O.R.X

Machine learning in operational risk

Making a business case for its practical implementation

White paper: September 2019

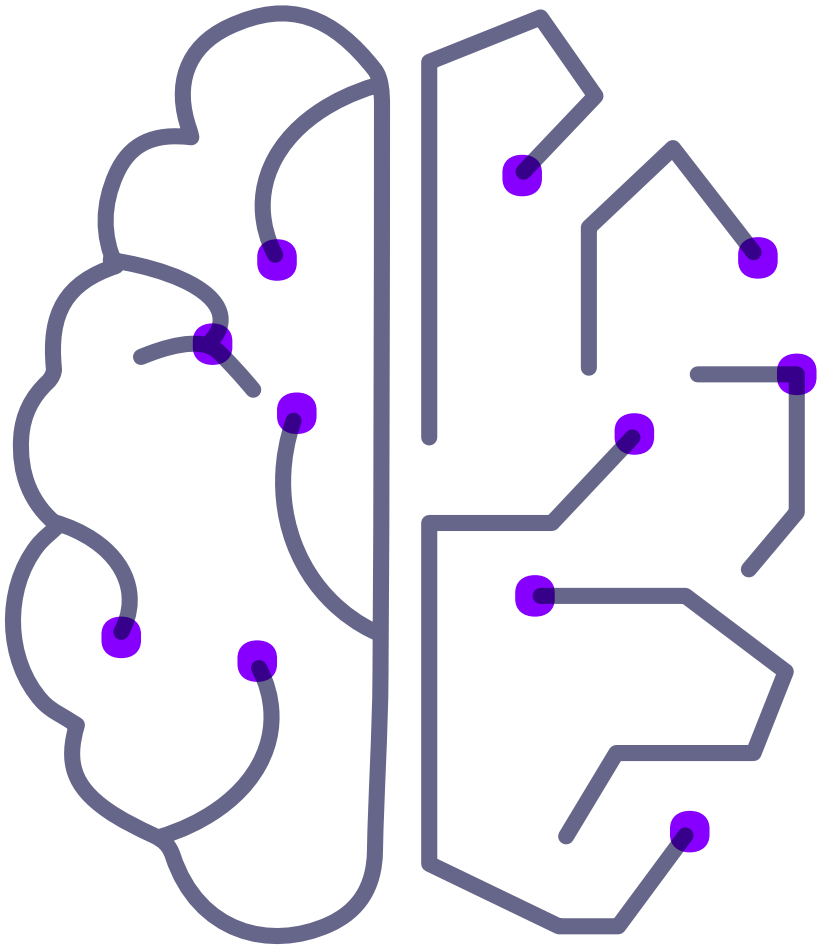

About this white paper

The purpose of this white paper is to make a case for the use of advanced analytics, and specifically machine learning techniques, for operational risk management in financial firms.

It lays out the opportunities that ORX believes lie in the use of these techniques, and provides information on how they can be integrated into day-to-day operational risk management activities.

"We believe that the application of advanced analytics, including machine learning and artificial intelligence (AI), will be a core part of any future strategy for the management of operational and non-financial risk. This paper focuses specifically on the opportunities that machine learning can offer."

Measurement & Data

Topics covered:

Machine learning | Advanced analytics | Artificial intelligence | Robotic process automation

This white paper is the result of work done as part of our research programme. Research is a core part of ORX. It helps us to support successful operational and non-financial risk management and measurement.

Taking part is free for ORX members, who also get access to all our published research reports. If your firm isn't a member of ORX, then you can still participate in our studies for a small fee.

For more information visit: www.orx.org

Connect with ORX

www.orx.org | @ORX_Association | @ORX_association

Introduction

In 2017, ORX conducted a study interviewing 15 Chief Risk Officers (CROs) on the future of operational risk. Many of the interviewees identified advanced analytics (beyond capital modelling and stress testing) as a potential game changer for operational risk management. Around half earmarked it as a priority for innovation and investment over the next five years.[1]

However, despite its huge potential machine learning has remained largely unexplored by operational risk. In a recent literature review[2] of the application of machine learning in risk management, only six out of 50 papers focused on operational risk management.

We believe that the application of advanced analytics, including machine learning and artificial intelligence (AI), will be a core part of any future strategy for the management of operational and non-financial risk. This paper focuses specifically on the opportunities that machine learning can offer. It is intended for operational risk functions who are beginning to explore using these techniques, and want to make a business case for their application.

