

Issue N. 29 - 2025

ARGO

New Frontiers in **Practical Risk Management**



This is a creation of **iason**.

The ideas and model frameworks described in this document are the result of the intellectual efforts and expertise of the people working at **iason**. It is forbidden to reproduce or transmit any part of this document in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose without the express written permission of a company in the **iason Group**.



Argo magazine

Year 2025 - Issue Number 29

Published in July 2025

First published in October 2013

Last published issues are available online:

<http://www.iasonltd.com>

Front Cover: **Atanasio Soldati**, *Untitled*, 1952.



ESSENTIAL SERVICES FOR
FINANCIAL INSTITUTIONS



Editors:

Antonio CASTAGNA (Managing Partner)

Luca OLIVO (Managing Director)

Executive Editor:

Giulia PERFETTI

Graphic Designer:

Lorena CORNA

Scientific Editorial Board:

Gianbattista ARESI

Michele BONOLLO

Alessandro CAPPO

Marco CARMINATI

Antonio CASTAGNA

Dario ESPOSITO

Massimo GUARNIERI

Antonio MENEGON

Luca OLIVO

Giulia PERFETTI

Massimiliano ZANONI

Francesco ZORZI

Milan Headquarter:

Corso Europa, 15

20122 Milan

Italy

London Headquarter:

3rd Floor, 120 Baker Street

W1U 6TU London

United Kingdom

Madrid Headquarter:

Calle Miguel Ángel, 16

28010 Madrid

Spain

Contact Information:

info@iasonltd.eu

www.iasonltd.com

iason is a registered trademark.

Articles submission guidelines

Argo welcomes the submission of articles on topical subjects related to the risk management. The articles can be indicatively, but not exhaustively, related to models and methodologies for market, credit, liquidity risk management, valuation of derivatives, asset management, trading strategies, statistical analysis of market data and technology in the financial industry. All articles should contain references to previous literature. The primary criteria for publishing a paper are its quality and importance to the field of finance, without undue regard to its technical difficulty. *Argo* is a single blind refereed magazine: articles are sent with author details to the Scientific Committee for peer review. The first editorial decision is rendered at the latest within 60 days after receipt of the submission. The author(s) may be requested to revise the article. The editors decide to reject or accept the submitted article. Submissions should be sent to the technical team (info@iasonltd.eu). \LaTeX or Word are the preferred format, but PDFs are accepted if submitted with \LaTeX code or a Word file of the text. There is no maximum limit, but recommended length is about 4,000 words. If needed, for editing considerations, the technical team may ask the author(s) to cut the article.

Table of Contents

Editorial	p. 5
------------------	------

Just in Time - iason Notes	p. 7
-----------------------------------	------

INSURANCE RISK

Risks Aggregation, Tail Dependence and Beyond	About the Authors	p. 13
<i>F. Bertonati, D. Bruttomesso, A. Cavazzana, Z. Karfa and G. Mori</i>	Introduction	p. 16
	Risk Aggregation: Regulatory Evolution	p. 16
	Risk Aggregation: Supervisory Standards	p. 19
	Measures of Dependency	p. 20
	Aggregation Methodologies	p. 22
	Tail Dependence Methodologies	p. 25
	Impact Analysis of Tail Dependence	p. 31
	Conclusions	p. 33
	References	p. 35

MARKET RISK

A Comparison of Advanced Methods for the Quantile Estimation in the Risk Management Field	About the Authors	p. 39
<i>M. Bonollo and L. Mastrototaro</i>	Introduction	p. 40
	Financial Context. The DRC Risk Measure	p. 40
	The Quantile Estimation	p. 41
	Comparison of the Different Estimators	p. 42
	Conclusions	p. 44
	References	p. 45

INNOVATION

AI Risk Management Frameworks	About the Authors	p. 47
<i>N. Mazzoni, S. Martucci and M. Ranieri</i>	Introduction	p. 49
	AI: a Brief View	p. 49
	Risks and the Increasing Need for Regulatory and Risk Management Frameworks	p. 53
	AI Regulatory Framework	p. 57
	AI Risk Management Frameworks	p. 60
	Conclusions	p. 66
	References	p. 67

DEAR READERS,

We are pleased to welcome you to this new edition of our Argo series, dedicated to current developments in the banking and insurance sectors.

This edition delves into one of the most pressing challenges of modern finance: understanding and managing risk in an increasingly complex, data-driven, and regulated environment. Our featured articles explore a spectrum of topics — from the statistical modeling of extreme financial events to the evolving regulatory landscape for AI technologies — with a shared emphasis on methodological precision, institutional resilience, and forward-looking governance.

