On the wicked problem of quantifying and managing non-financial risks and the role of digital technology in providing solutions

Received (in revised form): 30th September, 2023

Tom Butler

Professor, Business Information Systems, University College Cork, Ireland

Dr Tom Butler is a professor of information systems and regulatory technologies at University College Cork. Tom was the Principal Investigator of Ireland's GRC Technology Centre and led a multidisciplinary team of knowledge engineers, information systems and legal researchers conducting research and development on artificial intelligence technologies to enable financial institutions to manage better regulatory compliance and operational risk. Tom has authored 227 publications, including 82 full papers, 84 conference papers and 11 inventions.

Business Information Systems, University College Cork, Cork, Ireland
E-mail: tbutler@ucc.ie

Robert Brooks

European Managing Director, Accenture, Cyber Risk and Regulation, UK

Robert Brooks is a managing director at Accenture. He has spent the past 30 years working in financial services across credit and commodity operations, trading risk management, operational risk management and non-financial risk. His focus is technology-led business transformations. Currently, he advises large corporations, including financial institutions, on operational risk and digital resilience, with a particular emphasis on cybersecurity. He holds a degree in law, a masters in computing and qualifications in accounting (ACMA), risk management (PRM) and quantitative finance (CQF).

Accenture, Cyber Risk Regulation, Pulborough, UK
E-mail: robert.baker-brooks@accenture.com

Abstract The management of operational risk in financial institutions has all the characteristics of a 'wicked problem'. Certainly, the repeated efforts of the Bank of International Settlements, (BIS) Basel Committee on Banking Supervision (BCBS) to have banks control and mitigate their operational risks speak to the tractability of extant approaches to addressing them effectively. The original 'Principles for the Sound Management of Operational Risk'[1] and its recent revisions,[2] the BCBS 'Principles for Effective Risk Data Aggregation and Risk Reporting'[3] and the 'Principles for Operational Resilience',[4] collectively offer a sound foundation for addressing this enduring problem. Why then are solutions so elusive for banks to implement? This paper first outlines the institutional and social web of conditions and factors that contribute to the existence of this 'wicked problem'. It then identifies how AI-based digital technologies can once and for all effectively address the problem of operational risk in large banks. Nevertheless, as powerful as today's digital technologies are, they require an organising vision, particularly if they are to contribute to the management of operational risk. This paper informs such a vision and identifies a comprehensive artificial intelligence-based digital architecture to realise it.

Keywords: *operational risk, wicked problem, digital technology, artificial intelligence, enterprise data fabric*

INTRODUCTION

Following the financial crisis of 2007–8, the growth in the volume, velocity, variety and complexity of regulations, new business opportunities and the evolution of information and communication technologies (ICT), led to the emergence of FinTech, RegTech and RiskTech paradigms.[5,6] Much of the innovation in developing and applying novel digital technologies since 2008 has come from some large established banks, BigTech firms and innovative tech start-ups from Silicon Valley and other digital technopoles globally.

The same period saw significant technological advances in the fundamental digital technologies on which financial innovation and risk and compliance technologies could be based, whether it is artificial intelligence (AI), blockchain/distributed ledger technology, smart contracts, the Internet of Things or the Cloud.[7] The processing power, bandwidth and speed of operation in executing instructions of fundamental ICT such as computers (CPUs and memory) and internetworking technologies, also saw improvements in line with Moore's Law. In other words, in 2023, digital technologies enable effective and efficient digitalisation of every aspect of human endeavour, be it business or personal. Having conducted research on governance, risk and compliance (GRC) in the financial industry since 2012, the lead author found no technological barriers to the solution of the industry's enduring problems with operational risk — indeed, this has been the case since before the pandemic. However, the pace of change and innovation in this important area of practice was noted as glacial by the European Commission's Regulatory Obstacles to Financial Innovation Expert Group (ROFIEG).[8] As argued in 2023 by Peter Hughes, large financial institutions appear to be married to a pre-crisis, analogue-era model in managing and accounting for non-financial risk.[9] Hence, he provides a convincing case for a digitally-enabled risk accounting approach to properly measure risk exposure in quantitative and qualitative terms. We contend that with few exceptions, financial institutions have failed to adequately leverage the power of digital technology to properly control operational risk and make their business enterprises resilient.

Before identifying how non-financial risks can be managed using digital technologies, the true scale of what is a 'wicked problem' is discussed in the second section. The third section then summarises the causes of this problem, while the fourth identifies AI-enabled digital technologies that address the enduring problems of operational risk management. The final section then offers several conclusions.

OPERATIONAL RISK AS A 'WICKED PROBLEM': PRACTICES AND CONSEQUENCES

Researchers argue that financial institutions are beset by 'wicked problems'.[10] Millan and Overall cite the seminal work of Rittel and Webber[11] and the observation of C. West Churchman to define 'wicked problems' succinctly as 'a class of social system problems, which are ill-formulated; where the information is confusing; where there are many clients and decision makers with conflicting values; and where the ramifications in the whole system are thoroughly confusing'.[12] Echoing Churchman and others, Brian Head argues that uncertainty, complexity and diversity characterise wicked problems.[13] It does not require much reflection to conclude that while complexity and uncertainty are associated with all types of risk, the divergence of objectives, values, frames of reference, interests, commitments and viewpoints of banking and risk managers across three lines of defence, present particular challenges to addressing systematically the challenges posed by operational risk.

How accurate are the reported losses from operational risk?

The losses from operational risk in financial institutions are significant. However, recent research reports that operational risks have consequences for the entire financial system, threatening its stability and incurring losses beyond those associated with operational risks.[14] How big are the reported losses associated with operational risk events? The Basel III Monitoring Report 2023 states:

> In total, €522.6 billion of gross and €470.8 billion of net operational risk losses have been reported over the past 10 years. Operational risk gross losses were €70.5 billion in 2012 and peaked in 2014 at €81.0 billion. Since then, gross losses have decreased significantly to €29.5 billion in 2021, the lowest value of the past 10 years. This trend was observed in 2021 despite the COVID-19 pandemic.[15]

This downward trend reflects the magnitude of operational risk losses associated with the Global Financial Crisis and the fact that many of the losses reported up to 2014 reflect the time lag between discovery, reporting and settlement. In 2021, the Operational Risk Exchange (ORX) reported €20.3bn in losses due to operational risk events for its 82 members.[16] By 2023, ORX declared 'an average decrease of €1.6bn per year, and a drop of €8.2bn total gross loss between 2017 and 2022'.[17] Thus, the overall trend is reassuring indicating that banks are either doing something right in addressing the operational risk problem or are doing nothing wrong in the period observed. However, the 2023 ORX report indicates new threats as 'more sophisticated technology and low-risk fraud had led to a reported loss number of 76,620 events in 2022, a 26.4% rise of 16,020 from the previous year',[18] back to levels reported in 2017.

There is one anomaly between the ORX and the Basel III Monitoring Report that requires comment. While ORX members include many global systemically important banks (G-SIBs), a difference of just €9.5bn in 2021 to account for the operational risk losses in the remainder of the entire financial industry covered in the Basel III Monitoring Report is problematic and places a question mark over the accuracy of the reported figures. So, are the losses reported to the ORX more accurate than those stated in the Basel III Monitoring Report 2023? As one of the reviewers of this paper commented, perhaps the anomaly exists due to timing differences in reporting losses to the BCBS and ORX. That may be so, nevertheless, Mark Cooke, former Group Head of Operational Risk at HSBC and former Chairman of ORX, casts doubt on the losses currently reported by ORX. He argues 'that existing risk reporting systems are simply failing to cope with "the new normal" and that risk events are going unreported and — worse — undetected altogether'.[19] If Cooke is correct, then reported losses need to be viewed critically as they are extremely difficult to corroborate, as regulators and researchers alike note the difficulty in 'obtaining reliable data on operational risks and losses'.[20] In preparing this paper, it was difficult to obtain accurate, credible figures. Nevertheless, two research papers by regulators give pause for thought in terms of the levels of losses that can be experienced, which are certainly worrying given the increase in loss events reported in the 2023 ORX report.