The paper is divided into four sections. The first describes the role that we believe machine learning techniques will play in operational risk measurement and management. This is followed by the key opportunities that we see in their application. The third section contains example use cases that we have collected from the ORX Machine Learning Working Group (MLWG), a group of ORX member institutions that are already successfully applying these techniques (see page 3 for further information). The final section addresses what we see as the key considerations that financial firms, and the operational risk discipline more generally, will need to address in order to make greater use of machine learning.

Five building blocks to support a reinvented operational risk management (from the 2017 ORX and McKinsey report, The future of operational risk)

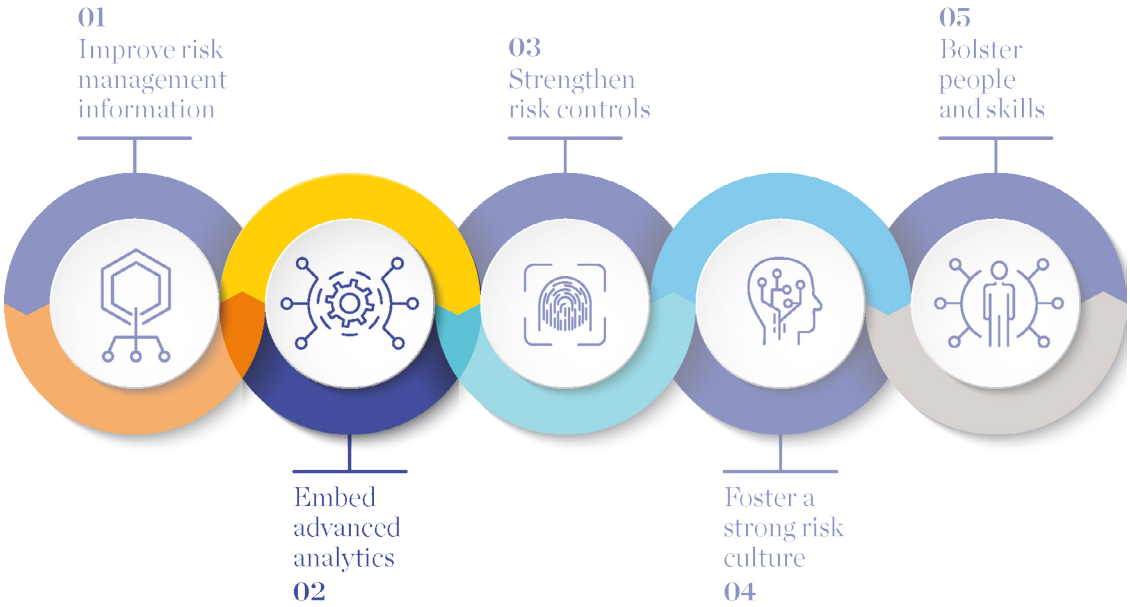

1 https://managingrisktogether.orx.org/research/future-operational-risk

2 Machine Learning in Banking Risk Management: A Literature Review. Leo, Sharma, Maddulety. Risks 2019, 7(1), 29; https://doi.org/10.3390/risks7010029

Advanced analytics – a potential game changer

The opportunities that machine learning techniques offer – from task optimisations and better use of resources, to cost savings and gaining deeper insights into data – are considerable, but often not used to their full potential. Operational risk is one of the disciplines that could considerably benefit from their application.

Perhaps more than any other risk discipline, operational risk is characterised by a wide variety of often idiosyncratic risks that fall under its scope. This is combined with an ever-increasing need for information to report on and manage these risks.

Adding further pressure, operational risk functions have been preparing for sweeping changes in regulatory capital requirements. The announcement of the Standardised Measurement Approach (SMA) has, in many organisations, resulted in a shift of focus away from capital calculation towards the more active management of operational risk.

With this new focus, there is an awareness that operational risk functions need to assume a more commercial and proactive role in supporting business activities. At the same time, businesses are facing new, previously unknown challenges in the changing operating environment introduced by the fourth industrial revolution.

In addition, thinking beyond the primary benefits of resource optimisation and cost savings, operational risk has a central role to play in the management of new risks that Artificial Intelligence and automation bring with them. This is an important topic, but beyond the central scope of this paper and will be explored in more detail in future ORX studies.