*We open this issue with “**Risks Aggregation, Tail Dependence and Beyond**” by G. Mori et al. This contribution examines the growing importance of accurately modeling risk dependencies, particularly tail dependencies, within financial institutions. Grounded in the context of the Solvency II framework and EIOPA’s supervisory role, the article explores how traditional aggregation methods may understate systemic exposures during extreme events. By integrating regulatory expectations with rigorous quantitative analysis, the authors provide a timely perspective on capital adequacy, stress testing, and risk model validation.*

*Continuing with the topic of quantitative rigor, “**A Comparison of Advanced Methods for Quantile Estimation in the Risk Management Field**” by M. Bonollo and L. Mastrototaro analyzes the challenges of modeling Default Risk Charge (DRC) under the new FRTB regime. The article addresses a critical but often underexplored issue: the reliability of extreme quantile estimation in finite sample settings. By benchmarking traditional approaches against more advanced statistical estimators on real-world data, the authors offer practical insights for modelers and risk managers navigating regulatory implementation under conditions of uncertainty.*

*Expanding our lens to the technological transformation of finance, “**AI Risk Management Frameworks**” by S. Martucci, N. Mazzoni and M. Ranieri tackles the growing integration of Artificial Intelligence, particularly Generative AI, into financial ecosystems. While AI adoption accelerates across sectors, it introduces complex risks related to transparency, accountability, and regulatory compliance. This article offers a comprehensive overview of both global and regional regulatory initiatives, with a focus on the EU AI Act. It also outlines structured risk management frameworks that institutions can deploy to address AI-specific challenges such as model bias, systemic instability, and governance gaps.*

We close this issue by recommending a visit to our online Research section, where you can also subscribe to our newsletter for monthly insights on practical risk management topics.

We wish you a happy reading!

Antonio Castagna

Luca Olivo

Giulia Perfetti



Just in Time

Jason Notes

Bank of Italy: Modelling Transition Risk-adjusted Probability of Default



The paper introduces a novel methodology to estimate the impact of climate-related transition risk on the one-year probability of default for Italian non-financial firms. To this end, the authors construct a detailed dataset that integrates information from the EU Emissions Trading System (EU-ETS) with market and corporate financial data. Within the EU-ETS framework, firms with emissions exceeding (or falling below) their free allowances face additional costs (or generate revenues), which can influence their creditworthiness.

[read more](#)

Date July 2025

EBA - Report on the 2024 Market Risk Benchmark Exercise



The EBA Report Results from the 2024 Market Risk Benchmarking Exercise presents the results of the 2024 supervisory benchmarking exercise pursuant to Article 78 of the Capital Requirements Directive (CRD) and the related regulatory and implementing technical standards (RTS and ITS) that define the scope, procedures and portfolios for benchmarking internal models for market risk (MR). The report summarizes the conclusions drawn from a hypothetical portfolio exercise (HPE) conducted by the EBA during 2023/24.

[read more](#)

Date May 2025

EBA - Report on the 2024 Credit Benchmark Exercise



The EBA's 2024 Credit Risk Benchmarking Exercise evaluates differences in risk-weighted assets (RWAs) among banks using the Internal Ratings-Based (IRB) approach. The assessment reviews progress in implementing the IRB roadmap and its impact on enhancing consistency across institutions. The results show notable progress, including a rise in approvals of significant model changes across various asset classes. However, full alignment remains a work in progress, as many institutions are still awaiting final model validations or are in the process of implementation. This limits the ability to observe consistent trends in RWA variability. Probability of Default (PD) variability has decreased significantly across most asset classes, largely due to regulatory efforts and model standardization.

[read more](#)

Date May 2025

Artificial Intelligence: Financial Industry Market Overview



Artificial Intelligence, and in particular GenAI, is increasingly being adopted across a wide range of industries, with significant capital investment and a growing number of business functions being reshaped. The financial industry has traditionally been among the earliest adopters of technological innovation, aiming to enhance productivity and improve operational efficiency. The same trend holds true for AI and GenAI, where the industry continues to be at the forefront of adoption.

[read more](#)

Date July 2025

ICT Risk: Focus on DORA



In 2025, the EU reached a key milestone in its digital finance strategy with the entry into force of the Digital Operational Resilience Act (DORA) and the Markets in Crypto-Assets Regulation (MiCAR). As key components of the Digital Finance Package, these regulations aim to create a harmonized framework that strengthens both the security and resilience of the EU financial markets while supporting its technological innovation. DORA establishes a unified framework to ensure that financial entities across the EU can effectively manage ICT risks, structured around five pillars.

[read more](#)

Date June 2025

CBDCs Worldwide Projects Overview



Following the COVID-19 era, the shift towards cashless payments has evolved even more rapidly than before. In this context, in order to avoid a reduction of their role within financial systems, central banks have started to analyze the integration of new financial innovations into their infrastructure. As a result, many central banks started assessing the issuance of central bank money in digital form. These are the so-called CBDCs, a type of digital money issued, complementary to cash, directly by a central authority. In recent years, several projects have been launched and are currently under assessment by central banks across the world, with some already deployed.

[read more](#)

Date June 2025

Analysis of ESG Disclosure on a Sample of Italian and European Banks



In December 2023, Bank of Italy published the main findings of its analysis of the "Accounting impacts and disclosures of ESG risks for a sample of Italian and European banks". Based on the European Commission's Sustainable Finance Action Plan, banking and financial supervisory authorities have, in fact, progressively intensified their efforts to assess the degree to which ESG risks are integrated into the business processes of financial intermediaries.

