business activity	14.8	15.2	2.28	9.23	24.3
Cloud service purchases	18.6	20.1	5.66	5.60	26.2
Binary variables at event level					
Malicious Indicator	0.437	0	0.496	0	1.00
Security Incidents	0.0815	0	0.274	0	1.00
Data Breaches	0.427	0	0.495	0	1.00
Phishing / Skimming	0.0494	0	0.217	0	1.00
Privacy Violations	0.436	0	0.496	0	1.00
Other incidents	0.00648	0	0.0802	0	1.00

Notes: [a] Zeros are a consequence of rounding accuracy. The top panel reports the variables from equation (1.1) that vary with each individual event in the sample. The second panel contains variables from equation (1.1) that vary by each firm contained in the sample. The third panel are the variables from equation (1.1) that vary at the sector level and obtained from the US census Bureau 2018 Annual Business Survey. The bottom panel are dummy variables that indicate the type of the incident.

Table 1.1: Summary of variables used in the regression

1.4.2 Baseline results

The results of the baseline regressions are presented in Table 1.2. The cost of a cyber attack is positively correlated with both firm size and the number of connected events. Columns I-III report the baseline regression with and without sector and year effects. We favour the regressions with their inclusion as the coefficients remain robust to unobserved heterogeneity across sectors and variation common to all firms (e.g., the macroeconomic environment).[22] The point estimate of firm size – the logarithm of firm revenues – in Column III is 0.231. A coefficient smaller than 1 suggests the marginal cost is decreasing with respect to revenues, i.e. costs don't increase linearly with the size of the firm.[23] The partial elasticity between firm size and costs implies that for a 1% increase in size there is an increase in the expected cost of 0.23%. Incidents that affect multiple firms – i.e. *connected events* – are similarly associated with higher expected costs: a unit increase in the number of affected firms translates to approximately a 2.6% increase in expected costs.[24] Finally, in column IV we show that these results are robust the inclusion of mare granular sector fixed effects.

Events with malicious intent are associated with a lower expected cost. Taking the estimate from the third column suggests that on average malicious events are associated with costs 66% lower than other event types.[25] This is perhaps surprising, given the significant press coverage that

[22] As discussed above the standard errors are clustered at the sector level. The results with Ecker-White errors and firm-level clustering are reported in the appendix. The magnitude of standard errors varies, although this has little impact on the precision of the estimates.

[23] An alternative way to see this is to correct for firm size on the costs variable, i.e. using a ratio of costs to firm size. We present the results of this regression in Table A.1.4 in the appendix. They confirm that the losses are not proportionate to firm size, and are decreasing relative to firm size.

[24] Revenues, like various measures of firm size, could be heterogeneous across sectors. Sector fixed effects should go some way into controlling for this. In untabulated results available upon request, we construct a dummy variable that equals 1 if firm revenues are above the median within that sector, such that we have a within-sector measure of small versus large firms. Including this in the regression confirms the robustness of our original result.

[25] The percentage change is calculated using the bias correction of Kennedy (1981), $g = \exp(\hat{\beta} - \frac{1}{2}V(\hat{\beta})) - 1$.

Dependent Variable: Log(Cost)				
Regressor	I	II	III	IV
log(Firm Size)	0.241***	0.220***	0.231***	0.220***
	(0.0300)	(0.0228)	(0.0234)	(0.0222)
Connected events	0.0176**	0.0257***	0.0257***	0.0238***
	(0.00740)	(0.00555)	(0.00548)	(0.00661)
Malicious	-1.31***	-1.33***	-1.20***	-1.09***
	(0.230)	(0.179)	(0.207)	(0.324)
Security Incident	11.0***	13.0***	13.6***	13.8***
	(0.338)	(0.631)	(0.676)	(0.632)
Data Breach	11.6***	14.1***	14.6***	14.8***
	(0.186)	(0.469)	(0.477)	(0.603)
Phishing/Skimming	12.7***	14.7***	15.1***	15.4***
	(0.486)	(0.573)	(0.554)	(0.683)
Privacy Violation	10.8***	13.2***	14.0***	14.3***
	(0.387)	(0.708)	(0.755)	(0.701)
Other	12.6***	14.4***	15.3***	15.2***
	(0.405)	(0.636)	(0.666)	(0.895)
Year Fixed Effects	N	Y	Y	Y
Sector Fixed Effects	N	N	Y	N
Sub Sector Fixed Effects	N	N	N	Y
R^2	0.11	0.19	0.21	0.25
N	3705	3705	3705	3705

Notes: Results from estimating equation (1.1). *, ** and *** denote significance at the 10, 5 and 1 percent level respectively. Standard errors (reported in parentheses) are clustered by sector. Column I is an OLS regression without controls for Year or Sector fixed effects. Column II is an OLS regression without Sector fixed effects. Column III is an OLS regression including both fixed effects. Column IV replaces the sector fixed effects with the finer categorisation of sub-sector.

Table 1.2: The drivers of cyber risk - baseline results

cyber attacks get and the concern expressed by multiple organisations.[26] Looking more closely

[26]A number of other factors may help explain this finding. For one, cyber security actions adopted by many firms protect them from the effects of malicious cyber incidents. There are various well-developed tools that are built to predict and manage cyber attacks, which may be less effective against events that occur as a result of human error

at the distribution of costs within each category can provide evidence as to what may be the key driver. In Figure A.1.3 of the annex, we show the distribution of costs per case type, malicious versus non-malicious. Security incidents and data breaches are the only case types with variation across both malicious and non-malicious events. Within security incidents in particular there is a stark contrast between the distribution of malicious and non-malicious events, with the latter being significantly more costly.[27] Finally, while on average the cost of malicious events may be lower, it may still be the case that when focusing at the worst type of events in terms of losses – i.e. when looking at the right tail of losses – malicious events regain prominence. Our finding should thus not be taken as a reason to gloss over the threat that is posed by malicious cyber attacks, as we show in the next section.

1.4.3 Beware of the tails

Losses stemming from operational and cyber incidents are typically characterised by a set of "heavy-tailed" distributions (Cohen et al., 2019). Therefore, it is reasonable to assume that the conditional distribution is not homogeneous across cost quantiles. Of particular interest in this context is the tail of this distribution, which characterises events of low frequency but high severity. Identifying the features of such events is important to policy-makers and supervisors as they carry the potential to generate substantial economic losses and systemic disruption.

Figure 1.4 displays the estimates of the coefficients of *firm size*, *connected* events, and *malicious* events at quantiles ranging between the 0 and 100th percentile. Estimates do vary at different

inside firms. Moreover, well coordinated cyber attacks can go undiscovered for a long time, in which case the cost of the attack can be difficult to estimate or even identify. Finally, some cyber attacks potentially carry large reputational costs that are hard to quantify and are hence not adequately reflected in loss data.

[27] Non-malicious security incidents include events such as network failures or software bugs that can be very costly. Network outages could be caused by operator errors, surge or usage spike, hardware infrastructure failure, or loss of electrical power. Firms could expect to face 1.6 hours of downtime every week, which has been estimated to cost them, on average, $5,600 per minute (Knobbe, 2020).

quantiles. The estimate of *firm size* has a lower magnitude at the both ends of the distribution, i.e. it shows an inverted U pattern. *Connected* events have larger estimates towards the upper end of the distribution. Most interestingly, we observe that as the *malicious* variable approaches the 90th percentile the coefficient trends upwards, towards zero, and eventually into positive territory. Malicious events thus do exhibit a significantly different behaviour at the tail-end of the distribution.

Turning the attention to the tail of the distribution, we next present the results of cost quantile regressions between the 95th and 99.5th percentiles. The specification of the regression is analogous to the baseline regression in Table 1.2 (Column III). *Firm size* and *connected* events are both lower than their mean estimates. The *malicious* indicator has a positive coefficient across all the upper quantiles. A significant effect is observed at the 99.5% level. This result suggests that, ceteris paribus, the tail of the loss distribution is more sensitive to shocks from malicious events. Well coordinated malicious attacks – that happen less frequently – are likely to exceed the costs of non-malicious cyber events. These estimates should be taken with some caution: with limited observations, estimates of what occurs in quantiles can be subject to bias (Chernozhukov and Umantsev, 2001). Nonetheless, uncovering this relationship reveals an important caveat of only studying the central measures of the distribution. While our benchmark regression may show that malicious events are less damaging, sophisticated hacks can actually exacerbate costs at the tail end of the distribution. Understanding the potential damage of high-frequency, low-probability events is paramount from a policy perspective.

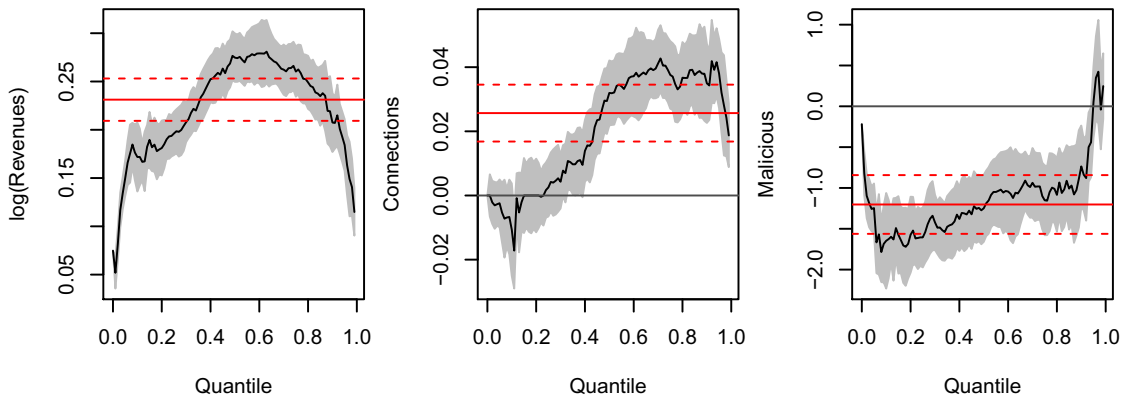

Notes: The plot shows the change in selected coefficients as cost quantiles vary. The specification of the regression is analogous to the baseline regression of Table 1.2 (Column III). The grey shading highlights the 90% confidence interval of the coefficient and red lines denote the estimate of the conditional mean by OLS.

Figure 1.4: Selected coefficients from cost quantile regressions

1.5 Digitalisation and cloud-based technologies

Until not so long ago, firms looking to adopt digital technology had to invest in their own data infrastructure and hardware. With the advent of cloud technologies, this has dramatically changed. Cloud technology enables firms to rent computing power and storage from service providers, turning some fixed costs into marginal costs and giving firms more flexibility in handling their operations in a potentially more protected environment. This can be particularly advantageous for smaller firms with fewer resources to spend on IT.[28] Cloud computing also exhibits positive externalities such as the reduction of energy consumption and carbon emissions (Etro, 2015). Evidence suggests that firms increasingly take advantage of these benefits, as adoption of digital technology continues to trend upwards (Chen and Srinivasan, 2019).

Digital technologies also pose risks and challenges. Networked production facilities, vehicles,

[28]However, firms are still responsible for the configuration of machines and safe storage of sensitive data while interfacing with external applications.

Dependent Variable: Log(Cost)				
Regressor	95%	97.5%	99%	99.5%
Panel A: Wild Bootstrap				
log(Firm size)	0.184***	0.149***	0.115***	0.153***
	(0.0170)	(0.0167)	(0.0225)	(0.0416)
Connected events	0.0381***	0.0241***	0.0187**	0.00301
	(0.00763)	(0.00756)	(0.00768)	(0.0179)
Malicious	0.0508	0.525*	0.247	0.342
	(0.309)	(0.316)	(0.347)	(1.11)
Panel B: Clustered Bootstrap				
log(Firm size)	0.184***	0.149***	0.115***	0.153***
	(0.0177)	(0.0140)	(0.0202)	(0.00303)
Connected events	0.0381***	0.0241***	0.0187***	0.00301***
	(0.00592)	(0.00370)	(0.00473)	(0.000870)
Malicious	0.0508	0.525	0.247	0.342***
	(0.239)	(0.335)	(0.454)	(0.102)
Year Fixed Effects	Y	Y	Y	Y
Sector Fixed Effects	Y	Y	Y	Y
Case Type Fixed Effects	Y	Y	Y	Y

Notes: Results from estimating equation (1.1) at different quantiles. *, ** and *** denote significance at the 10, 5 and 1 percent level, respectively. In Panel A, the standard errors (reported in parentheses) are calculated using the wild bootstrap method proposed by Feng et al. (2011). In Panel B, the standard errors are calculated using the method of Hagemann (2017). Both methods are computed using the R Package `quantreg`. For the definition of the regressors, see Table 1.1.

Table 1.3: Quantile regressions

transport infrastructure, and a host of other devices connected to the internet present new opportunities to cyber criminals. The growing complexity of digital infrastructures could increase the likelihood of failures and interruptions, as well as the attendant costs. Cloud service providers have recently drawn the attention of regulators due to the risks associated to their operations, not least given the high degree of concentration in the sector that increases the risk of single points of

failure. Tail-risks associated with an outage of a cloud service provider could lead to substantial losses and potentially bring the economy to a halt (Danielsson and Macrae, 2019).[29]

Dependent Variable: Log(Cost)					
Regressor	I	II	III	IV	V
log(Firm size)	0.222***	0.221***	0.220***	0.454***	0.450***
	(0.0224)	(0.0224)	(0.0228)	(0.128)	(0.126)
Connected events	0.0256***	0.0256***	0.0258***	0.0253***	0.0254***
	(0.00550)	(0.00549)	(0.00549)	(0.00523)	(0.00525)
Share of digital		-0.0142	0.0554	0.0461	0.114*
		(0.0445)	(0.0484)	(0.0631)	(0.0671)
log(Firm size) × Share of digital				-0.0156*	-0.0154*
				(0.00858)	(0.00846)
Share of cloud	-0.0211		-0.0378**		-0.0371*
	(0.0154)		(0.0188)		(0.0198)
Malicious	-1.33***	-1.33***	-1.33***	-1.34***	-1.34***
	(0.172)	(0.173)	(0.171)	(0.171)	(0.169)
Year Fixed Effects	Y	Y	Y	Y	Y
Sector Fixed Effects	N	N	N	N	N
Case Type Fixed Effects	Y	Y	Y	Y	Y
R^2	0.19	0.19	0.19	0.19	0.20
N	3705	3705	3705	3705	3705

Notes: Results from estimating equation (1.1). *, ** and *** denote significance at the 10, 5 and 1 percent level, respectively. All standard errors (reported in parentheses) are clustered by sector to account for the correlation in the sector-level variables. Firm size refers to the firm revenues; Share of digital is the percentage of firms per sector that keep more than 50% of information stored in digital format; Share of cloud is the percentage of firms per sector that keep more than 50% of their data in a dedicated cloud storage. Sector controls are dropped in all regressions due to multicollinearity with the sector-level variables.

Table 1.4: Regressions including the sector level cloud and digital storage variables

The jury is still out on whether the benefits outweigh the risks, or vice versa. To date, there is

[29]For a wider discussion of the benefits and risks of cloud computing, see for example Catteddu (2009) and Carr et al. (2019).

little empirical evidence to support either claim. We contribute to this discussion by extending our regression framework with variables that proxy for firms' exposure to digital and cloud services. In particular, we consider *share of digital* and *share of cloud*, which capture sector-level exposure to digital technologies and cloud technology, respectively. To avoid multicollinearity between sector-level variables and sector fixed effects, we drop the sector controls.

We present the results of the regressions in Table 1.4.[30] Our results suggest that firms in sectors with a higher exposure to cloud technologies benefit from a mitigating effect on expected costs stemming from cyber incidents. To interpret the magnitude of coefficients, recall how the digital and cloud variables are constructed: the proportion of firms that stated over 50% of their information was stored digitally and of their IT spending was on cloud technology, respectively. The measures loosely reflect the probability that any given firm within a sector has more of its data stored digitally and the probability that the firm has a higher spending on cloud services. The variables are recorded on a 0-100 percentage scale. Consider Column V in Table 1.4: a one percent increase in the probability of firms spending more than 50% of their IT budgets on cloud services is associated with a reduction of around 4% in expected costs. A stronger dependence on digital services appears to have no statistically significant effect, if considered in isolation (Column II). However, when we add an interaction term between firm size and digital dependence in Column V, we find a mitigating effect. Figure 1.5 plots the surface of this mitigating effect on costs as a function of firm size and digital dependence. As firm size increases, more exposure to digital technologies appears to have a mitigating effect on costs. A possible explanation for this relationship could be that as firms expand so do their resources and investment in personnel that ensure that digital technologies are safely maintained.

Overall, a higher exposure to digital and cloud infrastructures is associated with a mitigating

[30]The case type dummies displayed in previous output are subsumed into the Case Type Fixed Effects indicator.

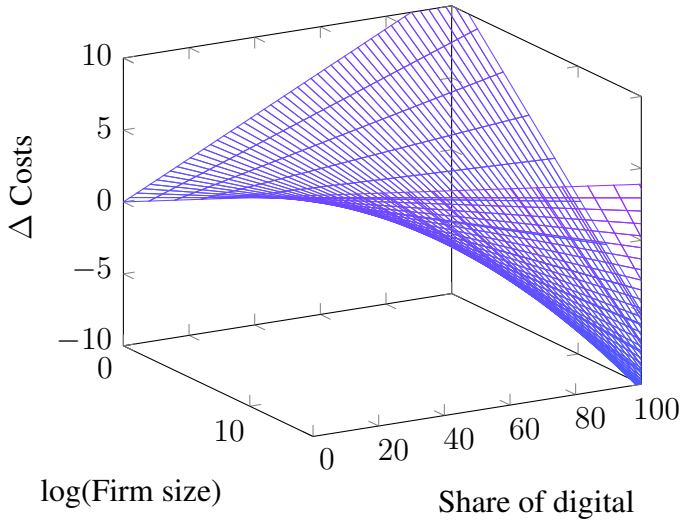

Figure 1.5: The expected joint effect of firm size and dependence on digital technology on firms costs

effect on costs, which is also firm-size dependent. Given data limitations, it is challenging to derive an exact identification strategy that enables us to define the causal relationship between firms' use of digital technologies and cyber risk. That said, our findings represent a first-pass analysis for a better understanding of the role that digital technologies and in particular cloud technology have to play in shaping cyber risk.

1.5.1 Revisiting the tails

In this section, we briefly revisit the behaviour of cyber costs at the tails of the loss distribution in relationship to our $digital$ and $cloud$ variables. Reliance on cloud services in particular has the potential to increase tail-risks (Danielsson and Macrae, 2019). If such risk would be present in our dataset, we would expect to observe a similar relationship to that seen for the $malicious$ variable. That is, the coefficients on $cloud$ and $digital$ become trend upwards as we move right in the distribution, eventually becoming positive.

Figure 1.6 shows the estimates for the $cloud$ (left) and $digital$ (right) variables across quantiles. The estimate for the $cloud$ variable shrinks from negative toward zero at the upper quantiles, but does not reach positive territory. The mitigating effect found in the baseline regression is reduced at the tails, but does not turn into a factor that could exacerbate tail risks. The $digital$ variable shows a positive relationship with cyber costs that tend to decline towards the upper quantiles (this confirms the sign of the interaction term in Column V of Table 1.4).

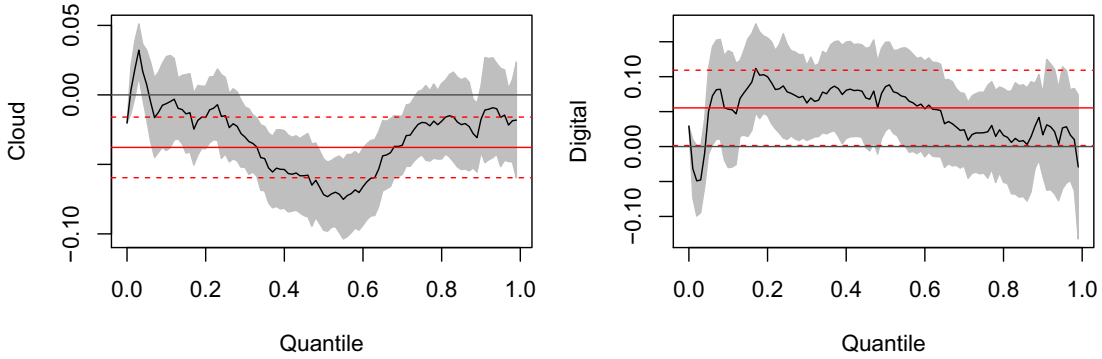

Notes: The graphs show the change in selected coefficients as quantiles vary. The specification of the regression is analogous to the baseline regression of Table 1.4 (Column III). The grey shading highlights the 90% confidence interval of the coefficients and red lines denote the estimate of the conditional mean by OLS.

Figure 1.6: Selected coefficients from quantile regressions

1.6 Dealing with cyber risk: is current IT investment enough?

In this section we analyse if investment in IT can help to mitigate costs from cyber incidents.[31] Investment in IT security arguably has obvious benefits, yet to date there appears to be little evidence

[31] According to Gartner (2021), worldwide IT spending is projected to total $4.2 trillion in 2021, with information security and risk management technology and services expected to grow 12.4% to reach $150.4 billion.

that the continually increasing size of IT budgets and spending are correlated with the mitigation of costs stemming from cyber-related incidents, despite IT security being considered an increasingly critical part of business continuity plans. Evidence that IT spending yields a beneficial return on investment could alter the perspective of firms that are reluctant to invest in IT security as it is viewed as a "sunk" cost. Roner et al. (2021) provide some evidence that cyber security investments can reduce losses arising from a cyber breach. We provide further support to this argument with alternative data.

We use a database constructed by Kennedy and Stratopoulos (2017) based on the InformationWeek IW500 survey. The survey gathers data on IT spending from 500 firms based in the US, across various sectors. The survey focuses on firms that are the most innovative IT users, i.e. to be included in the list a firm had to demonstrate sophisticated use and deployment of IT (Lim et al., 2011).[32] The IW500 dataset provides us with an estimate of firms' IT expenditures as a percentage of revenues. Figure A.1.5 in the appendix displays the trend in IT spending across sectors for the period 2002-2013. The finance and insurance sector is consistently one of the largest investors in IT, whereas construction and mining are at the lower end of the spectrum. Overall, the investment in IT tends to be relatively stable over time.

Table 1.5 summarises costs across sectors and the implied spending on IT. To compare the typical annual cost of cyber incidents to firms across sectors, we remove outliers using the interquartile range method (we are interested in these outliers and will return to them later in this section). On average, non-malicious incidents appear to have a higher average costs. Retail Trade and Finance and Insurance are the implied largest spenders on IT. These sectors along with Information and Manufacturing are inferred to spend upwards of a billion dollars annually on IT. In

[32]Previous studies have used the IW500 data to examine the relationship between IT expenditures and various aspects of firm activity and performance. However, there is little evidence on the impact that this investment has towards reducing the risk of losses.

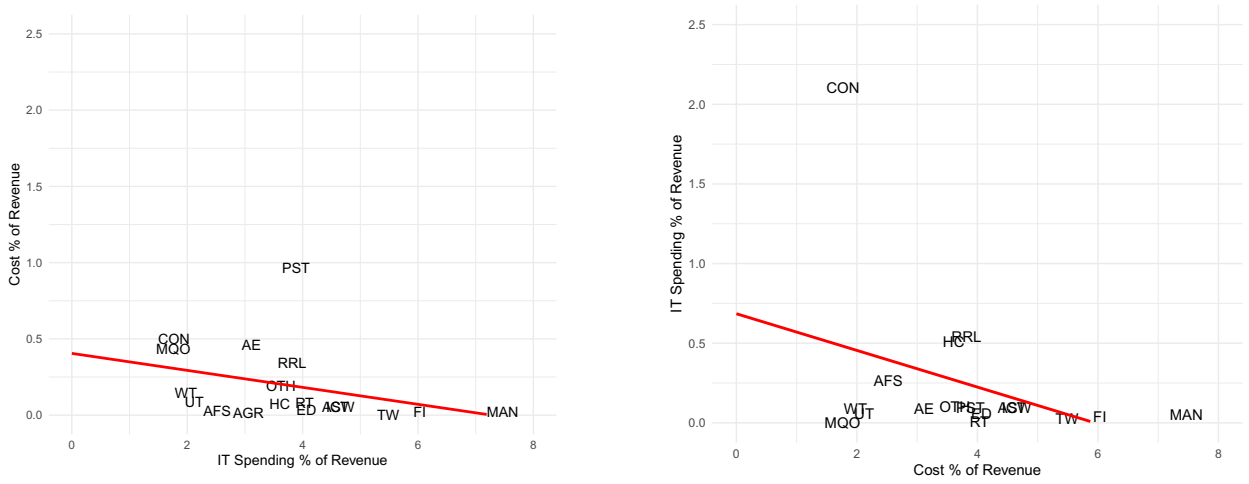

Notes: The panel on the left hand side shows the average annual costs for malicious incidents by sector normalised by the average revenues for firms within that sector on the y-axis. On the x-axis, we plot the spending as a percentage of revenues. The panel on the right hand side shows the average annual costs for non-malicious incidents by sector normalised by the average revenues for firms within that sector on the y-axis. On the x-axis, we plot the spending as a percentage of revenues.

Figure 1.7: Spending in IT relative to optimal and costs of cyber events per unit of revenue.

Figure 1.7 we show the relationship between spending and the typical annual costs for malicious and non-malicious incidents. For both, there is a weakly implied negative relationship, i.e. higher spending is correlated with lower costs when correcting for average firm size.[33] Sectors with an implied lower spending on IT appear to incur higher costs for malicious relative to non-malicious events. Thus there appears to be evidence that higher levels of IT spending are effective at protecting firms from non-malicious incidents, even if we cannot be certain about causality. For example, investment into new hardware may not be a direct investment in security, but may lead to fewer unintended failures that could be caused by old hardware.

One-off events that lead to significant damage and disruption are of particular interest. In Figure 1.8, we show a similar plot to the previous but using the 90th percentile of the distribution of

[33] Without normalising the costs by firm size (revenues) we would simply observe the fact that larger sectors have higher costs.

Sector	Revenues	Costs (all)	Costs (non-mal)	Costs (mal)	IT Spending
Accommodation and Food (AFS)	1,512.75	1.58	3.99	0.43	38.12
Admin., Support, WM (ASW)	659.86	0.65	0.62	0.36	30.48
Agriculture (AGR)	82.74	0.01	0.00	0.01	2.53
Arts and Entertainment (AE)	112.11	0.59	0.10	0.52	3.49
Construction (CON)	39.18	0.65	0.82	0.20	0.69
Educational Services (ED)	725.13	0.24	0.41	0.27	29.50
Finance and Insurance (FI)	21,362.75	5.52	7.69	4.18	1,288.88
Health Care (HC)	920.09	2.68	4.72	0.71	33.18
Information (ICT)	21,735.79	20.49	21.06	12.49	1,000.99
Manufacturing (MAN)	14,483.91	8.56	7.56	3.55	1,081.79
Mining (MQO)	1,621.28	2.45	0.05	7.00	28.53
Other Services (OTH)	31.73	0.04	0.03	0.06	1.15
Professional, Sci. and Tech. (PST)	388.15	0.97	0.36	3.76	15.09
Real Estate (RRL)	953.15	4.16	5.17	3.26	36.40
Retail Trade (RT)	33,325.43	12.70	1.74	27.63	1,344.87
Transportation and Warehousing (TW)	13,208.20	1.69	3.16	0.21	724.57
Utilities (UT)	1,892.43	1.37	1.04	1.69	40.14
Wholesale Trade (WT)	1,726.54	1.84	1.58	2.55	34.16
Total	6,376.73	3.88	3.34	3.83	318.59

Notes: The table summarises revenues, costs and IT spending across sectors. Revenues in the first column denote the average revenue of a firm within each sector. The three cost columns report the average annual cost of cyber incidents incurred by firms in the Advisen database. We remove outliers using the interquartile range method. We report the average for all incidents in the second column and then distinguish between non-malicious and malicious incidents in the third and fourth columns. The final column reports the implied total IT spending by sector (Revenues $\times$ IW500 measure). All figures are expressed in millions of US dollars. Sector abbreviations are denoted in parenthesis next to the sector.

Table 1.5: Summary of costs and spending by sector

costs across each sector for malicious and non malicious events. First we note a similar relationship, higher spending is associated with lower costs per unit of revenue at the 90th percentile. The estimated regression lines (in red) are similar. Finance and Insurance, Transportation and Warehousing and Manufacturing are some of the highest spenders as a percentage of revenue and appear to benefit from a reduction in relative costs at the mean and at the 90th percentile.

We next look at the difference between spending and costs at firm level. Figure A.1.6 in the

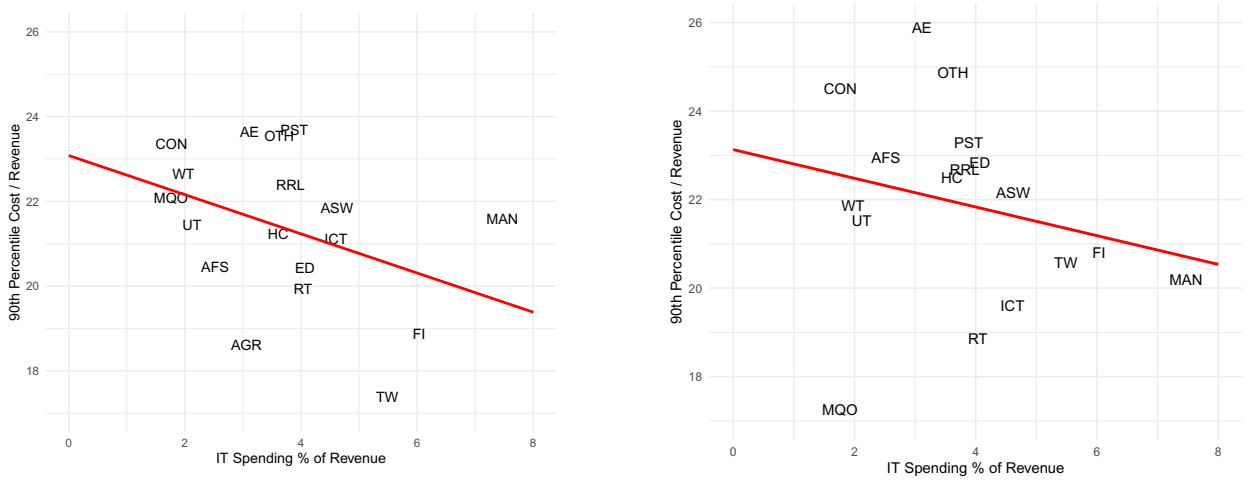

Notes: The panel on the left hand side shows the logarithm of the 90th percentile annual costs for malicious incidents by sector normalised by the average revenues for firms within that sector on the y-axis. On the x-axis, we plot the spending as a percentage of revenues. The panel on the right hand side shows the logarithm of the 90th percentile for non-malicious incidents by sector normalised by the average revenues for firms within that sector on the y-axis. On the x-axis, we plot the spending as a percentage of revenues.

Figure 1.8: Spending in IT relative to optimal and costs of cyber events per unit of revenue (90th percentile).

appendix shows the histogram of the difference in percentage of revenues spent on IT and the annual costs of cyber incidents as a percentage of revenues. The bulk of observations are centered around 0, indicating that spending and costs are approximately similar, as also noted by Romanosky (2016). However, we observe a longer left tail that point to the annual cost of cyber incidents sometimes exceeding investment in IT, and occasionally quite significantly. Of course, firms would not expect cyber incidents of such nature to occur with certainty every year and thus the optimal investment requires some balance over time. This leaves open the question of the optimal amount that firms should invest into IT security. The seminal work of Gordon and Loeb (2002) suggest that this should be where the marginal benefit of investment (reduction in costs) is equal to the marginal cost (dollars invested). Such work requires consideration of the distribution of cyber losses for firms and how to incorporate external information from databases similar to the

one used within this analysis. We leave this issue for future research.

