Operational and Cyber Risks in the Financial Sector*

Iñaki Aldasoro,[a] Leonardo Gambacorta,[a] Paolo Giudici,[b] and Thomas Leach[b]
[a]Bank for International Settlements
[b]University of Pavia

We use a unique cross-country data set at the loss event level to document the evolution and characteristics of banks' operational risk. Operational value-at-risk varies substantially—from 6 percent to 12 percent 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 recognized 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.

JEL Codes: D5, D62, D82, G2, H41.

1. Introduction

Operational risk emerged as a distinct risk category in the mid-1990s, following events such as Nick Leeson's "rogue" trader case at Barings Bank. Not long after, the Basel II standards introduced

*We thank two anonymous referees, Steve Bishop, Antoine Bouveret, Luke Carrivick, Alexander Demarco, Raymond Kleijmeer, Karin Reichardt, Ken Taniguchi, Philip Umande, Andreas Voegtli, and seminar participants at the 2020 CEBRA Annual Meeting, the Bank of Italy, Bank for International Settlements, Central Bank of Malta, IFI European Chapter Forum, and IFI Insurance Management Forum for helpful comments and suggestions. We are grateful to Oliver Denk and Gabriel Gomes for sharing their data on financial regulation. The views expressed here are those of the authors and do not necessarily represent those of the Bank for International Settlements or ORX. Author e-mails: Inaki.Aldasoro@bis.org; leonardo.gambacorta@bis.org; paolo.giudici@unipv.it; thomas.leach01@universitadipavia.it.

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 organizations, 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 analyzing a unique cross-country data set of operational losses. We present stylized 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 categorization 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 anonymized fashion in order to benchmark operational risk models. The sample we use contains over 500,000 operational loss events from 2002 until

[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 percent of risk-weighted assets can be attributed to operational risk in some jurisdictions (Liao, Ma, and Sands 2018).

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 characterize 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 recognized 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 characterized by a fat-tailed distribution.[3] Accordingly, estimates of 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 percent to 7.5 percent of gross income, against the 12 percent benchmark of the basic indicator approach. This finding may provide some support for the new regulatory framework that proposes the adoption of the standardized 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]

[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-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

Operational losses, on average, take over a year to be discovered and recognized 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 recognizing 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. These findings can inform policy discussions regarding the principles for executive compensation packages.

The stylized facts we present point to the existence of a link between operational losses and macroeconomic conditions. Abdymomunoy, Curti, and Mihov (2017) use data for U.S. 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 characterizing the run-up to the downturn. In other words, favorable 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 recognized 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

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.

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 the work of 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, Detragiache, and Tressel (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 organized as follows. The next section reviews the related literature. Section 3 describes the data and documents the duration between occurrence, discovery, and recognition of loss events. Section 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 5, whereas Section 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. Related Literature

Research on operational risk intensified after 2001, when the Basel Committee on Banking Supervision (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 conceptualize and quantify these risks (Cornalba and Giudici 2004; Power 2005; Chavez-Demoulin, Embrechts, and Nešlehorá 2006; Jarrow 2008; Antonini et al. 2009).

The literature points to links between the characteristics of financial institutions and operational risk. Curti, Frame, and Mihov (2019) and Shih, Samad-Khan, and Medapa (2000) find a positive relationship between size and operational losses. Chernobai, Jorion, and Yu (2011) use data for U.S. financial institutions and find 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 U.S. 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 sizable financial penalties. This can also affect bank returns (Byrne, Coughlan, and Tilley 2017; Köster and Pelster 2017).[5] Altunbaş, Thornton, and Uymaz (2018) find that banks are more likely to engage in misconduct when their CEOs have a long tenure. Eshraghi, Hagendorff, and Nguyen (2016) study regulatory enforcement actions issued against U.S. 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.

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

Operational risk could also be intertwined with business and financial cycles. Hess (2011) and Carrivick and Cope (2013) look at the consequences of the GFC on operational risk losses in the financial sector. Abdymomunov, Curti, and Mihov (2017) provide additional evidence of a relationship between operational losses in U.S. 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 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 are 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 categorization of different event types. Kaffenberger, Kopp, and Wilson (2017) examine the current regulatory framework and supervisory approaches, and identify 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) analyzes 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) analyze 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, Kovner, and Lee (2021) estimate that the impairment of any of the five most active U.S. banks could

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

result in significant spillovers to other banks, with 38 percent of the network affected on average.

3. Data

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 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 anonymized, 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, Schuermann, and Scott 2015). The full sample comprises over 700,000 observations of operational loss events occurring between 2002:Q1 and 2018:Q3. 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 2016:Q4, 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, Jorion, and Yu (2011) use the sample of data with 2,426 loss events reported by the Algo FIRST data set. Hess (2011) uses data reported

[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.

Table 1. Example of the Data Structure

RefID	Region	Business Line	Event Type	Gross Loss Amount	...	Loss Occurrence	Loss Discovery
123XYZ	Asia/Pacific	BL0101	EL0101	20,000	...	ddmmyyyy	ddmmyyyy
⋮	⋮	⋮	⋮	⋮	...	⋮	⋮

by SAS OpRisk Global Data with around 7,300 loss events from the banking industry. Cope, Piche, and Walter (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 monetary cost reflected in the books of the banks, above a minimum of EUR 20,000. After data anonymization by ORX, individual losses can only be identified by geography, business line, and event type. Table 1 provides an example of how the data are structured.

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 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 6.

Loss events are also associated with a *business line*. The business line classification, which again follows pre-specified standards,

[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/standard/OPE.htm.

Table 2. Overview of Event Types Based on the Operational Risk Reporting Standards of ORX

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.

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

comprises nine business lines, including asset management, clearing, retail banking, and trading and sales, among others. Table B.1 in Appendix B 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 4.

The data are also partitioned into macro-regions. These include North America, Latin America and the 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 in Appendix B for details). While data are collected so as to preserve bank anonymity,

Figure 1. Loss Timeline and Key Dates

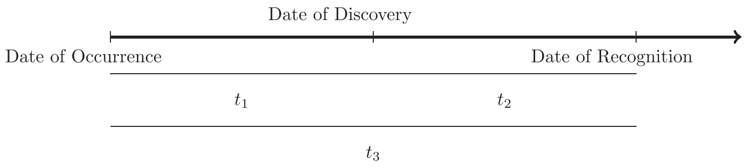

each loss event has a tag for bank size. This indicator variable divides financial institutions based on income into large, medium, and small.