[read more](#)

Date April 2025

iason Weekly Insights

Regulatory/Supervisory Pills



Among iason's various publications we also find the iason Pills.

With these daily Pills, iason aims to offer a summary on information, mostly, of the main regulatory and supervisory news in the banking and finance sector on both Pillar I and Pillar II risks of the Basel framework. The main purpose of these publications is to give the reader an effective, timely and brief overview of the main topics of the moment.

The author of the Iason Pills is Dario Esposito.

[read more](#)

Market View



Among iason's weekly insight you can also find the iason Market View, a weekly update on financial market by Sergio Grasso.

The author, with almost three decades of investment experience, presents an accurate analysis of market fluctuations of the week, giving a critical view of observed phenomenos and suggesting interesting correlations with the main world events.

[read more](#)

GOVERNANCE. METHODOLOGY. TECHNOLOGY.

iason is a company specialised in advanced solutions for the **Risk Management** of **Financial Institutions**.

We provide highly qualified **consulting services** in the **methodological** and **technological** fields, together with targeted support for **Data** and **Model Governance** projects in risk frameworks.

We strongly believe in **Research** because we want to guarantee our clients services and solutions that are always at the forefront of **Regulatory** and **Modelling** requirements.

ARGO MAGAZINE

Quarterly magazine on the new frontiers of Risk Management

RESEARCH PAPER SERIES

Research articles on innovative and advanced methodologies in the financial sector

JUST IN TIME

Real-time updates on regulatory changes

FOLLOW US!





Risks Aggregation, Tail Dependence and Beyond

About the Authors



Francesco Bertonati:

Quantitative Analyst

Graduated in Mathematics for Finance and Insurance, he currently works as a Quantitative Analyst and collaborates with the Risk Management team of one of the leading insurance companies in Italy. He is involved in the design and development of the internal model, with a particular focus on modeling the interdependencies between risk factors and drafting the model's technical documentation. Additionally, he contributes to the analysis and assessment of default risk and risks related to Underwriting activities in the Life insurance sector.



Daniel Bruttomesso:

Quantitative Analyst

He holds a degree in Mathematical Engineering and is currently cooperating as a Quantitative Analyst with the Risk Management team of one of the leading Italian insurance companies. He is involved in the development of the internal model, with a particular focus on modeling dependencies between different risk factors and preparing the related technical documentation. He is also involved in assessing underwriting risks related to Life insurance, as well as evaluating counterparty risk.





Amina Cavazzana:

Analyst

She holds a degree in Statistics for Economics and Business and is currently working as a Credit Risk Analyst, collaborating with the Risk Management team of one of the leading Italian insurance companies, supporting the development and implementation of the internal model. Her main activities include counterparty risk modeling, analysis of dependencies between risk factors, and the preparation of the official model documentation. She is also engaged in managing Underwriting/Life risks.



Zineb Karfa:

Analyst

She holds a Master's degree in Data Science for Economics and Finance and another in Economic Data Analytics. Currently working as an Analyst within the Risk Management team of a leading Italian firm, she focuses on developing risk models for asset management, with a particular emphasis on Value-at-Risk and time series analysis. She also contributes to the preparation of internal model documentation and supports activities related to counterparty risk, risk factor dependencies, and regulatory compliance in both banking and insurance contexts.





Gianmarco Mori:

Senior Quant Analyst

He holds a Bachelor in Banking and Finance and a Master degree in Statistics and Actuarial science. He obtained a specialization in Quantitative Finance at Politecnico di Milano and he had worked in a consulting firm for two year focused in Credit and Liquidity Risk Management. After these years, he moved to banking sector in Risk Management area and currently he has been working in Iason as financial engineer for some of the major italian players.



Risks Aggregation, Tail Dependence and Beyond

Francesco Bertoniati

Zineb Karfa

Daniel Bruttomesso

Gianmarco Mori

Amina Cavazzana

In an era marked by heightened uncertainty and intricate interconnections across financial systems, robust risk aggregation and dependence modelling have become critical pillars. Within the European insurance and banking sectors, capturing inter-risk dependencies, especially in the tails of distributions, is fundamental for assessing systemic vulnerabilities and ensuring solvency. This paper opens with an overview of the historical and regulatory evolution of risk aggregation practices in the financial sector, with a specific focus on insurance. Particular attention is given to the development of the Solvency II framework and the supervisory role of EIOPA, laying the groundwork for a more in-depth examination of dependence structures and the statistical tools used to model joint risk behaviour. A key focus is on tail dependence, a concept of paramount importance in the context of extreme loss events and capital adequacy. The discussion combines theoretical foundations with empirical illustrations to highlight how failing to properly account for tail dependence can lead to significant underestimation of aggregate exposures. The final section provides the methodologies to conduct a comparative impact analysis under the Solvency II capital requirements regime. By integrating regulatory context with quantitative rigour, the paper contributes to ongoing efforts to improve capital adequacy, model validation, and stress testing practices. This contribution is particularly relevant today, as both the banking and insurance industries face growing pressure to adopt models that better reflect real-world joint risk behaviour.