Dependent Variable: Log(Cost)					
Regressor	I	II	III	IV	V
log(Firm size)	0.228***	0.228***	0.228***	0.228***	0.230***
	(0.0232)	(0.0232)	(0.0197)	(0.0197)	(0.0190)
Connections	0.0159	0.0159	0.0245***	0.0245***	0.0251***
	(0.0101)	(0.0101)	(0.00679)	(0.00679)	(0.00708)
Malicious	-1.12***	-1.12***	-1.02***	-1.02***	
	(0.350)	(0.350)	(0.288)	(0.288)	
IT spending	-29.2	-29.2			
	(30.5)	(30.5)			
IT spending lag			-41.3*	-41.3*	-35.6*
			(22.6)	(22.6)	(21.5)
IT spending lag × Malicious					-22.9***
					(7.69)
Share of cloud		0.0217		-1.13***	-1.13***
		(0.158)		(0.135)	(0.128)
Share of digital		-0.286		1.34***	1.37***
		(0.188)		(0.152)	(0.153)
Year Fixed Effects	Y	Y	Y	Y	Y
Sector Fixed Effects	Y	Y	Y	Y	Y
Case Type Fixed Effects	Y	Y	Y	Y	Y
R^2	0.2	0.2	0.2	0.2	0.21
N	2611	2611	2953	2953	2953

Notes: Results from estimating equation (1.1). *, ** and *** denote significance at the 10, 5 and 1 percent level, respectively. All standard errors (reported in parentheses) are clustered by sector to account for the correlation in the sector-level variables. Firm size refers to the firm revenues; Share of digital is the percentage of firms per sector that keep more than 50% of information stored in digital format; Share of cloud is the percentage of firms per sector that keep more than 50% of their data in a dedicated cloud storage. IT spending denotes the percentage of revenue spent on IT within each sector per year.

Table 1.6: Regressions including the sector level cloud and digital storage variables

We attempt to identify if there is a significant effect of IT investment on reducing the costs of

a cyber incident.[34] We match our IW500 spending data with the sector and years of observations available in our sample. Data on spending ranges from 2002-2013 therefore we drop observations beyond 2013. We then include the level of spending as a variable in our regression. Table 1.6 displays the results. In columns I and II we match the spending year with the year in which the incident occurred. In III, IV, and V we match the year in which the incident with the lag of spending.[35]

The regressions using the contemporaneous spending measure indicate no effect of spending on the costs of cyber incidents. However, when including the lag there is a negative impact (i.e., mitigating effect). The magnitude of the estimate suggests that a 1% increase in IT spending is correlated with a decrease of 34% in costs the subsequent year. This is an important result and also echoes that of Roner et al. (2021). IT investments can take time to mature and the benefits may not be observed immediately, e.g. investment in staff training does not pay dividends until staff have acquired sufficient knowledge.

1.7 Conclusions

The digital revolution has increased the interconnectivity and complexity of the economic system. The use of technology and internet have improved firms' productivity, but also exposes them to cyber attacks. Moreover, the greater use of cloud services exposes further important economic sectors to common risks.

Despite the large and growing exposure to cyber risks, cyber costs are difficult to quantify. Using a unique database at the firm level for the US, we document the characteristics of cyber incidents and help quantify cyber risk. The average cost of cyber events has increased over the last

[34] Using UK based data Roner et al. (2021) show that investments in IT security lead to a reduction in the amount of a loss from a cyber incident.

[35] Here we gain observations from 2014 and lose those from 2002, hence the changing sample size in the regressions.

decade. These costs are higher for larger firms and more connected events, and relatively lower for cyber events with malicious intent (cyber attacks), but only if the attack is not conducted on a large scale: malicious events can be more costly in the upper tail of the distribution.

The financial sector experiences the highest number of cyber incidents (especially of a malicious type, privacy and lost data incidents). However, banks and insurance companies incur more limited losses relative to other sectors, likely due to the effects of regulation and higher investment in cyber security.

We document that developing technological skills helps firms mitigate the costs of cyber incidents, as does more reliance on cloud services. This last result should be taken with caution and qualified. As cloud connectivity increases and cloud providers become systemically important, cloud dependence is also likely to increase tail risks.

Finally, we document some evidence on the effect that spending on IT has on the costs of cyber incidents. We observe a negative relationship between spending in IT and the cost of cyber-events. This result provides some evidence of the ”unobserved” return on investment into IT and security and may encourage firms that are reluctant to invest into IT, as the returns on additional expenditures are hard to measure. While our analysis does not account for the systemic implications of failures in specific critical sectors, the results can inform policymakers as to where to direct their attention in order to improve the economy's overall cyber resilience.

CHAPTER 2

OPERATIONAL AND CYBER RISK IN THE FINANCIAL SECTOR

2.1 Introduction

Operational risk emerged as a distinct risk category in the mid 1990s, following events such as the Nick Leeson's "rogue" trader case at Barings bank. Not long after, the Basel II standards introduced operational risk capital requirements, with operational risk defined as "the risk of losses resulting from inadequate or failed internal processes, people, systems or from external events" (Basel Committee on Banking Supervision (2003)).[1]

Measuring and understanding operational risk is critical for both banks and public authorities. Operational risk currently represents a significant portion of banks' risk-weighted assets, second only to credit risk.[2] Regulators, central banks and international organisations, in turn, place the understanding and mitigation of operational risk – and subcomponents such as cyber risk – high in their agendas. While banks use internal data to determine their regulatory capital, there is limited work to identify the relationship between operational risk and the macroeconomic and supervisory environments – especially in an international context. Accordingly, policy discussions on the topic at the wider macroeconomic level tend to lack substantial empirical grounding. The prevalence of work-from-home arrangements in the wake of the Covid-19 pandemic only heightens the need to quantify and understand operational and cyber risks for financial institutions.

In this paper, we contribute to filling this gap by analysing a unique cross-country dataset of

[1] Before Basel II, losses stemming from operational risks were covered by capital provisions set aside from credit and market risk.

[2] Up to 40% of risk-weighted assets can be attributed to operational risk in some jurisdictions (Sands et al. (2018)).

operational losses. We present stylised facts on the evolution of operational losses since 2002; compute operational risk capital through different methods; use proportional hazards models to study the lag between occurrence, discovery and recognition of operational loss events; and link losses to the macroeconomic and supervisory environment. Finally, we construct a proxy for cyber losses using the event type categorisation of Basel II, document their evolution and compute an estimate of "cyber risk capital".

We use data at the loss event level from ORX, a consortium of financial institutions. The consortium was founded by banks with the aim of sharing operational loss risk data in an anonymised fashion in order to benchmark operational risk models. The sample we use contains over 500,000 operational loss events from 2002 until end-2016 for a group of 74 large banks across the globe. This makes our paper the most comprehensive in terms of its time series and, especially, cross-country coverage.

We document that, after a notable increase post-Great Financial Crisis (GFC), banks' operational risk losses have shown signs of decline since 2015. One category in particular is responsible for this pattern, namely "Clients, Products & Business Practices". It includes improper business practices like fiduciary breaches, aggressive sales, breaches of privacy, account churning and misuse of confidential information. These are the types of operational risks that characterise periods of financial excess, with mis-selling of mortgage-backed securities in the mid-2000s being a prime example. Towards the peak of the GFC there was a significant increase in the occurrence of this type of events (especially in North America), which were then recognised in the books of banks a few years later. Importantly, this pattern is observed only in terms of loss amounts and not in terms of frequency of occurrence.

Operational losses are characterised by a fat-tailed distribution.[3] Accordingly, estimates of

[3]In other words, there are a large number of inconsequential events from a cost perspective and a limited number of very costly events. The latter group in particular complicates the quantification of operational risks, as such low

operational risk capital can lead to notably different results depending on the method used and how well it captures what happens at extreme values of the distribution of operational losses. Indeed, our estimates for operational risk capital using methodologies from the Advanced Measurement Approach (AMA) range from 1% to 7.5% of gross income, against the 12% benchmark of the Basic Indicator Approach. This finding may provide some support for the new regulatory framework that proposes the adoption of the Standardised Measurement Approach (SMA) for all banks. This has two practical effects. First, it reduces heterogeneity in the application of different AMAs and the need for regulators to validate these models. Second, it simplifies the regulation, while at the same time preserving capital adequacy to cover operational risks.[4]

Operational losses, on average, take over a year to be discovered an recognised in banks' books. The time between occurrence, discovery and recognition, however, varies across event types, bank size and jurisdictions. From our summary statistics of duration times, we see that Internal fraud and Clients and business practices are the incidents that, on average, take the longest to be discovered and eventually accounted for. Two facts could explain this. First, perpetrators of internal fraud do their best to cover their tracks such that the event goes unnoticed for longer. Second, "business practices" events are often settled through lengthy legal proceedings that delay loss recognition. Large banks, in turn, tend to be slower in discovering and recognising operational losses in their books. Finally, we also find substantial heterogeneity across jurisdictions: banks in North America are the quickest to discover losses, whereas those in Eastern Europe are the slowest. Different approaches to regulation and supervision across jurisdictions may play a role in these results, and we note that a strengthening of quality in supervision is associated with shorter duration times.

frequency/high severity events are often cited as being "one-in-a-hundred years" events.

[4]That being said, it should be noted that the SMA may not entirely reduce the heterogeneity across estimates. Regulators across jurisdictions will have the option to apply a loss component to the calculation of the capital ratio, which, if applied, will rely on calculations based on previous losses. Thus, estimates across banks may still vary based on their internal historical losses.

These findings can inform policy discussions regarding the principles for executive compensation packages.

The stylised facts we present point to the existence of a link between operational losses and macroeconomic conditions. Abdymomunov et al. (2017) use data for US banks to document a contemporaneous correlation between macroeconomic conditions and operational risk losses, e.g. operational losses rise during economic downturns. We build on this idea and use a cross-country panel analysis to argue that the ultimate cause of the rising losses during economic downturns lies in the excesses characterising the run-up to the downturn. In other words, favourable conditions during periods of macroeconomic expansion and financial exuberance lead to the occurrence of events that are only discovered when the economic tide turns, and recognised in the books of banks even later.

Using deviations of policy rates from Taylor-rule implied benchmarks, we show that periods of accommodative monetary policy are followed by an increase in operational losses. This appears to be driven by the frequency rather than the severity of events. Periods of excessively accommodative monetary policy can lead to increased risk-taking by banks, which can boost the type of improper business practices that account for the lion's share of operational losses. Finally, in line with the work of De Nicolò and Lucchetta (2013), who find that banks in a higher competition environment increase monitoring efforts and reduce risks, and with Kim (2018) who finds that banks with lower market power take less liquidity risk, we find that periods of intense bank competition are also associated with lower operational losses.

Regulation can also play a role in moderating operational losses. The time pattern of losses stemming from internal fraud and improper business practices suggests that the quality of regulation and supervision can also be related to operational losses in the cross-section of countries. Indeed, we find that better regulation and supervision – as captured by the financial reform index

of Abiad et al. (2010) and Denk and Gomes (2017) – is associated with lower operational losses.

Finally, we provide estimates of cyber losses. Growing interconnectedness and reliance on technology has led to a growing focus and concerns regarding cyber and IT-related risks. These are most prominent for the financial system, given its critical role. We use the data to construct a proxy range of cyber losses (which are a subset of operational losses). We document that cyber losses, so far, represent a relatively small share of operational losses. In recent years, however, losses from cyber events saw a spike which aligns with the growing attention cyber risk has been receiving. Despite representing a relatively small share of operational losses, cyber risk capital can account for up to a third of total operational value-at-risk.

The paper is organised as follows. The next section reviews the related literature. Section 2.3 describes the data and documents the duration between occurrence, discovery and recognition of loss events. Section 2.4 uses the analytic and loss distribution approaches to estimate operational value-at-risk. The link between operational losses and the macroeconomic environment is the focus of Section 2.5, whereas Section 2.6 presents our estimate of cyber risks, a very important class of emerging risks in the financial sector. The last section discusses the main conclusions.

2.2 Related Literature

Research on operational risk intensified after 2001, when the BCBS introduced an amendment to the Basel Capital Accord to support operational risk with regulatory capital. Early work on the subject focused on issues related to how to conceptualise and quantify these risks (Power (2005), Cornalba and Giudici (2004), Chavez-Demoulin et al. (2006), Antonini et al. (2009), Jarrow (2008)).

The literature points to links between the characteristics of financial institutions and operational risk. Shih et al. (2000) and Curti et al. (2019b) find a positive relationship between size and operational losses. Chernobai et al. (2011) uses data for US financial institutions and finds that

most operational losses can be traced to a breakdown of internal controls. Firms suffering from these losses tend to be younger and more complex, and have higher credit risk, more anti-takeover provisions, and CEOs with higher stock option holdings and bonuses relative to salary. Operational losses can also pose risks for the financial system at large (i.e. systemic risks). Berger et al. (2018) find that operational risk at large US bank-holding companies is statistically and economically positively linked to standard measures of bank systemic risk.

Fraud and employee misconduct have contributed to operational losses and have come under scrutiny from regulators, often resulting in sizeable financial penalties. This can also affect bank returns (Byrne et al., 2017; Köster and Pelster, 2017).[5] Altunbaş et al. (2018) find that banks are more likely to engage in misconduct when their CEOs have a long tenure. Eshraghi et al. (2015) study regulatory enforcement actions issued against US banks to show that both board monitoring and advising are effective in preventing misconduct by banks. Fich and Shivdasani (2007) study whether external directors suffer reputational penalties if the firms they serve on were accused of financial fraud.

Operational risk could also be intertwined with business and financial cycles. Carrivick and Cope (2013) and Hess (2011) look at the consequences of the GFC on operational risk losses in the financial sector. Abdymomunov et al. (2017) provide additional evidence of a relationship between operational losses in US banks and macroeconomic conditions. We build on this literature and investigate why such relationships are observed. Sakalauskaite (2018) shows that banks' misconduct has been relevant over our sample period and that its intensity correlates with the business cycle. Interestingly, the study finds that misconduct initiation is related to bank remuneration schemes, increasing with CEO bonuses in periods of high economic growth and when bank leverage is high.

Growing concerns around the economic and social impact of cyber risk in financial institutions

[5] A related strand of literature investigates the link between operational losses and bank returns (Biell and Muller, 2013; Sturm, 2013; Gillet et al., 2010; Cummins et al., 2006; Allen and Bali, 2007).

contrasts with a relatively thin literature in the topic. Data on cyber incidents are scarce and thus quantitative analyses on the impact of cyber events is challenging. The absence of common agreed standards to record such events further complicates the analysis.[6] We devise a proxy for cyber-related incidents from the categorisation of different event types. Kopp et al. (2017) examines the current regulatory framework and supervisory approaches, and identifies information asymmetries and other inefficiencies that hamper the detection and management of systemic cyber risk. Kashyap and Wetherilt (2019) outline some principles for regulators to consider when regulating cyber risk in the financial sector. From a perspective of the wider economy, Romanosky (2016) analyse the characteristics of cyber incidents across different sectors.

Bouveret (2018) estimates that average losses due to cyber-attacks could amount to USD 97 billion or 9 percent of banks net income. Duffie and Younger (2019) analyse a sample of 12 systemically important U.S. financial institutions and suggest that these firms have sufficient stocks of high quality liquid assets to cover wholesale funding run-offs in a relatively extreme cyber event. However, Eisenbach et al. (2021) estimate that the impairment of any of the five most active U.S. banks could result in significant spillovers to other banks, with 38 percent of the network affected on average.

2.3 Data

2.3.1 Operational loss data

Our analysis is based on a database that collects operational losses reported by financial firms across the globe. The data are owned and managed by ORX, the largest operational risk association in the financial services sector. The association, established in 2002, is primarily a platform for

[6] Facchinetti et al. (2019) propose ordinal measures to evaluate cyber risk in the presence of lack of data regarding the severity of such events.

the secure and anonymous exchange of high-quality operational risk loss data, with the objective of improving the management and measurement of operational risk.[7]

Data on losses are submitted to ORX on a voluntary basis. Data are anonymised, so as to protect the identity of the institution which suffered the loss. This process removes the incentive for members to under-report their losses, a problem which affects public databases.[8] However, this comes at the cost of making the analysis of individual institutions more complicated (Ames et al., 2015). The full sample comprises over 700,000 observations of operational loss events occurring between Q1 2002 and Q3 2018. We will work predominantly with a sample of 521,082 incidents which is obtained after combining individual loss data with region and bank size data and truncating our data at Q4 2016, the reason for which we outline below. This is still considerably larger than other available data-sets on operational risk and has the added appeal – relative to detailed data-sets at the country level such as the one available to U.S. regulators – that it includes a cross-section of countries over a large period. Our sample size is substantially larger than in vendor data-sets reported by Algo FIRST and SAS OpRisk Global Data, which are commonly used in the literature. For example, Chernobai et al. (2011) use the sample of data with 2,426 loss event reported by the Algo FIRST dataset. Hess (2011) uses data reported by SAS OpRisk Global Data with around 7,300 loss events from the banking industry. Cope et al. (2012) also use the ORX Global Loss Data Database, which at the time had approximately 180,000 loss events.

Members report losses based on the operational risk reporting standards established by ORX. These standards follow the event type and business line classification defined in the operational risk framework of the BCBS.[9] To be included in the data, operational events need to have an associated

[7]For details on the ORX consortium, see: https://managingrisktogether.orx.org/about.

[8]Furthermore, as the ORX consortium was set-up by financial institutions themselves, it would run counter to the very initiative of being part of the consortium to under- or mis-report data.

[9]For details on the ORX reporting standards, see: https://managingrisktogether.orx.org/standards. For the BCBS classification, see: https://www.bis.org/basel_framework/chapter/OPE/30.htm.

monetary cost reflected in the books of the banks, above a minimum of EUR 20,000. After data anonymisation by ORX, individual losses can only be identified by geography, business line and event type. Table 2.1 provides an example of how the data are structured.

RefID	Region	Business Line	Event Type	Gross Loss Amount	...	Loss Occurrence	Loss Discovery
123XYZ	Asia/Pacific	BL0101	EL0101	20000	...	ddmmyyyy	ddmmyyyy
⋮	⋮	⋮	⋮	⋮	...	⋮	⋮

Table 2.1: Example of the data structure

Each loss event is associated with an *event type* category. In line with Basel II definitions, there are seven event type (level 1) loss categories. Table 2.2 provides an overview of these categories and their definition. They include a wide array of potential causes of operational losses, such as internal/external fraud, disasters, improper business practices related to either clients or products, IT related, etc. Most of our analysis will be done at the level 1 category. However, the data also include a subdivision of each loss into level 2 event types, allowing for even more granular analysis. We will make use of the level 2 event type information to proxy for cyber-related events in Section 2.6.

Loss events are also associated with a *business line*. The business line classification, which again follows pre-specified standards, comprises nine business lines, including asset management, clearing, retail banking and trading & sales, among others. Table B.2.1 in the Appendix provides a detailed description. The intersection between business line and event types is important for the calculation of operational risk capital, as discussed further in Section 2.4.

The data are also partitioned into macro-regions. These include North America, Latin America & Caribbean, Eastern Europe, Western Europe, Asia/Pacific and Africa. For some of the regions that are more densely populated in terms of bank coverage, a further division into sub-regions is possible (see Table B.2.2 in the Appendix for details). While data are collected so as to preserve

Event Type	Description
Internal fraud	Losses due to acts of a type intended to defraud, misappropriate property or circumvent regulations, the law or company policy, excluding diversity/ discrimination events, which involves at least one internal party.
External fraud	Losses due to acts of a type intended to defraud, misappropriate property or circumvent the law, by a third-party
Employee related	Losses arising from acts inconsistent with employment, health or safety laws or agreements, from payment of personal injury claims, or from diversity / discrimination events
Clients, products and business practices	Losses arising from an unintentional or negligent failure to meet a professional obligation to specific clients (including fiduciary and suitability requirements), or from the nature or design of a product.
Disasters	Losses arising from disruption of business or system failures.
Technology and infrastructure	System failures (hardware or software), disruption in telecommunication, and power failure can all result in interrupted business and financial loss.
Transactions and processing	Losses from failed transaction processing or process management, from relations with trade counterparties and vendors.

Notes: The definitions of event types used by ORX are mapped to those used under the Basel II framework.

Table 2.2: Overview of event types based on the operational risk reporting standards of ORX

bank anonymity, each loss event has a tag for bank size. This indicator variable divides financial institutions based on income into large, medium and small.

Finally, each loss event has three associated dates. The *date of occurrence* captures the date when the loss event was deemed to have taken place. The *date of discovery* captures the point in time at which staff became aware of the event that lead to the operational loss. Finally, the *date of recognition* represents the date when the loss was recorded in the accounts of the bank. Figure 2.1 depicts the timeline of a loss. We explore the factors that determine the duration of losses in Section 2.3.4. However, this also brings us to an important juncture regarding completeness of the data, which we discuss next.

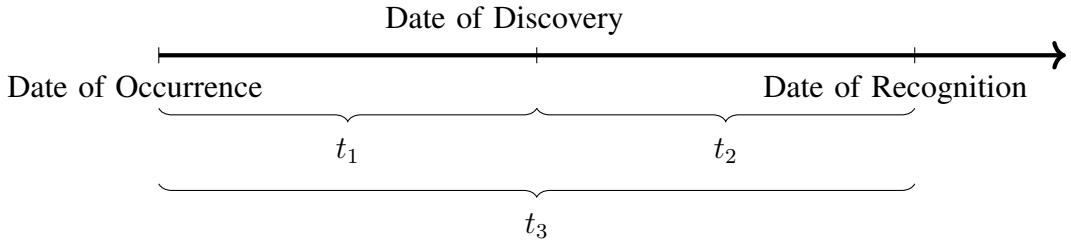

Figure 2.1: Loss timeline and key dates

2.3.2 Data bias and completeness

Given how data are collected, it is necessary to perform some adjustments to ensure that losses are comparable through time, especially when presenting aggregate figures. In particular, this refers to changes in the composition of the consortium membership and differences in the degree of completeness of the data across periods.

Figure B.2.1 in the online annex reports the evolution of the ORX consortium, in terms of total income and frequency of the reported losses. The number of banks in the consortium has grown over time, which could bias assessments of the evolution of operational losses when aggregating them over time. To account for this trend, when making comparisons over time, we divide gross losses and the frequency of events by the total income of the banks in the consortium for the given period. This adjusts for the growing number of banks in the sample, but also for their size. This second point is important, as simply dividing by the number of banks in the sample would fail to capture potential heterogeneity in banks' size.

In addition, Carrivick and Cope (2013) (herein CC) note that some losses are not reported to the consortium until long after the event has occurred. This is not related to wilful under-reporting of events, but is merely an artifact of the time it takes for events to be discovered and recognised. For example, legal proceedings can continue for years before a settlement is made. This is quite typical

for event types that include employment practices and workplace safety, and clients, products and business practices (see Section 2.3.4). While this issue affects in principle the whole sample (i.e. one cannot rule out that an event in, say, 2004, is yet to be discovered and recognised), it bites especially at the most recent end of the database. CC construct an approximate bias factor, which estimates the proportion of events that are unobserved in the data and use this to correct for the recent end of the sample. An alternative to this approach is to truncate the portion of the data that is most affected – a choice that can be underpinned by an analysis of how long it takes on average for events to be discovered and recognised in the books of banks. We follow this approach and in what follows consider observations until year end of 2016. We address this issue and our approach in more detail in Section 2.3.4.[10]

2.3.3 Additional data

For the analysis of the link between operational losses, macroeconomic conditions and regulatory characteristics, we complement the operational risk data with data from a variety of sources.

We proxy for the build-up of financial imbalances by using credit-to-GDP gap data from the Bank for International Settlements.[11] We obtain quarterly data for the credit-to-GDP gap across various regions from 2002Q1 until 2016Q4.

To capture competition in the banking sector, we use the Boone indicator (Boone, 2008), retrieved from the World Bank.[12] This measure proxies bank competition by the elasticity of profits to marginal costs. The elasticity is calculated by regressing the logarithm of profits on the logarithm of marginal costs.[13] The indicator is based on the premise that higher profits are achieved by

[10] In unreported results, available upon request, we also compute bias factors as in CC, and also confirm with aggregate data until December 2021 that the our choice of truncation gets rid of the period with the most pervasive under-reporting.

[11] See https://www.bis.org/statistics/c_gaps.htm.

[12] See https://datacatalog.worldbank.org/boone-indicator.

[13] The estimates of the Boone indicator in this database are based on the approach used by Čihák and Schaeck

more efficient banks, thus a more negative Boone indicator implies a higher degree of competition. We obtain annual data on the Boone indicator between 2002 and 2014 for various regions.

To measure the stance of monetary policy we use deviations of monetary policy rates from implied rates based on country-specific Taylor rules. The measure is constructed by subtracting the implied policy rate by the Taylor rule from the actual policy rate:

$$\tilde{\phi}_t = i_t - \phi_t \tag{2.1}$$

where i_t is the observed policy rate, ϕ_t denotes the rate implied by the Taylor rule, and $\tilde{\phi}$ denotes the deviation of the actual rate from the implied one. Central bank policy rates are sourced from the Bank for International Settlements and the implied Taylor rule rates are computed following Bogdanova and Hofmann (2012):

$$\phi = r^* + \pi^* + 1.5(\pi - \pi^*) + 0.5y \tag{2.2}$$

where, π denotes inflation, y captures the output gap, π^* is the inflation target and r^* is the long-run level of the real interest rate. We use quarterly data on deviations from the Taylor rule across various regions from 2002Q1 until 2016Q4.

Finally, to assess regulation and supervision in the cross-section of countries, we use an index of regulation and bank supervision, originally presented in Abiad et al. (2010) and extended in Denk and Gomes (2017). The full data-set is used to construct a measure of financial reforms across countries. To do so, various indicators are aggregated into a single index calculated as the simple average of the following seven dimensions: credit controls, interest rate controls, banking sector entry barriers, capital account controls, state ownership of banks, regulation of securities

(2010), but use marginal costs rather than average costs.

markets, and prudential regulation and bank supervision. The main variable of interest in our work is the measure of regulation and supervision. This variable takes into account the following four factors, i) Has a country adopted a capital adequacy ratio based on the latest Basel standard?; ii) Is the banking supervisory agency independent from executives' influence?; iii) Does the banking supervisory agency conduct effective supervision through on-site and off-site examinations? ; and, iv) Does a country's banking supervisory agency cover all financial institutions without exception? We use these questions to calculate an index at the regional level to be matched with the ORX data (an example of how this is done can be found in Section 2.5). The index runs from 0 to 1, whereby a score of 0 indicates a repressed regulatory and supervisory framework and a score of 1 a well-developed and liberalised framework. The series is provided annually from 2002 up to 2015. For further details, we refer the reader to Denk and Gomes (2017).

For each of these variables, we construct composite measures by weighting based on the banks in the sample.[14] For example and to fix ideas using the case of credit gaps, if the region Western Europe were made up of two UK banks, three German banks and four French banks, we would compute the statistic for the region as follows:

$$CreditGap_{WE} = \frac{2 \times CreditGap_{UK} + 3 \times CreditGap_{DE} + 4 \times CreditGap_{FR}}{9}$$

Against the background of limited data to underpin discussions of operational risk in the financial sector, we start by presenting stylised facts.

Table 2.3 displays summary statistics of operational risk losses by event type, region and bank size. A general observation is the large standard deviations in the data, an indicator of the heavy

[14]While we cannot associate a specific loss with any given bank, we know which banks comprise the sample at any given point in time.

tailed nature of the distribution of the data. From the perspective of event types (Panel A), on average the most costly events come from "Clients, products and business practices", which also contains the incident with the largest loss in the database. These types of events are "big ticket" items and, as we will see, are a common feature of losses stemming from the GFC. The largest losses tend to occur in Western Europe and North America (Panel B). Finally, the largest losses also appear to occur at larger banks, followed by small banks (Panel C).

	Mean	Std. Dev.	Max	Min
Panel A: By Event Type				
Internal fraud (EL01)	829,092	31,785,312	4,056,523,958	20,000
External Fraud (EL02)	148,384	2,537,026	500,000,000	20,000
Employee related (EL03)	121,266	1,075,761	174,382,494	20,000
Clients, products & business practices (EL04)	2,263,937	104,938,864	23,705,540,000	20,000
Disasters (EL05)	241,954	4,454,072	402,538,834	20,000
Technology and infrastructure (EL06)	623,200	21,597,528	2,224,579,168	20,000
Transactions and processing (EL07)	375,694	6,901,914	1,444,000,321	20,000
Panel B: By Region				
Africa	284,169	5,650,503	470,874,828	20,000
Asia / Pacific	491,863	7,202,508	814,464,293	20,000
Eastern Europe	498,447	6,782,258	500,000,000	20,000
Latin America & Caribbean	96,859	780,985	123,198,198	20,000
North America	990,052	59,807,389	20,180,094,936	20,000
Western Europe	747,111	54,789,230	23,705,540,000	20,000
Panel C: By Size				
Large	674,064	55,400,088	23,705,540,000	20,000
Medium	391,835	7,708,470	947,475,504	20,000
Small	519,024	15,605,569	2,744,201,136	20,000

Notes: The table presents summary statistics of losses by various categorisations. The summary statistics are based on 609,854 observations in total. We report information on the mean, standard deviation, maximum and minimum. Figures are in Euros.

Table 2.3: Summary statistics of loss events by categories

In Figure 2.2 we present evolution of the annual value and frequency where each year of losses is partitioned by event type (normalised by income, as per the discussion above). In terms of Date of Occurrence (upper panels), Clients, products and business practices clearly dominate in terms of loss amounts and featured heavily through the Great Financial Crisis. Transactions and process management in turn dominate in terms of frequency. This is consistent with the former being a high severity item, largely attributable to fines and regulatory actions, and the latter a high-frequency item, arising from thousands of daily operations taking place in banks. Contrasting the upper panels with the lower panels, which aggregate data based on the Date of Recognition, there is some initial evidence of a visible lag in the accumulation of losses. In the upper left panel, the peak arrives at around 2008 at the time of the GFC, whereas in the lower left panel the peak is in 2011. This lag is indicative of the fact that many losses in the Clients, products and business practices category face protracted legal proceedings before they are eventually settled and reflected in the accounts of the bank.

Figure 2.3 focuses on a geographic breakdown of loss events. North America and Western Europe clearly dominate in terms of the value of the losses. This is where the majority of the worlds' largest banks are headquartered, which were particularly affected by the events leading up to, and after, the GFC.