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 led to the operational loss. Finally, the *date of recognition* represents the date when the loss was recorded in the accounts of the bank. Figure 1 depicts the timeline of a loss. We explore the factors that determine the duration of losses in Section 3.5. However, this also brings us to an important juncture regarding completeness of the data, which we discuss next.

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.1 in Appendix B 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 willful under-reporting of events, but is merely an artifact of the time it takes for events to be discovered and recognized. 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 3.5). 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 recognized), 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 recognized 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 3.5.[10]

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 buildup 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 2002:Q1 until 2016:Q4.

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

[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.

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 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:

$$\tilde{\phi}_t = i_t - \phi_t, \tag{1}$$

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):

$$\phi = r^* + \pi^* + 1.5(\pi - \pi^*) + 0.5y, \tag{2}$$

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 2002:Q1 until 2016:Q4.

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, Detragiache, and Tressel (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 markets, and prudential regulation and bank supervision. The main variable

[13]The estimates of the Boone indicator in this database are based on the approach used by Čihák and Schaeck (2010) but use marginal costs rather than average costs.

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 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 liberalized 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 U.K. banks, three German banks, and four French banks, we would compute the statistic for the region as follows:

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

3.4 Stylized Facts

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

Table 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-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

[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.

Table 3. Summary Statistics of Loss Events by Categories

	Mean	Std. Dev.	Max.	Min.
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, and 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
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 and 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
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

Note: The table presents summary statistics of losses by various categorizations. 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.

Figure 2. Loss and Frequency of Operational Losses by Event Type

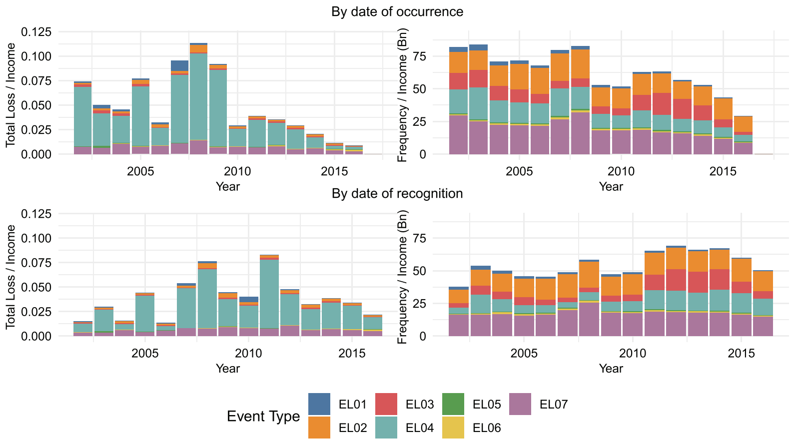

Note: On the left-hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right-hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of occurrence and the bottom panel by date of recognition. Each bar is partitioned by event type. EL01 = Internal Fraud; EL02 = External Fraud; EL03 = Employee Related; EL04 = Clients, Products and Business Practices; EL05 = Disasters; EL06 = Technology and Infrastructure; EL07 = Transactions and Processing.

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).

In Figure 2 we present evolution of the annual value and frequency where each year of losses is partitioned by event type (normalized 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,

Figure 3. Loss and Frequency of Operational Losses by Region

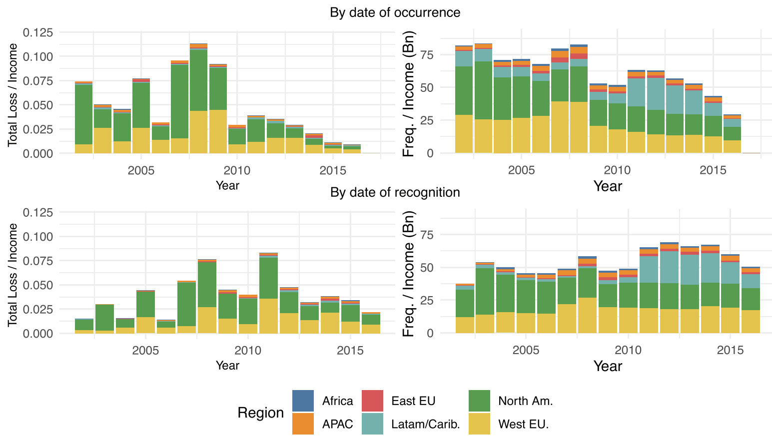

Note: On the left-hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right-hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of 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 the Caribbean; North Am: North America; and West EU: Western Europe.

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 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 world's largest banks are headquartered, which were particularly affected by the events leading up to, and after, the GFC.

Figure B.2 in Appendix B shows the losses and frequency but normalized 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 realized 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—D-SIBs and G-SIBs, respectively) for their role in events alleged to have taken place in the run-up to the crisis, such as the LIBOR (London interbank offered rate) scandal and the mis-selling of mortgage-backed securities.

3.5 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 suffer 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 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 Financial Stability Board'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 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.3 in Appendix B 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 led 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 on-site inspections are conducted, how efficiently the supervisor is 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 5, we look at the size of the bias by region. In Figure B.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 and 25 percent chance that an event is still unaccounted for after two years. To illustrate, if an event took place in a bank in Northern Europe on January 1, 2017, we estimate there is around 8 percent chance it has still not been accounted for in the books of the firm by January 1, 2019. As the curve in Figure B.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 2018:Q3. 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.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 percent, 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.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 seven quarters of data such that the series ends at 2016:Q4 (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.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 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 recognized 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 recognized 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:

$$\lambda(t_i|X_i) = \lambda_0(t) \exp(X_i\beta + FE), \tag{3}$$

where $\lambda_0(t)$ denotes the baseline hazard function, and 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 effect 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 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 with a hazard ratio 1.11 times higher than the baseline, i.e., the likelihood the event will be recognized 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 recognized, 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 recognize 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.

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

[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.

Table 4. Proportional Hazard Models with Supervisory Index

	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	

Note: 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.

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 standardized 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 percent 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, standardized 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.

4.1 Basic Indicator and Standardized 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 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 $\alpha = 0.15$.

Under the Basel II framework the TSA extends the BIA by adjusting the α terms for various bank business lines

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

(see Table B.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.

4.1.1 Basel III Standardized Approach

The standardized 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 euros, and five times average losses above 100 million euros. 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 2018b).

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 United States the focus was on internal and external loss data (Cruz, Peters, and Shevchenko 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 (LDA). 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.

4.2.1 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]

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

[17] Basel II rules require banks to calculate their regulatory capital requirement as the sum of expected and unexpected losses (i.e., the 99.9th 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.