IN recent years, the stability of financial systems has been increasingly challenged by global shocks, heightened complexity, and growing interconnections across institutions and markets.

These dynamics have elevated the importance of accurately modelling aggregated risk and capturing its dependencies, particularly in the tails of loss distributions, as a foundational element of both regulatory supervision and internal risk management. This is especially relevant in the insurance and banking sectors, where low-probability, high-severity events can compromise solvency and propagate systemic risk. In response to these challenges, regulators have responded with comprehensive frameworks aimed at ensuring resilience. The Solvency II directive, governing the European insurance industry, and the Basel III framework, guiding banking regulation, both emphasise the importance of solid capital requirements and sound risk aggregation practices. Central to these efforts is the challenge of capturing inter-risk dependence, particularly in the tails of loss distributions, where traditional correlation-based approaches often fall short.

This paper explores the historical and regulatory development of risk aggregation, with a specific emphasis on the insurance sector under Solvency II and the supervisory role of EIOPA. It proceeds to formalise a range of statistical dependence measures and to introduce various risk aggregation methodologies, emphasizing the limitations of linear correlation in favour of more reliable representations. A detailed treatment of tail dependence metrics follows, highlighting their significance in capturing extreme co-movements and informing stress testing and scenario analysis. The final section evaluates how alternative dependence structures and aggregation methods influence capital requirements under Solvency II and Basel III. This provides a framework for enhancing model validation practices and strengthening the resilience of financial institutions in the face of tail events and systemic shocks.

Risk Aggregation: Regulatory Evolution

The aggregation of risks has become a cornerstone of modern financial regulation, especially as financial systems grow more complex and interconnected. This chapter traces the evolution of regulatory approaches to risk aggregation, examining how major frameworks like Basel and Solvency have adapted to better capture dependencies and systemic vulnerabilities.

Evolution of Risk Aggregation and Regulatory Motivation

The evolution of regulatory frameworks for risk aggregation in the banking and insurance sectors has been driven by successive financial crises and the growing complexity of financial systems.

The European Union (EU) has progressively enhanced its regulatory framework to address systemic vulnerabilities, emphasizing the importance of understanding extreme co-movements in financial markets. Tail dependence, which captures the propensity of asset returns to exhibit extreme co-movements, poses significant challenges to traditional risk management approaches that often assume normality and independence.

In response to these challenges, European regulatory bodies have implemented measures to bolster financial stability. The European Central Bank (ECB), in its *Financial Stability Review - November 2024* [23], highlights the increasing interconnectedness of financial institutions and the potential for contagion during periods of market stress. Moreover, the adoption of advanced risk assessment tools, such as copula-based models, has been encouraged to better capture the complexities of joint extreme events.

Two dominant frameworks, the Basel Accords for banking

and the Solvency frameworks for insurance, have emerged as the cornerstones of global financial regulation. Though initially distinct, these regimes have progressively converged in their recognition of systemic risk, the need for sophisticated modelling techniques, and the integration of forward-looking stress scenarios.

Figures 1 and 2 present the key regulatory milestones in the Basel and Solvency frameworks, respectively. The timelines highlight the parallel evolution of banking and insurance regulation, from foundational directives to the most recent reforms, illustrating how both regimes have responded to financial crises and advanced toward risk-based capital standards.

The regulation of financial institutions has traditionally relied on models that assume linear correlations, normal distributions, and relative independence of risk factors. However, as financial markets became more interconnected and volatile, regulators began to acknowledge the shortcomings of these assumptions, particularly during crises, when asset returns tend to move together in extreme and non-linear ways. This phenomenon, known as tail dependence, reflects the heightened probability of joint extreme losses across institutions, asset classes, or geographies. Understanding and managing aggregated risks in this context has thus become essential to the stability of both banks and insurers.

To better understand the trajectory of regulatory development, we begin with the banking sector's response to risk aggregation challenges under the Basel framework.

Regulatory Frameworks in Banking: The Basel Accords

International Convergence of Capital Measurement and Capital Standards (1988) [6], also known as Basel I Accord, marked the first international attempt to standardise capital requirements, introducing a simple, additive system based on fixed risk weights. However, this model assumed linear relationships between exposures and failed to recognise correlations or tail events. As a result, it underestimated systemic vulnerabilities, particularly in times of financial stress.

In response to these shortcomings, *Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework* (2004) then revised in *International Convergence of Capital Measurement and Capital Standards: A Revised Framework - Comprehensive Version* (2006) [7] introduced a more nuanced and risk-sensitive approach. It leveraged internal ratings-based models that account for Probability of Default, Loss Given Default, and Exposure at Default. Although this marked significant progress, it did not directly address tail dependence. Instead, it relied on Gaussian copulas and normal distributions in portfolio models, which proved inadequate in capturing extreme joint losses, as revealed during the 2008 financial crisis.