A recent paper by regulators at the US Federal Reserve Bank finds that

> on average, the [banks] in our sample lose $185 million or the equivalent of 0.026% of their assets per quarter to operational risk . . . On average, 231 operational loss events with an average severity of $0.5 million occur at an institution over a given quarter.[21]

The data on which this analysis is based is historical, but the figures are significant. A more recent study by the US Federal Reserve Bank of New York of 35 large banks, including GSIBs, participating in the Dodd–Frank Act Stress Test found that they had 'severely adverse scenario projected operational risk losses [. . .] of $135 billion, or 23 percent of the $578 billion in aggregate losses projected for these firms over the nine quarters ending in March of 2020'.[22]

The failure and rescue of Credit Suisse (CS) in 2023 is an example of why complacency in addressing operational risk is foolish and the research by the US Federal Reserve Bank needs to be taken seriously. CSs operational risk failures included money laundering, misconduct involving corruption, tax evasion and corporate espionage. Losses included $1.1bn for stockholders as 'UBS stockholders received a wealth transfer from CS stockholder', with the cost to the Swiss taxpayer an estimated $6–7bn.[23] However, these are first-order losses; what of second-order effects and the losses then imposed on customers, suppliers and society? In Credit Suisse's case, staff lost over $400m of bonuses,[24] while 30,000–35,000 staff are being

made redundant by UBS, with the estimated costs of the merger at $17bn.[25] The conclusions of this research indicate that the size of a bank leads to greater operational risk and losses due to (1) risks related to the complexity of its structure and processes; (2) the effects of moral hazard related to the 'too-big-to-fail' perception and implicit government guarantees; and (3) innovation-related risks. Thus, it is argued that the actual, as opposed to the reported losses, may be much larger than assumed.

FAILED PEOPLE, PROCESSES AND SYSTEMS: THE INSTITUTIONAL RESPONSE TO OPERATIONAL RISK

The Basel Committee on Banking Supervision (BCBS) defines operational risk in terms of the financial losses arising from inadequate or failed internal processes, people and systems or from external events.[26] The current operational risk paradigm is heavily critiqued by Peter Hughes in his book 'Where Next for Operational Risk? A Guide for Risk Managers and Accountants'.[27] Paradoxically, Hughes' account of the extant approach to dealing with operational risk in banking involves the failure of people, processes and systems tasked with addressing the problem.[28]

Consequences of the complexity of IT systems and data architectures

In the period since operational risk was institutionalised in the 1990s,[29] the magnitude and complexity of this type of risk have increased as ICT has advanced and become more ubiquitous in automating and informating business activities within and across banks and financial markets. Today's G-SIBs grew largely through mergers and acquisitions, simply appending ICT infrastructures with the equivalent of technological duct tape and without federating and integrating data architectures. Thus, as processes and systems increased in complexity, so too did investment products of the 'exotic' category such as 'derivative', 'synthetic' and 'structured' instruments. According to Allan Grody and Peter Hughes:

> Substantial concentrations of risk are now a permanent feature of banks whose operating environments are invariably comprised of highly complex risk management ecosystems within similarly complex information technology infrastructures.[30]

The institution of the 'Principles for Effective Risk Data Aggregation and Risk Reporting'[31] by the Basel Committee on Banking Supervision was a regulatory response to have 34 G-SIBs address the ICT infrastructure and data architecture complexity problems from a risk management perspective. However, surveys published in 2015, 2018 and 2020 revealed that most G-SIBs were slow in implementing and achieving the principles.[32] In 2020, the survey concluded that 'As of the end of 2018, none of the banks are fully compliant with the BCBS 239 principles, as attaining the necessary data architecture and IT infrastructure remains a challenge for many'.[33]

In 2021 and 2022 the UK's Prudential Regulatory Authority (PRA) voiced its concerns regarding the reliability of the data used for regulatory reporting. In 2021, the 'Dear CEO' letter of the PRA states:

> Overall, we were disappointed to find significant deficiencies in a number of firms' processes used to deliver accurate and reliable regulatory returns . . . For some firms, there had been a historic lack of focus, prioritisation, and investment in this area.[34]

In 2022, the 'Dear CEO' letter underlined that 'deficiencies in banks' risk management governance and frameworks, many of which were symptoms of a broader root cause and manifestations of an inappropriate internal risk culture where lessons from the global financial crisis had not been sufficiently learnt'.[35]

The root cause of the risk and compliance reporting data management problem

In 2012, Andrew Haldane, former Chief Economist at the Bank of England, identified the central issues confronting G-SIBs' inability to manage risk. In his seminal 'Tower of Babel' paper he states that finance

> has no common language for communicating financial information. Most financial firms have competing in-house languages, with information systems silo-ed by business line. Across firms, it is even less likely that information systems have a common mother tongue. Today, the number of global financial languages very likely exceeds the number of global spoken languages. The economic costs of this linguistic diversity were brutally exposed by the financial crisis. Very few firms, possibly none, had the information systems necessary to aggregate quickly information on exposures and risks. This hindered effective consolidated risk management. For some of the world's biggest banks that proved terminal, as unforeseen risks swamped undermanned risk systems.[36]

This is fast becoming a digital 'Tower of Babel' which is impacting risk management and compliance reporting of all banks. Adequate risk governance is frustrated not only by a common language problem but also by the siloed nature of risk supervision and audit functions due to the nature of the three lines of defence model (3LoD).[37]

Deficiencies in operational risk management frameworks and information systems

In line with regulatory requirements for implementing a risk management framework (RMF), operational risk information systems are designed to serve the needs of the operational risk function in supporting periodic activities related to oversight, supervision and reporting. Any value-adding activities provided by operational risk functions tend to be *ad hoc* and focus on decision support around change management and new product implementation. The dominance of the extant operational risk paradigm is such that the digitalisation of financial services operations appears to have completely neglected the need to support the information needs of business managers and professionals in the first line of defence. So, what type of data and information are required?

Those critical of the existing paradigm believe that firms need to focus less on loss data and more on operational risk context and event information.[38] However, many practitioners focus also on near-miss events and learn from these. Nevertheless, the literature is critical of certain metrics and approaches employed in operational risk management frameworks. Firstly, risk and control self-assessments (RCSAs) are one of the cornerstones of the extant paradigm's risk management frameworks for risk identification. However, practitioners increasingly argue that they are prone to biases in supervised entities in business. Researchers point out that 'Whereas assessment based metrics can provide a vital source of risk intelligence at the operating unit level, they are inherently subjective and are not aggregatable or comparable along the vertical and horizontal dimensions of an enterprise'.[39]

Secondly, risk functions typically use a Red, Amber, Green (RAG) rating to categorise quantitative and qualitative unstructured data and assess threats and vulnerabilities. The RAG approach is widely criticised because it is

> inherently subjective and [colours] not aggregatable or comparable along the vertical and horizontal dimensions of an enterprise . . . For all practical purposes, the measurement of operational risk has been deferred by defining it in terms of a "qualitative" assessment process rather than a "quantitative" measurement process. This has left financial institutions to ponder how to link operational risk exposure to their frequency and severity measures of operational losses. If available (and not much is yet available) then operational risk loss data is rather inelegantly utilized to determine the parameters of a typically poorly articulated model for calculating the minimum capital requirement.[40]

Thirdly, another weapon in the RMF arsenal is scenario analysis, which is useful in the hands of the risk-aware professional. However, it suffers from the same issues as RCSAs because actors typically underestimate or misjudge risks due to availability bias; fear or over-optimism contaminates intuitive judgments of and risk perceptions based; stereotypical thinking and a failure to understand randomness; confirmatory bias and rejection of evidence that challenges their views; finally they also underestimate their ability to control risks.[41]

Finally, key risk indicators (KRIs) are another problematic mechanism as researchers argue that

financial institutions 'are generally not satisfied with their KRI frameworks. One reason for this could be the limited guidelines available on establishing a set of effective indicators that reflect and monitor operational risk exposure in an efficient manner'.[42]

CORE DIGITAL TECHNOLOGIES FOR NEXT-GENERATION OPERATIONAL RISK MANAGEMENT

Financial institutions have not, generally, implemented the type of digital risk information systems required to address the complete information needs of the first and second lines of defence in addressing operational risk. The information needs of the risk function are generally catered for in siloed GRC information systems either on-premises or as Software-as-a-Service in the Cloud. However, some researchers argue that first-line business managers are unaware of the risks they incur, whether through incompetence, error, misconduct or bias.[43] The exception here, is the IT function which manages risk and compliance of ICT operations as an integral activity in achieving its business objectives. However, these systems are, for the most part, siloed as well. Nevertheless, another significant difference is that the IT function operates with common languages defined in software, hardware and communication technology standards.