Figure B.2.2 in the appendix shows the losses and frequency but normalised by the income level of the bank (large, medium and small). The frequency of events tends to be quite stable across bank sizes. In terms of gross losses, there is much more variability, in particular in larger banks. Moreover, a large proportion of the losses that were realised around the crisis period can be attributed to large banks. This is in line with the increased scrutiny of large banks (including domestic and global systemically important banks – DSIBs and GSIBs respectively) for their role in events alleged to have taken place in the run-up to the crisis, such as the Libor scandal and the

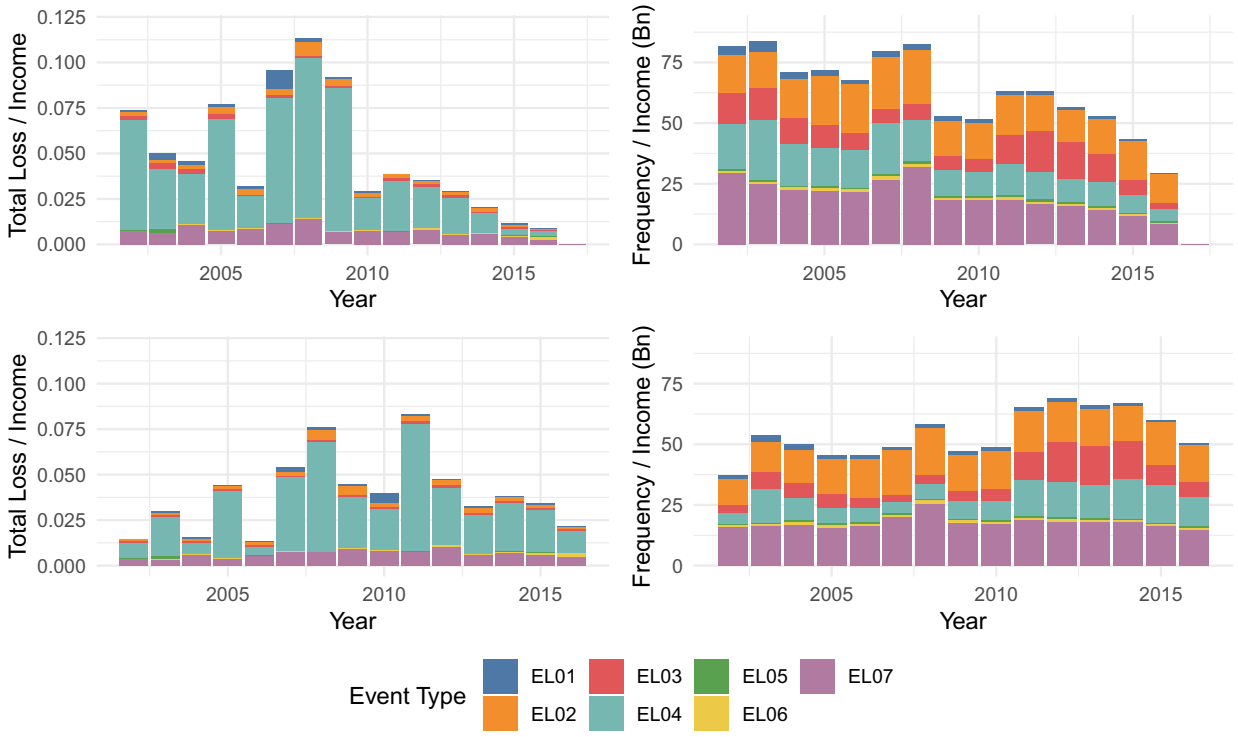

Notes: On the left hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right hand side we display the frequency divided by income (in billions). The upper panels of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by event type. EL01 = Internal fraud; EL02 = External fraud; EL03 = Employee related; EL04 = Clients, products & business practices; EL05 = Disasters; EL06 = Technology and infrastructure; EL07 = Transactions and processing.

Figure 2.2: Loss and frequency of operational losses by event type

mis-selling of mortgage-backed securities.

2.3.4 How long does it take for discovery and recognition of losses?

The time it takes for a loss to be discovered, reported and finally accounted for in banks' books can reveal important information regarding operational risks. Operational risk data suffers from an under-reporting bias, especially acute in more recent periods (Carrivick and Cope, 2013). That is, some events may have occurred but due to the fact they are not discovered or settled and accounted

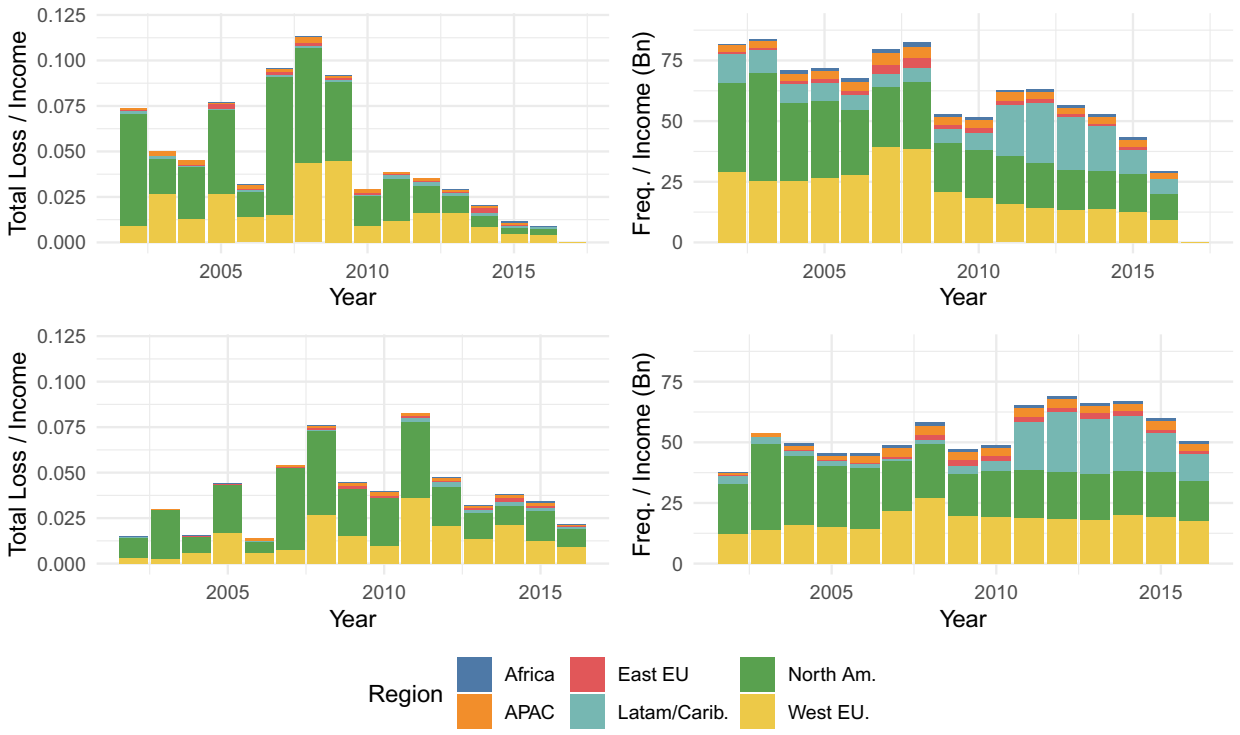

Notes: On the left hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by region. Abbreviations in the legend are defined as follows: APAC: Asia/Pacific; East EU: Eastern Europe; Latam/Carib: Latin America and Caribbean; North Am: North America; and West EU: Western Europe.

Figure 2.3: Loss and frequency of operational losses by event type

for, they are not observed in the database. Examples of such 'unobserved' incidents could be fraudulent activities that were well hidden by the perpetrator. In other cases, legal proceedings can take time to reach a settlement. The quantification of these lags is particularly relevant for CEO compensation and provides support for the introduction of the FSB's Principles and Standards on Sound Compensation (Cerasi et al., 2020). We follow up on this aspect below.

We study the duration of the three intervals defined in Figure 2.1, namely $t_1 = discovery - occurrence$, $t_2 = recognition - discovery$, and $t_3 = t_1 + t_2 = recognition - occurrence$. The

average duration of the three time intervals varies across different dimensions. Table B.2.3 in the appendix provides summary statistics for the duration of events by different categories.

In Panel A we show the breakdown by event types. Internal fraud and Clients and business practices are the incidents that, on average, take the longest to be discovered and eventually accounted for. This result is intuitive, as inside actors are likely to take steps to hide their illegal acts, which may be unearthed only when pressure from management and regulators intensifies. It is worth noting that t_1 and t_3 have a long tailed distribution: many incidents were discovered quickly, but a few extraordinary events which took a long time to be discovered and accounted for lead to a skew of the distribution. This is evident by the median often being well below the mean, as well as the high 95th quantile.

Panel B shows a summary by region. Regional differences could be driven by different regulatory approaches towards operational risk. This is more likely to manifest itself through Pillar II of the Basel capital framework, which leaves more room for supervisory discretion (i.e. how frequently are on-site inspections conducted, how efficiently is the supervisor communicating with banks). Moreover, different legal systems also affect the time to the booking of the loss in the bank's balance sheet. For example, on average, losses in North America are discovered more quickly than in Western Europe, possibly due to more pressure from supervisors and more direct supervision on operational loss problems after the GFC. However, on average, the time from discovery to recognition (t_2) is longer in North America than Western Europe, which may be an indication that the legal proceedings in North America are more protracted than those in Western Europe. Furthermore, banks of different size could face varying degrees of attention and scrutiny from regulators due to their different contribution to systemic risk. Panel C shows that, on average, larger banks face a longer duration of incidents.

Size of the data bias. As previously mentioned, the time to discovery and recognition has consequences for the completeness of the data reported. To obtain a proxy of how large the under-reporting bias might be, we can use the survival curve of the duration of time from occurrence to recognition. Since we focus on heterogeneity across regions in our regressions in Section 2.5, we look at the size of the bias by region. In Figure B.2.4 we show the survival probability by region by estimating the Kaplan-Meier curve from occurrence to recognition (t_3). This survival probability can be best interpreted as the probability of an event being accounted for after occurring. Estimating the Kaplan-Meier survival curve suggests that, depending on the region, there is approximately between 8-25% chance that an event is still unaccounted for after two years. To illustrate, Northern Europe, if an event took place in a bank in Northern Europe on the 1st January 2017, we estimate there is around 8% chance it has still not been accounted for in the books of the firm by 1st of January 2019. As the curve in Figure B.2.4 shows, this probability wanes over time.

To assess the implications of this for our data, we need to work backwards. Our granular loss data are in principle available until 2018Q3. Periods closer to this date will be associated with a higher incidence of events that have not been accounted for. By using the estimate of the survival curve, we can produce an approximate factor by which our sample could be biased. In Figure B.2.5 we show the bias factor proposed by Carrivick and Cope (2013), split by region. In the most recent year of the sample, the data could be underrepresented by around 30-100%, dependent on the region. We can apply this factor to our data by region to obtain an estimate of where the trend in frequency and losses should lie. In Figure B.2.6 we show how the annual trends in different regions might look with the correction factor.

Two approaches could remedy this problem. First, one could truncate the data to remove the years most affected by the bias. Regardless of where the database is truncated, there will be an under-reporting bias across all years, but by removing the most recent years we truncate the part of

the sample when the bias is most pervasive. Alternatively, one could use the bias factor to adjust the time series. This is not without its shortcomings, however. First, applying the correction factor may still underestimate or overestimate the actual size of unobserved losses. Moreover, it only tells us approximately how many incidents are unobserved but not much about the distribution of the losses associated with them in monetary terms. In light of this, in the next sections we opt for truncating the most recent 7 quarters of data such that the series ends at 2016Q4 (included). In this way we remove the years that are likely to misrepresent the actual losses and frequency. To maintain comparability across regressions we avoid using the correction factor.

The effect of supervision. Differences in the implementation of the Basel framework across regions could partly explain the heterogeneity in duration times in Panel A of Table B.2.3. To investigate this, we look at the cross-regional impact of regulation and supervision of banks on duration times, using the index of prudential regulation and bank supervision described in Section 2.3.

We model the duration of each t_i, accounting for the variation across these multiple dimensions, by employing a proportional hazards model as in Cox (1972). In a proportional hazards regression model, the measure of effect is the hazard rate, which is generally interpreted as the risk or probability of incurring the event of interest, conditional on the individual/entity of interest not having incurred the event up to a certain time. In our application, the hazard rate of each of the intervals can be interpreted as follows:

- $\lambda(t_1)$: probability of the loss being discovered at time t conditional on having occurred but being undiscovered until time $t_1 - 1$.
- $\lambda(t_2)$: probability of the loss being recognised in the books at time t, conditional on being discovered but not accounted for until time $t_2 - 1$.

- $\lambda(t_3)$: probability of the loss being recognised in the books at time t, conditional on having occurred and remaining unaccounted for until time $t_3 - 1$.

For each of the intervals defined above, we estimate the following equation,

$$\lambda(t_i|X_i) = \lambda_0(t) \exp(X_i\beta + FE) \tag{2.3}$$

where $\lambda_0(t)$ denotes the baseline hazard function, X_i is a vector of explanatory variables whose effect on the hazard is captured by the β coefficients. The explanatory variable in the vector X is our *supervisory index*. We include a yearly, regional, and event type fixed effects in the equation, denoted by FE. To construct X, we assign a score from the index to each observation given the year and the region in which it occurred. We have data on the supervisory index up until 2015, such that naturally losses beyond 2015 will be dropped from data used for analysis (hence the portion of our sample most affected by potential under-reporting bias is also not considered). We multiply the supervisory index by 100 to obtain a scale of 0-100, which makes the coefficients easier to interpret – a one unit increase in the supervisory index translates to a $\hat{\beta}$ increase in the likelihood of the event occurring.[15] We present the results of the regression in Table 2.4.

The estimated coefficients in the Cox proportional hazards regression model denote the change in the expected log of the hazard ratio relative to a one unit change in the independent variable, holding all other variables constant. Our results imply that increases in the supervisory index are associated with a rise in the likelihood of discovery and recognition of events. Focusing on the time from occurrence to recognition (t_3), a one unit increase in the supervisory index is associated

[15]To be clear, we do not uncover a causal relationship with this exercise. While we include various fixed effects to take into account unobserved factors that vary across years, bank size and regions, there are variables that we are not able to observe. For example, individual banks risk management and reporting practices. Moreover, we are not able to rule out reverse causality in the relationship between duration times and supervision. Supervision may become tougher if firms are lax with respect to reporting losses in a time frame deemed acceptable by supervisors.

	Dependent variable					
	t_1		t_2		t_3	
Regressor	$\hat{\beta}$	$exp(\hat{\beta})$	$\hat{\beta}$	$exp(\hat{\beta})$	$\hat{\beta}$	$exp(\hat{\beta})$
Supervisory Index	0.086*** (0.0011)	1.09	0.05*** (0.001)	1.05	0.1*** (0.0011)	1.11
Year FE	Y		Y		Y	
Region FE	Y		Y		Y	
Event Type FE	Y		Y		Y	
N	508,595		508,595		508,595	

Notes: The table contains the results of estimating a proportional hazards model. The dependent variables are the various duration measures. Standard errors are reported in parentheses. *, ** and *** denote significance at the 10, 5 and 1 percent level, respectively. We also show the exponent of the coefficient which denotes the hazard ratio.

Table 2.4: Proportional Hazard Models with Supervisory Index

with a hazard ratio 1.11 times higher than the baseline, i.e. the likelihood the event will be recognised at any date. This supports the guidance issued in Financial Stability Board (2014) regarding supervisors' interactions with financial institutions on the subject of risk culture. The report notes that since the GFC, supervisors are tending towards a more direct and intense approach to improve the resilience of the financial system. Our result supports the notion that this shift in approach should ensure that ex-post emerging risks are recognised, assessed, and addressed in a timely manner. This effect takes place not only over time, but also in the cross-section of regions. Financial institutions in regions with more effective supervisory frameworks are more likely to recognise and address operational risks in a timely manner. We note, however, that these results should be interpreted with caution, not least because we cannot claim a causal relationship given potential omitted variable and reverse causality bias.

2.4 Operational risk capital

The GFC laid bare two main shortcomings of the operational risk framework. Capital requirements for operational risk proved insufficient to cover operational risk losses incurred by some banks. Furthermore, the nature of these losses – covering events such as misconduct, and inadequate systems and controls – highlighted the difficulty associated with using internal models to estimate capital requirements for operational risk (Basel Committee on Banking Supervision, 2017).

The Basel II accord allowed three methods for calculating the capital charge assigned to operational risk. These are: i) the Basic Indicator Approach (BIA); ii) the Standardised Approach (TSA); and iii) the Advanced Measurement Approach (AMA). These methods vary in their increasing sophistication and risk sensitivity. Under the BIA, banks have simply to keep at least 15% of their gross income in the form of capital, averaged over the past three years. The TSA calculation is similar, but allows the percentage to vary according to different business lines. The AMA allows for a more sophisticated suite of methodologies to estimate the appropriate level of capital, often making use of historical loss data.

The approach in Basel III aims to streamline the operational risk framework. The three approaches in Basel II will be replaced with a single, risk-sensitive, standardised approach to be used by all banks. In this section, we outline the approaches to calculate operational risk and subsequently quantify and compare operational risk capital using the various approaches.

2.4.1 Basic indicator and standardised approaches

The simplest method that banks could use to calculate operational risk capital is the BIA. Banks that adopt the BIA must hold capital equivalent to the average over the past three years of a fixed percentage of gross income.[16] Formally, under the BIA, operational risk capital is calculated as

[16] Years of negative or zero income are excluded from the calculation.

follows,

$$K_{BIA} = \alpha \frac{1}{n} \sum_{j=1}^{3} \max(I_j, 0)$$

where I_j is the annual gross income, n is the number of previous years in which income is positive (expected to be three); and α = 0.15.

Under the Basel II framework the TSA extends the BIA by adjusting the α terms for various bank business lines (see in Table B.2.1). These are known as the β factors. Operational risk capital per business line is then calculated as follows:

$$K_{SA} = \frac{1}{3} \sum_{j=1}^{3} \max \left(\sum_{k=1}^{7} \beta_k I_{j,k}, 0 \right)$$

where k denotes the business line.

2.4.1 Basel III standardised approach

The standardised approach methodology aims to converge on a risk measure that combines the simplicity of the BIA and TSA, but also makes use of banks' historical loss information. The measure is based on the following components: (i) the Business Indicator (BI), a financial-statement-based proxy for operational risk; (ii) the Business Indicator Component (BIC), which is calculated by multiplying the BI by a set of regulatory determined marginal coefficients; and (iii) the Internal Loss Multiplier (ILM), which is a scaling factor that is based on a bank's average historical losses and the BIC. The final capital measure is calculated as,

$$K_{SMA} = BIC \times ILM$$

where the ILM is defined as:

$$ILM = \ln\left(\exp(1) - 1 + \frac{LC}{BI}\right)$$

and the Loss Component (LC) is calculated as the sum of seven times the average annual loss, seven times the average annual loss for events above 10 million Euro and 5 times average losses above 100 million Euro. The distinction in terms of various size losses aims to differentiate between banks with different loss distribution tails but with similar average loss totals (Basel Committee on Banking Supervision, 2018c).

2.4.2 Advanced measurement approaches

The AMA allows banks to use their own internal models to estimate the appropriate level of operational risk capital. Banks must demonstrate to regulators the accuracy of their internal models. Given the flexibility allowed by the AMA, the range of practices across banks has been quite broad. In Europe, the methodological focus of most banks was on using scenario analysis, while in the US the focus was on internal and external loss data (Cruz et al., 2015).

Three frameworks for calculating operational risk capital were proposed under the scope of AMA: i) Internal Measurement Approach (IMA); ii) Score Card Approach; and iii) Loss Distribution Approach. Below, we detail approaches i) and iii) to calculating operational risk from the available options under the AMA. We do not look at the Score Card Approach in great detail as it is based on subjective measures. In brief, the methodology takes a baseline level of capital which is modified based on a qualitative ranking or scoring various risks. We calculate operational risk capital based on an extension of the IMA and two LDA approaches, which we describe in detail in Appendix A. Below we describe the idea behind the LDA.

2.4.2 Loss Distribution Approach

The LDA aims to explicitly model the annual distribution of losses. In this framework, the frequency and severity of losses are each independently assumed to follow a statistical distribution, whose parameters are estimated directly from the data. The convolution of these two distributions is then used to compute the annual distribution of losses:

$$Z = \sum_{i=1}^{N} X$$

where Z denotes the annual loss, N the number of annual operational incidents and X the severity of losses. The operational risk capital is then defined as the 0.999 value-at-risk (VaR), which is the 99.9th quantile (q) of the distribution of the annual loss:[17]

$$K_{LDA} = \text{VaR}_q = \inf\{z \in \mathbb{R} : Pr[Z > z] \leq 1 - q\}$$

The VaR indicates the level of risk to which a firm, a portfolio or a single position may be exposed to over a given time period. Figure 2.4 displays an example of the distribution of annual losses, the relevant risk measures and their location on the distribution.[18]

2.4.3 Evaluating operational risk measures

As noted, at present (under Basel II), a variety of methodologies can be used to calculate a banks' operational risk capital. Under Basel III, these will be put aside in favour of a single standardised

[17] Basel II rules require banks to calculate their regulatory capital requirement as the sum of expected and unexpected losses (i.e. the 99.9$^{\text{th}}$ percentile). However, if a bank can demonstrate that it is adequately capturing expected losses in its internal business practices, it may base the minimum regulatory capital requirement on unexpected losses alone.

[18] VaR as an appropriate risk measure for capital has been challenged. Artzner et al. (1999) suggest that expected shortfall (ES) is better suited for risk management as it provides information not only about the probability of default but also about its severity. However, the use of VaR for capital allocation warrants justification from a regulator's point of view when considering minimisation of the possible shortfall and cost of capital (Cruz et al., 2015).

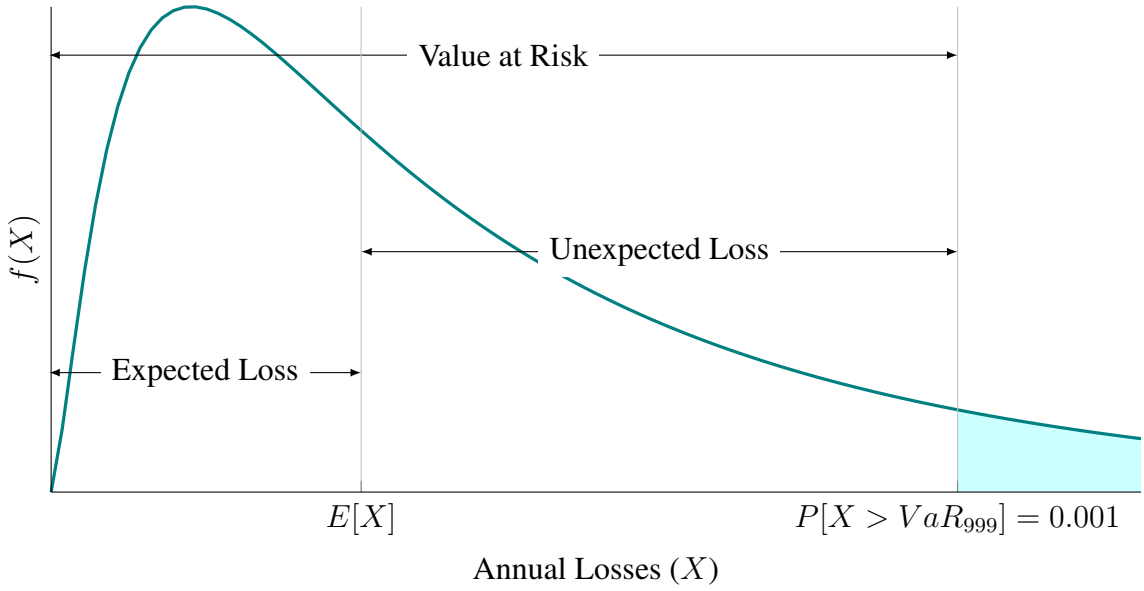

Figure 2.4: Distribution of losses and risk measures

measure. Proponents of such move suggest it will simplify the framework and provide adequately conservative measures that are not subject to gaming by participants (Tarullo, 2008; Admati, 2016). However, others suggest that the SMA may still be flawed. It is argued that practitioners would favour the granularity of the AMA approach, as without a clear regulatory requirement to keep collecting loss data at a detailed level, budgets to relevant departments could be at risk (Peters et al., 2016).

Migueis (2018) lays out some properties of an ideal approach to operational risk capital. These include, *conservatism of the measure*, *robustness to gaming*, *risk sensitivity*, *comparability*, *stability*, and *simplicity*. In this subsection, we perform a simple exercise to evaluate the different measures against these properties. Using a rolling window of 5 years of historical losses, we estimate the operational risk capital for our sample of banks based on various approaches. We then compare these estimates against the subsequent year's observed losses. Note that our estimates are not to be taken as a robust measure of operational capital. The objective here is simply to com-

pare the properties of the various estimates and we do not try to make any suggestion as to which measure is optimal for individual banks to adopt.

In Figure 2.5, we plot the estimated operational risk capital for each year versus the observed level of operational risk losses. We use five different measures of operational risk capital, the first two are the Basic Indicator Approach (detailed above) and a proxy of the Standardised Measurement approach.[19] We then use three models taken from the AMA framework. These include an analytical estimator proposed by Alexander (2008), a Monte Carlo approach to estimate the annual loss distribution, denoted as the Lognormal LDA, and a Bayesian approach to estimate the annual loss distribution, denoted as the Bayesian LDA. The details of the three approaches are contained in the online annex, in addition to the confidence intervals of each measure, where possible (see Figure B.2.3).

The BIA appears to be the most conservative estimate, as the observed losses never exceed the capital suggested by this measure. The SMA closely follows, with only a few spikes in losses exceeding the capital estimate. At the other end of the spectrum, the Lognormal LDA approach consistently underestimates a suitable level of capital. This is most likely due to a mis-specification of the severity distribution – the lognormal distribution fitted to the severity may not capture effectively the shape of the tail. The Bayesian LDA, which uses a generalised Pareto distribution, appears to explore more effectively the tail of the distribution and produces more conservative estimates. The analytical approach is reasonably conservative, although during the crisis period may have underestimated losses.

The degree of simplicity of measures varies significantly. Methodologies adopted under the AMA framework require significant statistical and mathematical expertise and are not straightfor-

[19]To calculate this proxy, we replace the Business Indicator with the BIA estimate, since we do not have granular information on the income components of banks in the sample. This puts the BI into the appropriate magnitude for computing the capital estimate. We also use a 5-year rolling window of losses rather than the proposed 10 in the Basel Committee on Banking Supervision (2018c) guidance.

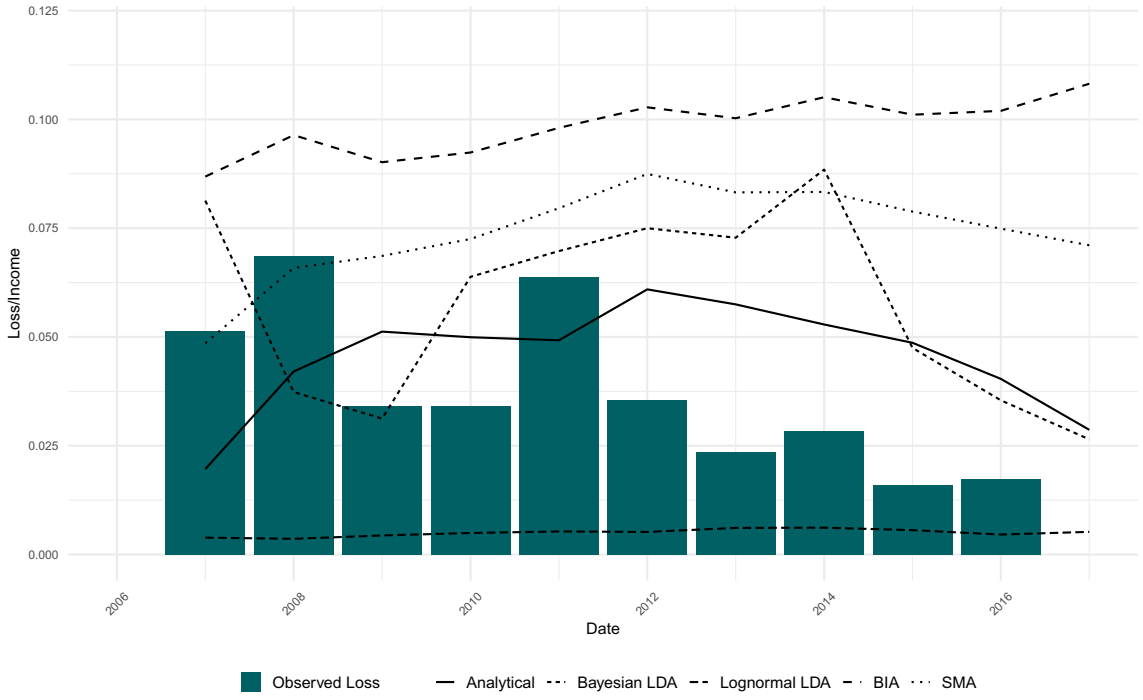

Notes: This plot shows the observed gross loss per year (bars) against the operational risk predicted by the various models. The analytical approach, the lognormal LDA, the Bayesian LDA, the Basic Indicator Approach and a proxy of the Standardised Measurement Approach. The y axis shows the observed or predicted loss - divided by the total income in the sample.

Figure 2.5: Implied capital by various approaches

ward to calculate. On the other hand, the BIA and TSA are much more clearly defined and are relatively easy to calculate. The SMA strikes a balance across the two. Simplicity also leaves banks' less scope to manipulate estimates to minimise their capital allocation. Moreover, simpler methodologies make for an easier comparability of estimates across institutions.

Risk sensitivity and the stability of capital requirements are closely related. Volatile estimates of capital can be costly for banks and arguably estimates should not be overly sensitive to risk, potentially leading to large swings in the allocation of capital (Heid, 2007). That said, capital should adjust appropriately to changes in the risk environment. As we see from our estimates, the

BIA remains relatively stable and has a coefficient of variation of 0.1, on par with the 0.13 of the SMA. In contrast, the Analytical, Lognormal and Bayesian estimates have coefficients of variation of 0.32, 0.18, 0.46, respectively. However, the cost of an overly conservative stable estimate is noted as methodologies from the LDA approach appear to adjust more appropriately with the decline in losses post 2012.

2.5 Operational losses and macroeconomic conditions

The increased risk-taking taking place during upswings in the financial cycle could be associated with operational losses surfacing down the line. Moreover, during these periods the operating environment and control structure of financial institutions could be weaker, and the implementation of controls could be viewed as restrictions to growth and entrepreneurship (European Systemic Risk Board, 2015). Abdymomunov et al. (2017) find evidence that operational losses for US banks are contemporaneously correlated with domestic macroeconomic conditions (i.e. operational losses increase in recessions). They argue that during economic downturns, banks are subject to pressures that translate into an increased likelihood of discovering losses that occurred in the past.

We extend their analysis by looking at the effect of lagged macroeconomic variables on the realisation of operational losses. We let the analysis in Section 2.3.4 guide our choice for the length of lags. In particular, we are interested in the time at which losses materialise in banks' balance sheets. However, given the significant lags between event occurrence and recognition seen in the previous section, we expect that the financial and economic environments that are conducive to risk-taking precede the actual financial impact. Looking at the survival curves for the duration between incident occurrence and recognition suggests that within two years, 87% of the incidents that occurred will have been accounted for (on average across regions). We therefore look at the cumulative effect of one and two year lags. Studying the intertemporal relationship between

operational losses and macroeconomic conditions strengthens the argument that it is in fact the excesses that take place during the upswing that lead to the occurrence of operational risk events with large associated costs, which only materialise in the books of banks a few years later.[20]

We use the lags of three different financial indicators and a supervisory index to study whether economic and financial conditions are correlated with future losses. Our variables are constructed as outlined in Section 2.3. We provide a summary of the variables in Table 2.5.

We use the credit-to-GDP gap as a measure of the build-up of financial imbalances, as also done for example in the context of the countercyclical capital buffer. The aim is to assess whether periods of excessive lending could be associated with a build up of operational risks. The average credit-to-GDP gap in our sample is around 3.04, which indicates that Credit-to-GDP ratio was, on average, above its long term trend across regions in our sample.

There has been a notable debate in the banking literature on the impact of bank competition on financial stability (Allen and Gale (2004)). We test this relationship by looking at whether periods of higher competitiveness in the banking sector are followed by periods of less/more frequent or severe operational losses. To this end, we use the Boone indicator – discussed in Section 2.3 – as the dependent variable. The average value of the Boone indicator is -0.087 with a standard deviation of 0.15.