Building on these developments and considering the global financial crisis *Revisions to the Basel II market risk framework* (2011) [8], known as Basel 2.5, introduced the Stressed Value-at-Risk (VaR) and the Incremental Risk Charge. *Basel III: A global regulatory framework for more resilient banks and banking systems* (2011) [4], which came into effect after 2013 with the implementation of the *Capital Requirements Regulation (EU) No 2013/575* [28], represented a substantial enhancement to regulatory rigor. It introduced the Liquidity Coverage Ratio, Net Stable Funding Ratio, capital conservation buffers, and leverage ratios. Furthermore, Basel III began a transition from VaR to Expected Shortfall (ES) under the *Fundamental review of the trading book: A revised market risk framework* [5], thereby improving the framework's ability to capture tail risk. Although tail dependence was not formally embedded, Basel III's stress testing frameworks simulate co-movements under systemic stress and promote more realistic modelling of non-linear dependencies.

Regulatory Frameworks in Insurance: The Solvency Regimes

In parallel with the evolution of banking regulations, the insurance sector also underwent a transformation through the Solvency frameworks.

Directive 2002/13/EC [29] regarding the solvency margin requirements for non-life insurance undertakings and *Directive 2002/83/EC* [30] concerning life assurance, introduced Solvency I. This framework was based on simplistic, factor-based formulas and assumed independence between risk types, ignoring potential diversification and systemic aggregation effects, particularly during market disruptions.

Directive 2015/35/EU [26], known as Solvency II, replaced this static regime with a more dynamic and risk-based system. It introduced the Solvency Capital Requirement (SCR), calculated using either a standard formula or internal models, to capture risk exposures over a one-year horizon with 99.5% confidence. Crucially, Solvency II allowed the use of copula functions, especially t-copulas and Clayton copulas, to model asymmetric and nonlinear dependencies between risks. In addition, it mandated regular stress testing and scenario analysis, extending to climate and systemic shocks, and provided mechanisms for group-level aggregation and oversight.

Methodological Advances and Future Directions

To effectively capture aggregated risks, both Basel III and Solvency II allow banks and insurance companies to use internal models, supported by advanced statistical techniques and solid governance frameworks.

These internal models are subject to supervisory approval and rigorous validation procedures under both Solvency II and Basel III frameworks. This includes requirements to demonstrate the models' reliability, appropriate use within risk management practices, and compliance with quantitative and qualitative standards set by regulators.

Copula-based modelling and heavy-tailed distributions are among the techniques now commonly used to quantify joint extreme losses across risk types and asset classes. Model validation includes benchmarking, back-testing and supervisory audits. Where empirical data are insufficient, Basel III and Solvency II allow expert judgment-provided it is transparent and justifiable.

Complementing these regulatory efforts, academic and actuarial communities have made significant contributions to the improvement of risk aggregation methods.

Researchers such as Bernard and Vanduffel (2016) [9] have argued that correlation-based models tend to overstate diversification benefits when tail dependence is present. Similarly, Bruneton (2011) [10] criticised Gaussian copulas for their inability to accurately capture extreme co-movements. In response to these limitations, Marri and Moutanabbir (2021) [41] proposed the use of Generalised Archimedean Copulas to enhance the precision of capital aggregation modelling.

Despite methodological advances, the adoption of internal models remains uneven across jurisdictions and sectors. In the EU, larger financial institutions, especially those classified as Group 1, are more likely to use internal models, given their greater resources and more complex risk profiles. In contrast, smaller entities typically rely on standard formulas due to the high costs and regulatory demands of model development and approval. This pattern is particularly evident in the insurance sector, where, as of 2020, only 14 insurers in Italy had received full or partial internal model approval [34]. Among banks, internal model use is also more common among larger institutions, although regulatory developments have increased scrutiny and pro-