Helping the industry explore advanced analytics

The ORX Machine Learning Working Group

In early 2019, ORX established a Machine Learning Working Group (MLWG). The group consists of representatives from nine ORX member institutions who are already using these techniques for operational risk management. They are dedicated to exploring the impact of machine learning on operational risk.

The purpose of the working group is to begin an industry dialogue about advanced analytics (AI and machine learning). The group meets on a regular basis to share common challenges, identify potential solutions, and explore the potential for collaboration. They also oversee and guide the direction of ORX's research on this topic.

At ORX, our vision is for the financial sector to work together to improve and innovate operational risk management. This is why we believe it is so important for our members to talk to each other about the operational risk issues facing the industry today. Our working groups are vital to helping us achieve this goal.

The opportunities for operational risk

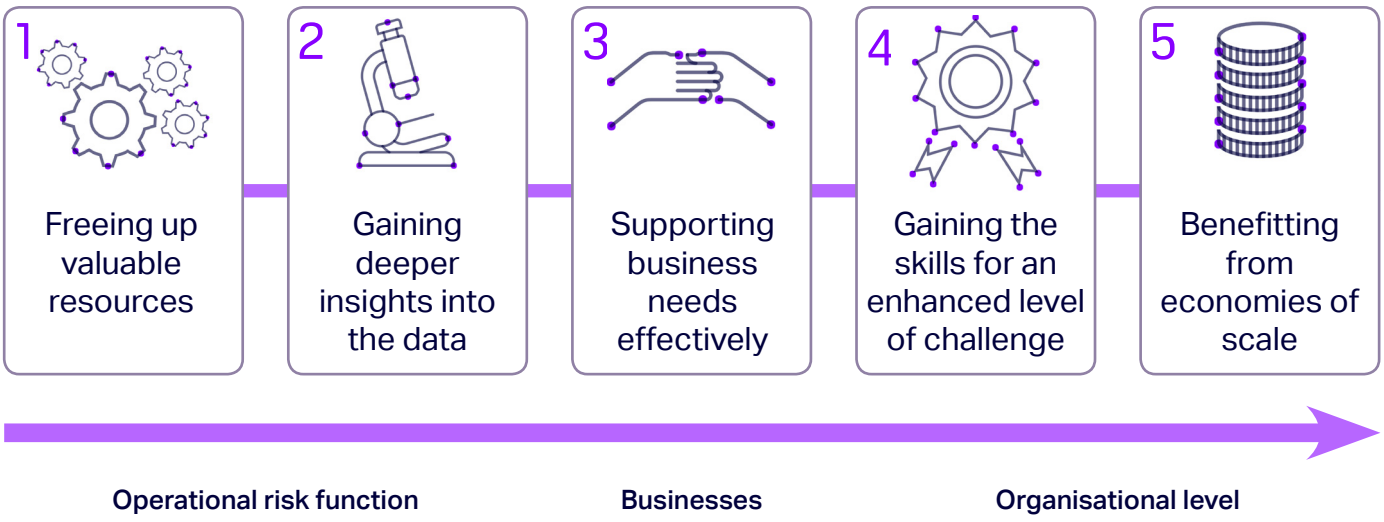

The five opportunities for operational risk from machine learning techniques

The five key areas

Based on conversations with ORX member institutions, we've identified five areas in which we believe operational risk functions will benefit from the use of machine learning:

1. Freeing up valuable resources
2. Gaining deeper insights into data
3. Supporting business needs effectively
4. Gaining the skills for an enhanced level of challenge
5. Benefitting from economies of scale

1. Freeing up valuable resources

Potentially, the biggest gains from the implementation of machine learning can be expected to come from the reduction or elimination of time-intensive and repetitive tasks, which take up valuable time of operational risk teams. These tasks are often related to the collection, handling and analysis of operational risk data, such as for quality assurance and regular reporting.

Many of these tasks lend themselves to robotic process automation, and this is where even simple applications of machine learning techniques can add considerable value. For example, these techniques can improve and speed up the rationalisation of control data by identifying equivalent controls and by inferring missing attributes in control libraries, based on free text descriptions of the control.

Moreover, already automated processes can be further improved with machine learning techniques. A classic example where advanced analytics has led to important improvements in efficiency is the screening of financial transactions to detect potential instances of fraud or money laundering. As one of the case studies on page 7 shows, some financial firms have already successfully applied machine learning to reduce the number of false positives and negatives that existing processes produce, thereby reducing the number of instances that need to be individually reviewed.