Low interest rate environments may also influence bank risk-taking via two channels. First, low interest rates affect banks measures of risk through valuations, incomes and cash flows. Second, low yields on risk-free assets may increase financial institutions' appetite for taking on more risk. Altunbaş et al. (2014) show that low levels of short-term interest rates over an extended period of time lead to an increase in bank risk. Against this backdrop, we evaluate to what extent the monetary policy stance may be linked with a build-up of operational risk losses. To do so, we use

[20]We corroborate the findings of Abdymomunov et al. (2017) and our own by running regressions to study the contemporaneous effect of macroeconomic variables on losses. These results are available upon request.

deviations of policy rates from implied Taylor rule rates as a proxy for periods in which monetary policy has been too accommodative. The mean of the deviations from the Taylor Rule is -1.29,i.e. for our sample monetary policy has been more accommodative than a Taylor rule would imply.

Bank supervision and regulation is an integral part of the Basel framework, which ultimately aims to minimise risk in the financial sector, including operational risk. We look at the cross-regional impact of regulation and supervision of banks on operational risk using an index of prudential regulation and bank supervision. We expect the effects of regulatory/supervisory reforms not to be observed immediately, as there is a period of adjustment for banks to comply with new standards.

	N	Mean	Std. Dev.	Min	Max	Start Date	End Date
Panel A: Quarterly Variables							
Loss per Income	598	0.0031	0.012	0	0.17	2002Q1	2016Q4
Frequency per Income (millions)	600	0.0045	0.0055	0	0.035	2002Q1	2016Q4
Credit to GDP Gap	584	3.04	11.3	-33.2	35	2002Q1	2016Q4
Deviations from Taylor Rule	600	-1.29	2.58	-15.28	13.65	2002Q1	2016Q4
Panel B: Yearly Variables							
Loss per Income	150	0.0031	0.0069	0	0.044	2002	2016
Frequency per Income (millions)	150	0.0045	0.0055	0	0.029	2002	2016
Boone Indicator	127	-0.087	0.16	-0.67	0.41	2002	2014
Supervisory Index	140	0.87	0.15	0.56	1	2002	2015

Notes: The table presents a summary of the variables used in our regressions. Panel A reports a summary of our quarterly variables and Panel B the yearly variables. For each series we report information on the total number of observations, mean, standard deviation, maximum and minimum. We also provide the start and end dates for which each series were used in our regressions.

Table 2.5: Summary of regression variables

We estimate several panel regressions at the quarterly frequency for the credit-to-GDP gap and the deviations from the Taylor rule, and at a yearly frequency for the Boone indicator and regulatory and supervisory index. The regressions take the following form:

$$\ln(Y_{it}) = \sum_k \beta_k X_{i,t-k} + \alpha_i + \gamma_t + \sum_k \epsilon_{i,t-k} \tag{2.4}$$

where Y_{it}, indicates the dependent variable in region i at time t, X_{it} denotes our main independent variable (either the credit-to-GDP gap, Boone indicator, deviations from the Taylor rule, or financial and supervisory index), α_i is a regional fixed effect and γ_t is a time fixed effect. We look at three dependent variables: namely the gross loss amount, the frequency of losses, and the severity of losses (which results from dividing gross losses by frequency), all normalised by gross income.

We start by looking at contemporaneous effects, before moving into the main regressions with lagged variables. We aggregate quarterly variables to their annual counterparts and combine them in a single regression. Table B.2.4 presents the results using models that include regional and time fixed effects. We consider the contemporaneous effect on losses aggregated at the recognition date (Panel A) and occurrence date (Panel B). Deviations from the Taylor rule have the most significant effect on operational losses, consistently across regressions. When the rule suggests monetary policy is too accommodative (too restrictive) there is an increase (decrease) in losses and the frequency of events. This holds both when doing the analysis by date of occurrence and date of recognition. Our results also suggest that more intense bank competition is associated with lower operational losses. Finally, the supervisory index is insignificant -– although this may be subject to reverse causality bias, as an increase in losses may prompt a tightening of supervisory measures.

Table 2.6 present the main results of this section, looking at the link between the lagged variables discussed above and operational losses.[21] When interpreting these results, it is important to bear in mind that they may be subject to omitted variable and reverse causality issues – hence

[21]The coefficients are the cumulative effect of the lagged dependent variables. The standard errors reported in parentheses are the standard error of the sum of the coefficients.

one should be careful not to give a causal interpretation. While we do control for instance for region and time fixed effects, as well as rely on lagged variables, this may not fully eliminate such concerns.

Gross losses and event frequencies are both positively correlated with the credit-to-GDP gap (Panel A), but not statistically significant. The results in Panel B suggest that more intense bank competition is associated with lower operational losses in subsequent periods. Recall that the more negative the Boone indicator the higher the competition in the banking sector, therefore a one standard deviation decrease in the Boone indicator (indicative of a more competitive market) is associated with a cumulative 29% decrease in annual operational losses as a fraction of income.

In Panel C, we see the results from the regressions including the deviations from the Taylor rule. The results suggest that following periods of overly accommodative monetary policy, operational losses increase in frequency and value. This provides support to the notion that risk-taking in low-yield environments can lead to a build-up of operational losses. A one standard deviation decrease in the Taylor gap is associated with a 20% increase of operational losses in the following four quarters and 28% after 8 quarters.

Panel D contains the results for the financial and supervisory index. Higher scores on the index are associated with lower gross amounts and frequency of operational losses per unit of income. The index ranges between 0.56 and 1 in the sample, and it is slow-moving because it depends on institutional characteristics. Operational losses are very sensitive to changes in the index: A 0.1 increase in the supervisory score is associated with a decrease in the gross loss (frequency) per unit of income of around 40% (26%) one year after. The cumulative effect of a 0.1 increase in two subsequent years rises in excess of 35% (20%) for gross loss (frequency) per unit of income. The severity of incidents also appears to fall after two years. Our results suggest that more stringent supervisory frameworks may help offset operational risks by reducing the frequency of

	Dependent variable		
	$\frac{TotalLoss}{Income}$	$\frac{Frequency}{Income}$	$\frac{Severity}{Income}$
Panel A			
Credit-GDP-Gap - 4 Lags	0.0094	0.010	-0.0011
	(0.011)	(0.011)	(0.0053)
Credit-GDP-Gap - 8 Lags	0.010	0.0086	0.0019
	(0.010)	(0.0096)	(0.0048)
Panel B			
Boone Ind. - 1 Lag	1.8*	1.1**	0.76
	(0.95)	(0.50)	(0.80)
Boone Ind. - 2 Lags	1.4	0.72	0.67
	(1.0)	(0.60)	(0.92)
Panel C			
Taylor Rule Dev. - 4 Lags	-0.079***	-0.086**	0.0070
	(0.022)	(0.035)	(0.026)
Taylor Rule Dev. - 8 Lags	-0.11***	-0.13*	0.018
	(0.036)	(0.070)	(0.043)
Panel D			
Supervision Index - 1 Lag	-3.9*	-2.6**	-1.3
	(2.0)	(1.1)	(0.99)
Supervision Index - 2 Lags	-3.5*	-2.0*	-1.6*
	(1.8)	(1.0)	(0.92)
Regional Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Notes: The table is divided into four panels summarising the results from 24 panel regressions. Each column denotes the dependent variables used, which are lagged. The coefficients shown are the sum of the lagged variables, i.e. the cumulative effect – for example at 4 lags the coefficient reported is, $\sum_{i=1}^{4} \hat{\beta}_i$. A robust sum of standard errors is reported in parenthesis. The sum of standard errors is calculated as $\sqrt{L'VL}$, where L is a (0,1) vector that denotes the linear combination of regressors and V is the estimated robust covariance matrix. We test that the sum of coefficients is significantly different from zero. The asterisks denote the significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions are two-way fixed effects models, including a regional and time effect. In Panels A and C the time unit is quarters, in Panels B and D the time unit is years.

Table 2.6: Operational losses, macroeconomic conditions and the regulatory environment

their occurrence, as presumably they lead banks to implement better risk management strategies.

2.6 Cyber risks in the financial sector

Cyber and related IT risks can be seen as a subset of operational risks and are frequently cited as a prominent threat to the financial system (Kopp et al., 2017; Kashyap and Wetherilt, 2019). In March 2017, the G20 Finance Ministers and Central Bank Governors noted that "the malicious use of information and communication technologies (ICT) could disrupt financial services crucial to both national and international financial systems, undermine security and confidence, and endanger financial stability". In December 2018 the Basel Committee on Banking Supervision published a report on the range of cyber-resilience practices (Basel Committee on Banking Supervision (2018b)). The Covid-19 pandemic may have opened up new possibilities for attacks. Given the widespread use of work-from-home arrangements, especially in the financial sector, threat actors are able to leverage operational uncertainty and the use of personal devices (Dingel and Neiman, 2020; Aldasoro et al., 2021).

An accurate quantification of cyber risks is challenging, as there is no precise definition of cyber events. This naturally also applies to the ORX database. We thus need to rely on a number of assumptions. In particular, we make use of event type definitions and consider as cyber events a subclass of operational risks events. Table 2.7 describes the event categories that are most likely to be associated with cyber events. As discussed above, we use the level 2 event type classification in order to compute a proxy range for cyber events. Given the nature of the classification, we are not able to accurately capture all events. Other categories not included could in principle have some cyber events within them. Similarly, some events included in the categories we consider might not be cyber-related, especially for the upper bound estimate. This approach is largely in line with the classification used by Curti et al. (2019a). We diverge slightly by not taking into account the Transactions and processing (EL07) category. This category is quite widely defined and it would

be very difficult to filter the non-cyber related incidents out. The full list presented in Table 2.7 (i.e. bold plus non-bold) constitutes our upper bound estimate for cyber events. We highlight in bold the event types we consider as a lower bound to approximate cyber events, after discussions with risk management experts acquainted with the event type categorisation.

Event Type Level 1	Event Type Level 2	Description
Internal Fraud	Unauthorised activity	e.g. Rogue trading, unreported transaction, mismarking positions
	Internal theft	e.g. Forgery, theft, extortion, embezzlement, bribes/kickbacks
	System security (internal)	Intentional damage to systems by internal staff
External Fraud	External theft and fraud	e.g. Robbery, Forgery, Cheque Kiting
	System security (external)	Wilful Damage e.g. Hardware/Software, Hacking Damage, Theft of Data
Technology and infrastructure failures	-	Losses arising from disruption of business or system failures

Notes: The table denotes the definitions of event types that could proxy for cyber related incidents. Taken together presents our upper bound on cyber risk and in bold are those that are used as our lower bound definition of cyber risk.

Table 2.7: Definitions of cyber event types

We first present summary statistics to provide a comparison of cyber losses with other operational losses. Table 2.8, presents statistics on the total number of incidents, mean, standard deviation and maximum values, by cyber and non cyber events. We provide summaries for both our lower bound and upper bound. There are 13,561 cyber events within the database according to our lower bound definition, which is a minor fraction of all losses, around 2%. The upper bound captures a much wider range of events and is roughly representative of a third of the incidents in the database. The true number of cyber incidents likely lies somewhere in between that range. When considering features of the distribution of cyber losses, the lower bound may be a better guide as the upper bound is likely to be populated with a significant amount of noise. Across both bounds we see a higher average cost for non-cyber events and also a larger standard deviation.

	N	Mean	Standard Deviation	Max
Panel A: Lower Bound				
Non-Cyber	596,293	627,195	49,844,267	23,705,540,000
Cyber	13,561	476,541	19,710,650	2,224,579,168
Panel B: Upper bound				
Non-Cyber	397,439	841,028	60,666,399	23,705,540,000
Cyber	212,415	217,484	10,615,633	4,056,523,958

Notes: The table presents summary statistics for losses by cyber and non-cyber events. Panel A presents summary statistics for the lower bound of cyber losses versus non-cyber losses. Panel B presents summary statistics for the upper bound of cyber losses versus non-cyber losses. With the exception of the first column, figures reported are in Euros.

Table 2.8: Cyber losses – summary statistics

We also present a time series of frequencies and amounts in Figure 2.6, as well a breakdown of losses and frequency by region, "cyber" event types and bank size, reported in Figures B.2.7-B.2.9 in the appendix.[22] The dominating event type is Technology & Infrastructure. Since "system security (external)" captures damage from hacking, we assume that these are typically failures that are out of the control of the firm – a typical example being a power outage. Damages from hacking appear to be low. In a companion paper, we show, using a different dataset which focuses only on cyber events, that the financial sector is relatively more resilient than other sectors in riding out attacks with malicious intent, most likely thanks to investments in security practices done by banks also under the auspices of regulators (Aldasoro et al. (2022)).

In terms of regions, Western Europe suffers more cyber losses than other regions, with the exception of 2016, when considerable cyber losses occurred in the U.S. When doing the split by bank size, in turn, the share across banks appears to be relatively stable. The peak in 2016, however, can be largely attributed to small and medium-sized banks. This could be an indicator that larger budgets and thus more investment in security pays dividends for larger banks.

[22]For the sake of space, we report these only for the lower bound estimate of cyber losses.

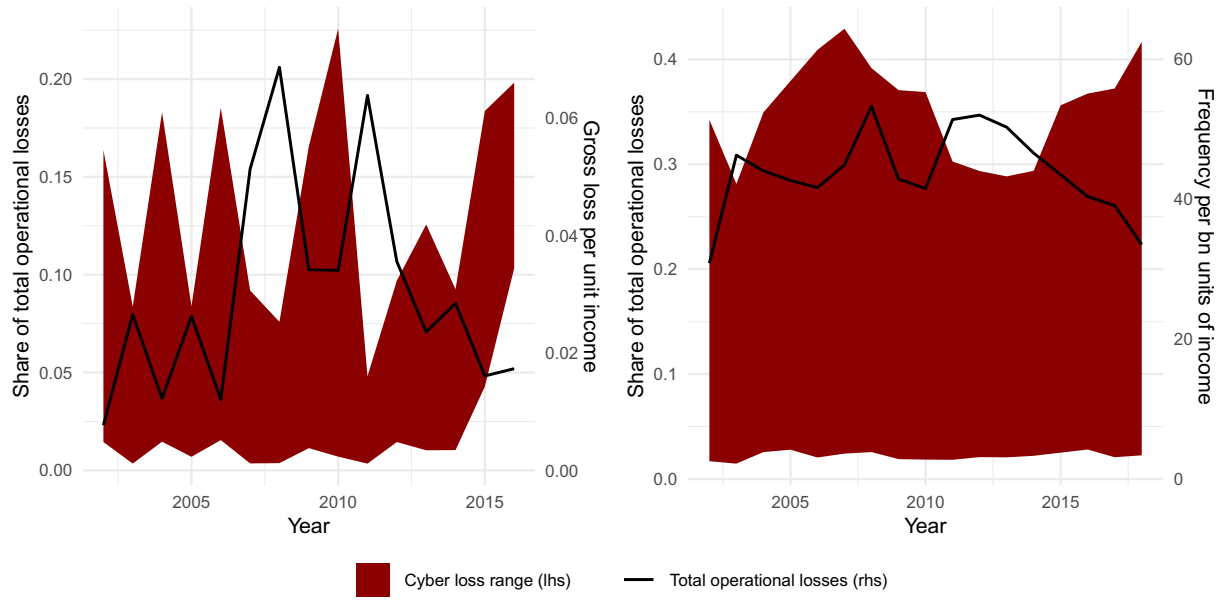

Notes: The left hand axis of the plots shows the estimated range of cyber losses across years as a share of all operational losses, which is shown by the red area in the graphs. The right y-axis in the left (right) panel gross losses (frequency) per unit of income. Events are aggregated by the date of recognition.

Figure 2.6: Operational and cyber events

2.6.1 Cyber risk capital

As a sub-component of operational risk, a proportion of capital should be allocated to account for losses stemming from cyber incidents. To complement the analysis in Section 2.4, we also compute estimates of cyber risk capital. We perform a similar exercise, by computing the cyber risk capital over time, but focusing solely on the Bayesian methodology. We compute estimates for both the lower and upper bounds as defined above. The results are summarised in Figure 2.7, which includes the estimate for total operational risk (red line) as a benchmark.

We use the Value-at-Risk (VaR) as the measure of appropriate capital from the estimated cyber loss distributions. The value of the VaR for the distribution of cyber losses is only a fraction of total operational VaR if the calculation is based on the analytical approach. At the lower bound,

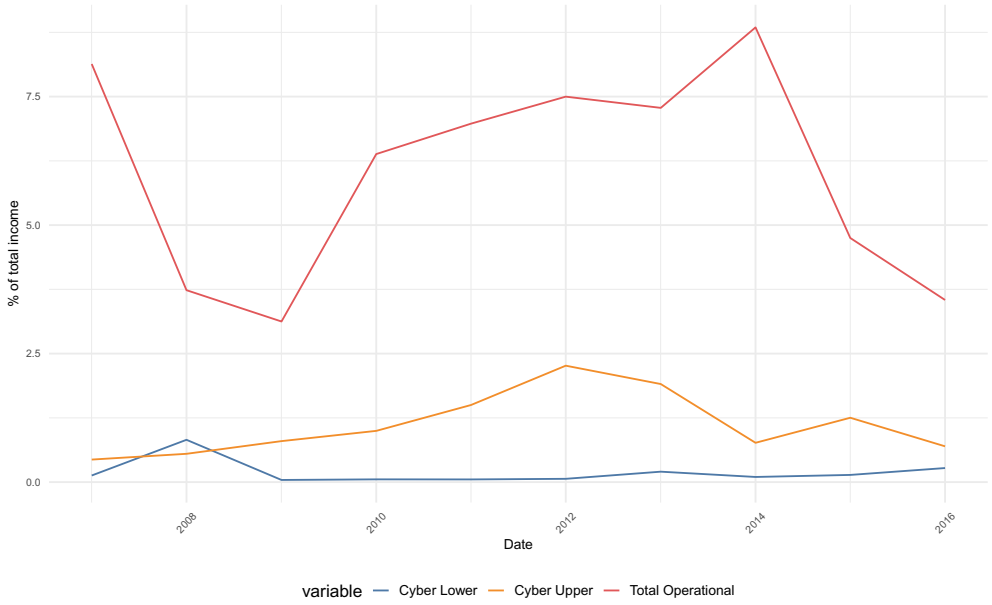

Figure 2.7: Operational and cyber value-at-risk

the value ranges between 0.04-0.8% of the gross income of the consortium, which corresponds to around 338 EUR millions and 7.8 EUR billions, respectively. At the upper bound this can jump up to around 2.5% at the peak in 2012. These figures reflect that cyber risk is a small fraction of total operational risks, as discussed above. At the time of the peak this would have represented around a third of total operational losses. These results should be interpreted with caution, not least in that they should be taken to underplay the threat of cyber risks. First, by construction the definition of operational risk is much broader and encapsulates cyber risk and thus will naturally be larger. Second, cyber is an emerging risk and reporting cyber-related losses is not always mandatory – thus their true distribution is very challenging to estimate. Accordingly, not all the costs of cyber events may be covered in our approximation. Third, our estimates group losses across the entire consortium and thus represent the total impact of cyber incidents on the financial sector as a whole. However, an isolated incident that leads to the business disruption of a large financial

institution and/or market infrastructure could have dire consequences for the institution and pose a significant systemic risk due to risk concentration and the lack of substitutes in the case of financial market infrastructures. Potential scenario analyses include cyber-attacks affecting the availability of a major payments system, or a breach that compromises the confidentiality of key financial or personal data, or corrupts the data of a major financial institution or data provider (Boer and Vazquez, 2017; Monetary Authority of Singapore, 2018; European Systemic Risk Board, 2020). One key finding is that intentional data manipulation could be especially damaging, as it may erode confidence, triggering feedback loops, and require a prolonged recovery period.

2.7 Conclusions

The GFC drew the attention of regulators and academics towards operational risk. Moreover, the shift to the new Standardised Approach in Basel III and especially the threat of cyber events feature prominently in policy debates around operational risk. We contribute to the debate by using a unique cross-country dataset at the operational loss event level for over 14 years and more than 70 large banks.

We provide stylised facts as a basis for discussions of operational risk in the financial sector. After a spike in operational losses in the immediate aftermath of the GFC, operational losses declined. The post-crisis spike is to a large extent accounted for by the severity of losses related to improper business practices that occurred in large banks in the run-up to the crisis, which materialised only later. An example of such event is the mis-selling of mortgage-backed securities that took place around 2005/2006 but was crystallised as a loss in the books of banks only a few years later, when heavy fines were imposed.

We compute operational value-at-risk and show it can vary substantially depending on the methodology. The average VaR for the financial institutions in the sample ranges from 1% to

7.5% of total gross income, depending on whether the method used is better able to capture the heavy-tailed nature of the data. These numbers are consistent with actual capital requirements, but notably smaller than the basic indicator approach. Our results provide some support for the shift to the standardised approach in Basel III. First, this would reduce heterogeneity of estimates across banks that come from various AMA methodologies. Moreover, the simplified approach could also free up resources at banks and supervisory authorities.

We document a substantial lag between the dates of discovery and recognition of loss events. On average, it exceeds one year, but it varies across regions, business lines, event types, and bank size. Internal fraud events and failures due to improper business practices are less likely to be discovered than other events, especially when the size of the financial firm is small. These findings can inform policy discussions on compensation practices.

We show that operational losses are higher after periods of excessively accommodative monetary policy. In other words, the link between monetary policy, and bank risk-taking found in the literature also extends to operational risk-taking. A higher quality of financial regulation and supervision is associated with lower operational risk losses. We also find that periods of increased bank competition correlate with future reductions in operational losses.

Finally, we use the categorisation of operational loss events to compute a proxy range of cyber events, a subset of operational events. Cyber losses represent a relatively small portion of overall operational risk losses, especially in terms of frequency. That said, recent years saw a notable increase in losses due to cyber events, with a strong peak in 2016. We note that a higher quality of financial regulation and supervision is also associated with lower cyber losses. Despite representing a relatively minor share of operational losses, cyber losses can account for up to a third of total operational risk capital. Better estimating the cost of cyber events for financial institutions is an important area for future research.

CHAPTER 3

LIBRA OR LIBRAE: DIGITAL CURRENCY BASKETS

3.1 Introduction

Carney (2019) posed the question of whether a Synthetic Hegemonic Currency (SHC) would be best provided by the public sector. The rationale behind is that a global currency, underpinned by a basket of reserve assets, could better support global outcomes. For example, a SHC could dampen the dominating influence of the US dollar on global trade, alleviate spillovers to exchange rates from shocks to the US economy, and trade across countries would become less dependent on the dollar.

Discussions around global currencies have been reignited in the overarching debate around central bank digital currency (CBDC) and stablecoins. From the private sector, Facebook, amongst numerous others, have announced plans for their own privately issued stablecoin that could emulate the characteristics of a SHC. In the most recent iteration of Facebooks proposition, the idea is to supply digital tokens that are pegged to major currencies, i.e. LibraUSD would be pegged to the US dollar.[1] In addition, there will be another token whose value is derived from a weighted basket of the currencies provided on the platform. The exact composition of this underlying basket and its targeted exchange rate is unspecified. In this paper, we assume that the objective is to devise a digital currency whose exchange rate fluctuations are minimised against several currencies of the worlds major currencies.

Facebook's plans have been met with some resistance from regulators and will face intense

[1]The project has since rebranded to the Diem Association: https://www.diem.com/en-us/association/

scrutiny before receiving any kind of regulatory approval. Why have regulators reacted with such caution to Facebook's plans to issue a stablecoin? Firstly, as a tech-giant Facebook can push Libra to its vast user-base, approximately 2.41 billion monthly active users.[2] To put this into perspective, currently it is estimated there are around 40 million bitcoin wallets and 1 million daily users.[3] Facebook would have to successfully penetrate 2% of its user base to match what is an upper bound on a proxy for the size of bitcoin's user base, the most frequently used cryptoasset. A subsequent concern was for the basket based currency to potentially interfere with monetary sovereignty and monetary policy, leading to a form of dollarisation.

Against this background, we investigate the empirical aspects of the design of a currency basket i.e. "Librae" in the sense that the value is composed of several currencies. First, we consider the optimal weights of the basket of underlying reference currencies, such as those included in the International Monetary fund Special Drawings Rights (SDR). After computing the optimal weights we construct the historical values of the designed stablecoin (SAC) and compare its volatility against a set of major currencies. For the optimal allocation of weights in the currency basket we follow Hovanov et al. (2004) to compute a currency invariant index. A particular advantage of this approach is that given a fixed set of currencies, the index of a currency will have the same value, regardless of the base currency choice. The index is determined by minimising the variance of a portfolio of currencies, expressed in Reduced Normalised Values (RNVALs). We construct a reference basket that contains the Dollar (USD), the Euro (EUR), the Yen (JPY), the Renminbi (CNY) and the Pound Sterling (GBP), the same currencies that are employed for the determination of the IMF's Special Drawing Rights (SDR) basket. We use daily foreign exchange rate data from January 2002 up until December 2020.

[2]https://newsroom.fb.com/company-info/

[3]The number of bitcoin wallets: https://www.statista.com/statistics/647374/worldwide-blockchain-wallet-users/ and the number of active wallets: https://coinmetrics.io/.

By construction, our basket based currency should have the lowest variance in comparison to those contained in the basket and our results confirm this. However, it is of interest to see how our basket fares against currencies outside of the basket, for example against the currencies of the most important remittance markets. We investigate whether the properties of the basket based stablecoin offer a hedge or alternative utility to foreign workers dependent on remittance markets. The association's white paper, Libra Association (2020), the use of the basket currency coin is motivated as follows:

> For countries that do not have a single-currency stablecoin on the Libra network, we believe LBR is a neutral and low-volatility alternative that could ensure users in such regions can benefit from accessing the network and increased financial inclusion. In this context, LBR could operate as a settlement coin in cross-border transactions, and people and businesses could convert the LBR they receive into local currency to spend on goods and services through third-party financial service providers. [...]

It is evident that the association has identified an opportunity for the basket-based currency to have a role to play in remittance markets. To provide an initial empirical study in evaluating the effectiveness of the basket to function as a tool in remittance markets, we first reconstruct the price of our SAC in base dollar. We then recompute the RNVALs, including the SAC and extending the currencies to those of major remittance markets, namely the Indian Rupee, the Mexican Peso, the Philippine Peso and the Nigerian Naira. Our empirical findings show that, overall, the stablecoin maintains the lowest volatility, thus could act as a hedge and store of value for overseas workers savings. We also make comparison to the IMF's SDRs, which performs almost as well.

To gain insight into the composition of the basket, we finally study the currencies which drive the volatility spillovers among exchange rates, using the framework of Diebold and Yılmaz (2014).

Specifically, we build a spillover network decomposition analysis of the currencies up to December 2020, thus including the period of the Covid-19 outbreak. Our spillover network decomposition shows that the USD is the currency whose dynamics has the largest impact on the others, especially in terms of exporting contagion, although in the latest period CNY has begun overtaking, and the Covid-19 crisis has disrupted the pre-existing equilibria. As a consequence, a shock to USD, expressed by a one standard deviation decrease in its normalised value with respect to the other currencies, causes a shock on all currencies and, through high order contagion, on the USD itself, leading to a new lower equilibrium. Differently, a shock in the value of the SAC, caused by a shock of a currency in the basket, is offset by the diversification effect and, therefore, the starting equilibrium is maintained. This implies that remittances converted in basket based stablecoin better maintain their value, with respect to those converted in dollars (or dollar based stablecoins).

3.2 Related Literature

Cryptocurrencies were primarily conceived under the advent of Bitcoin, outlined in Nakamoto (2008). This was the first decentralized payment system based on maintaining a public transaction ledger. Since then, as many as 5,500 cryptocurrencies exist as of 24 May 2020. Several authors have dealt with the description and functioning of cryptocurrencies (Segendorf, 2014; Dwyer, 2015). Legal concerns that have arisen through cryptocurrencies are discussed in Murphy et al. (2015). Cheng and Dai (2020) study the inflow capital control evasion phenomenon in cryptocurrencies and show that the relative CNY to USD bitcoin price, which indicates capital inflow volumes, of those with high excessive currency conversion profit reacts more negatively to carry trade returns.

Several studies have analysed cryptocurrencies statistical properties. Corbet et al. (2018) investigate the dynamic spillovers of cryptocurrencies with other financial assets, and find that the

two categories of financial instruments are isolated. Using a similar approach, Giudici and Pagnottoni (2019) and Giudici and Pagnottoni (2020) explore the dynamic relationship of Bitcoin exchanges and show their relative importance in transmitting information of fundamental Bitcoin price changes. The latter studies complement the findings achieved in the field of price discovery on Bitcoin (and cryptocurrency) exchanges – see Brandvold et al. (2015), Pagnottoni and Dimpfl (2019) and Dimpfl and Peter (2020). Katsiampa et al. (2019) examine the volatility interaction of eight cryptocurrencies through the diagonal BEKK and Asymmetric Diagonal BEKK methodologies and find that despite shocks in Bitcoin are most persistent, the cryptocurrency is not dominant. Bouri et al. (2019) and Resta et al. (2020) evaluate the effectiveness of several technical trading rules in cryptocurrency markets and provide support to the best performances of moving average based strategies. Other contributions in the field use network models and neural networks for cryptocurrency portfolio management and Bitcoin option pricing – see Giudici et al. (2020) and Pagnottoni (2019).

Research on stablecoins is a growing topic. The Financial Stability Board (2019) defines a 'stablecoin' as a crypto-asset designed to maintain a stable value relative to another asset (typically a unit of currency or commodity) or a basket of assets. Bullmann et al. (2019) make distinctions between stablecoins based on the collateral that underpins them, varying from cash in commercial bank accounts to the use cryptocurrencies like used in the MakerDao project. The Libra association has outlined plans to invest the funds that are received in return for stablecoins into "safe" assets e.g. sovereign debt. Stablecoins are close substitutes for cash, similarly to electronic money. This is not the first time that electronic money has been on the agenda for central banks and policy makers, after a flurry of innovations in this space, in 1996 and 1998 respectively the BIS and ECB published reports addressing the regulation of e-money and the implications for monetary policy (European Central Bank, 1998; Bank for International Settlements, 1996). However, these forms

of e-money never really gathered sufficient traction to trouble the initial concerns of policy makers of the time[4].

Keynes originally suggested the *bancor* as a unit of account of his proposed International Clearing Union, intended to fix to the dual dollar gold system. The solution was eventually conceived by the IMF who approved the SDR in 1967. The IMF's issuance of SDR could be seen as a supranational currency issued by central banks, although the SDR does not fulfil all functions of money. Whilst serving as a store of value and unit of account, SDRs are only used by some central banks and international institutions as a means of exchange to pay each other (Ocampo, 2019). For this, they may not be strictly considered as a"true" global currency.

A boost to the importance of SDRs was given in 2009, when China called for reforms to the international monetary system by adopting the SDR as a reserve asset. Against these developments, Humpage (2009) suggests that while the adoption of the SDR as a reserve asset is technically feasible, it would not reduce the dollar's role any time soon. Many foreign-exchange transactions, even excluding US residents, are denominated and settled in dollars. Producers typically invoice their products in dollars, which keeps their prices in line with their competitors and simplifies cross-border price comparisons among producers (Gopinath et al., 2016). Given the persistent importance of the US dollar, the question is whether this will remain so under the fintech transformation that is changing the financial world. And, in particular, whether a dollar based stablecoin is more likely to be adopted than a basket based one.

Flore (2018) recently notes the impact that blockchain could have on reducing costs in remittance markets. As previously mentioned a stablecoin backed by a basket of currencies could be an attractive asset for foreign workers that depend on remittances. Under the status quo, an appreciation in the value of the domestic currency can reduce the remittances ratio because workers

[4]For example, see Levene (2006)

want to to keep the additional earning from the appreciation of the currency. On the other hand, workers based in foreign countries, where the value of the domestic currency is declining, may remit money on an urgent basis. A basket based currency could dampen some of these effects as it is less susceptible to appreciation and depreciation of the domestic and foreign currencies. However, the effects are likely to be ambiguous and depend on how the stablecoin is used. If it gains acceptability in the home currency this could lead to new episodes of dollarisation, whereas if the currency is only used as a medium of exchange the effect could be negligible.