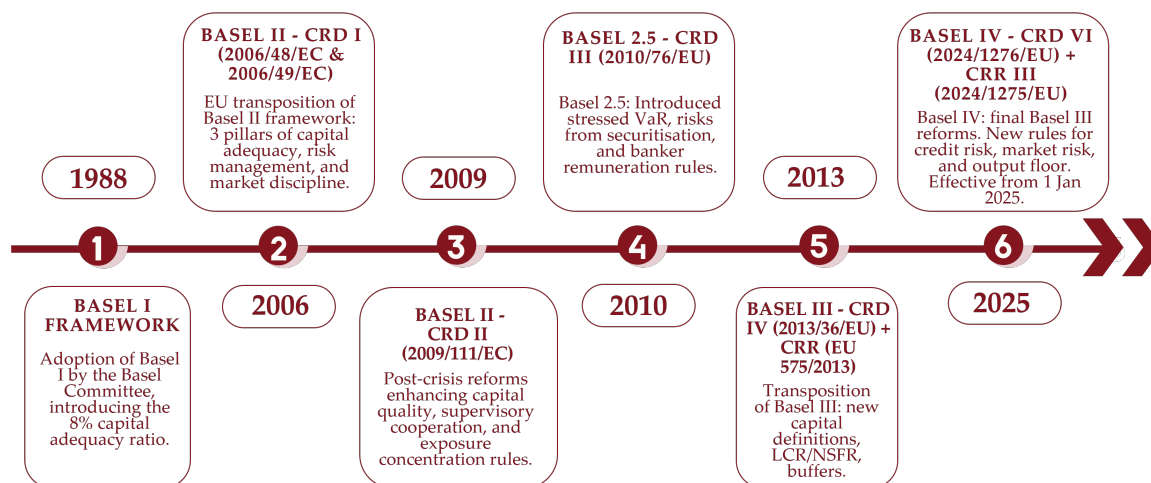


FIGURE 1: Evolution Timeline of Basel Regulatory Framework

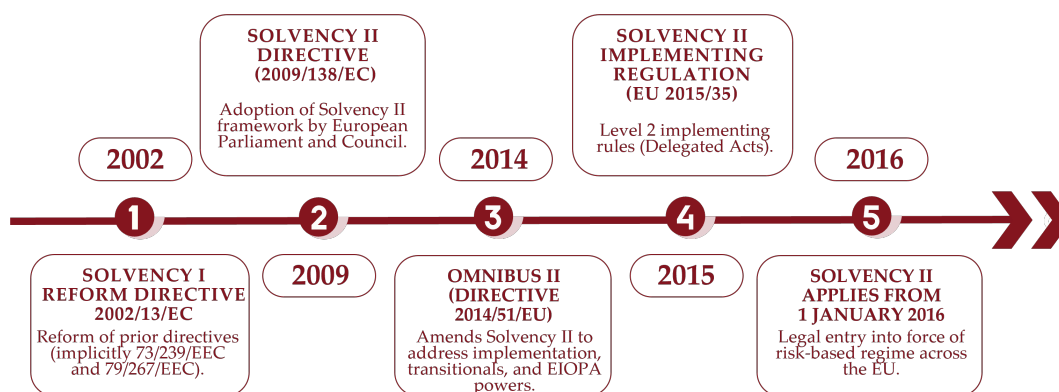


FIGURE 2: Evolution Timeline of Solvency Regulatory Framework

moted broader use of standardised approaches [20]. Although the Basel and Solvency frameworks serve different sectors, they are increasingly converging in their recognition of systemic risk, contagion, and tail dependence. The aftermath of the global financial crisis, and more recently the COVID-19 pandemic and a series of geopolitical shocks, have reinforced the need for regulatory models that move beyond average-case scenarios and address the complexities of extreme events.

Today, both frameworks are evolving to incorporate advanced methodologies that reflect these concerns. There is a growing emphasis on the use of copula models and heavy-tailed distributions within internal risk assessments, allowing institutions to more accurately capture dependencies and extreme co-movements among risk factors. In parallel, stress testing frameworks have become more sophisticated, aiming to model joint extreme losses across various business lines and risk types, thereby providing a more holistic view of institutional vulnerability.

Moreover, scenario aggregation techniques are being refined to better reflect the non-linearities in financial markets. These improvements help regulators and institutions alike to understand how risks can interact and amplify under adverse conditions, enhancing the reliability of risk management practices.

Looking ahead, the integration of emerging threats such as climate risks, cyber risks, and geopolitical instability is

expected to further influence regulatory models. This evolution will likely heighten the focus on non-linear tail dependencies and network-based contagion models, ultimately pushing both Basel and Solvency frameworks toward more dynamic, system-wide approaches to risk aggregation.

Environmental, Social, and Governance (ESG) factors are becoming increasingly central to regulatory policy and financial risk management. The United Nations Environment Programme Finance Initiative (UNEP FI) [52] has issued specific guidance on integrating ESG risks into insurance underwriting practices. Empirical evidence from Allianz Commercial [3] supports the predictive power of ESG metrics in assessing insurance risks and designing risk frameworks. Concurrently, artificial intelligence and machine learning (AI/ML) are reshaping internal modelling practices, introducing both new analytical capabilities and challenges related to governance and explainability. SupTech-technology-enabled supervision is being advanced by institutions like the ECB and the Cambridge SupTech Lab to enhance regulatory responsiveness and precision.

Another persistent challenge is model uncertainty, particularly under varying dependency structures. Embrechts et al. [19] introduced the concept of aggregation robustness to address this issue, while Cambou and Filipovic (2017) [12] provides further insights into improving the resilience of capital models under diverse assumptions.

Country/Region	Regulatory Framework	Supervisory Authority
European Union	Solvency II	European Insurance and Occupational Pensions Authority (EIOPA)
Switzerland	Swiss Solvency Test (SST)	Swiss Financial Market Supervisory Authority (FINMA)
United States	NAIC Risk-Based Capital (RBC)	National Association of Insurance Commissioners (NAIC)
Japan	Solvency Margin Ratio (moving to EV-based regime in 2025)	Financial Services Agency (FSA)
Singapore	RBC 2	Monetary Authority of Singapore (MAS)
China	C-ROSS / C-ROSS II	China Banking and Insurance Regulatory Commission (CBIRC)
Canada	LICAT (formerly MCCR for life, CCIR for P&C)	Office of the Superintendent of Financial Institutions (OSFI)

TABLE 1: Regulations and authorities in the World

Global Comparison

When comparing international approaches, it becomes clear that regulatory philosophies and capabilities vary widely. The following examples focus on the insurance sector, where differences in regulatory philosophy, model usage, and risk aggregation practices are particularly pronounced across jurisdictions.