Figure 4. Distribution of Losses and Risk Measures

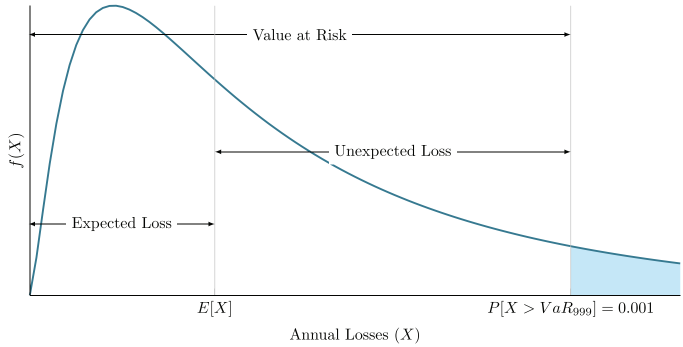

The VaR indicates the level of risk to which a firm, a portfolio, or a single position may be exposed over a given time period. Figure 4 displays an example of the distribution of annual losses, the relevant risk measures, and their location on the distribution.[18]

4.3 Evaluating Operational Risk Measures

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

[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 minimization of the possible shortfall and cost of capital (Cruz, Peters, and Shevchenko 2015).

Figure 5. Implied Capital by Various Approaches

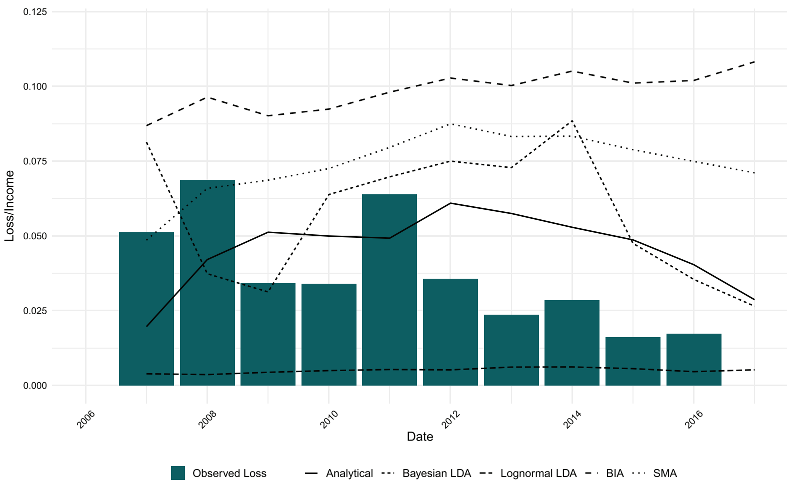

Note: This plot shows the observed gross loss per year (bars) against the operational risk predicted by the various models: the analytical approach, the lognormal LDA, the Bayesian LDA, the basic indicator approach, and a proxy of the standardized measurement approach. The y-axis shows the observed or predicted loss—divided by the total income in the sample.

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 five 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 compare 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 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 standardized 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 appendices, in addition to the confidence intervals of each measure, where possible (see Figure B.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 generalized 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 straightforward 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 minimize 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,

[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 (2018b) guidance.

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, and 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.

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, Curti, and Mihov (2017) find evidence that operational losses for U.S. 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 realization of operational losses. We let the analysis in Section 3.5 guide our choice for the length of lags. In particular, we are interested in the time at which losses materialize 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 percent 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 materialize 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 3. We provide a summary of the variables in Table 5.

We use the credit-to-GDP gap as a measure of the buildup 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 buildup 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 3—as an independent 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, Gambacorta, and Marques-Ibanez (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 buildup of operational risk losses. To do so, we use deviations of policy rates from implied Taylor-rule rates as

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

Table 5. Summary of Regression Variables

	N	Mean	Std. Dev.	Min.	Max.	Start Date	End Date
A. Quarterly Variables							
Loss per Income	598	0.0031	0.012	0	0.17	2002:Q1	2016:Q4
Frequency per Income (Millions)	600	0.0045	0.0055	0	0.035	2002:Q1	2016:Q4
Credit-to-GDP Gap	584	3.04	11.3	–33.2	35	2002:Q1	2016:Q4
Deviations from Taylor Rule	600	–1.29	2.58	–15.28	13.65	2002:Q1	2016:Q4
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

Note: 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.

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 minimize 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.

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:

$$\ln(Y_{it}) = \sum_{k} \beta_k X_{i,t-k} + \alpha_i + \gamma_t + \sum_{k} \epsilon_{i,t-k}, \qquad (4)$$

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 normalized 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.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 by 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 6 presents 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 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 percent 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 buildup of operational losses. A one-standard-deviation decrease in the Taylor gap is associated with a 20 percent increase of operational losses in the following four quarters and 28 percent after eight 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

[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.

Table 6. Operational Losses, Macroeconomic Conditions, and the Regulatory Environment

	Dependent Variable		
	$\frac{TotalLoss}{Income}$	$\frac{Frequency}{Income}$	$\frac{Severity}{Income}$
Panel A			
Credit-to-GDP Gap — 4 Lags	0.0094	0.010	–0.0011
	(0.011)	(0.011)	(0.0053)
Credit-to-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

Note: The table is divided into four panels summarizing 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 four lags the coefficient reported is $\sum_{i=1}^{4} \hat{\beta}_i$. A robust sum of standard errors is reported in parentheses. 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.

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 percent (26 percent) one year after. The cumulative effect of a 0.1 increase in two subsequent years rises in excess of 35 percent (20 percent) 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 their occurrence, as presumably they lead banks to implement better risk-management strategies.

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 (Kaffenberger, Kopp, and Wilson 2017; Kashyap and Wetherilt 2019). In March 2017, the G-20 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 2018a). 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 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

Table 7. Definitions of Cyber Event Types

Event Type Level 1	Event Type Level 2	Description
Internal Fraud	Unauthorized Activity	Example: rogue trading, unreported transaction, mis-marking positions
	Internal Theft	Example: forgery, theft, extortion, embezzlement, bribes/kickbacks
	System Security (Internal)	Intentional damage to systems by internal staff
External Fraud	External Theft and Fraud	Example: robbery, forgery, check kiting
	System Security (External)	Willful damage, e.g., hardware/software, hacking damage, theft of data
Technology and Infrastructure Failures		Losses arising from disruption of business or system failures

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

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. (2019). 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 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 categorization.

We first present summary statistics to provide a comparison of cyber losses with other operational losses. Table 8 presents statistics on the total number of incidents, mean, standard deviation, and maximum values, by cyber and non-cyber events. We provide

Table 8. Cyber Losses: Summary Statistics

	N	**Mean**	**Standard Deviation**	**Max.**
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
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

Note: 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.