One specific challenge for countries that face large inflows of worker remittances could lead to the emergence of "Dutch disease," that is, remittance inflows could result in an appreciation of the equilibrium real exchange rate that would tend to undermine the international competitiveness of domestic production, particularly that of nontraditional exports. Barajas et al. (2011) note that reasonable modifications in the modelling of the factors driving remittances, or in the various macroeconomic roles that remittances could moderate or even reverse the expected impact of remittance flows on the equilibrium value of the real exchange rate. Acosta et al. (2009) discuss two mechanisms by which this occurs, the first mechanism is demonstrated in the Salter-Swan-Conder-Dornbusch model, which points to a "spending effect," by which the increase in wealth following higher capital inflows from remittances, combined with exogenous tradable prices, causes the prices of nontradable goods and services to rise. The second mechanism is that remittances tend to increase household aggregate wealth. An increase in household wealth may lead to a decrease in labor supply as households substitute more leisure for work. A shrinking labor supply, in turn, puts upward pressure on wages. Rising wages raise production costs, and higher production costs can lead to a further contraction of the tradable sector.

3.3 Methodology

In this section we outline the methodologies employed in our empirical application. Firstly, we describe the optimal control problem which yields the optimal stablecoin weights. Secondly, we introduce our VAR model and, based upon which, we study the spillover effects across the currencies in the basket to determine their interconnectedness to understand which are the most relevant ones in terms of shock transmission.

3.3.1 Optimal control problem

We aim to build a basket of predetermined (reference) currencies with optimal weights, namely, weights which minimise the variability of a basket based stablecoin. This translates into an optimal control problem which takes as an objective function minimising the variance of the value of the basket by finding the optimal combination of weights assigned to the basket.

Hovanov et al. (2004) show that the values of any given currency depends on the base currency chosen. The latter fact creates ambiguity in evaluating the currency itself and its dynamics. To overcome this issue, Hovanov et al. (2004) proposed a reduced (to the moment t_0) normalized value in exchange (RNVAL) of the i-th currency:

$$\mathrm{RNVAL}_i\left(t/t_0\right) = \frac{c_{ij}(t)}{\sqrt[n]{\prod_{k=1}^{n} c_{kj}(t)}} \Big/ \frac{c_{ij}\left(t_0\right)}{\sqrt[n]{\prod_{k=1}^{n} c_{kj}\left(t_0\right)}} = \sqrt[n]{\prod_{k=1}^{n} \frac{c_{ik}(t)}{c_{ik}\left(t_0\right)}} \tag{3.1}$$

where $c_{ij}(t)$ denotes the exchange rate between currencies i, j at time t, with $i; j = 1, ..., n$ (where n denotes the number of currencies). By reducing to the moment t_0 and normalizing each currency observation by the geometric average of the other currencies at that specific point in time, the RNVAL allows the computation of a unique optimal, minimum variance currency basket, despite the base currency choice. The minimum variance currency basket is derived by searching

the optimal weight vector $w^* = (w_1^*, ..., w_n^*)$ which solves the following optimal control problem:

$$\text{Min}\left(S^2(w) = \sum_{i,j=1}^{n} w_i w_j \,\text{cov}(i,j) = \sum_{i=1}^{n} w_i^2 s_i^2 + 2\sum_{i,j=1}^{n} w_i w_j \,\text{cov}(i,j)\right) \tag{3.2}$$

subject to

$$\begin{cases} \sum_{i=1}^{n} w_i = 1 \\ \qquad w_i \geq 0 \end{cases}$$

where $cov(i,j)$ is the covariance of time series $RNVAL_i(t = t_0)$ and $RNVAL_j(t = t_0)$, $s_i^2(w)$ is the variance of the time series $RNVAL_i(t = t_0)$, with $i; j = 1, ..., n$, $t = 1, ..., T$.

The solution to the constrained optimisation problem in Equation (3.2) yields to the minimum variance weights which enable us to construct the stablecoin value.

3.3.2 VAR models and spillover analysis

We evaluate spillovers using the methodology of Diebold and Yilmaz (2012), which has been widely employed in the literature with the aim of measuring return and volatility spillovers – see, for instance, Abosedra et al. (2020). As in their seminal paper, we start from estimating a Vector AutoRegressive (VAR) model, that is:

$$x_t = \sum_{i=1}^{k} \Phi_i x_{t-i} + \varepsilon_t \tag{3.3}$$

where x_t is the $(n \times 1)$ vector of first differences in RNVALs at time t, Φ_i the $(n \times n)$ VAR parameter matrices, k the autoregressive order, ε_t a zero-mean white noise process having variance-covariance matrix Σ_ε, with n being the number of currencies considered in order to build the basket.

Note that the VAR model is built on the variables' first differences, as this ensures the stationarity of the analyzed time series.

The VAR in Equation 3.3 may also be rewritten in its corresponding vector moving average (VMA) representation, that is

$$x_t = \varepsilon_t + \Psi_1 \varepsilon_{t-1} + \Psi_2 \varepsilon_{t-2} + \cdots \tag{3.4}$$

where $\Psi_1, \Psi_2, ...$ the $(n \times n)$ are the matrices of VMA coefficients. The VMA coefficients are recursively computed as $\Psi_i = \Phi_1 \Psi_{i-1} + \Phi_2 \Psi_{i-2} + ... + \Phi_i \Psi_1$, having $\Psi_i = 0 \; \forall i < 0$ and $\Psi_1 = I_n$.

As it is widely accepted in the financial econometric literature, the variance decomposition tools are used to evaluate the impact of shocks in one system variable on the others. Strictly speaking, variance decompositions decompose the H-step-ahead error variance in forecasting x_i which is due to shocks to x_j, $\forall j \neq i$ and $\forall i = 1, ..., n$.

Diebold and Yilmaz (2012) founded their methodology on the H-step ahead forecast error variance decomposition. Considering two generic variables x_i and x_j, they define the own variance shares as the proportion of the H-step ahead error variance in predicting x_i due to shocks in x_i itself, $\forall i = 1, ..., n$. On the other hand, the cross variance shares (spillovers) are defined as the H-step ahead error variance in forecasting x_i due to shocks in x_j, $\forall i = 1, \cdots, n$ with $j \neq i$. In other words, denoting as $\theta_{ij}^g(H)$ the KPPS H-step forecast error variance decompositions, with $h = 1, \cdots, H$, we have:

$$\theta_{ij}^g(H) = \frac{\sigma_{jj}^{-1} \sum_{h=0}^{H-1} (e_i' \Psi_h \Sigma e_j)^2}{\sum_{h=0}^{H-1} (e_i' \Psi_h \Sigma \Psi_h' e_i)} \tag{3.5}$$

with σ_{jj} being the standard deviation of the innovation for equation j and e_i the selection vector, i.e. a vector having one as i^{th} element and zeros elsewhere. Intuitively, the own variance shares and cross variance shares (spillovers) measure the contribution of each variable to the forecast error variance of itself and the other variables in the system, respectively, thus giving a measure of the importance of each variable in predicting the others.

Note that the row sum of the generalized variance decomposition is not equal to 1, meaning $\sum_{h=0}^{H-1} \theta_{ij}^{g}(H) \neq 1$. Diebold and Yilmaz (2012) circumvent this problem by normalizing each entry of the variance decomposition matrix by its own row sum, i.e.

$$\tilde{\theta}_{ij}^{g}(H) = \frac{\theta_{ij}^{g}(H)}{\sum_{j=1}^{n} \theta_{ij}^{g}(H)} \tag{3.6}$$

This tackles the above mentioned issue and yields to $\sum_{j=1}^{n} \tilde{\theta}_{ij}^{g}(H) = 1$, and $\sum_{j,i=1}^{n} \tilde{\theta}_{ij}^{g}(H) = n$. As a measure of the fraction of forecast error variance coming from spillovers, Diebold and Yilmaz (2012) define the total spillover index (TSI):

$$TSI(H) = \frac{\sum_{\substack{j=1 \\ j \neq i}}^{n} \tilde{\theta}_{ij}^{g}(H)}{\sum_{j,i=1}^{n} \tilde{\theta}_{ij}^{g}(H)} \cdot 100 = \frac{\sum_{\substack{j=1 \\ j \neq i}}^{n} \tilde{\theta}_{ij}^{g}(H)}{n} \cdot 100 \tag{3.7}$$

Moreover, we also make use of directional spillovers indexes (DSI) to measure, respectively through equations (3.8) and (3.9), the spillover from exchange i to all other exchanges J (cfr. Eq. 3.8) and the spillover from all exchanges J to exchange i (cfr. Eq. 3.9) as:

$$DSI_{J\leftarrow i}(H) = \frac{\sum_{\substack{j=1 \\ j\neq i}}^{n} \tilde{\theta}^{g}_{ji}(H)}{\sum_{j,i=1}^{n} \tilde{\theta}^{g}_{ij}(H)} \cdot 100 \tag{3.8}$$

$$DSI_{i\leftarrow J}(H) = \frac{\sum_{\substack{j=1 \\ j\neq i}}^{n} \tilde{\theta}^{g}_{ij}(H)}{\sum_{j,i=1}^{n} \tilde{\theta}^{g}_{ij}(H)} \cdot 100 \tag{3.9}$$

Directional spillovers may be conceived as providing a decomposition of total spillovers into those coming from – or to – a particular variable. In other words, they measure the fraction of forecast error variance which comes from (or to) one of the variables included in the system - and, hence, the importance of the variable itself in forecasting the others. From the definitions of directional spillover indexes, it is natural to build a net contribution measure, impounded in the net spillover index (NSI) from market i to all other markets J, namely:

$$NSI_i(H) = DSI_{J\leftarrow i}(H) - DSI_{i\leftarrow J}(H) \tag{3.10}$$

Another very important metric to measure the difference between the gross shocks transmitted from market i to j and gross shocks transmitted from j to i is the net pairwise spillover (NPS), defined as:

$$NPS_{ij}(H) = \left(\frac{\tilde{\theta}^{g}_{ij}(H)}{\sum_{q=1}^{n} \tilde{\theta}^{g}_{iq}(H)} - \frac{\tilde{\theta}^{g}_{ji}(H)}{\sum_{q=1}^{n} \tilde{\theta}^{g}_{jq}(H)} \right) \cdot 100 \tag{3.11}$$

All the metrics discussed above are able to yield insights regarding the mechanisms of market exchange spillovers both from a system-wide and a net pairwise point of view. Furthermore, performing the analyses on rolling windows we are able to study the dynamics of spillover indexes over time.

3.4 Data and empirical findings

3.4.1 Data

To test our proposal we make use of historical data to perform a retrospective analysis. We use daily foreign exchange rate data spanning the period January 2002 - November 2019. We then extend the analysis until December 2020, to take into account variation related to the Covid-19 crisis period. To build our optimal basket of currencies, we collect data on the foreign exchange pairs between the currencies that are included in the IMF's Special Drawings Rights: the US dollar, the Chinese Renminbi, the Euro, the British pound and the Japanese Yen. According to our research assumption, we will assume that the obtained basket of currencies correspond to a stablecoin which can be exchanged and compared with a single currency based stablecoin. To understand the relationship between major currencies and remittances we also collect data on the largest remittance markets - namely, the Indian Rupee, Mexican Peso, Philippines Peso, Nigerian Naira. Moreover, for what concerns the volatility analysis, we divide the sample into subsets which define the pre-crisis period (2002-2008), crisis period (2009-2011) and post-crisis period (2012-2019). Finally, for the sake of comparison with a widely known basket-based currency such as the IMF SDR, we also collect data relative to the foreign exchange pair of the dollar with the IMF Special Drawing Rights.

3.4.2 Optimal basket and stability analysis

We present the results of the optimisation problem presented in Section 3.3 in Table 3.1. The weights of the IMF's SDR are also given for comparison. From the table, note that our method yields weights which are spread relatively equal across the currencies, in fact each are approximately a fifth, with a slightly heavier weighting on the EUR. The weights are quite different from the weights of the IMF SDR, which tend to be more concentrated on the USD dollar.

Currency	USD	CNY	EUR	GBP	JPY
Optimal Weights	0.17	0.2	0.23	0.19	0.21
IMF SDR Weights	0.42	0.11	0.31	0.08	0.08

Table 3.1: Optimal weights of the currency basket versus SDR weights

The IMF's precise methodology for determining weights remains undisclosed. The weighting scheme incorporates information from other economic variables such as trade flows. The most attractive feature of our weighting methodology is that it remains apolitical in nature rather than adopting the arbitrary inclusion of economic variables that become a matter of contention. The valuation of the SDR in terms of a currency basket had been among the most controversial decisions in the IMF history Mandeng (2019). This could certainly be of appeal to a private company considering issuing such an asset, as it keeps the methodology a purely statistical debate rather than a political one (Pontines, 2009). Furthermore, it also allows the computation of a unique optimal, minimum variance currency basket regardless of base currency choice and can be extended to any number of currencies. The rationale behind the SDR is to minimise transaction costs in international exchanges yet remaining independent on the monetary or exchange rate policy objectives of any single country. Many foreign-exchange transactions, even excluding US residents, are denominated and settled in dollars. Producers typically invoice their products in dollars, which keeps

their prices in line with their competitors and simplifies cross-border price comparisons among producers (Gopinath et al., 2016). This status quo tends to lend itself to a currency that has less variation against the dollar and consequently a weighting that favours the dollar such that there are stronger correlations with dollar movements. On the other hand Libra is focused on retail consumers, in particular those dependent on remittances. During bouts of depreciation, for example, in 2010, when the Fed embarked on QE2, the dollar hit all-time lows against several currencies. In this instance a Libra coin with less emphasis on the dollar could have sheltered holders from the depreciation of the dollar.

As noted above, fluctuations of SDRs will strongly be correlated with fluctuations in USD and EUR. Our proposed stablecoin (SAC) distributes the weights more evenly across the basket to minimise the variations in fluctuations. Since the basket is comprised of hard currencies the diversification tends to work, as the currencies move systematically over time relative to one another. In other words, if the value a particular currency depreciates relative to the SAC, but simultaneously there is an appreciation of another currency, their movements would balance each other, all else the same. Note that China has managed a floating peg against the USD and hence these two currencies are likely to be strongly linked. In the SDR these two currencies make up 53% of the basket compared to 37% in the SAC, arguably indicating more diversification is needed to offset movements in the dollar. To better interpret the results, Figure 1 represents the time series of the RNVALs of all considered currencies in the basket, along with our basket based stablecoin, in the considered period.

From Figure 1 note that, after an initial period of turbulence, the time series start to diverge roughly from the beginning of 2006 onwards. From that point in time, two clusters seem to emerge from the graph: the first one includes USD and CNY, while the second one pertains EUR, GBP and JPY. This is arguably due to the fact that, for many years, the CNY value was pegged to the dollar

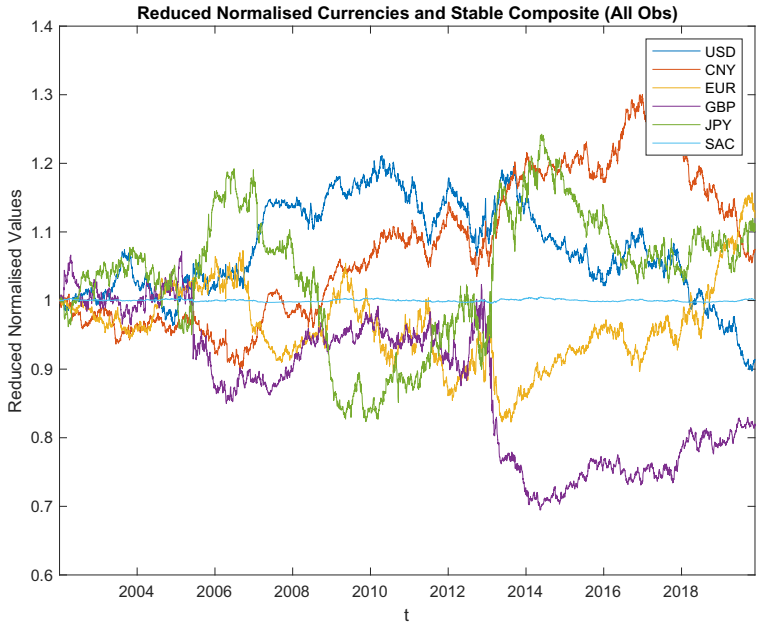

Figure 3.1: RNVALs of the basket currencies

Notes: The figure shows the time evolution of the RNVALs of the basket currencies over the sample period.

and, therefore, its dynamics over time shows quite similar patterns to that of the USD. As expected by construction, the RNVAL of the basket based stablecoin lies in the middle, "mediating" between the different currencies, and compensating single deviations with diversification benefits.

To understand the dynamics between currency and whether the basket based currency is more stable, Table 3.2 presents their volatilities, measured by their standard deviations, over the considered time period. The table presents also the correlations between the currencies, which help in the interpretation of the results. Table 3.2 shows, as far as correlations are concerned, that USD and CNY exhibit relatively strong negative or little correlation with others currencies in the basket, but are weakly positive between themselves, consistently with what observed in Figure 1. Moreover, one can clearly notice that the EUR acts as a good diversifier, as its pairwise correlations are quite low if compared to those between other currencies. More importantly, from the correlation matrix we can deduce that the stablecoin shows correlations with the other currencies whose values are

very close to zero. Low correlations with the other currencies is a clear sign of the ability of our stablecoin to remain orthogonal to the other fiat currencies' dynamics and, therefore, arguably stable. In terms of variability, the standard deviations show that the most volatile currency is CNY, followed by JPY and USD. Our stablecoin exhibits a standard deviation magnitude which is much lower than those of the other currencies and about ten times lower than that of the least volatile one, namely EUR. This is a clear sign of stability of the proposed stablecoin, as opposed to an hypothetical stablecoin pegged to one single currency.

	USD	CNY	EUR	GBP	JPY	SAC
USD	1	0.14	-0.68	0.01	-0.41	0.02
CNY	0.14	1	-0.4	-0.8	0.17	0.02
EUR	-0.68	-0.4	1	0.26	-0.09	0.03
GBP	0.01	-0.8	0.26	1	-0.64	0.02
JPY	-0.41	0.17	-0.09	-0.64	1	0.02
SAC	0.02	0.02	0.03	0.02	0.02	1
σ	0.07	0.1	0.06	0.1	0.09	0.002

Notes: The table presents the correlations between the RNVALs of the basket currencies based on the time series depicted in Figure 1. The bottom row gives the deviation measured over the whole period.

Table 3.2: Volatility and correlation of the RNVALs

To analyse the world's emerging market currencies with the highest portions of remittances, we first reconstruct a historical exchange rate for the SAC. We then recompute the RNVALs extending the basket to include the SAC, SDR and the emerging market currencies. The corresponding graphical representation is contained in Figure 3.2. Figure 3.2 shows that emerging market currencies such as the Mexican Peso (MXN) and the Nigerian Naira (NGN) have appreciated consistently with respect to the other fiat currencies in the basket over time. All the other currencies, instead, seem to belong to another cluster, in the sense that they do not follow an upward trend as the previous ones, but rather fluctuate below the value of 1, with different patterns. The only exception

is the Indian rupee (INR), whose value grows over time, although not with the same magnitude as MXN and NGN do. Note that both basket based stablecoins lie in the middle, similarly as in Figure 1, although their Reduced Normalised values fluctuate. This because the baskets are built using only five currencies, but are normalised with respect to all nine included in Figure 3.2.

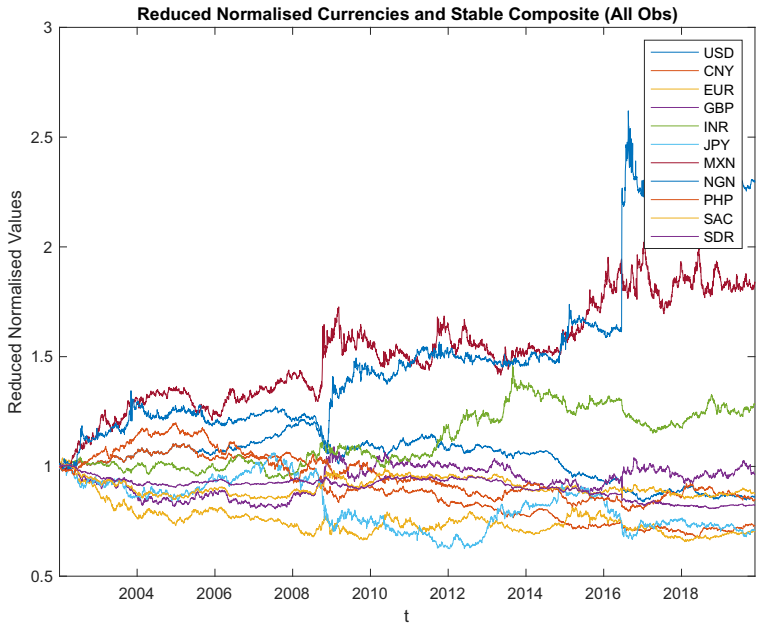

Figure 3.2: RNVALs of the basket and largest remittance currencies

Notes: The figure shows the time evolution of the RNVALs of the basket currencies including largest remittance currencies of the basket currencies over the sample period.

To determine whether a basket-based stablecoin would be a more valuable and more stable alternative than a stablecoin pegged to a single currency, especially for remittances, we can, in analogy with 3.2, compare the volatility of our stablecoin with that of a SDR based basket, and with the currencies of the most important emerging markets in terms of remittances. Table 3.3 contains the comparison, in terms of standard deviations, over the whole period and also in three distinct periods, corresponding to the pre-crisis period, the crisis period and the post-crisis period. From the top row of Table 3.3, it is clear that overall the stablecoin exhibits lower values of volatility,

when compared to the other traditional fiat currencies. The other rows in the table show that this is often the case, although especially during pre-crisis and crisis period few currencies exhibit slightly lower volatilities, depending on how and when they were affected by the global financial crisis. However, we can notice that the stablecoin's volatility is much lower than that of the other currencies which, although for some period slightly lower, show quite relevant jumps in magnitude. Moreover, the proposed stablecoin exhibit lower volatilities over the whole time period if compared to the single currencies in the basket and to the single emerging market currencies. This can be read as a strength of our stablecoin, as it could function as a better medium of exchange than a country's single currency, in particular as far as remittances are concerned. Note also that the SDR appears to be equally stable and is a valid alternative to our stablecoin.

	USD	CNY	EUR	GBP	INR	JPY	MXN	NGN	PHP	SAC	SDR
σ_{all}	0.09	0.14	0.07	0.06	0.13	0.11	0.23	0.41	0.10	0.04	0.05
σ_{pre}	0.04	0.03	0.08	0.05	0.02	0.05	0.10	0.07	0.06	0.04	0.02
σ_{cri}	0.05	0.05	0.03	0.06	0.03	0.10	0.09	0.10	0.03	0.03	0.02
σ_{post}	0.1	0.06	0.03	0.04	0.08	0.07	0.16	0.39	0.04	0.03	0.05

Notes: The table presents the standard deviations of the RNVALs of the basket currencies, of the emerging market currencies, our stablecoin and the SDR. The top row gives the deviation measured over the whole period (all), the second row measured during the pre-crisis period (pre), the third during the crisis period (cri) and the fourth during the post-crisis period (post).

Table 3.3: Measuring the volatility of the RNVALs

3.4.3 Spillover network analysis

We now consider spillovers between foreign exchange rates to evaluate the connectedness of the currencies composing the basket, and to understand which is the relative importance of each of the currencies in transmitting shocks. In this way, we are also able to determine which currencies potentially cause strong (or weak) changes in our proposed stablecoin value.

As far as specifications are concerned, VAR models are built on changes in reduced normalised values (RNVALs). We use a VAR lag determined by a Bayes-Schwarz information criterion (BIC) that penalises over-parameterisation compared to other widely employed information criteria. We use a $H = 100$ step-ahead forecast horizons for forward iteration of the system. Additionally, dynamic spillovers use a rolling estimation window of length 100 observations. Firstly, we provide an analysis of unconditional spillovers, that are spillovers evaluated on the whole sample period. The results are shown in Table 3.4.

	USD	CNY	EUR	GBP	JPY	FROM
USD	44.94	35.33	13.02	6.67	0.04	11.01
CNY	34.49	49.40	10.76	5.34	0.00	10.12
EUR	15.81	15.22	62.29	6.48	0.19	7.54
GBP	11.4	10.21	6.28	69.58	2.53	6.08
JPY	0.41	0.14	0.01	3.94	95.51	0.90
TO	12.42	12.18	6.01	4.49	0.55	35.66

Notes: The table presents the volatility spillovers between currencies. Reading the table horizontally shows the spillovers each currency receives from others, reading vertically shows the spillovers that each currency gives to the others.

Table 3.4: Currency spillovers

From Table 3.4 note that the USD and CNY are highly interconnected with the others, whereas EUR, GBP and in particular JPY are more isolated in terms of connectedness. USD and CNY are significantly dominant, and their contributions in terms of spillover towards other currencies are much higher than those of the remaining currencies in the basket.

The analysis of dynamic spillovers is able to clarify the results obtained in the unconditional spillover analysis by means of observing the evolution of spillovers over time. Figure 3.3 depicts the overall dynamic spillover plotted over the sample period. The overall spillover within the basket ranges from a minimum of 17.87% to a maximum of 80.00%. It seems that the overall spillover follows a generally decreasing trend, as it starts from 54.51% at the beginning of the

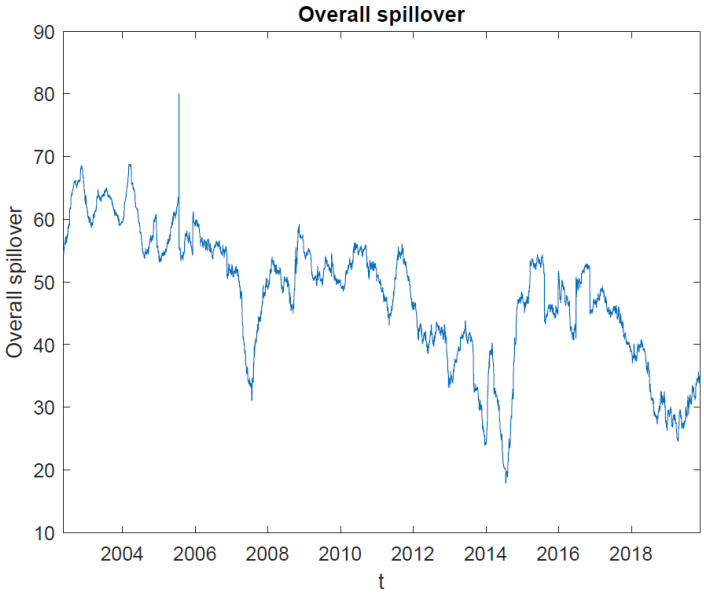

Figure 3.3: Overall spillovers

Notes: The figure represents the dynamic overall spillover index of the basket currencies over the sample period.

sample period, while it diminishes to 34.43% at the end of the studied time frame.

Dynamic directional spillovers can shed light on which of the currencies transmit price change spillovers to others and which of them receive spillovers from others. We plot directional from, to and net spillovers in Figures 3.4, 3.5 and 3.6, respectively.

From the joint analysis of Figures 3.4, 3.5 and 3.6 we can observe that USD is the most influential currency in terms of spillovers. Indeed, the magnitude of spillovers received from others is weak compared to that transmitted to others. Moreover, the net spillover dynamics summarises the dominant position of the USD, being it always positive and taking relatively high values over the sample period. However, the magnitude of spillovers transmitted by USD follows a decreasing trend over time, indicating the currency is gradually losing its potential to contribute to the evolution of the others, perhaps due to the affirmation of emerging economies in the latter period, especially after the 2009 crisis. CNY is indeed emerging as a dominant currency during the recent

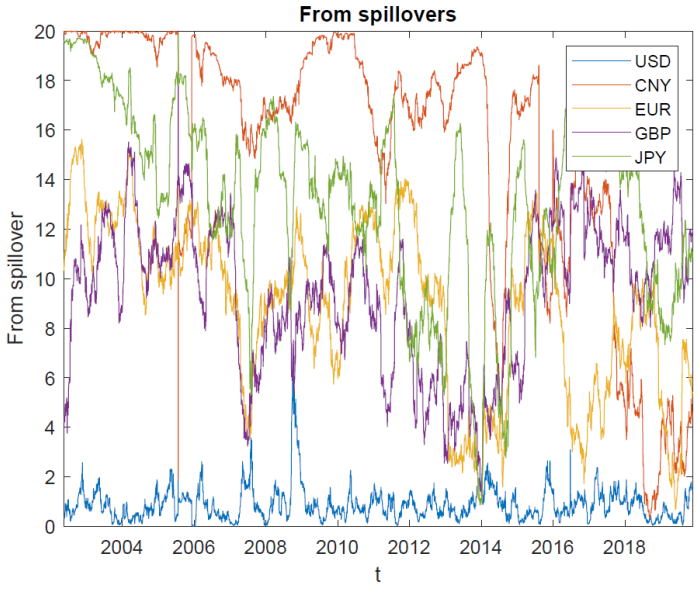

Figure 3.4: From spillovers

Notes: The figure represents the dynamic directional "From" spillovers of the basket currencies over the sample period.

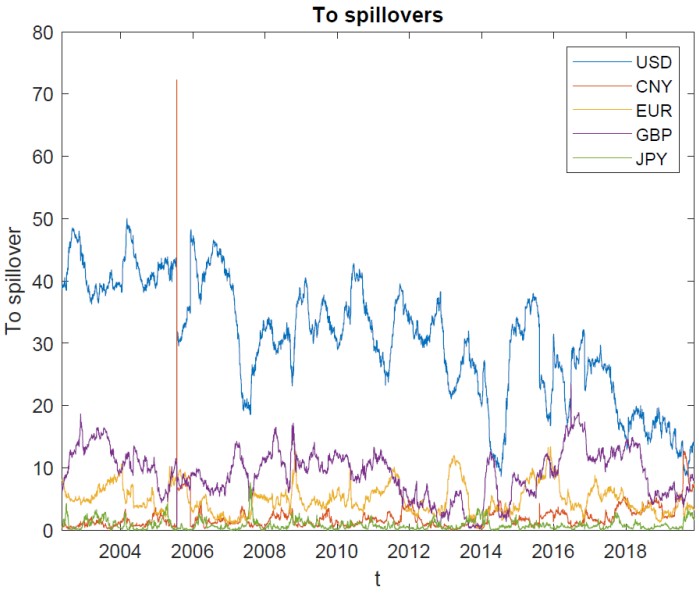

Figure 3.5: To spillovers

Notes: The figure represents the dynamic directional "To" spillovers of the basket currencies over the sample period.

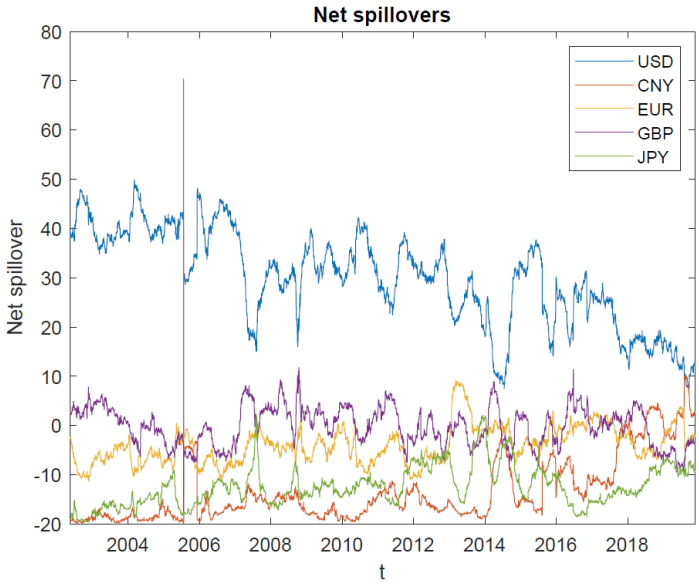

Figure 3.6: Net spillovers

Notes: The figure represents the dynamic directional net spillovers of the basket currencies over the sample period.

times. Despite that, the latter considerations are in line with the full sample results obtained above, which point to the dominance of USD as a spillover transmitting currency.