As depicted in Table 1, the EU and Switzerland have embraced model-intensive, principle-based frameworks that formally recognise tail dependence and allow for group-level aggregation. In contrast, the United States' NAIC Risk-Based Capital (RBC) framework relies on rule-based formulas with limited flexibility for internal models or inter-entity diversification. In Asia, regulatory maturity is mixed: while Japan and Singapore move closer to Solvency II, China's C-ROSS system reflects early steps toward internal model alignment. Canada's MCCR framework, meanwhile, limits inter-risk diversification and enforces capital additivity across risk types.

Switzerland's Swiss Solvency Test (SST) stands out for its balance of flexibility and oversight. It allows both standard and internal model use and supports holistic diversification through cluster modelling and intra-group transactions subject to regulatory approval.

In conclusion, the future of risk aggregation regulation lies in continued adaptation and convergence. As global financial systems become more interconnected and exposed to complex systemic risks, from climate change and cyber threats to geopolitical instability, regulatory frameworks must evolve to embrace data-driven modelling, stronger governance mechanisms, and coordinated global oversight.

Risk Aggregation: Supervisory Standards

The practical implementation of dependence modelling and risk aggregation in the insurance and banking sectors rests on a sophisticated regulatory framework. Unlike the historical and structural perspectives outlined previously, this chapter identifies and discusses the most influential and current regulatory and technical sources that govern how institutions are expected to manage, validate, and disclose their aggregation logic. These sources include legal instruments, supervisory guidelines, technical papers and professional standards, which together form the operational scaffolding of modern solvency and risk capital regimes.

Regulatory Instruments and Supervisory Guidelines in Insurance

A foundational reference for the insurance sector is *Commission Delegated Regulation (EU) 2015/35* [26], which supplements Solvency II and establishes the correlation structure within the standard formula. Regulation prescribes correlation coefficients across risk modules and presumes a linear dependency framework that determines diversification benefits when aggregating market, credit, life, health, and non-life risks. This linear correlation method offers computational simplicity and ease of supervisory comparison but has been widely criticised for failing to capture the nonlinear and tail-dependent relationships observed during stress events. Alternative techniques, such as copula-based models, have been proposed in literature to address these shortcomings, although they introduce additional complexity and model risk. IVASS, Italy's insurance regulator, echoes and operationalises these Solvency II principles through its *Regolamento n. 32 IVASS del 9 novembre 2016* [36], which requires Italian insurers to demonstrate that their internal aggregation methods are comprehensive and forward-looking. The regulation emphasises that ORSA processes should explicitly incorporate inter-risk dependencies, including under adverse scenarios, and that these dependencies be regularly reviewed as part of the risk governance process.

In addition, *IVASS Regolamento n. 20/2016* [35] allows the use of independent experts to evaluate internal models, especially in case of complex aggregation methodologies that cannot be validated using standard tools. These external professionals, often drawn from actuarial, quantitative finance or audit backgrounds, provide technical assurance to IVASS on the soundness of models dealing with stochastic dependencies, copula calibration, or tail aggregation methods. This complements broader European-level initiatives led by EIOPA, including the *Guidelines on ORSA (EIOPA-BoS-14/259)* [24], which provide supervisory expectations for documenting and validating internal aggregation logs. *EIOPA's 2020 Study on Diversification in Internal Models* [25] offers further insights. The report investigates how insurers across Europe quantify diversification, noting wide variability in the implementation of copulas and correlation matrices. The study highlights that even within a harmonised framework like Solvency II, substantial discretion exists in calibrating dependency structures, which may result in materially different capital outcomes. Such discretion underscores the importance of model governance and supervisory benchmarking.

Regulatory Instruments and Supervisory Guidelines in Banking

For banking institutions, the *Capital Requirements Regulation (CRR, Regulation EU No. 575/2013)* [28] and the *Capital*