2. Gaining deeper insights into data

Over decades, operational risk functions have accumulated large datasets of event data for internal, regulatory compliance and capital calculation purposes. These datasets often contain free text descriptions which can potentially contain a wealth of information that is currently untapped.

While reviewing them individually is time intensive to the point where it becomes economically unviable, machine learning techniques for text-mining can be used to gain deeper insight into the loss history, e.g. on root causes and risk drivers.

This is just one example of how machine learning can provide deeper insights into data, and it is also potentially one of the most straightforward uses on an existing data resource. For this reason, ORX has begun a research study focusing on this area of machine learning application.

Automated categorisation

In a related application, free-text descriptions can be used further to classify loss data where existing categories have been added partially, inconsistently or when underlying taxonomies have been subsequently changed.

3. Supporting business needs effectively

To assume a more commercial role that supports the front line in day-to-day business activities, operational risk needs to be able to provide more and more easily consumable information. Gaining deeper insights into existing data will not only help operational risk functions, but will also support the front line in their operational risk management activities.

In addition to providing deeper insights, machine learning techniques will allow a wider range of information to be taken into account, and condense the information to facilitate decision-making processes. It is important to keep in mind that the user will typically require at least a basic understanding of the application to develop a certain level of trust into this kind of data augmentation. As in the case of other types of analytical and IT applications, employee buy-in is key for their successful implementation.

4. Gaining the skills for an enhanced level of challenge

With the increasingly widespread use of machine learning and AI, their application can be expected to become more of a regulatory focus. As part of this, operational risk professionals will need the skills to be able to challenge the use of these techniques with the business. Having experience with their different uses will provide operational risk teams with the necessary understanding to enable them to provide to do this effectively.

It will also allow them to assume a leadership role in conversations around transparency, and provide the necessary impulse for the financial institutions to actively engage with customers and regulators over their internal use of machine learning techniques.

5. Benefitting from economies of scale

The use of machine learning techniques for making data collection, handling and analysis for operational risk more efficient will also bring wider benefits to the whole firm.

Where these techniques are used to improve existing preventative controls, such as fraud detection, considerable gains can be made even from incremental increases in predictive accuracy. Moreover, if firms can reduce the time needed for mundane tasks, such as loss data entry and categorisation, this will benefit not just operational risk functions, but also first line staff. In these cases, organisations will benefit from the economies of scale realised through even small improvements.

"Potentially, the biggest gains from the implementation of machine learning can be expected to come from the reduction and elimination of time-intensive and repetitive tasks, which take up valuable time of operational risk teams."

Use cases

In early 2019, ORX established a Machine Learning Working Group dedicated to exploring the impact of machine learning on operational risk. The group consists of representatives from nine ORX member institutions who are already using machine learning techniques for operational risk and are seeking to further increase the scope of their uses.

As part of this work, we collected examples of use cases describing these applications. Our analysis of the use cases showed that there are three main areas in which machine learning techniques lend themselves to support operational risk:

1. Fraud detection
2. Text-mining for data augmentation
3. Data quality assurance

Three areas where machine learning can support operational risk

Fraud detection

- Improve processes
- Reduce false positives/ negatives

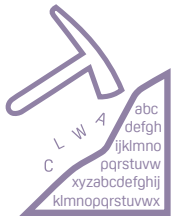

Text-mining

- Deal with large amounts of data
- Analyse free-text descriptions

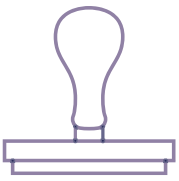

Data quality assurance

- Reduce effort
- Identify duplicated entries
- Accurately identify outliers

ORX Text-mining research study

The purpose of this research study is to make a case for the use of text-mining machine learning techniques on operational risk loss descriptions.

Historical event descriptions are an unstructured data resource that can contain valuable information not formally captured elsewhere. This information might be inferred from free text descriptions, but it may not be economical to manually review them.

Machine learning techniques can be used to create clusters of terms with similar meanings to establish whether this can provide deeper insight into the data.

If you would like more information about this research study, please contact Annika Westphal: annika.westphal@orx.org

Machine learning in operational risk:
Key uses and case studies

1. Fraud detection

The use of machine learning techniques for fraud and anti-money laundering detection is a well-known, almost "classic" application of machine learning to manage these operational risks.[3] The labelling of financial transactions as either suspicious or innocuous is one of the prime examples of how advanced analytics has sped up processes and reduced resource requirements.