In order to understand better how digital technologies can support efficient, effective and timely operational risk management, following the European Central Bank (ECB) it is proposed that the focus should be on digital technologies that underpin core capabilities such as data architecture governance and quality, risk data aggregation, risk management (identification, assessment, measurement, monitoring and control)[44] and reporting capabilities.[45] These 'are deemed essential preconditions for proper risk governance and sound risk-based decision-making and necessitate state-of-the-art IT infrastructure'.[46]

Solving the 'wicked problems' of managing operational risk using digital technologies that are 'state-of-the-art' would not be easy for many organisations but may be quite difficult for G-SIBs given the unquestioning commitments to the extant risk management paradigm.

Comprehending, adopting and implementing digital technology: The role of an organising vision

In 2019, Accenture's McIntyre and Skan reported that over the three years to 2019, the 161 largest retail and commercial banks spent an estimated $1tn on digital technology.[47] Obtaining business value or a satisfactory return on investment from digital initiatives eludes many organisations. McIntyre and Skan provide valuable insights into the financial industry's performance in this regard and estimate that just 12 per cent of these financial institutions were leveraging the benefits of digital transformation at scale. According to Werth *et al.*, the digital transformation of much of the financial industry is evolutionary rather than disruptive.[48] Furthermore, the pace of digital transformation in the financial industry is glacial, as indicated by the European Commission Expert Group.[49] Something is clearly amiss here: In a risk context, take the glacial progress also in implementing digital technology to address the risk data aggregation problem.

Given the significant operational losses indicated earlier and the problematic labour-intensive processes in managing operational risk, implementing controls and regulatory compliance reporting, it may be concluded that banks have yet to develop appropriate organising visions to manage operational risk using digital technology and make crucial operations resilient. In their seminal paper in *Organisation Science*, Swanson and Ramiller define 'an organizing vision is a focal community idea for the application of information technology in organizations'.[50] It is a shared, cognitive view of how digital technologies enable success in information systems innovation in firms. It helps frame organisational innovation based on digital technologies by providing a focus for its interpretation and related sensemaking, thereby influencing decision-makers in the adoption process and helping to legitimise the implementation and assimilation of the technology. Finally, it facilitates commitment and the mobilisation of necessary resources. In the authors' experience, G-SIBs generally speaking, and with few exceptions, do not possess, let alone apply, an organising vision for enterprise digital transformation that incorporates operational risk management needs.

What state-of-the-art information technologies can banks use to digitise and enhance operational risk management capabilities?

The European Commission Expert Group on Regulatory Obstacles to Financial Innovation (ROFIEG) reviewed and analysed state-of-the-art information technologies to identify those that had the greatest potential for innovation in financial services, including risk management and compliance reporting.[51] The first author of this paper was a member of this group. This group of regulatory, industry and academic experts identified AI, blockchain/distributed ledger technology, smart contracts, Internet of Things, quantum computing and the Cloud as technologies of innovation. Building on the report and insider insights, academic researchers evaluated these technologies in terms of the operational risks they also present as the downside of their adoption and implementation.[52] This paper argues that AI and the Cloud offer immediate promise, the IoT and, particularly, quantum computing are future game changers bringing benefits but greater threats to banking operations. First, however, a short overview of the areas associated with operational risk events and where these might be applied is given.

Major categories of operational risk and implications for digital risk management

Operational risk may be categorised into seven event types: internal fraud (IF); external fraud (EF); employment practices and workplace safety (EPWS); clients, products and business practices (CPBP); damage to physical assets (DPA); business disruption and system failures (BDSF) and execution, delivery and process management (EDPM).

Figure 1 presents the allocation of losses, the percentage of total losses and loss amounts for seven operational risk event type categories associated with 303,562 operational losses incurred by 34 large US banks from 2001 to 2016.[53] It is notable that the EDPM category is populated by high volumes of events of relatively short duration and low-value losses, while the CPBP category is characterised by low volumes of extended duration and high-value losses. Significantly, EDPM is characterised by short lags between discovery and recognition, while

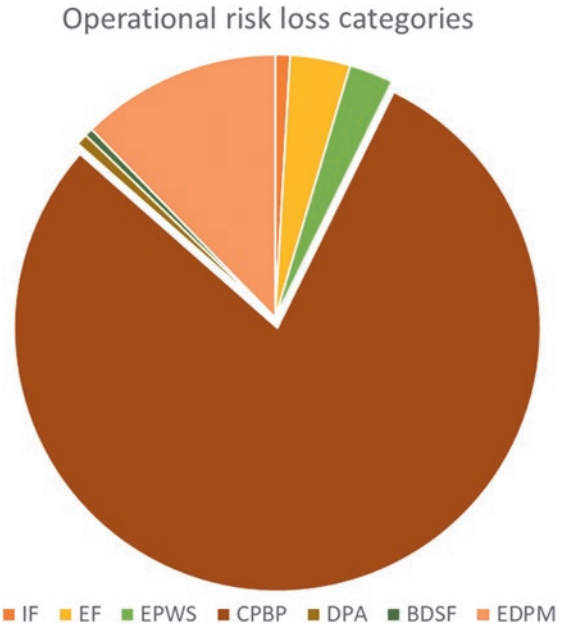

Risk category	Percent	Losses
IF	0.9%	$ 2.20bn
EF	3.7%	$ 9.10bn
EPWS	2.7%	$ 6.50bn
CPBP	79.20%	$ 192.10bn
DPA	0.70%	$ 1.60bn
BDSF	0.50%	$ 1.20bn
EDPM	12.30%	$ 29.90bn

Figure 1: Categories of operational risk by event frequency
Source: Curti, Frame and Mihov, 2022

Table 1: Average durations by region and event type (in days)

	Time to discovery	Time to recognition	Total days
Region			
North America	146	150	296
West Europe	403	110	513
Risk category			
Transactions and process management	254	143	397
Business practices	566	261	827

CPBP has significant lags between discovery and recognition. Table 1 draws on a study published by the Bank of International Settlements[54] to illustrate the time to discovery and recognition in two main regions by major risk category.

While ICT directly contributes just 0.49 per cent to operational losses due to disruption or failure, it must be remembered that the two main categories that contribute to 90 per cent of operational losses, CBPB and EDPM, are increasingly digitised and delivered using ICT with varying degrees of indirect or second-order contributions to risk events caused by poor systems design, mis-operation, absence of controls, information or cybersecurity issues, poor data governance and management, etc. Thus, the above categories are misleading and direct attention away from important sources of operational risk.

Two important points arise regarding the general application of digital technologies: (1) errors and misconduct can be controlled and or detected by digital technologies; and (2) risk managers can use the data produced by such systems to anticipate, detect, identify, assess and mitigate operational risks and make business activities resilient. As banks digitise more and more of their operations, it is the volume, variety and complexity of ICT and the quality, accuracy, availability and transparency of the data they produce that are important determinants of operational risk, these features also provide the solution to those self-same problems. Thus, in digitally transforming their operations, or upgrading/replacing existing IT systems, banks need to identify opportunities to minimise or eliminate human 'touch points' across processes in the front, middle and back-office to control for risks.[55] If this is not possible, then artificial intelligence technologies should be implemented ideally in a private Cloud platform or distributed in a hybrid cloud mesh.[56]

Applying AI to manage non-financial risks

AI provides financial institutions with powerful capabilities to address the 'wicked problem' that is operational risk in banking. It is assumed here that an appropriate organisational vision is developed and socialised and the values and commitments of c-suite, board, business, risk and IT functions are aligned. First, this paper gives a short overview of AI technologies and how they need to be implemented holistically and integrated rather than applied piecemeal, as may be the case. Researchers argue that while

> AI technologies offer great potential for financial institutions and supervisory authorities to automate and informate business processes, compliance and risk management activities, and regulatory processes, the benefits will not be fully realised unless an integrative combinatorial approach is adopted involving the three key AI approaches [of] knowledge representation, natural language processing, and machine learning. There is therefore an imperative to understand the strengths and limitations, the hype and reality of AI.[57]

Knowledge representation using semantic technologies

Semantic technologies such as ontologies solve the problem of providing a human (conceptual) and machine-readable (in a semantic, rules and logic-based data representation language, for example) common language for business, by categorising and defining unambiguously business concepts (using descriptions, relationships, axioms and rules) such as people/roles, activities/processes, services/products and their relationships. They are in widespread use in defence and intelligence organisations, the life sciences and in major industries, including the financial industry. However, they are just catching

on in the financial industry. In fact, just about every 'thing' that a business needs to represent digitally, manage, query and analyse can be represented. The ubiquitous business model canvas began as an ontology.