Differently from what emerged in the full sample analysis, instead, the dynamic analysis shows that CNY has not been such a leading currency in transmitting price change shocks from an historical viewpoint. Indeed, the full sample result is arguably driven by a noticeable spike which occurred on 21 July 2005. Indeed, during that day the Chinese Central Bank officially announced the abandonment of the eleven-year-old peg to the dollar and pegged the CNY to a basket of currencies whose composition was not disclosed. This caused a prompt revaluation to CNY 8.11 per USD, as well as to 10.07 CNY per euro. However, the peg to the dollar was reinstated as the financial crisis strengthened in July 2008. These results indicate that CNY does not particularly contribute to the price change evolution of the other currencies in the basket, although it can exert shocks through sudden policy decisions.

The dominance of the USD and, to a lesser extent, of CNY emerges also when analysing

the directed network structure of the currencies in terms of net pairwise spillovers represented in Figure 3.7. In this context, the network edges are represented by the currencies in the basket, whereas links represent the magnitude of net pairwise spillovers for each currency pair.

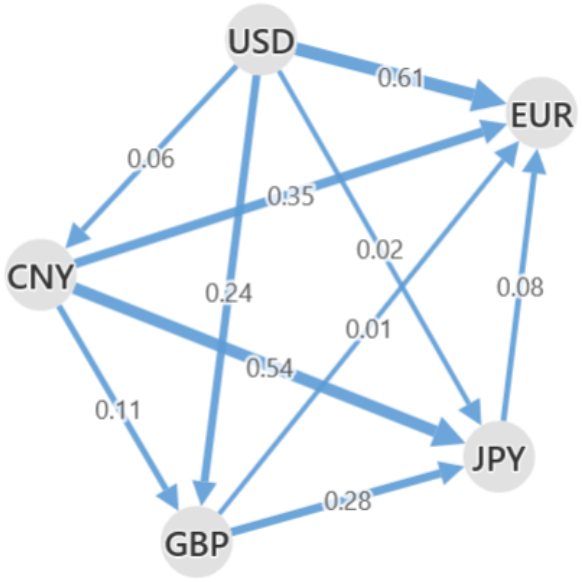

Figure 3.7: Spillover network (full sample)

Notes: The figure represents the spillover network of the currencies included in the basket over the full sample period. The nodes are represented by the currencies included in the basket. The magnitude of the links is represented by the net pairwise spillovers between each currency pair.

To verify the loss of importance of the dollar, we extend the spillover network analysis to cover the Covid-19 crisis period. Specifically, we analyse two subsamples with the year 2020 as the cutoff point, to detect major changes in country forex spillover dynamics.

The spillover network gives a ranking in terms of spillover transmission capacity and, therefore, price discovery. The most influential currency in terms of price change shock transmission is USD, followed by CNY and, to a lesser extent, GBP. The two receivers are instead JPY and, at most, EUR. The highest influence is given by USD towards EUR, followed by CNY to JPY. This

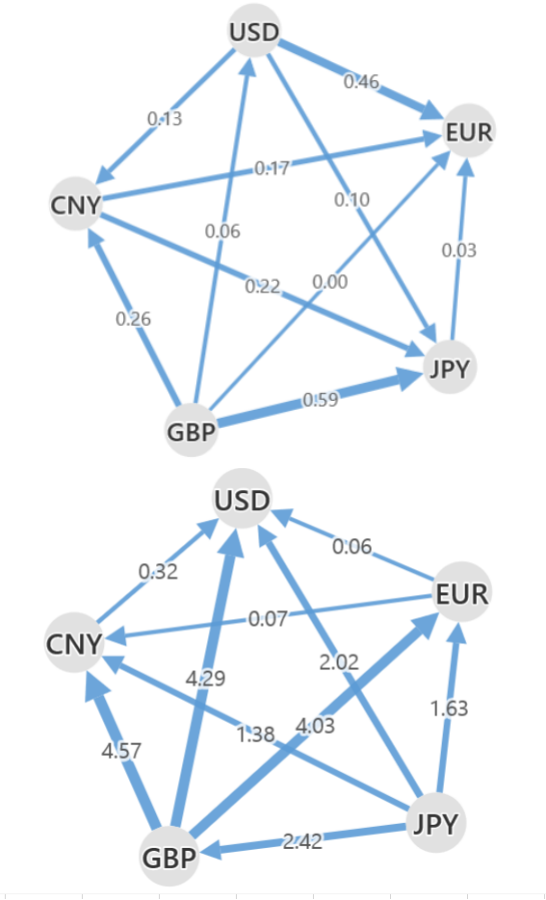

Figure 3.8: Spillover network (sub-samples)

Notes: The figure represents the spillover network of the currencies included in the basket. The first representation corresponds to the period April 2017 to December 2019 (left panel), while the second one from January 2020 to December 2020 (right panel). The nodes are represented by the currencies included in the basket. The magnitude of the links is represented by the net pairwise spillovers between each currency pair.

suggests that the contagion occurs to a greater extent within Asian currencies and across American and European ones.

However, the picture is different when analysing spillovers during two distinct sub-samples: one ranging from April 2017 to December 2019, and another one from January 2020 to December 2020, both depicted in Figure 3.8. Indeed, overall interconnectedness has increased in the second sub-sample, likely due to the Covid-19 outbreak; thus markets move more similarly as a conse-

quence of the epidemic. This is equivalent to say there are more contagion dynamics occurring since the Covid-19 outbreak and that the magnitude of information transmitted from the currencies sharply increased after the Covid-19 crisis. Moreover, empirical outcomes highlight that the contagion dynamics starkly different from that of the pre-crisis period. USD is no longer dominant, and it becomes a spillover receiver during the pandemic period. Currencies which before were mostly receivers, instead, started becoming transmitter of shocks, such as JPY, GBP and, to a lesser extent, EUR. This suggests that the Covid-19 pandemic, while inducing a more coordinated response of the forex market to shocks, has literally disrupted the equilibria existing prior the crisis. It also highlights the importance of monitoring the spillover dynamics in the basket both to have a systemic risk indicator and to determine lead-lag relationships among currencies in the basket when designing basket-based stablecoins.

3.5 Conclusion

In the paper we present a methodology to build a basket based stablecoin whose weights can maximise stability over a long time period. The weights have been calculated, retrospectively, from 2002 to 2019, and are distributed more evenly across the currency basket than those suggested by the SDR.

The proposed stablecoin (Librae) appears to be less volatile than single currencies and, therefore, with respect to single currency stablecoins (Libra). It can thus constitute a valuable proposal especially for workers who live abroad and make remittances to their own country, a market segment with a high potential of being attracted by payments in stablecoins.

We have also proposed a variance decomposition technique based on a VAR model aimed at showing which currencies mostly impact the Foreign Exchange market and whether a single currency or a basket based stablecoin is more resilient to currency shocks. Our results show that

the dollar is the currency which mostly impact the market, and that a basket based coin is better than a dollar based one, from a stability and value maintenance viewpoint. However, CNY is taking over as spillover transmitter and USD is gradually losing its influence over time, especially with regards to the latest period, characterised by the Covid-19 otubreak.

With a basket based stablecoin it is possible to offset the risk of currencies shocks. This is of relevance for different policy purposes and, in particular, for emerging markets and countries having high remittances. Indeed, by holding stablecoins rather than single currencies the risks associated to currency shocks are mitigated and stablecoin holders can count on a currency whose value is less volatile than traditional fiat currencies and, thereby, more reliable. The latter fact has also positive consequences on the cross-border payments side, provided that the stability of the stablecoin mitigates the foreign exchange risk, thus contributing to the fact that buyers and sellers give or receive an amount of money whose value is less sensitive to variations over time.

Future research may consider basket that dynamically evolve over time, although these are bound to be more difficult to achieve consensus. Furthermore, currency volumes in circulation may be taken to account, along with the technical characteristics of the coins (for example: cybersecurity, redeemability, reliability), from a different, more theoretical, viewpoint. Future works might also consider a basket composed of different currencies: for example, without the Chinese Renminbi, and in line with the developments of the Diem coin, recently announced by the Libra foundation.

BIBLIOGRAPHY

Abdymomunov A, Curti F, Mihov A. 2017. US banking sector operational losses and the macroeconomic environment. Technical report, Available at SSRN 2738485.

Abiad A, Detragiache E, Tressel T. 2010. A New Database of Financial Reforms. *IMF Staff Papers* **57**: 281–302.

Abosedra S, Arayssi M, Sita BB, Mutshinda C. 2020. Exploring gdp growth volatility spillovers across countries. *Economic modelling* **89**: 577–589.

Acosta PA, Baerg NR, Mandelman FS. 2009. Financial development, remittances, and real exchange rate appreciation. *Economic Review* **94**.

Admati AR. 2016. The missed opportunity and challenge of capital regulation. *National Institute Economic Review* **235**: R4–R14.

Aldasoro I, Gambacorta L, Frost J, Whyte D. 2021. Covid-19 and cyber risk in the financial sector. *BIS Bulletin* Forthcoming.

Aldasoro I, Gambacorta L, Giudici P, Leach T. 2020. Operational risks in the financial sector. BIS Working Papers 840, Bank for International Settlements.

Aldasoro I, Gambacorta L, Giudici P, Leach T. 2022. The drivers of cyber risk. *Journal of Financial Stability* Forthcoming.

Alexander C. 2008. *Statistical Models of Operational Loss*, chapter 11. Handbook of Finance.

Allen F, Gale D. 2004. Competition and Financial Stability. *Journal of Money, Credit and Banking* **36**: 453–480.

Allen L, Bali TG. 2007. Cyclicality in catastrophic and operational risk measurements. *Journal of Banking & Finance* **31**: 1191–1235.

Altunbaş Y, Gambacorta L, Marques-Ibanez D. 2014. Does monetary policy affect bank risk-taking? *International Journal of Central Banking* **10**: 95–135.

Altunbaş Y, Thornton J, Uymaz Y. 2018. CEO tenure and corporate misconduct: Evidence from US banks. *Finance Research Letters* **26**: 1–8.

Ames M, Schuermann T, Scott HS. 2015. Bank capital for operational risk: A tale of fragility and instability. *Journal of Risk Management in Financial Institutions* **8**: 227–243.

Anderson R, Barton C, Böhme R, Clayton R, Hernandez Ganan C, Grasso T, Levi M, Moore T, Vasek M. 2019. Measuring the changing cost of cybercrime. In *The 2019 Workshop on the Economics of Information Security (WEIS 2019)*.

Antonini G, Cope EW, Mignola G, Ugoccioni R. 2009. Challenges in measuring operational risk from loss data. *Journal of Operational Risk* **4**: 3–27.

Artzner P, Delbaen F, Eber JM, Heath D. 1999. Coherent measures of risk. *Mathematical finance* **9**: 203–228.

Baldwin A, Gheyas I, Ioannidis C, Pym D, Williams J. 2017. Contagion in cyber security attacks. *Journal of the Operational Research Society* **68**: 780–791.

Bank for International Settlements. 1996. Implications for central banks of the development of electronic money.

Barajas A, Chami R, Hakura D, Montiel P, Tressel T. 2011. Workers' remittances and the equilibrium real exchange rate: Theory and evidence [with comment]. *Economía* **11**: 45–99. ISSN 15297470.
URL http://www.jstor.org/stable/41343450

Basel Committee on Banking Supervision. 2003. Basel II: The new Basel capital accord. Consultation, Basel, Switzerland.

Basel Committee on Banking Supervision. 2017. High-level summary of Basel III reforms. Technical report, Bank for International Settlements.

Basel Committee on Banking Supervision. 2018a. Cyber resilience: Range of practices. Technical report, Bank for International Settlements.

Basel Committee on Banking Supervision. 2018b. Cyber-resilience: range of practices. Other, Basel, Switzerland.

Basel Committee on Banking Supervision. 2018c. Standardised measurement approach for operational risk. Consultative document, Basel, Switzerland.

Berger AN, Curti F, Mihov A, Sedunov J. 2018. Operational risk is more systemic than you think: Evidence from US bank holding companies. *Available at SSRN 3210808* .

Biell L, Muller A. 2013. Sudden crash or long torture: the timing of market reactions to operational loss events. *Journal of Banking & Finance* **37**: 2628–2638.

Biener C, Eling M, Wirfs JH. 2015. Insurability of cyber risk: An empirical analysis. *The Geneva Papers on Risk and Insurance-Issues and Practice* **40**: 131–158.

Boer M, Vazquez J. 2017. Cyber security and financial stability: How cyber-attacks could materially impact the global financial system. Technical report, Institute of International finance.

Bogdanova B, Hofmann B. 2012. Taylor rules and monetary policy: a global 'great deviation'? *BIS quarterly review, September* .

Boone J. 2008. A new way to measure competition. *The Economic Journal* **118**: 1245–1261.

Bouri E, Lau CKM, Lucey B, Roubaud D. 2019. Trading volume and the predictability of return and volatility in the cryptocurrency market. *Finance Research Letters* **29**: 340–346.

Bouveret A. 2018. Cyber risk for the financial sector: A framework for quantitative assessment. IMF Working Paper WP/18/143, International Monetary Fund.

Brandvold M, Molnár P, Vagstad K, Valstad OCA. 2015. Price discovery on bitcoin exchanges. *Journal of International Financial Markets, Institutions and Money* **36**: 18–35.

Brenner J. 2017. Keeping America safe: Toward more secure networks for critical sectors. Report on a series of mit workshops, MIT Internet Policy Research Initiative.

Bullmann D, Klemm J, Pinna A, et al. 2019. In search for stability in crypto-assets: are stablecoins the solution? Technical report, European Central Bank.

Byrne B, Coughlan J, Tilley SV. 2017. An empirical analysis of the impact of fines on bank reputation in the US and UK. *Available at SSRN 2980352* .

Cabras S, Castellanos ME. 2007. A default bayesian procedure for the generalized pareto distribution. *Journal of Statistical Planning and Inference* **137**: 473–483.

Cameron AC, Miller DL. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources* **50**: 317–372.

Cameron AC, Trivedi PK. 2005. *Microeconometrics: methods and applications*. Cambridge university press.

Carney M. 2019. The growing challenges for monetary policy in the current international monetary and financial system. In *Remarks at the Jackson Hole Symposium*.

Carr B, Pujazon D, Vazquez J. 2019. Cloud service providers and criticality: Potential treatments and solutions. Technical report, Institute of International Finance.

Carrivick L, Cope EW. 2013. Effects of the financial crisis on banking operational losses. *The Journal of Operational Risk* **8**: 3.

Catteddu D. 2009. Cloud computing: benefits, risks and recommendations for information security. In *Iberic Web Application Security Conference*. Springer, 17–17.

Cerasi V, Deininger S, Gambacorta L, Oliviero T. 2020. How post-crisis regulation has affected bank ceo compensation. *Journal of International Money and Finance,* Forthcoming.

Chande N, Yanchus D. 2019. The cyber incident landscape. Working Paper 32, Bank of Canada.

Chavez-Demoulin V, Embrechts P, Nešlehová J. 2006. Quantitative models for operational risk: extremes, dependence and aggregation. *Journal of Banking & Finance* **30**: 2635–2658.

Chen W, Srinivasan S. 2019. Going digital: Implications for firm value and performance. Working Paper 19-117, Harvard Business School.

Cheng J, Dai Y. 2020. Is bitcoin a channel of capital inflow? evidence from carry trade activity. *International Review of Economics & Finance* **66**: 261–278.

Chernobai A, Jorion P, Yu F. 2011. The determinants of operational risk in US financial institutions. *Journal of Financial and Quantitative Analysis* **46**: 1683–1725.

Chernozhukov V, Umantsev L. 2001. Conditional value-at-risk: Aspects of modeling and estimation. *Empirical Economics* **26**: 271–292.

Čihák M, Schaeck K. 2010. Competition, efficiency, and soundness in banking: An industrial organization perspective. *European Banking Centre Discussion Paper* .

Cohen RD, Humphries J, Veau S, Francis R. 2019. An investigation of cyber loss data and its links to operational risk. *Journal of Operational Risk* **14**: 1–25.

Cope EW, Piche MT, Walter JS. 2012. Macroenvironmental determinants of operational loss severity. *Journal of Banking & Finance* **36**: 1362–1380.

Corbet S, Meegan A, Larkin C, Lucey B, Yarovaya L. 2018. Exploring the dynamic relationships between cryptocurrencies and other financial assets. *Economics Letters* **165**: 28–34.

Cornalba C, Giudici P. 2004. Statistical models for operational risk management. *Physica A: Statistical Mechanics and its applications* **338**: 166–172.

Cox DR. 1972. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)* **34**: 187–202.

Crosignani M, Macchiavelli M, Silva AF. 2020. Pirates without borders: The propagation of cyberattacks through firms' supply chains. Staff Report 937, Federal Reserve Bank of New York.

Cruz MG, Peters GW, Shevchenko PV. 2015. *Fundamental aspects of operational risk and insurance analytics: A handbook of operational risk*. John Wiley & Sons.

Cummins JD, Lewis CM, Wei R. 2006. The market value impact of operational loss events for US banks and insurers. *Journal of Banking & Finance* **30**: 2605–2634.

Curti F, Gerlach J, Kazinnik S, Lee M, Mihov A. 2019a. Cyber risk definition and classification for financial risk management. *Federal Reserve Bank of St Louis, August, mimeo* .

Curti F, Mihov A, Frame WS. 2019b. Are the largest banking organisations operationally more risky?
URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3210206

Dalla Valle L, Giudici P. 2008. A bayesian approach to estimate the marginal loss distributions in operational risk management. *Computational Statistics & Data Analysis* **52**: 3107–3127.

Danielsson J, Macrae R. 2019. Systemic consequences of outsourcing to the cloud. VoxEU, CEPR.

De Nicolò G, Lucchetta M. 2013. Bank Competition and Financial Stability: A General Equilibrium Exposition. CESifo Working Paper Series 4123, CESifo Group Munich.

Denk O, Gomes G. 2017. Financial re-regulation since the global crisis? OECD Economics Department Working Papers 1396, Organisation for Economic Co-operation and Development.

Diebold FX, Yilmaz K. 2012. Better to give than to receive: Predictive directional measurement of volatility spillovers. *International Journal of Forecasting* **28**: 57–66.

Diebold FX, Yılmaz K. 2014. On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics* **182**: 119–134.

Dimpfl T, Peter FJ. 2020. Nothing but noise? price discovery across cryptocurrency exchanges. *Journal of Financial Markets* : 100584.

Dingel JI, Neiman B. 2020. How many jobs can be done at home? Working Paper 26948, National Bureau of Economic Research.
URL http://www.nber.org/papers/w26948

Duffie D, Younger J. 2019. Cyber runs. Hutchins Center Working Paper 51, Brookings Institution.

Dwyer GP. 2015. The economics of bitcoin and similar private digital currencies. *Journal of Financial Stability* **17**: 81–91.

Eisenbach TM, Kovner A, Lee MJ. 2021. Cyber risk and the U.S. financial system: A pre-mortem analysis. *Journal of Financial Economics* ISSN 0304-405X.
URL https://www.sciencedirect.com/science/article/pii/S0304405X21004578

Eshraghi A, Hagendorff J, Nguyen DD. 2015. Can bank boards prevent misconduct? *Review of Finance* **20**: 1–36.

Etro F. 2015. The economics of cloud computing. In *Cloud Technology: Concepts, Methodologies, Tools, and Applications*. IGI Global, 2135–2148.

Etzioni A. 2011. Cybersecurity in the private sector. *Issues in Science and Technology* **28**: 58–62.

European Central Bank. 1998. Report on electronic money.

European Systemic Risk Board. 2015. Report on misconduct risk in the banking sector. ESRB reports, European Systemic Risk Board, Frankfurt.

European Systemic Risk Board. 2020. Systemic cyber risk. ESRB reports, European Systemic Risk Board, Frankfurt.

Facchinetti S, Giudici P, Osmetti SA. 2019. Cyber risk measurement with ordinal data. *Statistical Methods & Applications* : 1–13.

Feng X, He X, Hu J. 2011. Wild bootstrap for quantile regression. *Biometrika* **98**: 995–999.

Fich EM, Shivdasani A. 2007. Financial fraud, director reputation, and shareholder wealth. *Journal of Financial Economics* **86**: 306–336.

Figini S, Gao L, Giudici P. 2015. Bayesian operational risk models. *Journal of Operational Risk* **10**.

Financial Stability Board. 2014. Guidance on supervisory interaction with financial institutions on risk culture: a framework for assessing risk culture. Compendium of standards, Financial Stability Board.

Financial Stability Board. 2019. Regulatory issues of stablecoins.

Financial Stability Board. 2019. Third-party dependencies in cloud services: Considerations on financial stability implications. Technical report, Financial Stability Board.

Florakis C, Louca C, Michaely R, Weber M. 2020. Cybersecurity risk. Working Paper 28196, National Bureau of Economic Research.
URL http://www.nber.org/papers/w28196

Flore M. 2018. How blockchain-based technology is disrupting migrants' remittances: a preliminary assessment. *Luxembourg, EUR* **29492**.

Gartner. 2021. Gartner forecasts worldwide security and risk management spending to exceed $150 billion in 2021.
URL https://www.gartner.com/en/newsroom/press-releases/2021-05-17-gartner-forecasts-worldwide-security-and-risk-managem

Gillet R, Hübner G, Plunus S. 2010. Operational risk and reputation in the financial industry. *Journal of Banking & Finance* **34**: 224–235.

Giudici P, Pagnottoni P. 2019. High frequency price change spillovers in bitcoin markets. *Risks* **7**: 111.

Giudici P, Pagnottoni P. 2020. Vector error correction models to measure connectedness of bitcoin exchange markets. *Applied Stochastic Models in Business and Industry* **36**: 95–109.

Giudici P, Pagnottoni P, Polinesi G. 2020. Network models to enhance automated cryptocurrency portfolio management. *Frontiers Artif. Intell.* **3**: 22.

Goldstein J, Chernobai A, Benaroch M. 2011. An event study analysis of the economic impact of IT operational risk and its subcategories. *Journal of the Association for Information Systems* **12**: 1.

Gopinath G, Boz E, Casas C, Díez FJ, Gourinchas PO, Plagborg-Møller M. 2016. Dominant currency paradigm. Working Paper 22943, National Bureau of Economic Research.
URL `http://www.nber.org/papers/w22943`

Gordon LA, Loeb MP. 2002. The economics of information security investment. *ACM Transactions on Information and System Security (TISSEC)* **5**: 438–457.

Hagemann A. 2017. Cluster-robust bootstrap inference in quantile regression models. *Journal of the American Statistical Association* **112**: 446–456.

Heid F. 2007. The cyclical effects of the basel ii capital requirements. *Journal of Banking & Finance* **31**: 3885–3900.

Hess C. 2011. The impact of the financial crisis on operational risk in the financial services industry: empirical evidence. *The Journal of Operational Risk* **6**: 23.

Hovanov NV, Kolari JW, Sokolov MV. 2004. Computing currency invariant indices with an application to minimum variance currency baskets. *Journal of Economic Dynamics and Control* **28**: 1481–1504.

Humpage O. 2009. Will special drawing rights supplant the dollar? https://voxeu.org/article/sdr-vs-what-it-takes-be-international-reserve-currency.

Jarrow RA. 2008. Operational risk. *Journal of Banking & Finance* **32**: 870–879.

Kamiya S, Kang JK, Kim J, Milidonis A, Stulz RM. 2018. What is the impact of successful cyberattacks on target firms? NBER Working Paper 24409.

Kamiya S, Kang JK, Kim J, Milidonis A, Stulz RM. 2021. Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics* **139**: 719–749. ISSN 0304-405X.
URL https://www.sciencedirect.com/science/article/pii/S0304405X20300143

Kashyap AK, Wetherilt A. 2019. Some principles for regulating cyber risk. In *AEA Papers and Proceedings*, volume 109. 482–87.

Katsiampa P, Corbet S, Lucey B. 2019. High frequency volatility co-movements in cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money* **62**: 35–52.

Kennedy DB, Stratopoulos TC. 2017. Mapping it spending across industry classifications: An open source dataset.
URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3073236

Kennedy PE. 1981. Estimation with correctly interpreted dummy variables in semilogarithmic equations. *The American Economic Review* **71**: 801–801. ISSN 00028282.
URL http://www.jstor.org/stable/1806207

Kim J. 2018. Bank competition and financial stability: Liquidity risk perspective. *Contemporary Economic Policy* **36**: 337–362.

Knobbe D. 2020. Network outages: Do they cost more than you think?
URL https://info.pivitglobal.com/blog/cost-of-downtime

Kopp E, Kaffenberger L, Jenkinson N. 2017. Cyber risk, market failures, and financial stability. IMF Working Paper WP/17/185, International Monetary Fund.

Köster H, Pelster M. 2017. Financial penalties and bank performance. *Journal of Banking & Finance* **79**: 57–73.

Levene T. 2006. 'Cashless society' card that flopped. *The Guardian* .
URL https://www.theguardian.com/money/2006/apr/15/moneysupplement1

Libra Association. 2020. Libra white paper.

Lim JH, Stratopoulos TC, Wirjanto TS. 2011. Path dependence of dynamic information technology capability: An empirical investigation. *Journal of Management Information Systems* **28**: 45–84.

Makridis C. 2020. Do data breaches damage reputation? evidence from 43 companies between 2002 and 2018.
URL https://ssrn.com/abstract=3596933

Makridis C, Dean B. 2018. Measuring the economic effects of data breaches on firm outcomes: Challenges and opportunities. *Journal of Economic and Social Measurement* **43**: 59–83.

Makridis C, Liu T. 2021. Abnormal returns and dispersion in cybersecurity exposure. URL https://ssrn.com/abstract=3746589

Mandeng OJ. 2019. The sdr – a blueprint for libra. Technical report, LSE Institute of Global Affairs.

Migueis M. 2018. Forward-looking and incentive-compatible operational risk capital framework. *Journal of Operational Risk* **13**.

Monetary Authority of Singapore. 2018. Financial stability review. Technical report, Monetary Authority of Singapore.

Murphy E, Murphy M, Seitzinger M. 2015. Bitcoin: Questions, answers, and analysis of legal issues. *Congressional Research Service* .

Nakamoto S. 2008. Bitcoin: A peer-to-peer electronic cash system. Accessed at: https://bitcoin.org/bitcoin.pdf.

Ocampo JA. 2019. Is it time for a 'true global currency'? https://www.weforum.org/agenda/2019/04/is-it-time-for-a-true-global-currency.

Pagnottoni P. 2019. Neural network models for bitcoin option pricing. *Frontiers in Artificial Intelligence* **2**: 5.

Pagnottoni P, Dimpfl T. 2019. Price discovery on bitcoin markets. *Digital Finance* **1**: 139–161.

Peters G, Shevchenko PV, Hassani B, Chapelle A. 2016. Should the advanced measurement approach be replaced with the standardized measurement approach for operational risk? *Journal of Operational Risk* **11**.

Pontines V. 2009. Optimal common currency basket in east asia. *Applied Economics Letters* **16**: 1139–1141.
URL https://doi.org/10.1080/13504850701335392

Power M. 2005. The invention of operational risk. *Review of International Political Economy* **12**: 577–599.

Resta M, Pagnottoni P, De Giuli ME. 2020. Technical analysis on the bitcoin market: Trading opportunities or investors' pitfall? *Risks* **8**: 44.

Romanosky S. 2016. Examining the costs and causes of cyber incidents. *Journal of Cybersecurity* **2**: 121–135.

Roner C, Caterina CD, Ferrari D. 2021. Exponential Tilting for Zero-inflated Interval Regression with Applications to Cyber Security Survey Data. BEMPS - Bozen Economics & Management Paper Series BEMPS85, Faculty of Economics and Management at the Free University of Bozen.
URL https://ideas.repec.org/p/bzn/wpaper/bemps85.html

Rowe B. 2007. Will outsourcing IT security lead to a higher social level of security? In *Workshop on Economics of Information Security*. Pittsburgh.

Rowe BR, Gallaher MP. 2006. Private sector cyber security investment strategies: An empirical analysis. In *The fifth workshop on the economics of information security*. The Workshop on the Economics of Information Security.

Sakalauskaite I. 2018. Bank risk-taking and misconduct. *University of Amsterdam and the Tinbergen Institute Working Paper* .

Sands P, Liao G, Ma Y. 2018. Rethinking operational risk capital requirements. *Journal of Financial Regulation* **4**: 1–34.

Segendorf B. 2014. What is bitcoin. *Sveriges Riksbank Economic Review* **2**: 71–87.

Shih J, Samad-Khan A, Medapa P. 2000. Is the size of an operational loss related to firm size. *Operational Risk* **2**: 21–22.

Sturm P. 2013. Operational and reputational risk in the European banking industry: The market reaction to operational risk events. *Journal of Economic Behavior & Organization* **85**: 191–206.

Tarullo DK. 2008. *Banking on Basel: the future of international financial regulation*. Peterson Institute.

Wolff J, Lehr W. 2017. Degrees of ignorance about the costs of data breaches: What policymakers can and can't do about the lack of good empirical data. Technical report, Available at SSRN 2943867.

APPENDIX A

THE DRIVERS OF CYBER RISK

A.1 Additional material and robustness checks

	Mean	Median	Std. dev.	Minimum	Maximum
Variables varying at individual event level					
Costs ($ mil)	8.72	0	58.6	0	5,000
connections	2.00	2.00	7.09	0	161.
Variables varying at firm level					
Firm size (Revenues $ mil)	6,910	27.7	31,000	0[a]	521,000
Variables varying at sector level					
Digital share of business activity	14.6	15.2	2.06	9.23	24.3
Cloud service purchases	18.3	20.1	5.04	5.60	26.2
Binary variables at event level					
Malicious	0.307	0	0.461	0	1.00
Security incident	0.0368	0	0.188	0	1.00
Data breach	0.400	0	0.490	0	1.00
Phishing/skimming	0.0366	0	0.188	0	1.00
Privacy violation	0.522	1.00	0.500	0	1.00
Other	0.00441	0	0.0663	0	1.00

Notes: [a] Zeros are a consequence of rounding accuracy. This table summarises all variables from the full set of US-based observations (120,184) for comparison with the sub-sample used for regressions. The top panel reports the variables from equation (1.1) that vary with each individual event in the sample. The second panel contains variables from equation (1.1) that vary by each firm contained in the sample. The third panel are the variables from equation (1.1) that vary at the sector level and obtained from the US census Bureau 2018 Annual Business Survey. The bottom panel are dummy variables that indicate the type of the incident.