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 percent. 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.

We also present a time series of frequencies and amounts in Figure 6, as well a breakdown of losses and frequency by region, "cyber" event types, and bank size, reported in Figures B.7–B.9 in Appendix B.[22] The dominating event type is "Technology and 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

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

Figure 6. Operational and Cyber Events

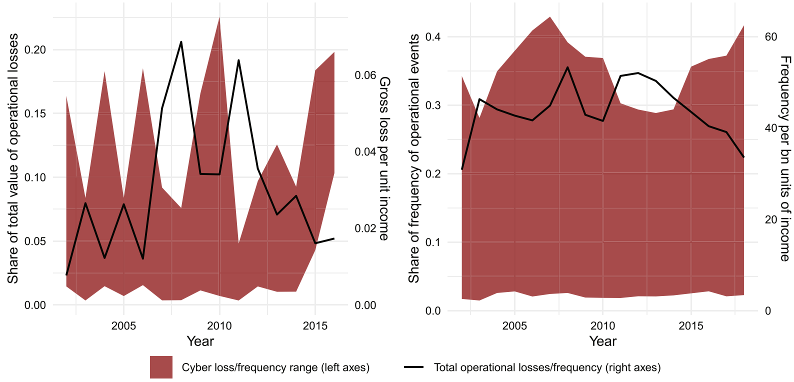

Note: In the left panel, the left-hand axis of the plot shows the estimated range of cyber losses across years as a share of all operational losses, shown by the red area in the graph. In the right panel, the left-hand axis of the plot shows the estimated range of the frequency of cyber incidents as a share of all operational loss incidents, shown by the red area in the graph. The right axis in the left (right) panel shows gross losses (frequency) per unit of income. Events are aggregated by the date of recognition.

paper, we show, using a different data set 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 United States. 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.

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 4, we also compute

Figure 7. Operational and Cyber Value-at-Risk

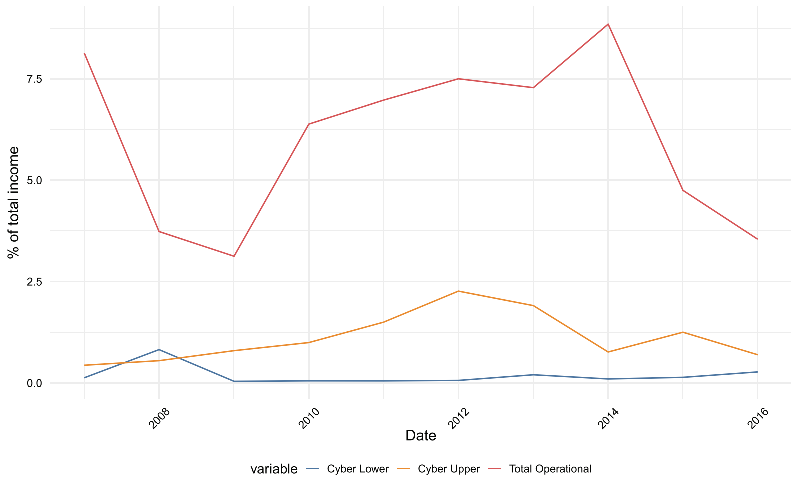

estimates of cyber risk capital. We perform a similar exercise, by computing the cyber risk capital over time, but focusing solely on the Bayesian methodology. We compute estimates for both the lower and upper bounds as defined above. The results are summarized in Figure 7, which includes the estimate 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, the value ranges between 0.04 and 0.8 percent of the gross income of the consortium, which corresponds to around 338 EUR million and 7.8 EUR billion, respectively. At the upper bound this can jump up to around 2.5 percent 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.

7. Conclusions

The GFC drew the attention of regulators and academics towards operational risk. Moreover, the shift to the new standardized 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 data set at the operational loss event level for over 14 years and more than 70 large banks.

We provide stylized 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 materialized only later. An example of such event is the mis-selling of mortgage-backed securities that took place around 2005/06 but was crystallized 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 percent to 7.5 percent 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 standardized 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 categorization 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.

Appendix A. Description of the Calculation of Capital

A.1 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 4. Alexander's alternative γ factors, labeled 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:

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

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

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

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 (A.1).

A.2 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 Markov chain Monte Carlo (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, Peters, and Shevchenko (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, Gao, and Giudici (2015) to estimate the annual loss distribution, considering a convolution between a generalized Pareto distribution for the mean loss (severity), with a Poisson distribution for the number of loss events (frequency), as in Chavez-Demoulin, Embrechts, and Nešlehová (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,

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

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} denotes 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):

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

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.

$$\lambda_i \sim \Gamma(\alpha, \beta) \qquad \text{(A.5)}$$

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

$$F_i \sim GPD(\mu, \xi, \sigma). \qquad \text{(A.6)}$$

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.

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

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.

Appendix B. Additional Tables and Figures

Table B.1. Overview of Business Lines Based on the Operational Risk Reporting Standards of ORX

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 and estates, investment advice, cards—credit and 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
Agency 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 that can only be categorized at corporate level

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

Table B.2. Overview of Regions and Sub-regions

Region	Sub-regions
North America	United States, Canada
Latin America and Caribbean	—
Eastern Europe	—
Western Europe	Southern Europe, Northern Europe, United Kingdom, Western Europe
Asia/Pacific	—
Africa	—

Table B.3. Summary of Durations by Region, Event Type, and Size (in days)

	t_1				t_2				t_3			
	Mean	**Std. Dev.**	**Median**	Q_{95}	**Mean**	**Std. Dev.**	**Median**	Q_{95}	**Mean**	**Std. Dev.**	**Median**	Q_{95}
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 and 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 and Process Management	255	648	2	1,689	144	396	7	793	399	780	49	2,182
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 and 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
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

Note: The table shows the summary of the different definitions of duration in our data. We present the mean, standard deviation, median, and the 95$^{\text{th}}$ quantile, by various breakdowns: in panel A by event type, panel B by region, and panel C by bank size.

Table B.4. Panel Regression of Contemporaneous Variables

Regressor	Dependent Variables		
	$\frac{Loss_{it}}{Income_{it}}$	$\frac{Freq_{it}}{Income_{it}}$	$\frac{Severity_{it}}{Income_{it}}$
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-to-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
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-to-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

Note: 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 normalized 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.5. Operational Losses and the Macroeconomic Environment, with Bias Adjustment

	Dependent Variable		
	$\frac{\textit{TotalLoss}}{\textit{Income}}$	$\frac{\textit{Frequency}}{\textit{Income}}$	$\frac{\textit{Severity}}{\textit{Income}}$
Panel A			
Credit-to-GDP Gap — 4 Lags	0.0061	0.0061	0.00064
	(0.0096)	(0.0096)	(0.0047)
Credit-to-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

Note: The table is divided into two panels summarizing 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 four lags the coefficient reported is $\sum_{i=1}^{4} \hat{\beta}_i$. A robust sum of standard errors is reported in parentheses. 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.