Ontologies vary in expressivity and may consist of conceptual models, vocabularies, taxonomies and concepts and relationships expressed in first-order and description logics. Most importantly, ontologies are essentially metadata models that help integrate, federate and virtualise data and their representations across heterogeneous siloed structured databases (SQL and NoSQL) and unstructured data stores (eg in XML), documents (Word, PDF etc), spreadsheets, webpages and websites.

At the Global Standards for Granular Data Working Conference at the European Central Bank in March 2017, the Bank of England's John Palmer argued that ontologies and related semantic technologies, such as graph databases, provide a solution to the problems of siloed risk data and its aggregation for risk management and regulatory compliance reporting. Why? Ontologies offer the solution for the problem of data synonyms and homonyms across siloed databases and other types of data stores and documents. Take, for example, data objects/entities or fields that have the same name/labels but refer to different concepts and data types; or data objects/entities and fields that have different names/labels but refer to the same type of concept representing business data. This is one reason why risk data aggregation is so problematic. Thus, ontologies enable the integration of models (including meta-data models), vocabularies, taxonomies of business objectives, functional activities, products and services with related risk taxonomies up, down, across and between organisations and public data sources.

Figure 2 represents an upper-level conceptual ontology of the financial domain that incorporates key concepts relevant to this paper. Domain and operational ontologies that expand key concepts such as regulations, risk, losses, business activity, systems, IT assets, data assets etc, will need to be developed and expressed formally in a machine language. However, organisations will have most of that information at hand, if not already expressed in taxonomies, such as risk and control taxonomies. One would also expect IT functions to have quite mature capabilities in capturing metadata on all IT systems and related assets. The importance and relevance of this approach is that it provides the first and second line of defence as business activities with the ability to answer complex questions, such as which systems' assets support specific critical operational activities and the IT risks (threats, vulnerabilities, impacts, including losses) to which a business is exposed. The model is sophisticated enough to also map operational losses at a granular level across people, processes and systems dimensions to identify where controls are failing etc. The power of ontologies is, however, explained in the remainder of this paper.

Related AI technologies, which include machine learning and natural language processing (NLP), when used in conjunction with ontologies, can better help filter through the noise in complex systems of system and data architectures and environments helping organisations achieve resilience. When business ontologies are employed with an enterprise knowledge base such as a graph database (eg Neo4J and RDF Triple Stores such as Jena TDB, StarDog and MarkLogic), these can be used to solve a variety of data analytics problems. For example, ontologies and knowledge graphs are used together to enhance the detection of operational risks, such as suspicious activities, fraud and money laundering and in the area of cyber risk. (See Palantir Foundry Ontology[58] and AML as one example.)

An enterprise graph database stores real-world data on the physical instances of the business concepts and relationships represented in an ontology, which is also referred to as a graph data model. For example, while an ontology would capture the relationship 'employee' *has* 'e-mail' in a subject-predicate-object semantic triple, the physical instance data 'John Doe' *has* 'jdoe@gmail.com' and 'John Doe' *has* 'jdoe@bank.com' could be stored in the enterprise knowledge base. In this case, the ontology would be used to integrate data from a bank's various databases for analysis. The power of a graph database is that any information about John Doe in any format anywhere, whether in e-mail, social media, webpages or operational databases can be captured in the graph database, building a powerful picture of John Doe