Table A.1.1: Summary of variables from full sample

	N	Mean	Standard deviation	σ_k	ρ_k
Accommodation and Food Services	125	6,189,843	22,366,049	9.95	0.0469
Admin. and Support and Waste Management	462	9,658,941	75,249,945	9.74	0.00181
Agriculture, Forestry, Fishing and Hunting	2	1,355,000	1,902,117	-	-
Arts, Entertainment, and Recreation	31	4,891,540	8,626,132	17.2	0.155
Construction	48	819,930	2,300,230	8.65	-0.0171
Educational Services	113	1,145,089	3,600,769	6.52	-0.00700
Finance and Insurance	901	6,157,930	25,769,555	10.5	0.00704
Health Care and Social Assistance	273	2,346,022	13,610,937	5.55	0.0293
Information	553	14,464,877	214,374,866	11.0	0.0209
Management of Companies and Enterprises	11	515,806	1,556,415	26.4	0.0187
Manufacturing	148	15,823,066	63,595,183	10.3	0.0853
Mining, Quarrying, and Oil and Gas Extraction	4	2,112,400	3,289,039	-	-
Other Services (except Public Administration)	80	3,409,466	14,926,987	12.6	-0.0111
Professional, Scientific, and Technical Services	338	17,746,993	221,520,625	10.8	0.00518
Real Estate and Rental and Leasing	54	2,521,698	5,819,034	8.42	0.0759
Retail Trade	359	11,666,479	42,163,609	12.3	-0.00267
Transportation and Warehousing	54	29,170,690	84,187,668	11.9	0.0427
Utilities	39	1,512,738	2,405,099	16.9	0.00792
Wholesale Trade	110	38,134,792	268,239,144	10.7	0.00748
Total	3,705	10,398,854	121,742,478	-	-

Notes: The first column reports the number of incidents in each sector. The second and third columns show the mean and standard deviation of the dependent variable (the cost of a cyber incident in US$) prior to a log transformation. The fourth and fifth column denote the within sector variation and within sector correlation calculated using the residuals of the regression in column I of Table 1.2 according to the method of (Cameron and Trivedi, 2005, p. 835). Missing entries could not be calculated due to a lack of observations in the clusters.

Table A.1.2: Summary statistics by sector and within cluster correlation

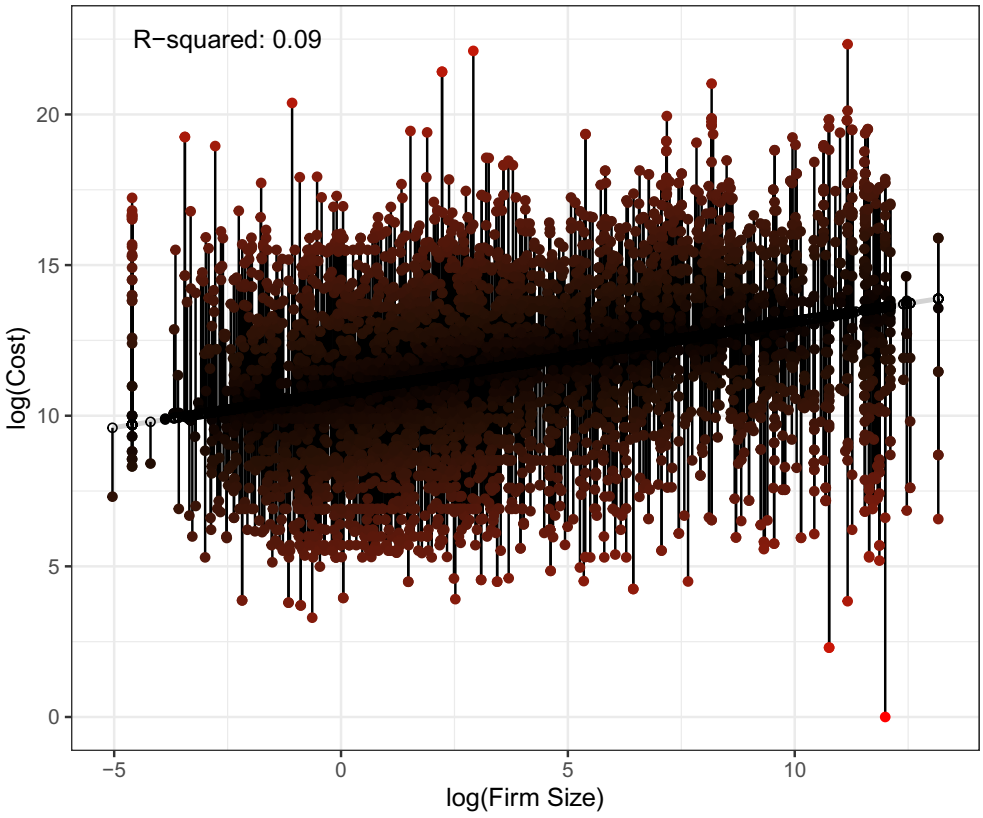

Figure A.1.1: Residuals of the estimation of firm revenues on the cost of cyber incidents

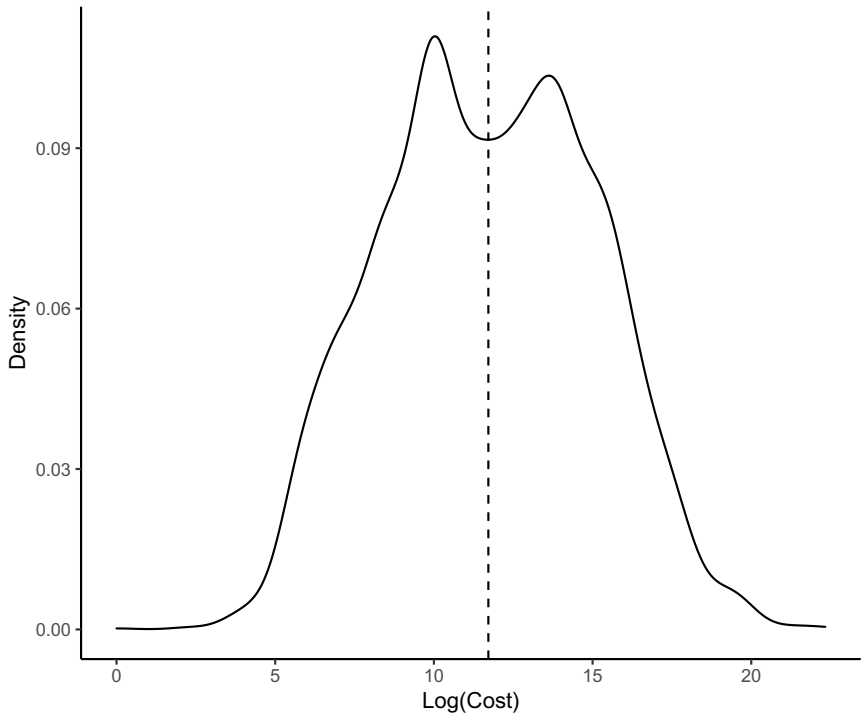

Figure A.1.2: The density of costs after the log transformation.

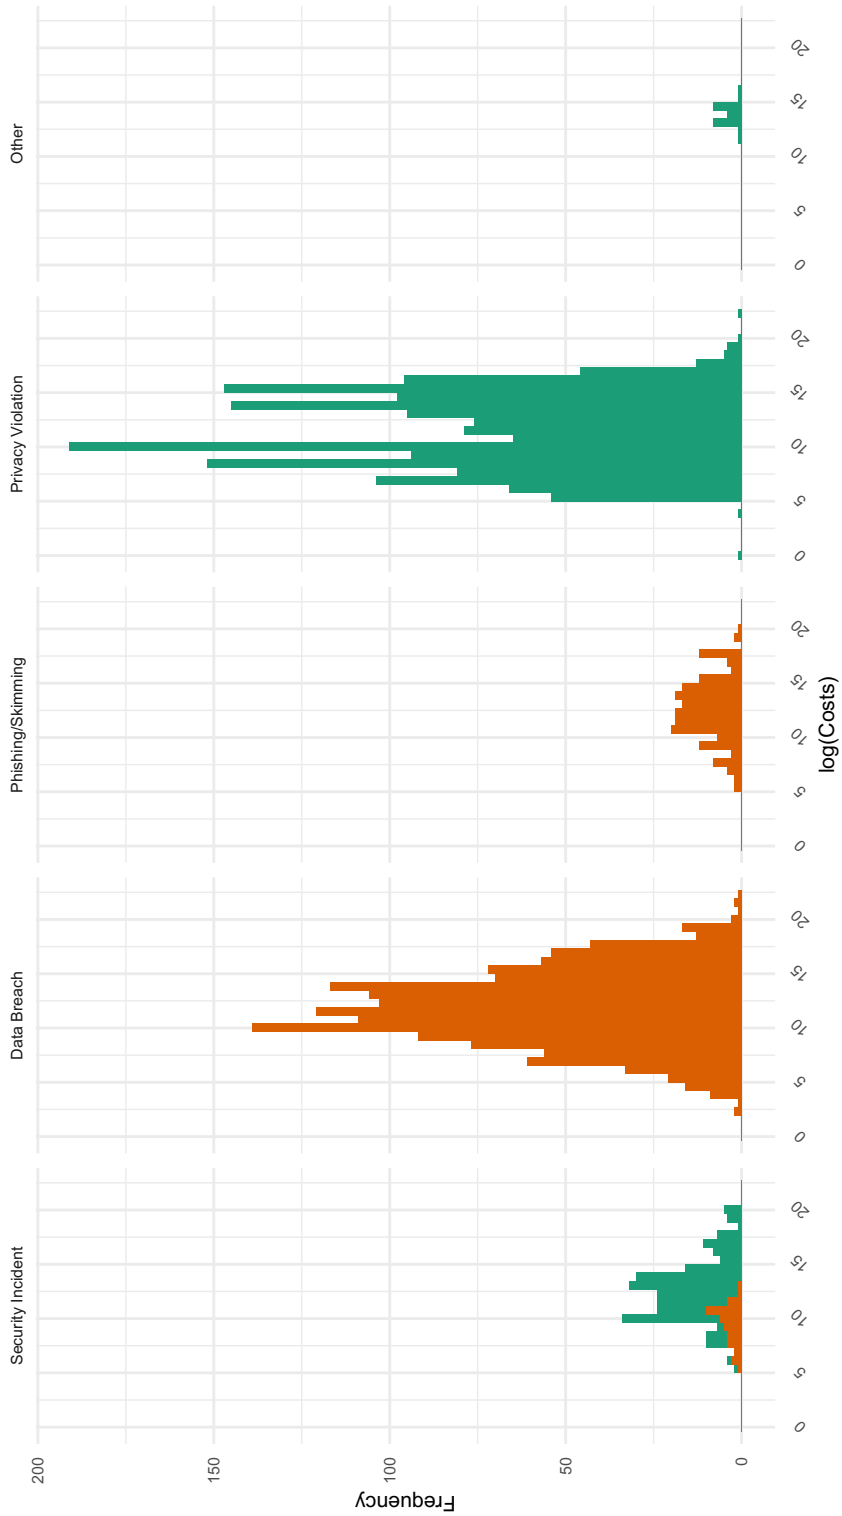

Figure A.1.3: Distribution of costs by case types

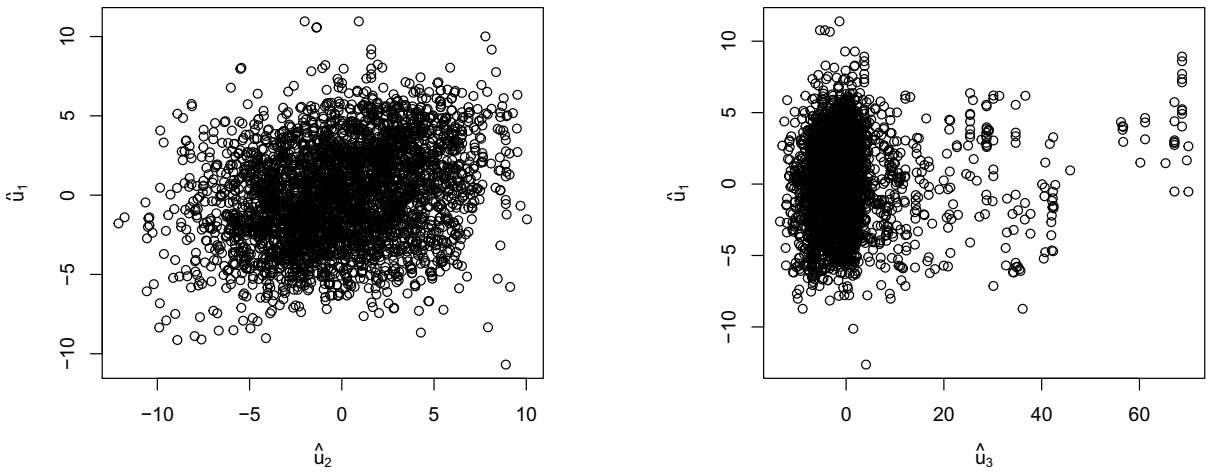

Notes: The two figures show the partial residual plots used to observe the existence of a second order relationship between costs and firm size, as well as costs and connections. The panel on the left hand side shows the residuals obtained by regressing costs on the variables from our baseline equation with the omission of firm size, denoted by $\hat{u}_1$, on the y-axis. On the x-axis are the residuals obtained by regressing firm size on the variables from the baseline equation, denoted $\hat{u}_2$. The right hand panel plots the partial residuals between costs and connections variable where the residuals from the regression of connections on controls is denoted by $\hat{u}_3$. The plots do not show any evidence of the presence of a second order term, we thus do not include any second order terms in our regressions.

Figure A.1.4: Partial residual plots to identify second order relationships

Dependent Variable: Log(Cost)					
Regressor	I	II	III	IV	V
Panel A: Ecker White Errors					
log(Firm Size)	0.241***	0.220***	0.231***	0.170***	0.234***
	(0.0134)	(0.0132)	(0.0140)	(0.0306)	(0.0140)
Connected events	0.0176***	0.0257***	0.0257***	0.0246***	-0.0251
	(0.00664)	(0.00638)	(0.00645)	(0.00653)	(0.0164)
log(Firm Size)^2				0.00698**	
				(0.00315)	
(Connected events)^2					0.000882***
					(0.000249)
Malicious	-1.31***	-1.33***	-1.20***	-1.20***	-1.17***
	(0.171)	(0.186)	(0.191)	(0.190)	(0.191)
Panel B: Cluster Robust Error by Firm					
log(Firm Size)	0.241***	0.220***	0.231***	0.170***	0.234***
	(0.0229)	(0.0192)	(0.0196)	(0.0448)	(0.0194)
Connected events	0.0176	0.0257**	0.0257**	0.0246**	-0.0251
	(0.0109)	(0.0114)	(0.0114)	(0.0117)	(0.0185)
log(Firm Size)^2				0.00698	
				(0.00505)	
(Connected events)^2					0.000882***
					(0.000304)
Malicious	-1.31***	-1.33***	-1.20***	-1.20***	-1.17***
	(0.219)	(0.218)	(0.223)	(0.222)	(0.223)
Year Fixed Effects	N	Y	Y	Y	Y
Sector Fixed Effects	N	N	Y	Y	Y
R^2	0.11	0.19	0.21	0.21	0.21
N	3705	3705	3705	3705	3705

Notes: Results form estimating equation (1.1). *, ** and *** denote significance at the 10, 5 and 1 percent level respectively. Panel A reports the estimates with Ecker-White standard errors (in parentheses). Panel B reports the estimates with cluster robust errors at the firm level (in parentheses). Definitions of the regressors are reported in Table 1.1. The regressions are analogous to those in Table 1.2.

Table A.1.3: Baseline model with alternative error clustering

Dependent Variable: Log(Cost per Unit Revenue)					
Regressor	I	II	III	IV	V
log(Firm size)	-0.820***	-0.843***	-0.825***	-0.844***	-0.822***
	(0.0295)	(0.0264)	(0.0247)	(0.0942)	(0.0244)
Connected events	0.0144	0.0221*	0.0243**	0.0241*	-0.0195
	(0.0105)	(0.0112)	(0.0113)	(0.0115)	(0.0195)
log(Firm size)^2				0.00146	
				(0.00782)	
(Connected events)^2					0.000762**
					(0.000332)
Malicious	-1.37***	-1.40***	-1.21***	-1.21***	-1.18***
	(0.241)	(0.248)	(0.250)	(0.250)	(0.250)
Security Incident	12.5***	14.9***	16.6***	16.7***	16.8***
	(0.293)	(0.600)	(0.688)	(0.718)	(0.698)
Data Breach	13.4***	16.2***	17.9***	17.9***	18.0***
	(0.256)	(0.600)	(0.693)	(0.717)	(0.702)
Phishing/Skimming	14.5***	16.9***	18.3***	18.4***	18.5***
	(0.470)	(0.703)	(0.765)	(0.793)	(0.771)
Privacy Violation	12.5***	15.3***	17.3***	17.3***	17.4***
	(0.177)	(0.550)	(0.648)	(0.675)	(0.656)
Other	13.4***	15.6***	17.9***	17.9***	18.2***
	(0.826)	(1.02)	(1.08)	(1.11)	(1.10)
Year Fixed Effects	N	Y	Y	Y	Y
Sector Fixed Effects	N	N	Y	Y	Y
R^2	0.45	0.5	0.52	0.52	0.52
N	3699	3699	3699	3699	3699

Notes: Results form estimating equation (1.1) using the cost per unit revenue as an alternate dependent variable. *, ** and *** denote significance at the 10, 5 and 1 percent level respectively. All standard errors (reported in parentheses) are clustered by sector. The definition of the regressors are reported in Table 1.1, however in this specification Firm size refers to the number of employees. Column I is an OLS regression without controls for year and sector fixed effects. Column II is an OLS regression without year fixed effects. Column III is an OLS regression including both fixed effects. Column IV tests for the presence of a non-linear relationship between firm revenues and the costs. Column V tests for the non-linear relationship between connections and costs.

Table A.1.4: Baseline model using scaled dependent variable

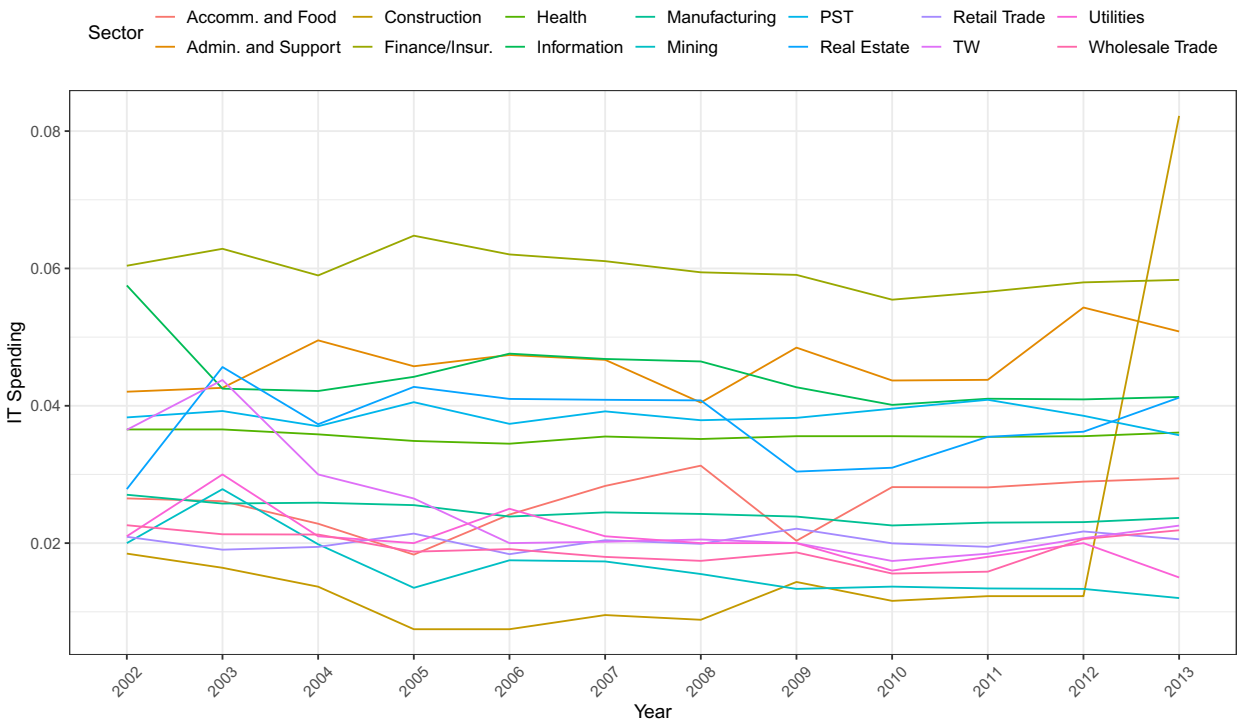

Notes: The plot shows the change in estimated IT spending by sector as a percentage of revenues using the methodology in Kennedy and Stratopoulos (2017). Data is available over the period 2002-2013 and for 14 sectors. The following sectors are abbreviated in the legend: Professional, Science and Technology (PST) and Transport and Warehousing (TW).

Figure A.1.5: IT expenditures across sectors

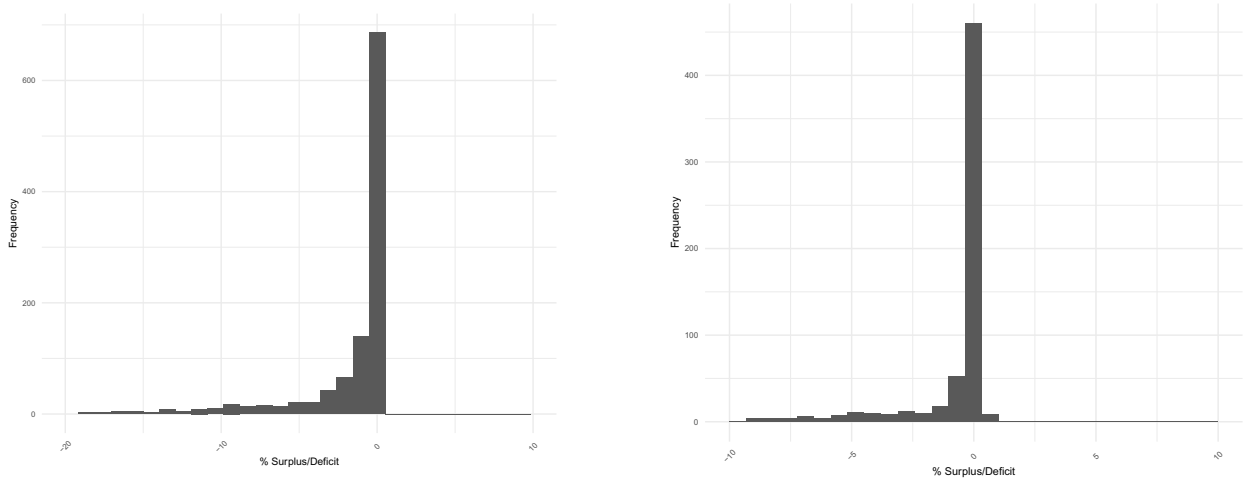

Notes: The panel on the left hand side shows the histogram of the difference between annual spending as a percentage of revenue minus the annual cost of cyber incidents suffered by firms that were the result of non-malicious incidents. The panel on the right reports the equivalent histogram using malicious incidents.

Figure A.1.6: Firms over / under spending.

APPENDIX B

OPERATIONAL AND CYBER RISK IN THE FINANCIAL SECTOR

B.1 Description of the calculation of capital

Extension of the Internal Measurement Approach.

As done with all frameworks under the Advanced Measurement Approach, the Internal Measurement Approach partitions a bank's operational risk exposures into a series of business lines and operational risk event types. Each intersection of business line and event type is known as a cell. For each cell, a separate expected loss figure is calculated. Due to data limitations, we use solely business lines as individual cells rather than the intersection of business lines and event types. A γ factor is then used to translate the expected loss into a capital charge. Alexander (2008) proposes a method to determine the γ factors that translate into observable quantities in the loss frequency distribution, and therefore the parameter can be calibrated based on operational risk data.

The basic idea is to map the expected loss to a level of capital that covers the unexpected annual loss, defined as the 99.9th percentile of annual loss net of mean annual loss (shown in Figure 2.4. Alexander's alternative γ factors, labelled as ϕ are thus defined as follows:

$$\phi = (99.9^{\text{th}}\text{percentile} - \text{mean})/\text{standard deviation}$$

where mean and standard deviation refer to the measures of the annual loss distribution. Under the assumption that loss severity is random, Alexander's approach suggests ϕ is calculated as follows:

$$\phi = (99.9^{\text{th}}\text{percentile} - \lambda\mu_L)/\sqrt{\lambda(\mu_L^2 + \sigma_L^2)} \tag{B.1}$$

where σ_L is the standard deviation of annual losses, μ_L is the mean of annual losses and λ is the mean frequency of losses under the assumption they follow a Poisson distribution. The calculation of operational risk capital then becomes:

$$K_{IMA} = \phi \times \mu_L \times \sqrt{\lambda} \times \sqrt{1 + \left(\frac{\sigma_L}{\mu_L}\right)^2} \tag{B.2}$$

the term $\sqrt{1 + \left(\frac{\sigma_L}{\mu_L}\right)^2}$ is included to account for the uncertainty in loss severity. Note that higher variation leads to a greater the capital charge. To calculate the operational risk capital based on this approach, we first obtain the mean, μ_L, and standard deviation, σ_L, of annual losses from the ORX database. For each business line, i, we use maximum-likelihood estimation to fit $\hat{\lambda}_i$ and then compute the estimate of $\hat{\phi}_i$ from equation B.1.

LDA and Bayesian methodology.

The LDA gives great flexibility to banks with respect to estimating the capital necessary to cover operational losses. In our analysis we use two models from the LDA suite for capital calculation. In this section we focus on the Bayesian approach, and note that the alternate MCMC methodology also used in our analysis follows a similar logic. More details on this approach and LDA more widely can be found in Cruz et al. (2015).

Various methodologies can be used to estimate the frequency and severity distributions and subsequently perform the convolution of the two. Here, we detail a Bayesian approach to estimating the annual loss distribution, which tends to give greater flexibility and avoids estimation problems typically encountered when working with extreme value distributions. We consider non-

informative priors for which Bayesian estimates converge to maximum likelihood ones. We follow the approach used by Figini et al. (2015) to estimate the annual loss distribution, considering a convolution between a Generalised Pareto distribution for the mean loss (severity), with a Poisson distribution for the number of loss events (frequency), as in Chavez-Demoulin et al. (2006).

The annual losses can be written as a product of Frequency (the number of loss events during a certain time period) and Severity (the mean impact of the event, in terms of financial losses, in the same period). In particular,

$$L_{it} = s_{it} \times n_{it} \tag{B.3}$$

where for the business line/event type intersection i and for t time periods available, L_{it} denotes the annual operational loss, s_{it} denotes the severity and n_{it} the frequency. As noted above, we aggregate over business lines rather than the intersection of business lines and event types. Following the operational risk literature, we consider the following three general assumptions: i) within each intersection i, and each time period t, the distribution of the frequency n_{it} is independent of the distribution of the severity s_{it}; ii) for any given time period t, the losses occurring in different intersections, i, are independent of each other; iii) for any given intersection, i, losses occurring in different time periods, t, are independent of each other.

Let $f(s_t|\theta)$ and $f(n_t|\lambda)$, denote the likelihood functions of the severity and frequency respectively, where θ denotes the parameter vector of the severity distribution and λ denotes the parameter vector of the frequency distribution, we have that, according to assumptions i)-iii):

$$L(s, n|\theta, \lambda) = \prod_{t=1}^{T} f(n_t|\lambda) f(s_t|\theta) \tag{B.4}$$

While expert input can be useful to construct informative priors, we use uninformative priors

with high variance, as in Dalla Valle and Giudici (2008). For the frequency, we use the conjugate gamma distribution.

$$\lambda_i \sim \Gamma(\alpha, \beta) \tag{B.5}$$

We choose $\alpha = 0.01$ and $\beta = 0.01$. The severity is assumed to follow a general Pareto distribution:

$$F_i \sim GPD(\mu, \xi, \sigma) \tag{B.6}$$

First, we assume the location parameter, $\mu = 0$. We then follow Cabras and Castellanos (2007) and use an uninformative prior for ξ and σ of the severity distribution.

$$\pi(\xi, \sigma) \propto \sigma^{-1}(1+\xi)^{-1}(1+2\xi)^{-1/2}, \quad \xi > -0.5, \sigma > 0 \tag{B.7}$$

Since there are no analytical solutions to this problem, we use the Metropolis-Hastings algorithm to estimate the posterior distributions of the annual frequency and severity. We then take the convolution of the two distributions to obtain the annual loss distribution.

B.2 Tables and Figures

Business Line[a]	**Description**
Corporate finance	Structuring, issuance or placement of securities and similar instruments, not just for capital raising
Trading and sales	Products / Positions held in the Trading Book of the firm and Corporate Investments.
Retail banking	Retail loans, Retail deposits, Banking services, Trusts & estates, Investment advice, Cards - Credit & Debit
Commercial banking	Project finance, Real estate finance, Export finance, Trade finance, Factoring, Leasing, Loans guarantees, Bills of exchange
Clearing	Financing and related services
Agency services	Bank account, deposit services, "plain vanilla" investment products
Asset management	Management of individual assets invested in financial instruments on behalf of others (i.e. not in the bank's own name for its own account) in which the bank has the power to make investment decisions. This includes activities where each customer's assets are held in a separate portfolio, as well as those where the assets of different customers are pooled in one portfolio.
Retail brokerage	Various services related to administration and management of estates, trusts, assets, portfolios etc.
Private banking	Limited category for items than can only be categorised at corporate level

Notes: [a]The definitions of business lines used by ORX are mapped to those used under the Basel II framework.

Table B.2.1: Overview of business lines based on the operational risk reporting standards of ORX

Region	**Sub-regions**
North America	US, Canada
Latin America & Caribbean	-
Eastern Europe	-
Western Europe	Southern Europe, Northern Europe, United Kingdom, Western Europe
Asia / Pacific	-
Africa	-

Table B.2.2: Overview of regions and sub-regions

	t_1				t_2				t_3			
	Mean	Std. Dev.	Median	Q_{95}	Mean	Std. Dev.	Median	Q_{95}	Mean	Std. Dev.	Median	Q_{95}
Panel A - By Event												
Internal fraud	292	602	13	1,558	145	365	15	773	437	713	121	1,975
External fraud	116	381	0	651	81	271	6	324	197	490	39	1,061
Employee-related	163	557	0	1,077	448	846	26	2,382	611	956	85	2,771
Clients & business practices	556	1,044	0	3,182	255	547	24	1,346	810	1,133	224	3,415
Disasters	60	237	0	348	133	298	27	666	192	384	53	902
Technology	76	283	0	394	63	192	3	324	139	359	17	723
Transactions & process management	255	648	2	1,689	144	396	7	793	399	780	49	2,182
Panel B - By region												
Africa	189	495	6	1,168	125	336	12	654	314	617	63	1,622
Asia/Pacific	234	586	6	1,521	80	234	5	412	314	642	40	1,764
Canada	114	379	0	713	132	345	21	697	246	518	49	1,319
Eastern Europe	483	775	92	2,176	189	448	12	1,140	672	886	266	2,626
Latin America & Caribbean	163	567	0	1,085	437	887	9	2,510	600	1,006	49	2,922
Northern Europe	163	440	6	938	74	212	12	318	237	508	46	1,156
Southern Europe	707	1,141	80	3,472	179	456	8	1,056	886	1,238	228	3,680
United Kingdom	210	579	5	1,430	66	211	0	345	276	635	28	1,679
United States	146	474	0	1,002	148	325	29	769	295	610	61	1,599
Western Europe	416	886	12	2,748	111	312	7	589	526	933	83	2,886
Panel C - By size												
Large	252	689	0	1,820	194	516	12	1,136	446	855	63	2,504
Medium	233	616	1	1,538	140	383	4	833	372	724	50	2,025
Small	245	601	4	1,606	149	395	14	858	394	738	61	2,172

Notes: The table shows the summary of the different definitions of duration in our data. We present the Mean, Standard Deviation, Median and the 95th quantile, by various breakdowns. In Panel A by event type, Panel B by region and Panel C by Banksize

Table B.2.3: Summary of durations by region, event type and size (in days)

Regressor	Dependent Variables		
	$\frac{Loss_{it}}{Income_{it}}$	$\frac{Freq_{it}}{Income_{it}}$	$\frac{Severity_{it}}{Income_{it}}$
Panel A: Recognition Date			
Taylor Rule	-0.0829**	-0.0708*	-0.0120
	(0.0359)	(0.0397)	(0.0349)
Boone Indicator	0.870	0.711	0.159
	(0.619)	(0.481)	(0.677)
Credit GDP Gap	0.00594	0.0124	-0.00645
	(0.0124)	(0.0130)	(0.00571)
Supervisory Index	-3.17	-2.44	-0.723
	(2.64)	(1.95)	(1.08)
R^2	0.1	0.19	0.19
N	123	123	123
Panel B: Occurrence Date			
Taylor Rule	-0.0578**	-0.0491**	-0.00870
	(0.0239)	(0.0241)	(0.0201)
Boone Indicator	1.01	0.889**	0.117
	(0.620)	(0.379)	(0.653)
Credit GDP Gap	0.0128	0.0124*	0.00033
	(0.00852)	(0.00681)	(0.00531)
Supervisory Index	-0.0972	-0.439	0.342
	(1.83)	(0.819)	(1.20)
R^2	0.12	0.29	0.29
N	123	123	123
Time FE	Y	Y	Y
Region FE	Y	Y	Y

Notes: The table contains the results of a panel regression with all macroeconomic variables. The dependent variables are the logarithm of the loss, frequency and severity normalised by income. Standard errors are robust with small sample correction. *, ** and *** denote significance at the 10, 5 and 1 percent level, respectively. All standard errors are robust to small sample. Panel A shows the coefficients when aggregating by recognition date and Panel B by occurrence date.