Figure B.1. Sample Size and Frequency of Events

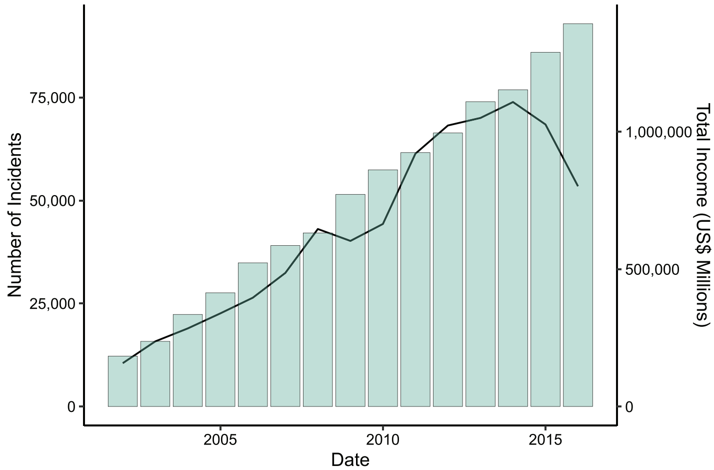

Note: The plot shows the total number of incidents per year alongside the total income of the consortium. The bars denote the total income (right axis) and the line denotes the frequency of incidents (left axis).

Figure B.2. Loss and Frequency over Time Partitioned by Bank Size

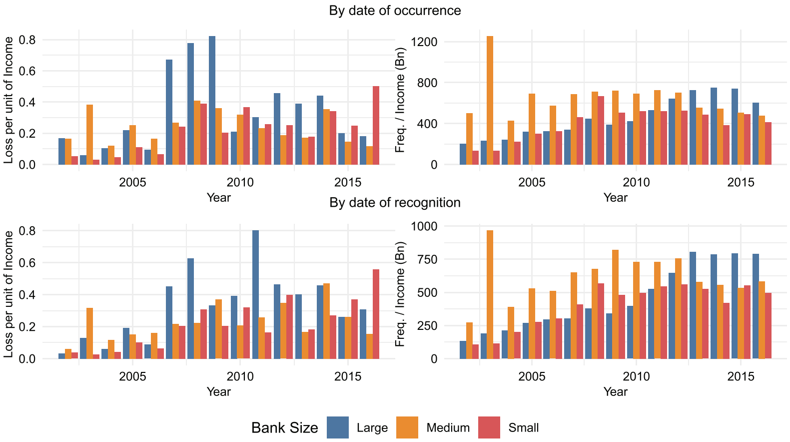

Figure B.3. Confidence Intervals for VaR

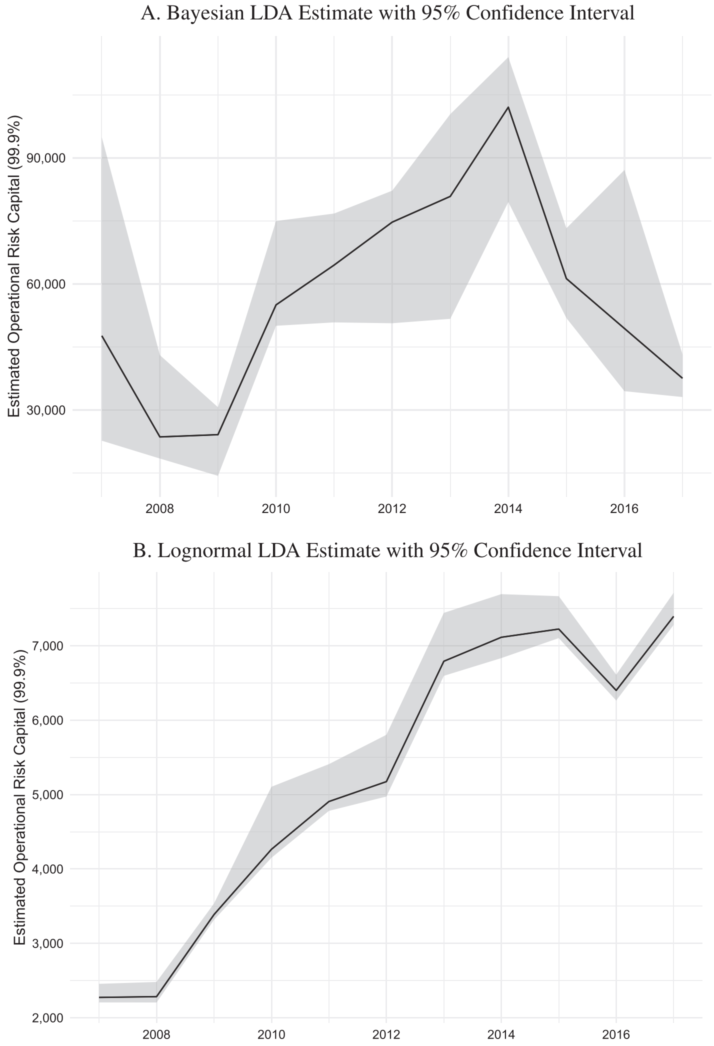

Note: The plot shows the estimated operational risk capital by two different methodologies and the 95 percent confidence interval for the location of the 99/9 percent quantile of the annual loss distributions. These are calculated by using the approximation put forward in Cruz, Peters, and Shevchenko (2015). The upper or lower bound can be calculated as $B = K\alpha \pm F_N^{-1}\sqrt{(K\alpha(1-\alpha))}$, 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.4. Estimated Survival Curves by Region