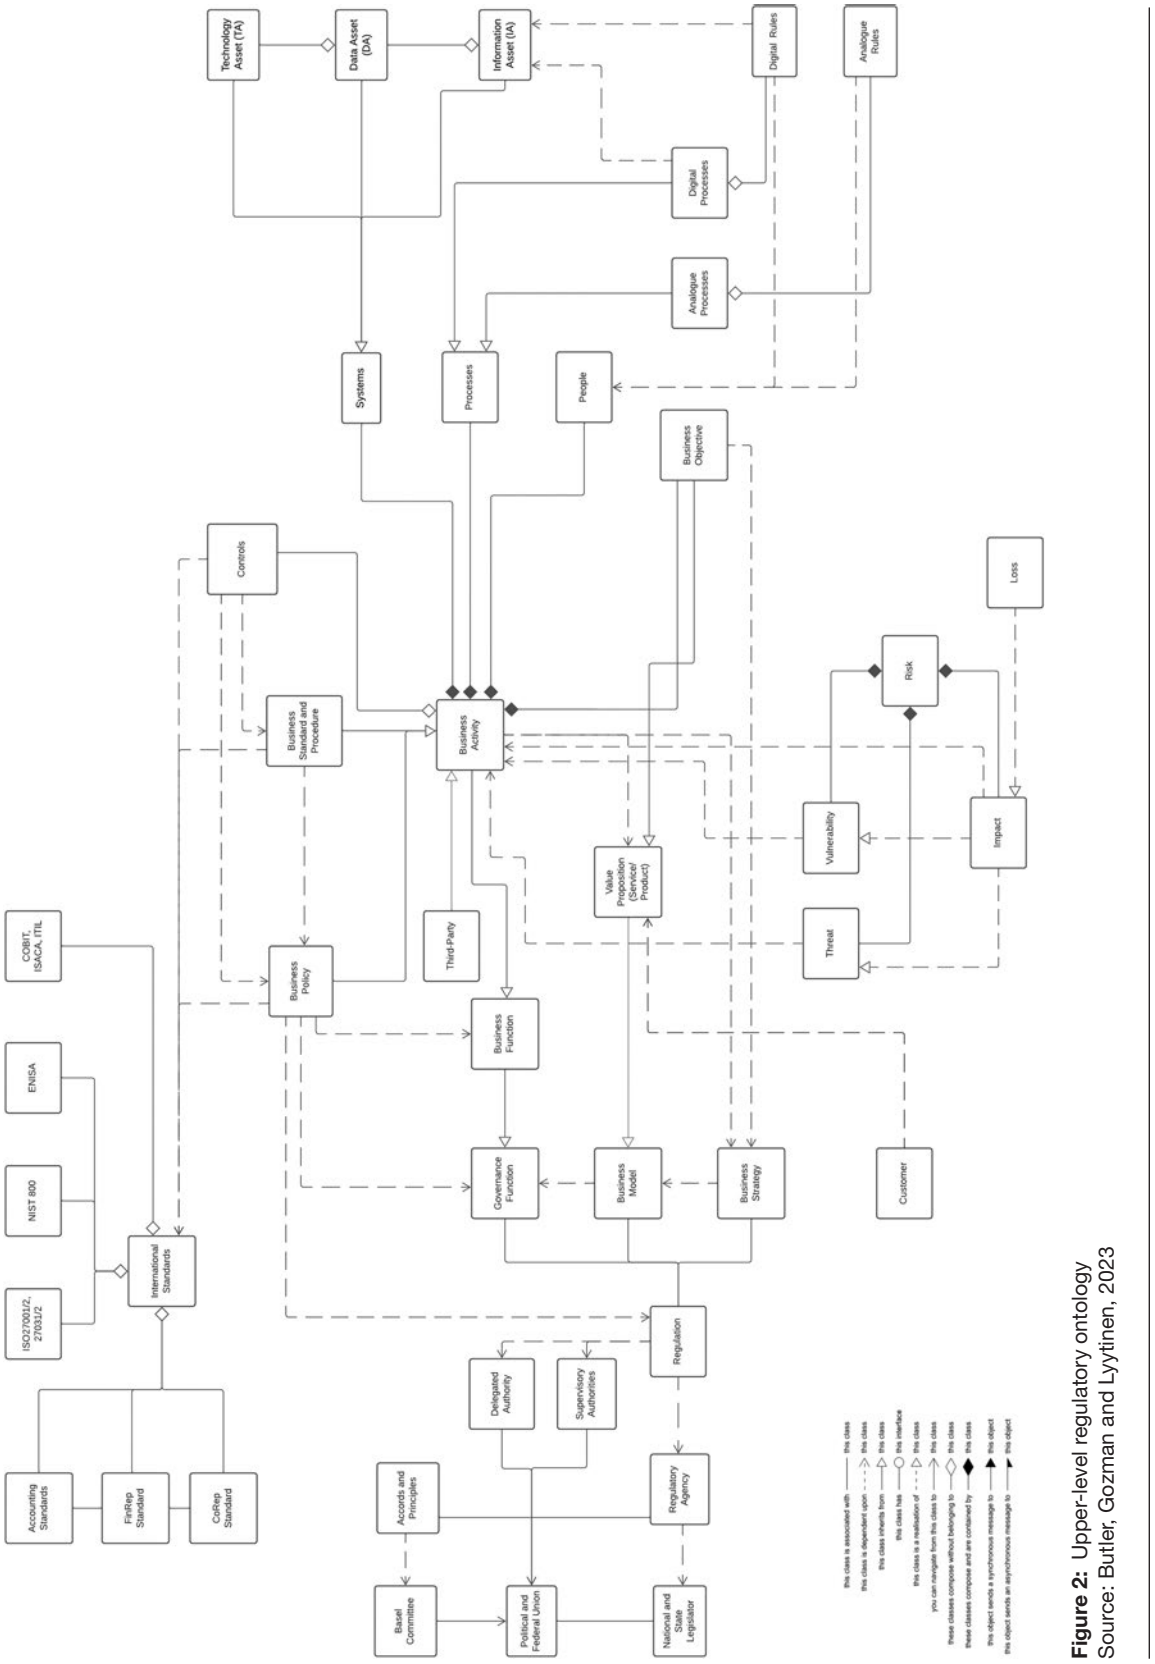

Figure 2: Upper-level regulatory ontology
Source: Butler, Gozman and Lyytinen, 2023

and his activities not captured in the bank's relational databases, thereby providing additional capabilities to help identify potential misconduct. JP Morgan's Workplace Activity Data Utility (WADU) system is an example of the potential of employee surveillance technologies. The reason intelligence services rely on ontologies and graph databases is that they can identify suspicious activities through linked data globally. Similar approaches are used to detect money laundering and fraud in several vendor solutions.

Creating AI-enabled enterprise knowledge fabrics

Figure 3 illustrates the relationship between the various elements of an enterprise data fabric. The enterprise ontology and graph database are *sine qua non* of an effective enterprise data fabric. These semantic technologies act together to integrate, federate or virtualise operational and other data sources, both internal and external (including those of third parties and market data). Standards-based semantic technologies such as the World Wide Web Consortium's (W3C) SPARQL (and the shapes constraint query language SHACL) provide services called endpoints to access data stores underpinning firm-specific and third party critical operations that deliver financial services. These are used to field semantic queries over data in multiple databases and other sources for operational and risk data integration, aggregation, federation or virtualisations. In the latter SPARQL–SHACL can implement compliance rules, as demonstrated in the Bank of England–FCA Digital Regulatory Reporting (DRR) project.

Thus, the enterprise data fabric described here solves both the 'common language' and the 'risk data aggregation' problems. Semantic reasoners are technologies that perform inferencing and enable the identification of new knowledge about the relationships between data. Semantic technologies such as these are powerful tools in the management of financial and non-financial risks. However, as argued elsewhere,[60] knowledge representation tools such as ontologies, reasoners, inference engines and graph databases confer even greater capabilities on machine learning and NLP technologies when used as part of enterprise data fabrics.[61] However, as all IT has downside risks and consequences,[62] the proposed

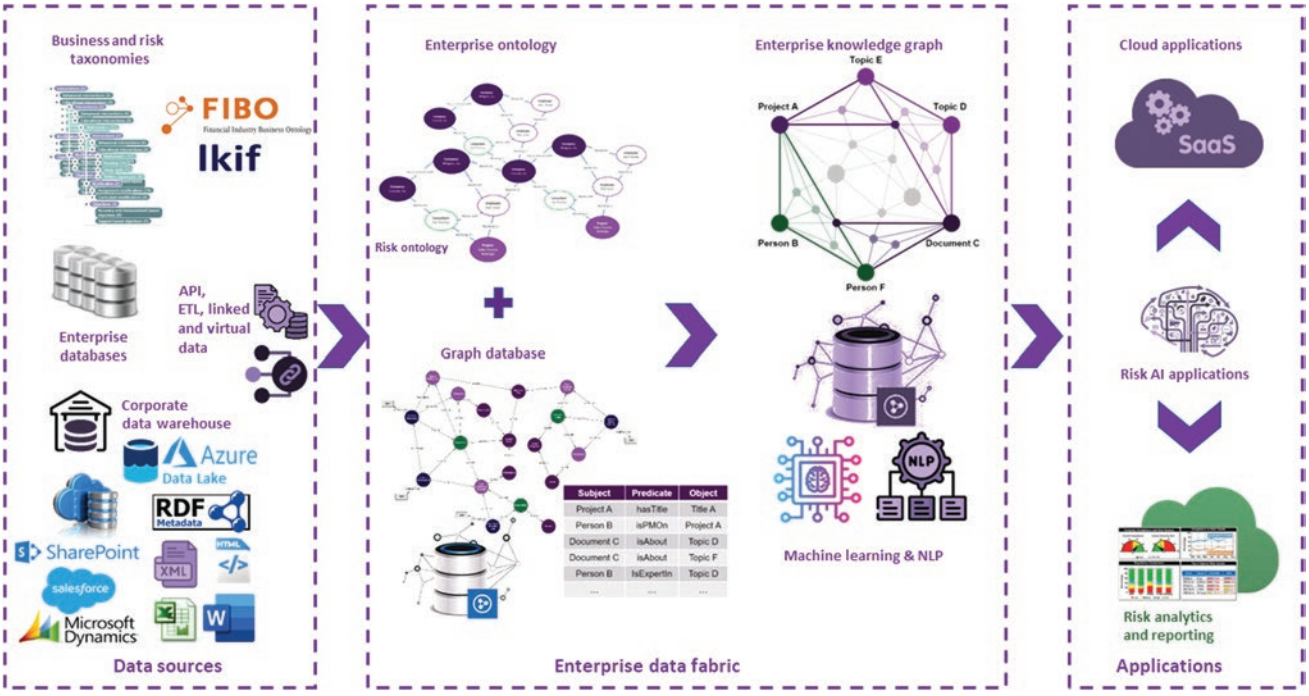

Figure 3: Understanding the power of AI-enabled enterprise knowledge graphs and data fabric
Source: Enterprise Knowledge, 2018[59]

architecture would also negatively change a bank's operational risk profile in terms of third-party dependencies, AI-related risks and the risks of data leakage etc.

AI-based enterprise data fabrics enable semantic, cognitive and perceptual computing for operational risk

The most powerful AI platforms perform *semantic*, *cognitive* and *perceptual computing*. *Semantic computing* has as its cornerstone knowledge representation (capturing conceptual and data semantics in models such as ontologies), but also employs NLP which is also founded on linguistic semantics. *Cognitive computing* has machine learning (ML) as its cornerstone and incorporates ontologies as declarative knowledge (semantics and rules) about the domain. ML approaches here also include deep learning (DL) which is a form of ML using artificial neural networks (ANN). *Perceptual computing's* cornerstone is in NLP and this is partnered with ML which identifies patterns of meaning in the data.