Machine learning techniques can help improve these processes further and reduce false-positives and false-negatives (alerts of fraud or money laundering that turn out to be false alarms or instances where fraudulent transactions are overlooked).

Example use case

Card not presented fraud model

The firm developed a model to detect fraudulent transactions in real time where the card was not presented at the time of transaction.

This was an internal proof of concept model developed to replace a vendor model.

The motivation for using machine learning techniques was the extremely large data volumes that needed to be handled, which meant even small improvements in model prediction had a high potential for cost savings.

3 Zareapoor, Shamsolmoali. 2015. Application of Credit Card Fraud Detection: Based on Bagging Ensemble Classifier. Procedia Computer Science 48: 679–86 https://www.sciencedirect.com/science/article/pii/S1877050915007103?via%3Dihub

2. Text mining for data augmentation

Effective operational risk management relies on large amounts of data for risk quantification and risk management. This includes internal loss data, control libraries, internal risk indicators and other internal data, as well as external loss data and macroeconomic data.

Maintaining these datasets can be time intensive, especially where it involves the categorisation of individual entries. We have seen financial firms using machine learning techniques to analyse free-text descriptions, for example of loss events, to complete the categorisation of data entries.

Example use case

NLP tagging of loss incidents

The firm used NLP tagging of losses to infer root causes from already existing tags and free text descriptions. Within the internal dataset of loss incidents, some attributes or dimensions of the incidents had been added as additional requirements over time, and therefore had only been collected from a certain time onwards.

Understanding root causes of incidents enables better risk mitigation strategies, but analysing root causes individually from free text is a time intensive, manual and somewhat subjective process.

3. Data quality assurance

With growing datasets, the data and quality assurance become more time intensive. Machine learning techniques can help reduce this effort by identifying duplicated entries and identifying outliers more accurately.

Example use case

Idiosyncratic event identification

The firm used an isolation forest model as part of its capital allocation model. The model was designed to identify events unlikely to occur again because of enhanced governance or controls. The model was still being implemented at the time this report was written.

How to pre-empt potential challenges

To implement machine learning applications effectively, operational risk functions should consider the following challenges that may arise in order to address and pre-empt them early.

Uncertainty about the returns on investment

While financial firms are discovering new areas of application for machine learning and AI, senior management can remain sceptical about the expected returns on investment. Experience with these applications is often still limited – especially in operational risk. This makes it all the more important to clearly define the scope of their application and their expected benefits. Having a compelling proof of concept and a project that is clearly defined is likely to have more resonance with senior management, and will lead to increased long-term benefits.

Small-scale applications to reduce the time spent on mundane tasks and make existing processes more efficient are good starting points. Moreover, off-the-shelf applications and freely available machine learning libraries offer a low-cost way of helping staff to gain experience with these techniques.

Gaining employee buy-in

No discussion of the implementation of machine learning techniques is complete without a discussion of the importance of employee buy-in. Gaining employee and user buy-in will mean that the implemented technology is used to its full potential, and that any queries and issues are raised quickly. It will also mean that supervised learning applications, which rely on user input, are adequately trained to support their successful implementation.

Depending on the type of application, fear among employees of becoming obsolete needs to be addressed during this phase. In this case, it is important to clearly communicate that any process automation will enable them to apply their skills in other areas that require cognitive abilities, and to spend their time doing more engaging and valuable work.

It'll also mean that front-line employees using, for example, the outputs from models augmenting existing data for decision making will feel comfortable with using these. All of this will ensure expected ROIs are realised and will help to make the case for widening the scope of proof of concepts and pilot projects.

Challenges of implementing machine learning techniques

Uncertainty about ROIs	Scarcity of machine learning skills	Gaining employee buy-in	Ethical considerations	Regulatory focus
				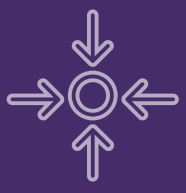

Solutions and answers

Uncertainty about ROIs	Scarcity of machine learning skills	Gaining employee buy-in	Ethical considerations	Regulatory focus
Clearly defined pilot projects Small-scale proof of concepts Off-the-shelf applications	Operational risk has the quantitative skills to build on Developing ML skills in house to explore use of techniques	Meeting potential concerns early Making sure users feel comfortable and understand the applications	Gaining first-hand knowledge about ML (and AI) Starting the conversation early	Proactively engaging about potential areas of concern Pooling industry knowledge

Internal ⟷ External

Scarcity of machine learning skills

With the increased attention on advanced analytics, more industries are trying to benefit from opportunities that machine learning has to offer. This also means that employers are competing more than ever to attract the people with the right quantitative skill set and knowledge of how to implement the techniques.