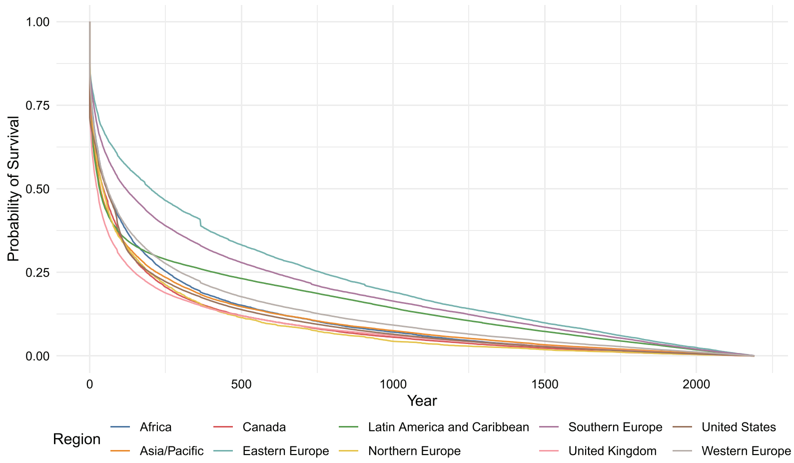

Figure B.5. Estimated Bias Factor by Region

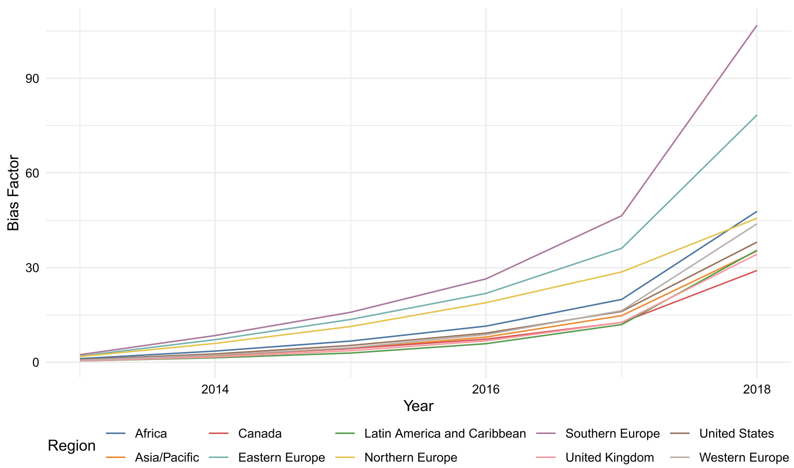

Figure B.6. Annual Frequencies Adjusted for Data Bias by Region

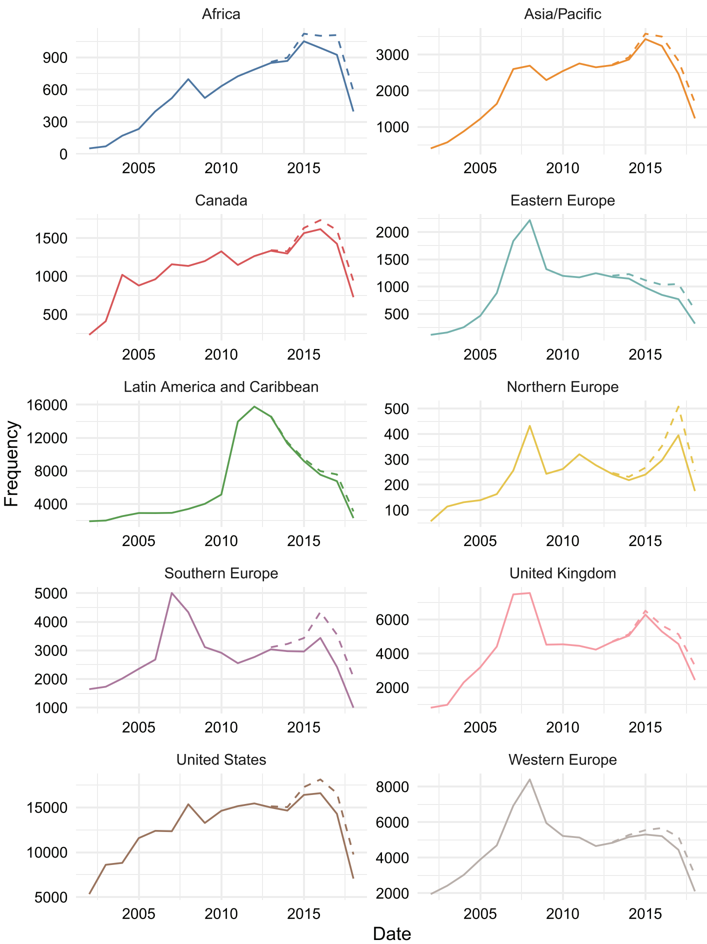

Figure B.7. Loss and Frequency of Cyber Losses by Event Type

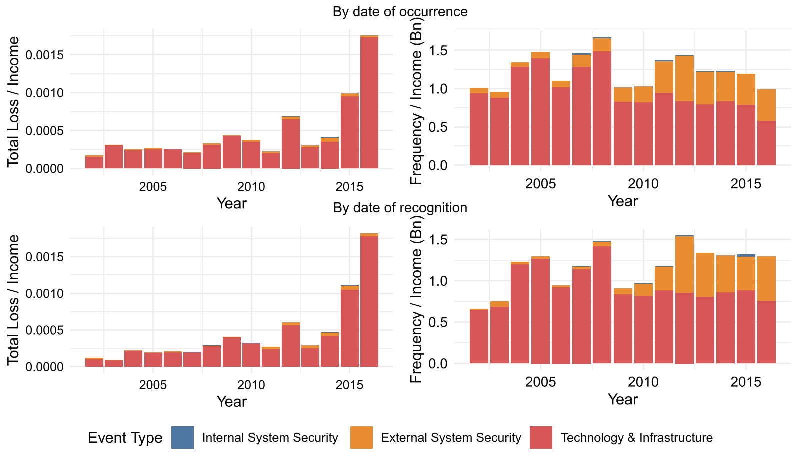

Note: On the left-hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right-hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of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

Figure B.8. Loss and Frequency of Cyber Losses by Region

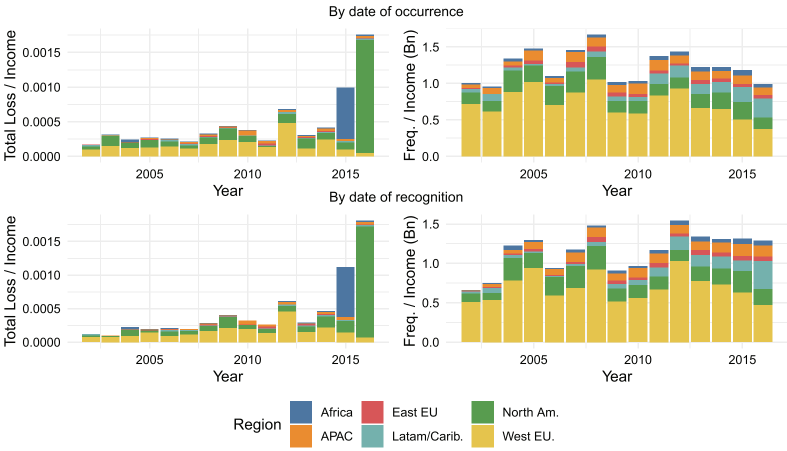

Note: On the left-hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right-hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of 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 the Caribbean; North Am: North America; and West EU: Western Europe.