Machine learning, which includes by definition, DL/ANN approaches, provides the capabilities for *descriptive, diagnostic, predictive* and *prespictive analytics*. These techniques are particularly useful in risk identification, assessment, measurement, monitoring and control specification. In all cases, ML and DL/ANN are empowered and enhanced by data access and semantic enrichment provided by knowledge-based enterprise data fabrics. It is vital to understand the distinction between these four approaches to machine learning due to their implications for the future digitalisation of operational risk management.

Descriptive analytics provide historical insights into risk data using statistics. Its main contribution is in the area of risk intelligence. Key to this are data visualisation inputs into dashboards, graphs, charts and information for risk, controls and compliance reports. For example, it might be used to detect historical suspicious activities about fraud or money laundering, market misconduct, surveillance of communications or to present patterns of historical operational risk losses.

Diagnostic analytics extend descriptive data analysis to identify causal and correlative patterns. Diagnostic analysis is referred to as root cause analysis (RCA) to determine the origins of operational risk events (eg cybersecurity risks) and related losses using data discovery, mining, inferencing and drill down and drill through the large volumes of linked data in enterprise data fabrics.

Predictive analytics include ML algorithms that apply classification, regression and clustering techniques to forecast the probabilities of future states of the world, such as risk event outcomes. Essentially, ML algorithms perform pattern matching and predict the likely path and probability of outcomes of detected patterns. DL and ANN are particularly suited to uncovering previously unknown relationships between large, related data sets and identifying probable future outcomes. *Predictive analytics* are particularly suited to predicting market manipulation, account churning, improper trading and risk related to product defects, for example.

Prescriptive analytics builds on information provided by descriptive, diagnostic and predictive analytics. This ML approach typically employs recurrent neural networks (RNN) for predictive analytics and multi-objective reinforcement learning (MORL) for prescriptive analytics over operational data.[63] It may be particularly suited to offering recommendations for effective risk mitigation measures and controls.

NLP enables semantic and perceptual computing by first processing digitalised visual, audio and text-based into machine-readable and computable format and then adding meaning through further processing to inform decisions, guide actions of human or digital agents and/or generate responses which can be understood by humans. To date, NLP use cases include unpacking regulatory provisions in legislation and related instruments, in supervisors' rulebooks and interpreting and tagging them,[64] processing unstructured anti-money laundering (AML)/know your customer (KYC) data internally and externally from social media and websites etc and assessing conduct risk events by identifying anomalies in unstructured data such as e-mails, correspondence, telephone and video calls and business reports. The power of NLP's perceptual computing is demonstrated by ChatGPT which applies a sophisticated NLP

large-language model (LLM) and employs Transformer, a deep neural network architecture, and an unsupervised learning approach in generating human-readable answers in response to specific questions. As promising as this might sound for risk management in financial institutions, the data integration, aggregation and virtualisation issues must be first addressed.

There is more to AI than perceptual computing's deep learning with neural networks, semantic computing and ontologies are required to fill significant knowledge gaps. In the financial industry, ChatGPT is incapable of addressing the many different operational risk management tasks that need to be solved over heterogeneous siloed, structured and unstructured textual data. While there are mountains of structured data, they lack semantic context and integration and the volume of textual data is typically not large enough for an LLM. The risk information operational risk managers need requires precise, explainable answers from RiskTech algorithms — LLMs like ChatGPT lack precision and provide answers that are often incorrect. Furthermore, they cannot explain their reasoning. Risk management requires current data as well as historical data if risks are to be anticipated and controlled. Historical operational risk event loss data is often without context and is, therefore, a poor predictor of future events. Finance is highly specialised and for LLM to successfully operate data from all G-SIBs might be required to make it work successfully.

AI use cases in banking and insurance

The three AI computing paradigms — knowledge representation using ontologies, ML (DL/ANN) and NLP — when used together over relevant, accurate, complete, integrated and/or aggregated operational and risk data, can enable digital risk management, specifically digital operational risk management and operational resilience. The authors' research identifies the following areas as examples:

- compliance and risk horizon scanning;
- risk (threat and vulnerability) intelligence across all operational risk categories enhancing and validating RCSAs and providing the ability for data-driven KRIs;
- enhanced risk-based problem-solving and decision support;
- risk alerts and compliance monitoring;
- automated customer advice using chatbots and virtual assistants to minimise conduct risk and human errors;
- robotic process automation in front, middle and back-office digitisation to eliminate human touchpoints and people risks;
- unpacking and mapping regulatory provisions to risks, controls and compliance reports;
- enhanced credit and risk underwriting;
- enabling smart contracts;
- KYC and AML risk management;
- internal and external fraud detection and prevention;
- consumer risk assessment;
- data and information asset risk;
- biometrics and identity risk;
- cybersecurity risk;
- risk control and compliance workflows;
- diligence, vendor and third party risk surveillance;
- employee and trader surveillance across front, middle and back-office.

Money laundering is one operational risk category that is well served by AI-based technologies. Several third party vendor offerings are available on the market. The risk in adopting such applications is that without an enterprise data fabric, financial institutions will end up with point solutions for each use case. For AI to achieve its full potential, domain ontologies are required. Vendors such as Palantir offer customisable out-of-the-box solutions. The first capability of such a solution is to perform anomaly detection and behaviour analysis over transaction data to identify suspicious activities or unusual patterns of transactions, including cash deposits and transfers between customer accounts or cross-border payments. Machine learning models and algorithms may be applied to assess, evaluate and risk-weight transactions for further investigations. A second capability is customer due diligence, where customer onboarding can be automated and KYC enhanced using AI's ability to integrate, make inferences from and reason over data from disparate internal and external unstructured data sources, including international watchlists and media platforms, to verify customer identity and detect any change in risk profiles. AI's

machine learning capabilities can also identify and learn from new data patterns and adapt as new approaches to money laundering or terrorist financing are applied by criminals or terrorists. AI's enhanced suspicious activity identification capabilities not only eliminate false negatives but significantly reduce false positives and associated costly investigations by money laundering reporting officers (MLROs).

Market manipulation is a significant category of people risk. AI's capabilities at pattern and trend recognition also come into play. Ontologies can integrate structured and unstructured data from heterogeneous sources, including market data from order books, trading books and external media sources. Inferencing and reasoning engines can detect suspicious activities and market movements. This is reinforced by behavioural analysis of the historical activities of traders and other participants to establish trading profile deviations to generate alerts for risk and compliance officers. As with employee surveillance, AI can monitor communication and network relationships between traders and counterparties to identify, assess and issue alerts on possible collusion in market manipulation. The patterns and trends identified by AI can also help predict and/or identify emerging manipulation threats.

The AI capabilities described above can also be employed in identifying misconduct through enhanced internal and external employee monitoring and surveillance. Ontologies, ML and NLP provide powerful tools to federate and integrate unstructured internal and external communication data (eg e-mails, messaging and attachments) for analysis to identify unethical conduct or collusion across the front, middle and back-office, whether to manipulate and defraud customers, or engage in rogue trading, etc. Patterns of keywords and phrases associated with misconduct in all electronic communications, including voice, suspicious or anomalous communication patterns across media, provide indicators of potential misconduct, as do communications with known or suspected bad actors. When combined with the approach described in market manipulation, integration of communication, transaction and other data provides the basis for enhanced inferencing, reasoning and analytics to forensically investigate misconduct.

Financial risk management is also clearly enhanced by the AI approach described herein. In fact, financial risk management requires digital risk management capabilities based on real-time AI to address the many human biases that underpin much of the decision-making in financial institutions.

CONCLUSIONS

This paper addresses the problems that perplex banks — large and small — and give committed, operational risk managers pause for thought if not sleepless nights. Operational risk in large banks is conceptualised as a 'wicked problem', one that eludes proper resolution. The role of the extant operational risk paradigm in contributing to this is summarised and an estimate given of the true first-order losses: However, these estimates exclude 'opportunity costs, forgone revenue, and costs related to risk management and control enhancements implemented to prevent future operational losses'.[65] Thus, the actual scale of losses is unaccounted for and therefore potentially unmanaged. The Risk Accounting Standards Board (RASB) provides a systematic approach to accounting for losses associated with operational risk. As argued by Peter Hughes, risk accounting not only quantifies non-financial risks using risk units (RU), which enables the reporting of risk exposures, it also places a financial value on operational risks at a granular level: Finally, the approach can be digitised and integrated with existing systems of record and digital initiatives.[66] This paper concludes that investment in the digitalisation of operational risk management would pay handsome dividends in enabling the management of the unmanageable, thereby reducing the future costs of operational risk losses.