With this already scarce talent being increasingly more in demand, many organisations will have to develop these skills in house. Thanks to its historically strong focus on risk quantification, operational risk teams already have the necessary quantitative skills that they can build on and develop further. This is an opportunity operational risk functions should be conscious of in order to avoid being left behind.

Ethical considerations

Machine learning applications typically require large amounts of data, which leads to two potential areas of concern. On the one hand, there is a question around the data that can be used to feed the algorithms. On the other hand, when it comes to customer data, there is a general awareness among financial firms that they need to tread carefully. This has become an even more pressing matter since the introduction of the EU's General Data Protection Regulation (GDPR).

Moreover, users of the applications need to have an understanding of how the algorithms work. Not only does this enable users to trust the applications, it also allows them to detect outcomes that may not be in line with what they should expect.

Machine learning techniques can lead to new insights on customer behaviour and help to improve operational risk management – but each dataset needs to be carefully assessed as to its use. Adequate training on the use of the application is important to help users gain the necessary understanding of the application.

The ethical use of AI is also an area where we can expect increased regulatory scrutiny in the future.[4]

Pre-empting regulatory challenge

With the general expectation that machine learning and other advanced analytical techniques will be an increased regulatory focus, proactively engaging with regulators early over the areas of application is essential. To play an active role in these conversations, operational risk as a discipline can benefit from pooling resources and engaging with peers over important aspects of machine learning. ORX is assisting financial institutions in this collaborative effort.

Moreover, having hands-on experience with machine learning techniques will allow operational risk to successfully assume a leadership role in the industry and drive conversations around topics such as data usage, potential biases in machine learning algorithms, and transparency of the models used.

Did you know?

ORX Membership gives you access to:

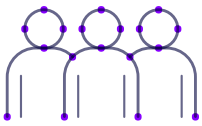

98
Member firms

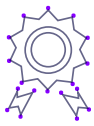

17
years' experience

700k+
loss events in our database

Find out more at www.orx.org

4 https://www.mas.gov.sg/~/media/MAS/News%20and%20Publications/Monographs%20and%20Information%20Papers/FEAT%20Principles%20Final.pdf

Conclusion

This white paper lays out the main opportunities that ORX believes machine learning techniques have to offer for operational risk. It is intended as a useful resource for operational risk functions who are interested in making use of these techniques.

As addressed at the beginning of this white paper, we believe that machine learning will be a central element to any future strategy for managing operational and non-financial risks. Beyond this, there will be the further impacts that advanced operational analytics will have on business strategy.

Operational risk is created through failed processes, systems and human error, and across many contexts. The application of machine learning and other advanced analytics holds promise to enable operational risk functions to help businesses implement strategies more intelligently, and to achieve their goals in a sustainable manner.

In addition, using machine learning and AI in various areas will create a new type of operational risk exposure, sometimes shifting concentrations of these exposures to areas within the business that might not have been initially anticipated. Operational risk teams having experience with these techniques will be able to assist businesses in capturing a global view of where these risks are and enable them to operate effectively in the changing realities created by the fourth industrial revolution.

"The opportunities that machine learning techniques offer – from task optimisations and better use of resources, to cost savings and gaining deeper insights into data – are considerable, but often not used to their full potential."

"In addition, thinking beyond the primary benefits of resource optimisation and cost savings, operational risk has a central role to play in the management of new risks that Artificial Intelligence and automation bring with them."

www.orx.org
+44 (0)1225 904 397

Managing risk together

ORX believes many heads are better than one. We're here to bring the best minds of the international operational risk community together. By pooling our resources, sharing ideas, information and experiences, we can learn how best to manage, understand and measure operational risk and become less vulnerable to losses.

We work closely with over 90 member firms to develop a deeper understanding of the discipline and practical tools. We set the agenda, maintain industry standards, and garner fresh insights.

ORX is owned and controlled on an equal basis by its members.

For more information about ORX, visit our website at **www.orx.org**

Authors

Dr Luke Carrivick
Head of Research and Information
luke.carrivick@orx.org

Annika Westphal
Statistical Assistant Manager
annika.westphal@orx.org

www.orx.org

@ORX_Association

@ORX_association