Figure B.9. Loss and Frequency of Cyber Losses by Bank Size

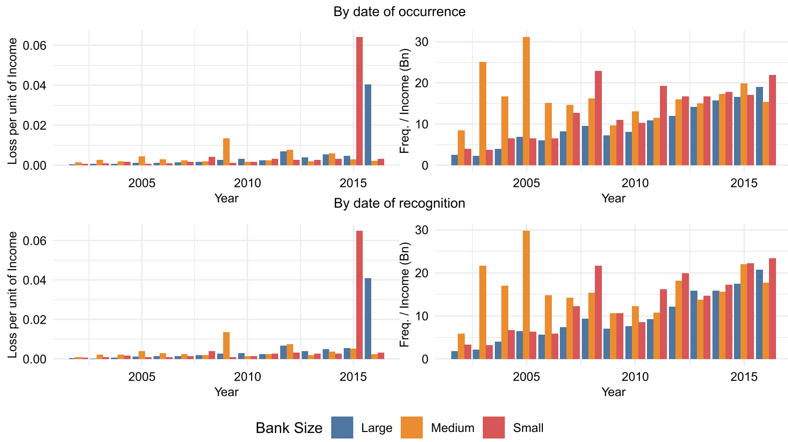

Note: On the left-hand side of the quadrant of plots we show the total value of losses per year divided by the total consortium annual income. On the right-hand side we display the frequency divided by income (in billions). The upper panel of the quadrant of plots shows incidents aggregated by date of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

References

Abdymomunov, A., F. Curti, and A. Mihov. 2017. "US Banking Sector Operational Losses and the Macroeconomic Environment." Technical Report. Available at SSRN 2738485.

Abiad, A., E. Detragiache, and T. Tressel. 2010. "A New Database of Financial Reforms." *IMF Staff Papers* 57 (2): 281–302.

Admati, A. R. 2016. "The Missed Opportunity and Challenge of Capital Regulation." *National Institute Economic Review* 235 (February): R4–R14.

Aldasoro, I., L. Gambacorta, J. Frost, and D. Whyte. 2021. "Covid-19 and Cyber Risk in the Financial Sector." BIS Bulletin No. 37 (January 14).

Aldasoro, I., L. Gambacorta, P. Giudici, and T. Leach. 2022. "The Drivers of Cyber Risk." *Journal of Financial Stability* 60 (June): Article 100989.

Alexander, C. 2008. "Statistical Models of Operational Loss." In *Handbook of Finance*, Vol. III, ed. F. J. Fabozzi. Wiley.

Allen, F., and D. Gale. 2004. "Competition and Financial Stability." *Journal of Money, Credit and Banking* 36 (3): 453–80.

Allen, L., and T. G. Bali. 2007. "Cyclicality in Catastrophic and Operational Risk Measurements." *Journal of Banking and Finance* 31 (4): 1191–1235.

Altunbaş, Y., L. Gambacorta, and D. Marques-Ibanez. 2014. "Does Monetary Policy Affect Bank Risk?" *International Journal of Central Banking* 10 (1, March): 95–135.

Altunbaş, Y., J. Thornton, and Y. Uymaz. 2018. "CEO Tenure and Corporate Misconduct: Evidence from US Banks." *Finance Research Letters* 26 (September): 1–8.

Ames, M., T. Schuermann, and H. S. Scott. 2015. "Bank Capital for Operational Risk: A Tale of Fragility and Instability." *Journal of Risk Management in Financial Institutions* 8 (3): 227–43.

Antonini, G., E. W. Cope, G. Mignola, and R. Ugoccioni. 2009. "Challenges and Pitfalls in Measuring Operational Risk from Loss Data." *Journal of Operational Risk* 4 (4): 3–27.

Artzner, P., F. Delbaen, J. M. Eber, and D. Heath. 1999. "Coherent Measures of Risk." *Mathematical Finance* 9 (3): 203–28.

Basel Committee on Banking Supervision. 2003. "Basel II: The New Basel Capital Accord." Consultative Document, Bank for International Settlements.

———. 2017. "High-level Summary of Basel III Reforms." Technical Report, Bank for International Settlements.

———. 2018a. "Cyber-resilience: Range of Practices." Bank for International Settlements.

———. 2018b. "Standardised Measurement Approach for Operational Risk." Consultative Document, Bank for International Settlements.

Berger, A. N., F. Curti, A. Mihov, and J. Sedunov. 2018. "Operational Risk is More Systemic than You Think: Evidence from US Bank Holding Companies." Available at SSRN 3210808.

Biell, L., and A. Muller. 2013. "Sudden Crash or Long Torture: The Timing of Market Reactions to Operational Loss Events." *Journal of Banking and Finance* 37 (7): 2628–38.

Boer, M., and J. Vazquez. 2017. "Cyber Security and Financial Stability: How Cyber-attacks Could Materially Impact the Global

Financial System." Technical Report, Institute of International Finance.

Bogdanova, B., and B. Hofmann. 2012. "Taylor Rules and Monetary Policy: A Global 'Great Deviation'?" *BIS Quarterly Review* (September): 37–49.

Boone, J. 2008. "A New Way to Measure Competition." *Economic Journal* 118 (531): 1245–61.

Bouveret, A. 2018. "Cyber Risk for the Financial Sector: A Framework for Quantitative Assessment." IMF Working Paper No. 18/143.

Byrne, B., J. Coughlan, and S. V. Tilley. 2017. "An Empirical Analysis of the Impact of Fines on Bank Reputation in the US and UK." Available at SSRN 2980352.

Cabras, S., and M. E. Castellanos. 2007. "A Default Bayesian Procedure for the Generalized Pareto Distribution." *Journal of Statistical Planning and Inference* 137 (2): 473–83.

Carrivick, L., and E. W. Cope. 2013. "Effects of the Financial Crisis on Banking Operational Losses." *Journal of Operational Risk* 8 (3): 3–29.

Cerasi, V., S. M. Deininger, L. Gambacorta, and T. Oliviero. 2020. "How Post-crisis Regulation has Affected Bank CEO Compensation." *Journal of International Money and Finance* 104 (June): 102–53.