It should be pointed out that financial institutions need not start from scratch in developing ontologies. Several ontologies focus on the financial and regulatory domains. Perhaps the most comprehensive is FIBO, which is an industry-developed family of ontologies that provide a common language for and describes the semantics of financial business entities, processes, products agents, people and systems, etc across the financial industry. Financial Industry Business Ontology (FIBO) was developed by the Enterprise Data Management Council (EDM

Council) and the software engineering standards organisation, the Object Management Group (OMG). EDM Council members such as G-SIBs have contributed to its development since 2008.[67] The Legal Knowledge Interchange Format (LKIF) ontology from the European Union[68] models the legal rules in legislation and regulations, while the suite of Financial Regulation Ontologies (FRO) combines FIBO and Legal Knowledge Interchange Format (LKIF) to model US financial regulations.[69] These three ontologies inform the development of domain-specific and operational ontologies to tackle specific use cases, such as automating the analysis of suspicious activities in financial institutions.

The design and development of an AI-enabled enterprise data fabric and knowledge base can be iterative, so there is no excuse for inaction as banks can focus on critical operations, however, an organising vision is a prerequisite, as is a roadmap to completion. Following, the BCBS 'Principles for Operational Resilience',[70] it is the authors' recommendation that banks focus on the firm-specific and third party critical operations that deliver financial services, which have the greatest historical loss and build out and up from there. Creating enterprise ontologies can be off-putting for the C-suite, as they are a long-term investment and not a quick technological fix. This is one reason for the continuing existence of this 'wicked problem'. Thus, a targeted approach that can be integrated with existing digital transformation initiatives by incorporating and applying AI technologies to predict, identify, assess and/or mitigate operational risks in the clients, products and business practices (CPBP) category is important. However, as indicated by the ORX recently, 'Generative AI has contributed to the highest number of operational risk "incidents" since 2017'[71] as it enables sophisticated external fraud. However, just as AI creates a problem, it also holds the key to a solution.

References

1 Basel Committee on Banking Supervision (BCBS) (June, 2011) 'Principles for the Sound Management of Operational Risk', Bank for International Settlements (BIS), available at https://www.bis.org/publ/bcbs195.pdf (accessed 30th September, 2023).

2 Basel Committee on Banking Supervision (BCBS) (2020) 'Revisions to the Principles for the Sound Management of Operational Risk', Bank for International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d508.pdf (accessed 30th September, 2023).

3 Basel Committee on Banking Supervision (BCBS) (January, 2013) 'Principles for Effective Risk Data Aggregation and Risk Reporting', Bank for International Settlements (BIS), available at https://www.bis.org/publ/bcbs239.pdf (accessed 30th September, 2023).

4 Basel Committee on Banking Supervision (BCBS) (March, 2021) 'Principles for Operational Resilience', Bank for International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d516.pdf (accessed 30th September, 2023).

5 Arner, D. W., Barberis, J. and Buckley, R. P. (2015) 'The Evolution of Fintech: A New Post-Crisis Paradigm', *Georgetown Journal of International Law*, Vol. 47, p. 1271.

6 Gomber, P., Koch, J. A. and Siering, M. (2017) 'Digital Finance and FinTech: Current Research and Future Research Directions', *Journal of Business Economics*, Vol. 87, pp. 537–80.

7 Butler, T., Gozman, D. and Lyytinen, K. (2023) 'The Regulation of and Through Information Technology: Towards a Conceptual Ontology for IS Research', *Journal of Information Technology*, Vol. 38, No. 2, pp. 86–107.

8 Paech, P. T., Butler, T., *et al.* (December, 2019) 'Thirty Recommendations on Regulation, Innovation and Finance', Expert Group on Regulatory Obstacles to Financial Innovation (ROFIEG), available at https://finance.ec.europa.eu/publications/final-report-expert-group-regulatory-obstacles-financial-innovation-30-recommendations-regulation_en (accessed 30th September, 2023).

9 Hughes, P. J. (2023) 'Risk Accounting: The Complete Guide to Quantifying and Accounting for Non-financial Risks', Grosvenor House Publishing Ltd, Surbiton.

10 McMillan, C. and Overall, J. (2016) 'Wicked Problems: Turning Strategic Management Upside Down', *Journal of Business Strategy*, Vol. 37, No. 1, pp. 34–43.

11 Rittel, H. W. and Webber, M. M. (1973) 'Dilemmas in a General Theory of Planning', *Policy Sciences*, Vol. 4, No. 2, pp. 155–69.
12 Churchman, C. W. (1967) 'Wicked Problems', *Management Science*, Vol. 14, No. 4, pp. B141–2.
13 Head, B. W. (2022) 'Wicked Problems in Public Policy: Understanding and Responding to Complex Challenges', Springer Nature, Cham.
14 Berger, A. N., Curti, F., Mihov, A. and Sedunov, J. (2022) 'Operational Risk is More Systemic Than You Think: Evidence from US Bank Holding Companies', *Journal of Banking & Finance*, Vol. 143, 106619.
15 Basel Committee on Banking Supervision (BCBS) (February, 2023) 'Basel III Monitoring Report', Bank for International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d546.pdf (accessed 30th September, 2023), p. 82.
16 Hoffman, C. (11th July, 2022) 'Global Banks Report Declining Operational Risk Gross Losses in 2021 Despite Market Turbulence', Trade Finance Global, available at https://www.tradefinanceglobal.com/wire/global-banks-report-declining-operational-risk-gross-losses-in-2021-despite-market-turbulence/#:~:text=In%202021%2C%2085%2C585%20operational%20risk,10bn%20between%202016%20and%202021 (accessed 30th September, 2023).
17 Zampano, G. (15th August, 2023) 'Operational Risk 'Incidents' Surge 26% because of AI and Low-risk Fraud', Banking Risk and Regulation, available at https://www.bankingriskandregulation.com/operational-risk-incidents-surge-26-because-of-ai-and-low-risk-fraud/ (accessed 30th September, 2023).
18 *Ibid.*
19 Hoefer, E., Cooke, M. and Curry, T. (8th September, 2020) 'Three Lines of Defense-Failed Promises and What Comes Next', Financial Regulatory Forum, Reuters, available at https://www.reuters.com/article/bc-finreg-risk-management-three-lines-of-idUSKBN25Z2FN (accessed 30th September, 2023).
20 Berger *et al.* ref 14 above, p. 4.
21 Curti, F., Frame, W. S. and Mihov, A. (2022) 'Are the Largest Banking Organizations Operationally More Risky?', *Journal of Money, Credit and Banking*, Vol. 54, No. 5, pp. 1232–59.
22 Afonso, G., Curti, F. and Mihov, A. (2019) 'Coming to Terms with Operational Risk, (No. 20190107)', Federal Reserve Bank of New York, available at https://libertystreeteconomics.newyorkfed.org/2019/01/coming-to-terms-with-operational-risk/ (accessed 30th September, 2023).
23 Böni, P., Kroencke, T. A. and Vasvari, F. P. (2023) 'The UBS–Credit Suisse Merger: Helvetia's Gift', available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4486417 (accessed 30th September, 2023), SSRN 4486417, p. 1.
24 Walker, O., Morris, S. and Smith, R. (May 2023) 'Credit Suisse Staff Prepare to Sue Regulator Finma Over Lost AT1 Bonuses', Financial Times, available at https://www.ft.com/content/b93216e7-b54b-42a1-a283-ce0dc5b476b4 (accessed 30th September, 2023).
25 Inman, P. (28th June, 2023) 'UBS "Preparing to Cut more than Half of Inherited Credit Suisse Workforce"', The Guardian, available at https://www.theguardian.com/business/2023/jun/28/ubs-preparing-to-cut-more-than-half-of-inherited-credit-suisse-workforce (accessed 30th September, 2023).
26 Basel Committee on Banking Supervision (BCBS) (June, 2006) 'Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework', Bank for International Settlements (BIS), available at https://www.bis.org/publ/bcbs128.htm (accessed 30th September, 2023).
27 Hughes, P. J. (2021) 'Where Next for Operational Risk? A Guide for Risk Managers and Accountants', Grosvenor House Publishing Ltd, Surbiton.
28 Butler, T. (2022) 'Book Review: Where Next for Operational Risk? A Guide for Risk Managers and Accountants', *Journal of Risk Management in Financial Institutions*, Vol. 15, No. 3, pp. 301–4.
29 Power, M. (2005) 'The Invention of Operational Risk', *Review of International Political Economy*, Vol. 12, No. 4, pp. 577–99.
30 Grody, A. D. and Hughes, P. J. (2016) 'Risk Accounting-Part 1: The Risk Data Aggregation