Table B.2.4: Panel Regression of Contemporaneous Variables

	Dependent variable		
	$\frac{\tilde{TotalLoss}}{Income}$	$\frac{\tilde{Frequency}}{Income}$	$(\frac{Severity}{Income})$
Panel A			
Credit-GDP-Gap - 4 Lags	0.0061	0.0061	0.00064
	(0.0096)	(0.0096)	(0.0047)
Credit-GDP-Gap - 8 Lags	0.0071	0.0071	0.0039
	(0.0095)	(0.0095)	(0.0040)
Panel B			
Taylor Rule Dev. - 4 Lags	-0.046**	-0.061**	0.015
	(0.022)	(0.028)	(0.021)
Taylor Rule Dev. - 8 Lags	-0.069**	-0.095*	0.026
	(0.031)	(0.055)	(0.036)
Regional Fixed Effects	Y	Y	Y
Time Fixed Effects	Y	Y	Y

Notes: The table is divided into two panels summarising the results from 12 panel regressions. Each column denotes the dependent variables used, which are logged and corrected for an underreporting bias. For these regressions we extend our data collection of the credit-to-GDP gap and deviations from the Taylor rule to match the full database at 2018 Q3. Each panel distinguishes between the dependent variables used. The coefficients shown are the sum of the lagged variables i.e. the cumulative effect, for example at 4 lags the coefficient reported is, $\sum_{i=1}^{4} \hat{\beta}_i$. A robust sum of standard errors is reported in parenthesis. The sum of standard errors is calculated as $\sqrt{L'VL}$, where L is a 0,1 vector that denotes the linear combination of regressors and V is the estimated robust covariance matrix. We test that the sum of coefficients is significantly different to 0. The asterisks denote the significance as follows: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions are two way fixed effects models, including a regional and time effect.

Table B.2.5: Operational losses and the macroeconomic environment, with bias adjustment

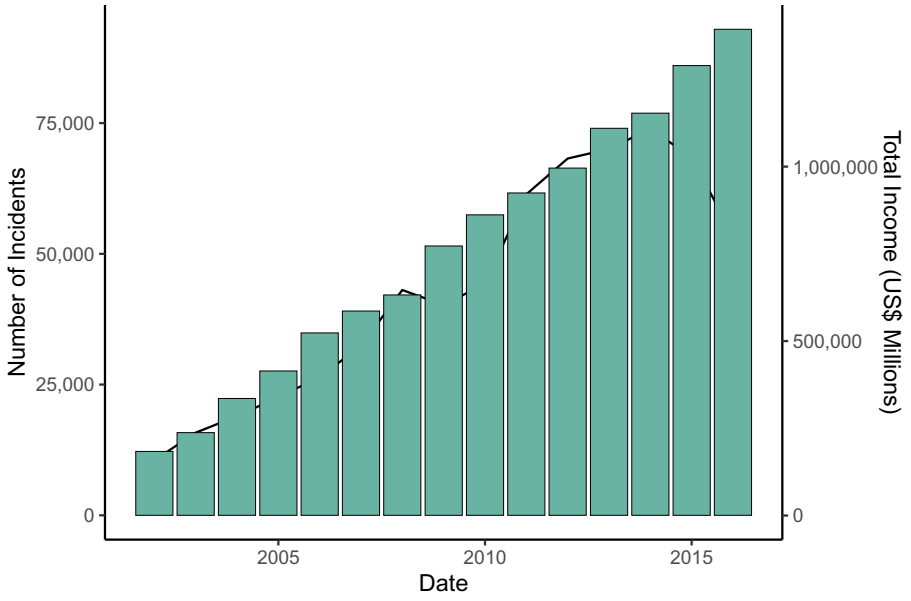

Notes: The plot shows the total number of incidents per year alongside the total income of the consortium. The bars denote the total income (right axis) and the line denotes the frequency of incidents (left axis).

Figure B.2.1: Sample size and frequency of events

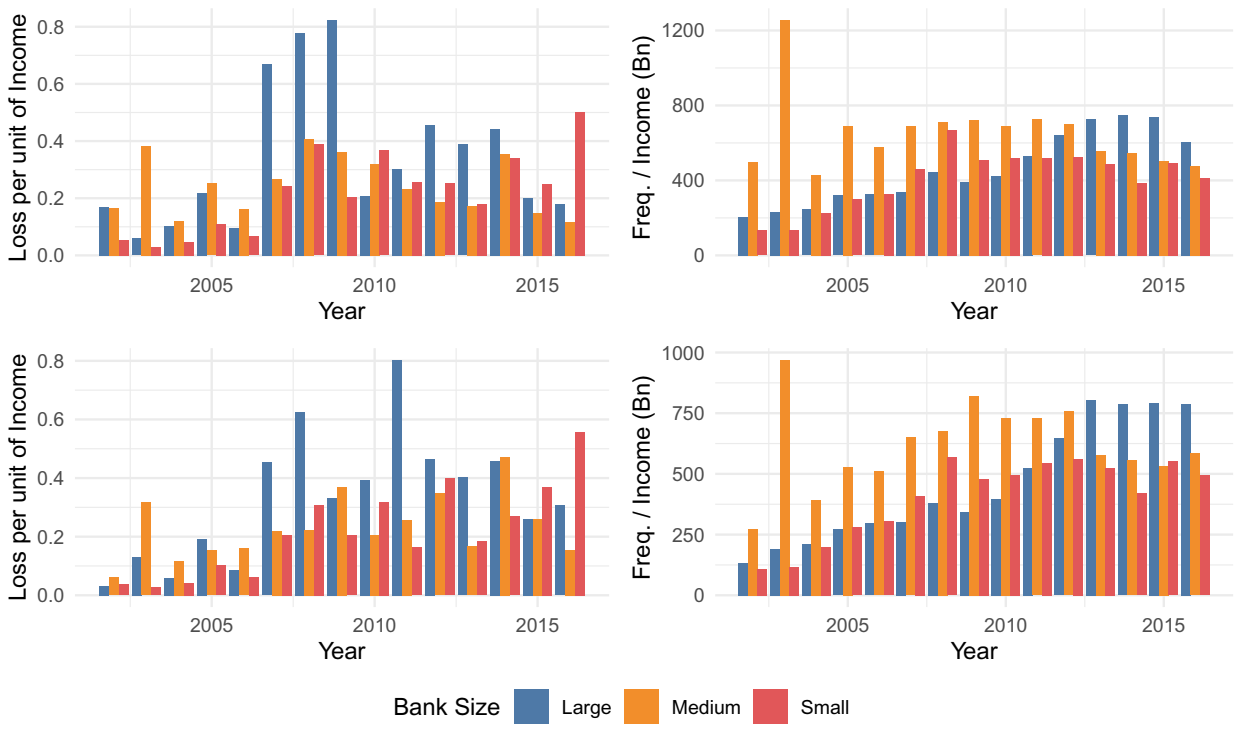

Figure B.2.2: Loss and frequency over time partitioned by bank size

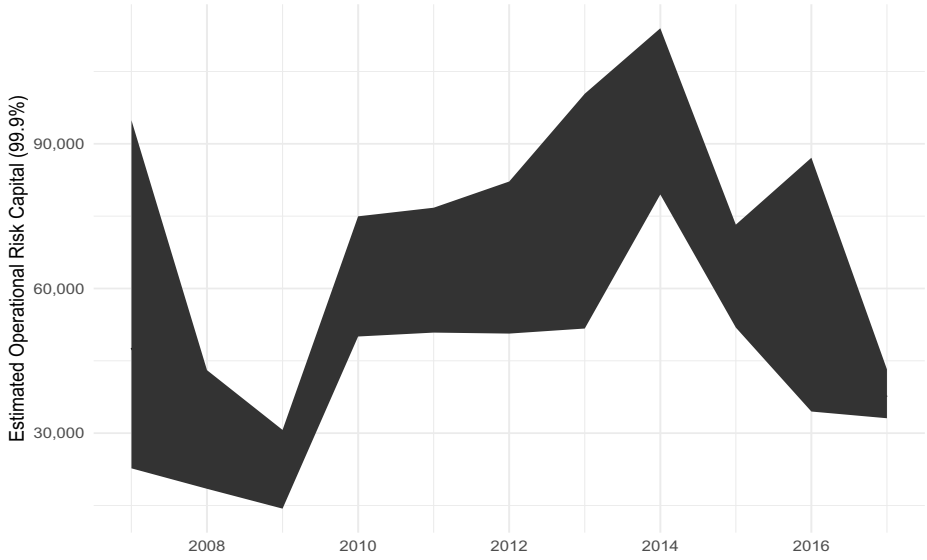

(a) Bayesian LDA Estimate with 95% Confidence Interval

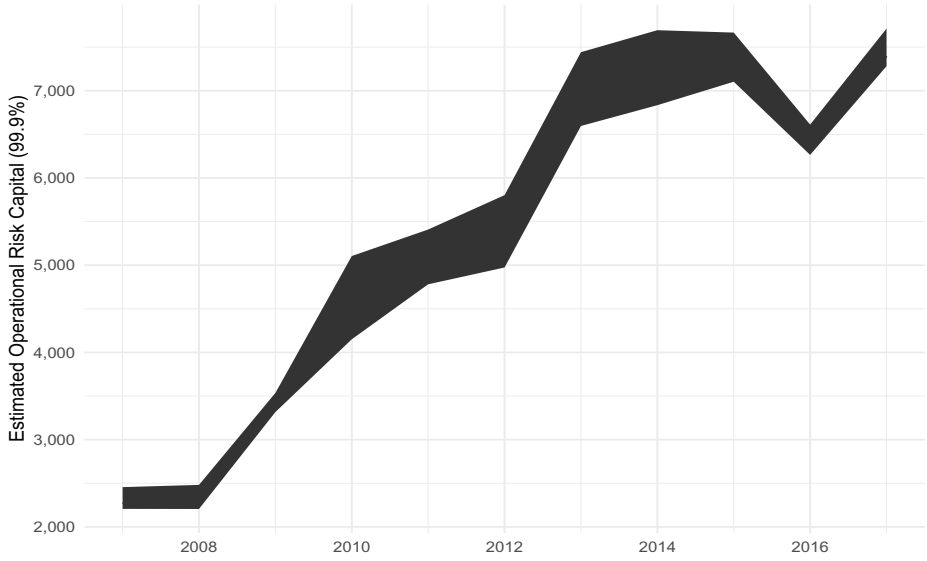

(b) Lognormal LDA Estimate with 95% Confidence Interval

Notes: The plot shows the estimated operational risk capital by two different methodologies and the 95% confidence interval for the location of the 99/9% quantile of the annual loss distributions. These are calculated by using the approximation put forward in Cruz et al. (2015). The upper or lower bound can be calculated as $B = K\alpha \pm F_N^{-1}\left(\frac{1+\gamma}{2}\right)\sqrt{(K\alpha(1-\alpha))}$. Where, K denotes the number of Monte Carlo random draws of the annual losses, F_N^{-1} the inverse of the standard normal distribution, γ the desired confidence interval and α the chosen quantile.

Figure B.2.3: Confidence intervals for VaR

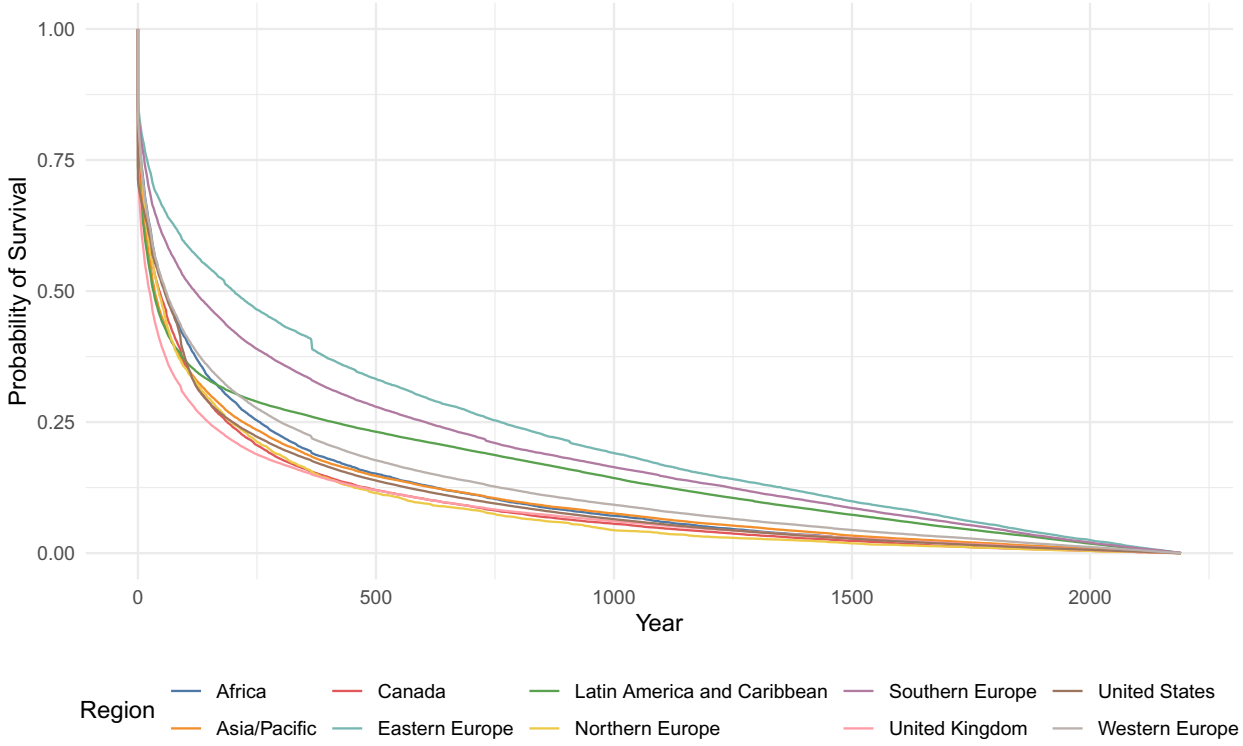

Figure B.2.4: Estimated survival curves by region

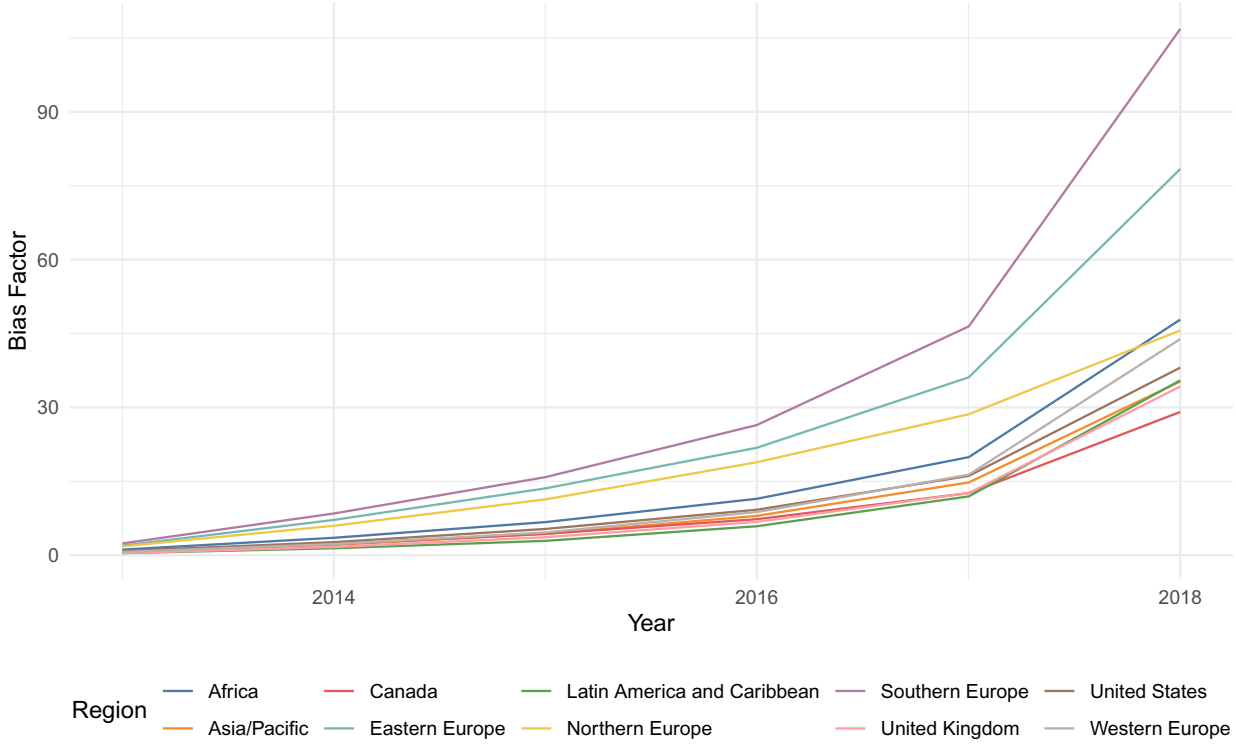

Figure B.2.5: Estimated bias factor by region

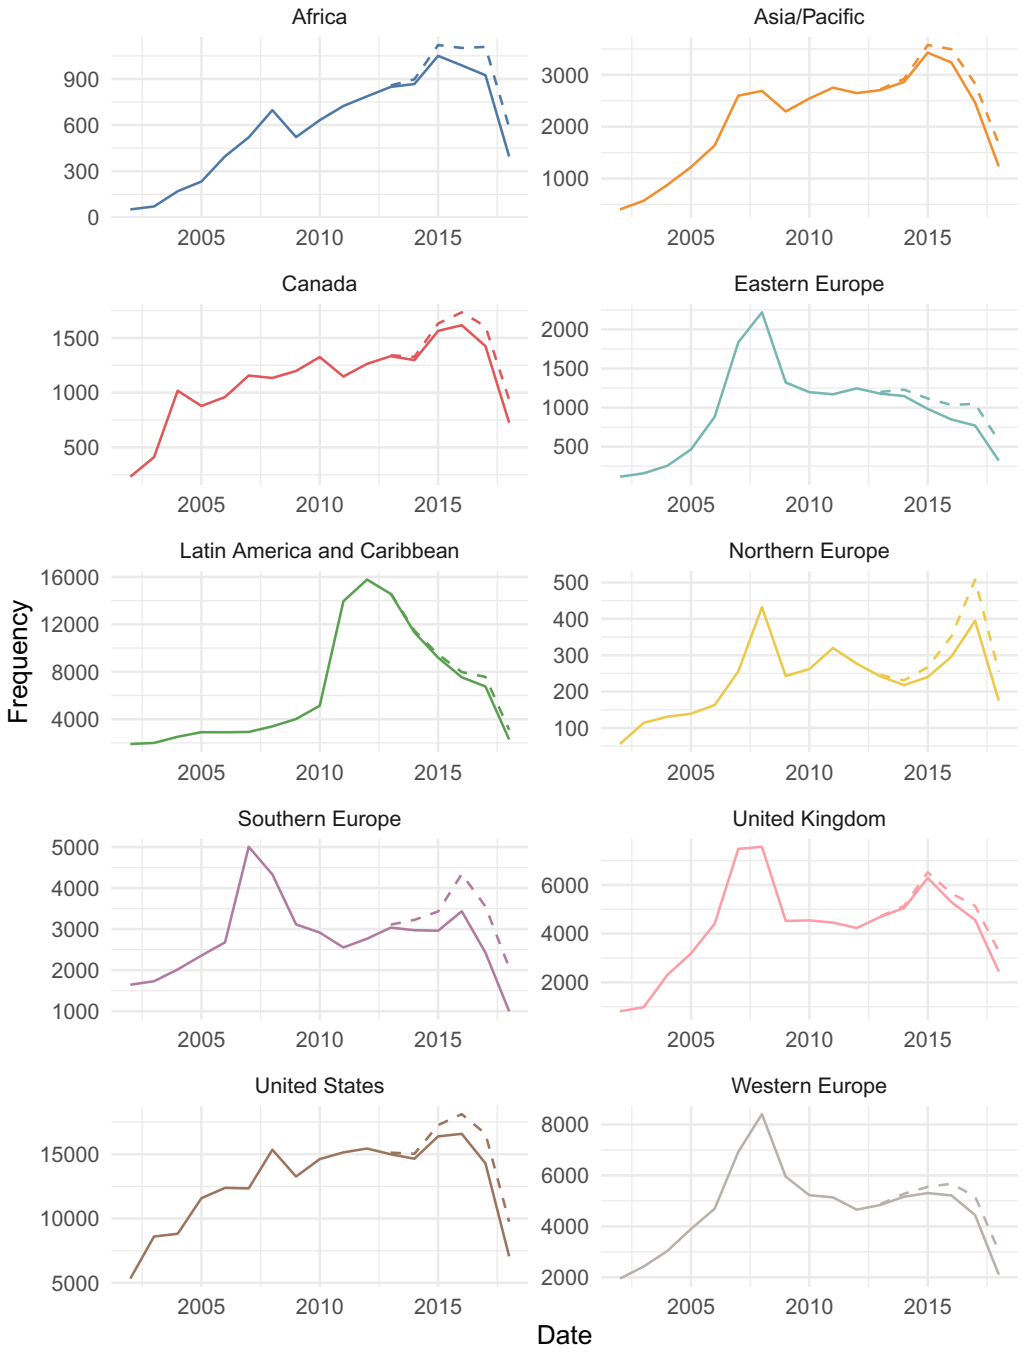

Figure B.2.6: Annual frequencies adjusted for data bias by region

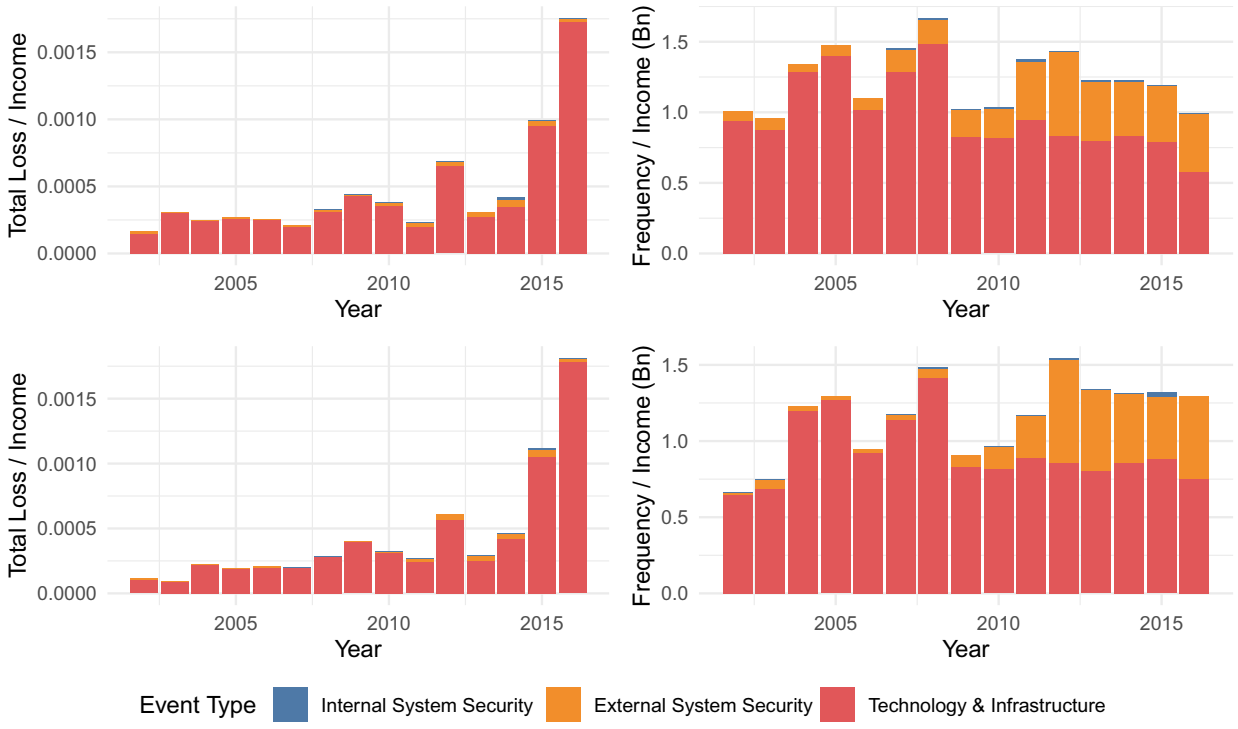

Notes: On the left hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by cyber event type.

Figure B.2.7: Loss and frequency of operational losses by event type

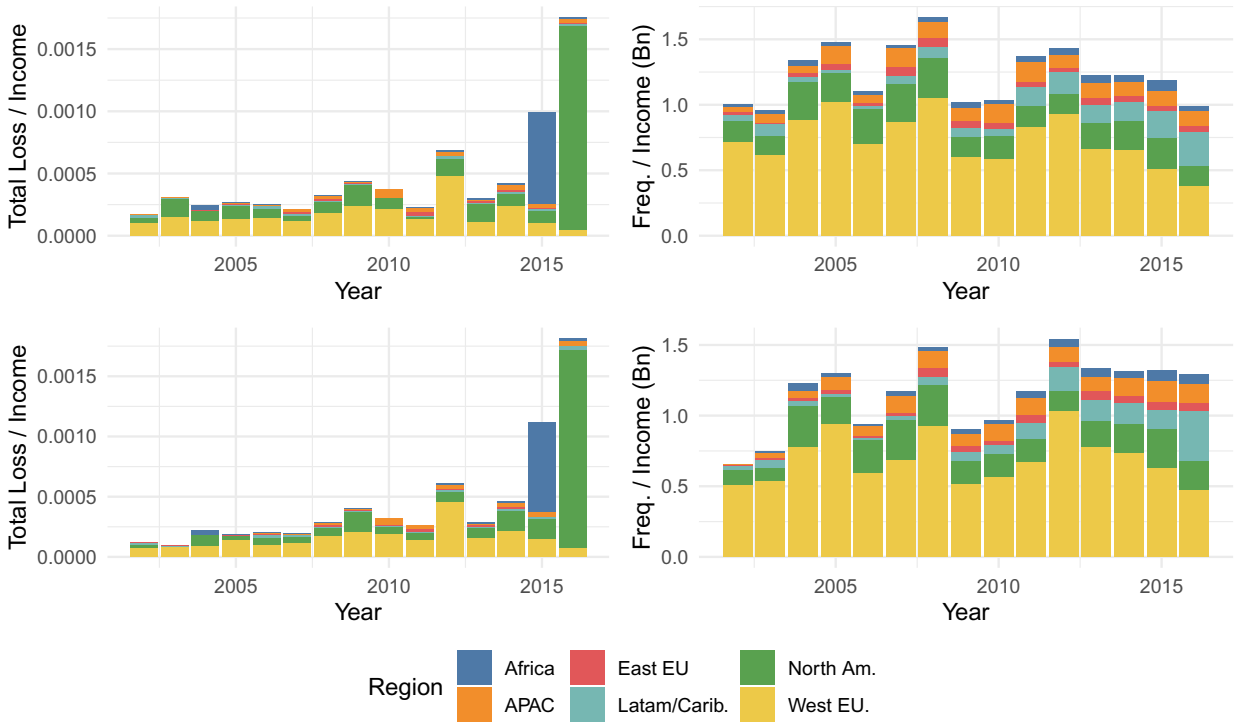

Notes: On the left hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by region. Abbreviations in the legend are defined as follows: APAC: Asia/Pacific; East EU: Eastern Europe; Latam/Carib: Latin America and Caribbean; North Am: North America; and West EU: Western Europe.

Figure B.2.8: Loss and frequency of operational losses by event type

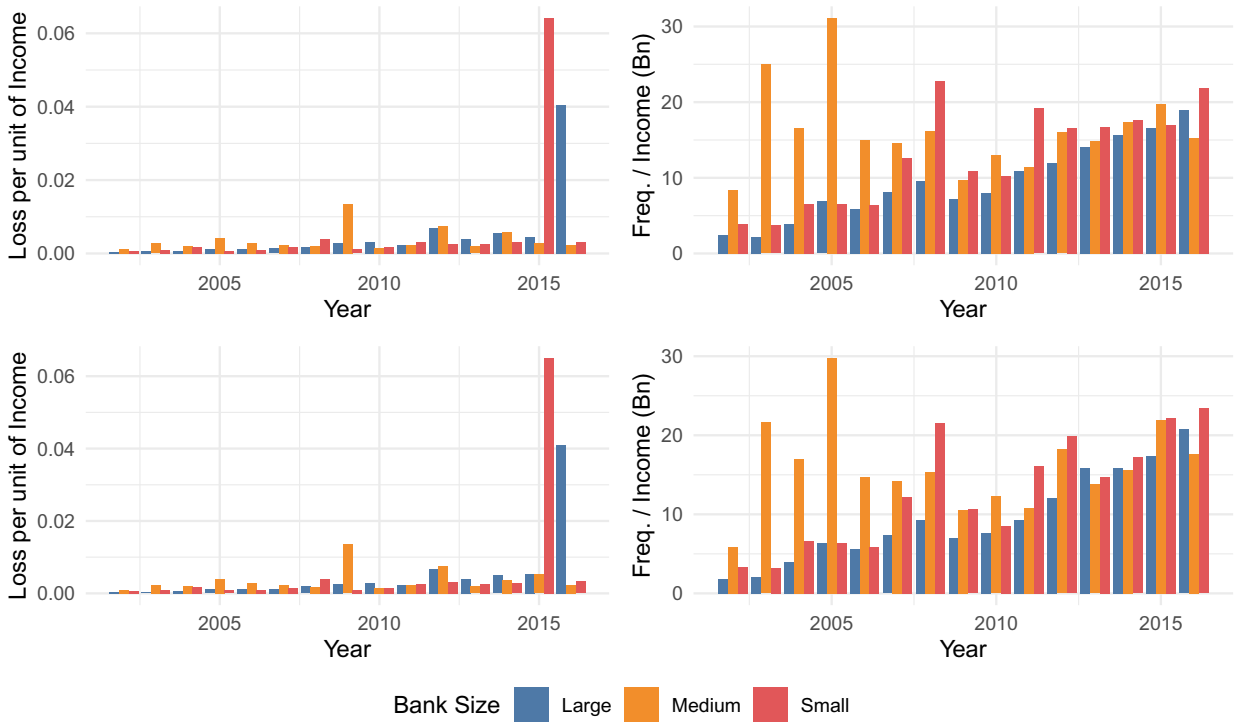

Notes: On the left hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by bank size.

Figure B.2.9: Loss and frequency of operational losses by event type