Chavez-Demoulin, V., P. Embrechts, and J. Nešlehová. 2006. "Quantitative Models for Operational Risk: Extremes, Dependence and Aggregation." *Journal of Banking and Finance* 30 (10): 2635–58.

Chernobai, A., P. Jorion, and F. Yu. 2011. "The Determinants of Operational Risk in US Financial Institutions." *Journal of Financial and Quantitative Analysis* 46 (6): 1683–1725.

Čihák, M., and K. Schaeck. 2010. "Competition, Efficiency, and Soundness in Banking: An Industrial Organization Perspective." Discussion Paper No. 2010-205, European Banking Centre.

Cope, E. W., M. T. Piche, and J. S. Walter. 2012. "Macroenvironmental Determinants of Operational Loss Severity." *Journal of Banking and Finance* 36 (5): 1362–80.

Cornalba, C., and P. Giudici. 2004. "Statistical Models for Operational Risk Management." *Physica A: Statistical Mechanics and Its Applications* 338 (1–2): 166–72.

Cox, D. R. 1972. “Regression Models and Life-tables.” *Journal of the Royal Statistical Society: Series B (Methodological)* 34 (2): 187–202.

Cruz, M. G., G. W. Peters, and P. V. Shevchenko. 2015. *Fundamental Aspects of Operational Risk and Insurance Analytics: A Handbook of Operational Risk.* John Wiley & Sons.

Cummins, J. D., C. M. Lewis, and R. Wei. 2006. “The Market Value Impact of Operational Loss Events for US Banks and Insurers.” *Journal of Banking and Finance* 30 (10): 2605–34.

Curti, F., W. S. Frame, and A. Mihov. 2019. “Are the Largest Banking Organizations Operationally More Risky?” Available at SSRN 3210206.

Curti, F., J. Gerlach, S. Kazinnik, M. Lee, and A. Mihov. 2019. “Cyber Risk Definition and Classification for Financial Risk Management.” Mimeo, Federal Reserve Bank of St. Louis (August).

Dalla Valle, L., and P. Giudici. 2008. “A Bayesian Approach to Estimate the Marginal Loss Distributions in Operational Risk Management.” *Computational Statistics and Data Analysis* 52 (6): 3107–27.

De Nicolò, G., and M. Lucchetta. 2013. “Bank Competition and Financial Stability: A General Equilibrium Exposition.” CESifo Working Paper No. 4123.

Denk, O., and G. Gomes. 2017. “Financial Re-regulation Since the Global Crisis?” OECD Economics Department Working Paper No. 1396, Organisation for Economic Co-operation and Development.

Dingel, J. I., and B. Neiman. 2020. “How Many Jobs Can Be Done at Home?” Working Paper No. 26948, National Bureau of Economic Research.

Duffie, D., and J. Younger. 2019. “Cyber Runs.” Hutchins Center Working Paper No. 51, Brookings Institution.

Eisenbach, T. M., A. Kovner, and M. J. Lee. 2021. “Cyber Risk and the U.S. Financial System: A Pre-mortem Analysis.” *Journal of Financial Economics* 145 (3): 802–26.

Eshraghi, A., J. Hagendorff, and D. D. Nguyen. 2016. “Can Bank Boards Prevent Misconduct?” *Review of Finance* 20 (1): 1–36.

European Systemic Risk Board. 2015. "Report on Misconduct Risk in the Banking Sector." Report, European System of Financial Supervision.

———. 2020. "Systemic Cyber Risk." Report, European System of Financial Supervision.

Facchinetti, S., P. Giudici, and S. A. Osmetti. 2020. "Cyber Risk Measurement with Ordinal Data." *Statistical Methods and Applications* 29 (1): 173–85.

Fich, E. M., and A. Shivdasani. 2007. "Financial Fraud, Director Reputation, and Shareholder Wealth." *Journal of Financial Economics* 86 (2): 306–36.

Figini, S., L. Gao, and P. Giudici. 2015. "Bayesian Operational Risk Models." *Journal of Operational Risk* 10 (2).

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.

Gillet, R., G. Hübner, and S. Plunus. 2010. "Operational Risk and Reputation in the Financial Industry." *Journal of Banking and Finance* 34 (1): 224–35.

Heid, F. 2007. "The Cyclical Effects of the Basel II Capital Requirements." *Journal of Banking and Finance* 31 (12): 3885–3900.

Hess, C. 2011. "The Impact of the Financial Crisis on Operational Risk in the Financial Services Industry: Empirical Evidence." *Journal of Operational Risk* 6 (1): 23.

Jarrow, R. A. 2008. "Operational Risk." *Journal of Banking and Finance* 32 (5): 870–79.

Kaffenberger, L., E. Kopp, and C. Wilson. 2017. "Cyber Risk, Market Failures, and Financial Stability." IMF Working Paper No. 17/185.

Kashyap, A. K., and A. Wetherilt. 2019. "Some Principles for Regulating Cyber Risk." *AEA Papers and Proceedings* 109 (May): 482–87.

Kim, J. 2018. "Bank Competition and Financial Stability: Liquidity Risk Perspective." *Contemporary Economic Policy* 36 (2): 337–62.

Köster, H., and M. Pelster. 2017. "Financial Penalties and Bank Performance." *Journal of Banking and Finance* 79 (June): 57–73.

Liao, G., Y. Ma, and P. Sands. 2018. "Rethinking Operational Risk Capital Requirements." *Journal of Financial Regulation* 4 (1): 1–34.

Migueis, M. 2018. "Forward-Looking and Incentive-Compatible Operational Risk Capital Framework." *Journal of Operational Risk* 13 (3).

Monetary Authority of Singapore. 2018. "Financial Stability Review." Technical Report, Monetary Authority of Singapore.

Peters, G., P. V. Shevchenko, B. Hassani, and A. Chapelle. 2016. "Should the Advanced Measurement Approach Be Replaced with the Standardized Measurement Approach for Operational Risk?" *Journal of Operational Risk* 11 (3): 1–49.

Power, M. 2005. "The Invention of Operational Risk." *Review of International Political Economy* 12 (4): 577–99.

Romanosky, S. 2016. "Examining the Costs and Causes of Cyber Incidents." *Journal of Cybersecurity* 2 (2): 121–35.

Sakalauskaite, I. 2018. "Bank Risk-taking and Misconduct." Working Paper, University of Amsterdam and the Tinbergen Institute.

Shih, J., A. Samad-Khan, and P. Medapa. 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 and Organization* 85 (January): 191–206.

Tarullo, D. K. 2008. "Banking on Basel: The Future of International Financial Regulation." Peterson Institute for International Economics.