and Risk Reporting (BCBS 239) Foundation of Enterprise Risk Management (ERM) and Risk Governance', *Journal of Risk Management in Financial Institutions*, Vol. 9, No. 2, pp. 130–46.
31 Basel Committee on Banking Supervision (BCBS), ref 3 above.
32 Basel Committee on Banking Supervision (BCBS) (2015) 'Progress in Adopting the Principles for Effective Risk Data Aggregation and Risk Reporting', Bank of International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d501.htm (accessed 30th September, 2023); BCBS (June, 2018) 'Progress in Adopting the Principles for Effective Risk Data Aggregation and Risk Reporting', Bank of International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d443.pdf; BCBS (2020) 'Progress in Adopting the Principles for Effective Risk Data Aggregation and Risk Reporting', Bank of International Settlements (BIS), available at https://www.bis.org/bcbs/publ/d501.htm (accessed 30th September, 2023).
33 *Ibid.* BCBS (2020).
34 Bailey, D. and Jackson, R. (10th September, 2021) 'Dear CEO Letter', Bank of England, Prudential Regulation Authority, available at https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2021/september/thematic-findings-on-the-reliability-of-regulatory-returns.pdf (accessed 30th September, 2023).
35 Benjamin, N. and Jackson, R. (12th January, 2022), 'Dear CEO Letter', Bank of England, Prudential Regulation Authority, available at https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/letter/2022/january/artis-2022-priorities.pdf (accessed 30th September, 2023).
36 Haldane, A. G. (14th March, 2012) 'Towards a Common Financial Language', Securities Industry and Financial Markets Association (SIFMA) 'Building a Global Legal Entity Identifier Framework' Symposium, New York, 2012, available at http://www.bis.org/review/r120315g.pdf (accessed 30th September, 2023).
37 Hughes, ref 27 above.
38 Hughes, ref 9 above; Hughes, ref 27 above; Grody and Hughes, ref 30 above.
39 Grody and Hughes, ref 30 above, p. 14.
40 *Ibid.*
41 Shefrin, H. (2016) 'Behavioral Risk Management: Managing the Psychology that Drives Decisions and Influences Operational Risk', Palgrave Macmillan, New York.
42 Andersen, L. B., Häger, D. and Vormeland, H. B. (2016) 'Causal Analysis of Operational Risk for Deriving Effective Key Risk Indicators', *Journal of Risk Management in Financial Institutions*, Vol. 9, No. 3, pp. 289–304.
43 See Curti *et al.*, ref 21 above; Hughes, ref 27 above.
44 Grody and Hughes, ref 30 above.
45 European Central Bank (ECB) (May, 2018) 'Report on the Thematic Review on Effective Risk Data Aggregation and Risk Reporting', European Central Bank, available at https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.BCBS_239_report_201805.pdf (accessed 30th September, 2023).
46 *Ibid.* p. 1.
47 McIntyre, A., Skan, J., Leruste, C. A. and Caminiti, F. (2019) 'Caterpillars, Butterflies, and Unicorns: Does Digital Leadership in Banking Really Matter?', Accenture, pp. 1–15, available at https://www.accenture.com/_acnmedia/pdf-102/accenture-banking-does-digital-leadership-matter.pdf (accessed 1st June, 2022).
48 Werth, O., Schwarzbach, C., Cardona, D. R., Breitner, M. H. and von der Schulenburg, J. M. (2020) 'Influencing Factors for the Digital Transformation in the Financial Services Sector', *Zeitschrift für die gesamte Versicherungswissenschaft*, Vol. 109, No. 2, pp. 155–79, https://doi.org/10.1007/s12297-020-00486-6.
49 Paech *et al.*, ref 8 above.
50 Swanson, E. B. and Ramiller, N. C. (1997) 'The Organizing Vision in Information Systems Innovation', *Organization Science*, Vol. 8, No. 5, pp. 458–74, available at https://doi.org/10.1287/orsc.8.5.458.
51 *Ibid.*
52 Butler, T. and O'Brien, L. (2019) 'Artificial Intelligence for Regulatory Compliance: Are We There Yet?', *Journal of Financial Compliance*, Vol. 3, No. 1, pp. 44–59.
53 Swanson and Ramiller, ref 50 above.

54 Aldasoro, I., Gambacorta, L., Giudici, P. and Leach, T. (February, 2020) 'BIS Working Papers No. 840, Operational and Cyber Risks in the Financial Sector', Bank for International Settlements, available at https://www.bis.org/publ/work840.pdf (accessed 30th September, 2023).
55 Hughes, ref 27 above.
56 IBM Hybrid Cloud Mesh, available at https://www.ibm.com/products/hybrid-cloud-mesh (accessed 30th September, 2023).
57 Butler and O'Brien, see ref 52 above.
58 Palantir 'The Foundry Ontology', available at https://www.palantir.com/platforms/foundry/foundry-ontology/ (accessed 30th September, 2023); Palantir 'Anti-Money Laundering', available at https://www.palantir.com/offerings/anti-money-laundering/ (accessed 30th September, 2023).
59 Significantly adapted from EK Team (1st November, 2018) 'What is an Enterprise Knowledge Graph and Why Do I Want One?', Enterprise Knowledge, available at https://enterprise-knowledge.com/what-is-an-enterprise-knowledge-graph-and-why-do-i-want-one/ (accessed 31st October, 2023).
60 Butler and O'Brien, ref 52 above.
61 *Ibid.*
62 Butler, T., Gozman, D. and Lyytinen, K. (2023) 'The Regulation of and Through Information Technology: Towards a Conceptual Ontology for IS Research', *Journal of Information Technology*, Vol. 38, No. 2, pp. 86–107.
63 Lepenioti, K., Pertselakis, M., Bousdekis, A., Louca, A., Lampathaki, F., Apostolou, D., Mentzas, G. and Anastasiou, S. (2020) 'Machine Learning for Predictive and Prescriptive Analytics of Operational Data in Smart Manufacturing', in Dupuy-Chessa, S. and Proper, H. A. (eds), *Advanced Information Systems Engineering Workshops*, Vol. 29, No. 382, pp. 5–16, available at https://doi.org/10.1007/978-3-030-49165-9_1 (accessed 31st October, 2023).
64 Butler, T. and O'Brien, L. (2019) 'Understanding RegTech for Digital Regulatory Compliance', in Lynn, T., Mooney, J. G., Rosati, P. and Cummins, M. (eds), 'Disrupting Finance: FinTech and Strategy in the 21st Century', Pagrave Pivot, Cham, pp. 85–102.
65 Curti *et al.*, ref 21 above.
66 Hughes, ref 9 above.
67 Bennett, M. (2016) 'An Industry Ontology for Risk Data Aggregation Reporting', *Journal of Securities Operations & Custody*, Vol. 8, No. 2, pp. 132–45.
68 ESTRELLA Project (2008) 'LKIF Core Ontology — Standardized Transparent Representations to Extend Legal Accesability', available at https://github.com/RinkeHoekstra/lkif-core (accessed 9th October, 2023).
69 Ziemer, J. (2016) 'The Financial Regulation Ontologies', available at http://finregont.com/ (accessed 30th September, 2023).
70 Basel Committee on Banking Supervision (BCBS), ref 4 above.
71 Zampano, ref 17 above.