Requirements Directive (CRD IV, Directive 2013/36/EU) [27] transpose the Basel III framework into EU law. While the CRR defines Pillar 1 capital requirements using separate formulas for credit, market and operational risk, it does not allow cross-risk diversification. Aggregation, therefore, becomes essential under Pillar 2 via the Internal Capital Adequacy Assessment Process (ICAAP), as mandated by CRD IV. Supervisory authorities use the SREP (Supervisory Review and Evaluation Process) to assess how banks aggregate risks, evaluate diversification assumptions and determine if additional capital buffers are required. In practice, some supervisors have imposed Pillar 2 Requirements (P2R) or Pillar 2 Guidance (P2G) to counteract what they view as unjustified capital relief resulting from optimistic internal aggregation. The European Banking Authority (EBA) further reinforces this framework through its *2016 Guidelines on ICAAP and ILAAP* [22], which obligate banks to justify their risk aggregation techniques and quantify any diversification effects claimed. Banks must demonstrate through data, stress testing and scenario analysis that their dependencies are realistic and consistent with past risk behaviour. On a global level, the Basel Committee's 2017 finalisation of Basel III introduced an output floor that limits the benefits derived from internal model variability, thereby indirectly constraining the effects of optimistic dependency modelling. While the final floor is set at 72.5% of standardised capital requirements, it is being phased in from 50% in 2025, gradually increasing to 72.5% by 2030. This imposes a ceiling on the extent to which model-derived diversification can reduce capital and aims to ensure comparability across institutions. In the insurance domain, the *IAIS's Insurance Core Principle 16* [33], stipulates that internal models used for solvency assessments must integrate all material risks and dependencies, and their assumptions must be transparent, documented, and validated. Supervisors are encouraged to assess not only the quantitative integrity of dependency modelling but also the appropriateness of its use within enterprise risk management frameworks.

Divergent Approaches to Cross-Risk Diversification in Banking and Insurance

One of the most fundamental distinctions between the banking and insurance regulatory frameworks lies in their treatment of cross-risk diversification in capital requirement calculations. While both sectors aim to ensure financial resilience, their respective regulatory philosophies diverge sharply in terms of how risk aggregation is handled. The banking sector, under the Basel III framework, particularly within Pillar I, takes a notably conservative stance. Regulatory capital requirements are calculated for credit, market, and operational risks separately, and these are simply added together without allowance for any diversification effects between them. Even internal models are developed and validated separately, without an integrated cross-risk modelling approach. By treating each risk as independent in the regulatory capital calculation, Basel reinforces a preference for transparency, comparability, and operational robustness.

By contrast, Solvency II explicitly allows for diversification across risk modules. Recognising that insurance portfolios often involve a broad spectrum of interconnected risks, from market and credit to life, health, and non-life, the Solvency II standard formula explicitly incorporates a diversification framework.

Although Pillar I of the Basel framework disallows diversification effects, Pillar II, through the Internal Capital Adequacy Assessment Process, permits banks to adopt more nuanced aggregation approaches. However, the embrace of diversification is conditional and often heavily constrained.

Supervisors tend to be cautious, demanding strong evidence of diversification effects, especially under stress scenarios as emphasised by the European Banking Authority, in its *Supervisory Review and Evaluation Process Guideline* (2022) [21].

The differing treatment of diversification reflects fundamental differences in sectoral risk characteristics. The Basel framework prioritises comparability and avoids complex modelling of heterogeneous risk interactions, favouring simplicity and robustness. Solvency II, having developed later, takes a more integrated and calibrated approach to the insurer's portfolio of risks, recognising inter-risk diversification both in its standard formula and in internal model approvals.

Standards, Academic Contributions and Emerging Challenges

Regulatory discourse has also been enriched by professional standards and academic contributions. The Actuarial Standards Board in its *Actuarial Standard of Practice No. 55: Capital Adequacy Assessment*. [2] provides actuarial guidance on evaluating capital adequacy and explicitly instructs actuaries to consider risk correlations and dependencies in their solvency assessments. The standard further requires disclosure of the methods, data sources, and limitations used in dependency modelling, thereby enhancing the transparency of actuarial evaluations. Academic studies, such as those by Bernard and Vanduffel (2016) [9], critique existing models for neglecting tail dependence and advocate for multivariate frameworks using copulas. These authors also propose techniques to quantify the model risk inherent in selecting a particular dependency structure.

Together, these regulatory instruments, supervisory expectations and academic insights establish the baseline against which institutions must design and validate their risk aggregation methodologies. They also offer a comparative lens through which to evaluate the interplay between regulatory conservatism and modelling sophistication. The challenge for firms is to balance compliance with the prudential need for resilience, while preserving the economic benefits of legitimate diversification. This ongoing tension continues to shape the evolution of supervisory practices and the future direction of capital regulation.

Measures of Dependency

Before delving into specific measures, it is essential to understand why dependency matters and how it influences the overall risk profile of a financial institution. This chapter explores various tools used to measure and represent dependencies, ranging from traditional correlation metrics to more advanced structures suited for capturing complex and tail-dependent behaviours.

Statistical Meaning and Importance

In the context of risk modelling, particularly in insurance and banking, understanding and accurately representing dependencies among risks is of paramount importance. Risk modelling typically involves two foundational components:

- The marginal distribution of each individual risk;
- and the dependency structure that links these marginal distributions.

While estimating marginal risk distributions is already a complex task, modelling dependencies presents an even greater challenge. This is because dependencies encapsulate

the systemic, often hidden, relationships between risks, that can significantly alter the overall risk profile of an enterprise [19].

Dependencies can emerge due to macroeconomic factors such as inflation, interest rates, and exchange rates, which affect both assets and liabilities. For instance, inflation can erode the real value of assets while increasing the cost of claims, especially in long-tail insurance lines, thereby impacting both sides of the balance sheet. Likewise, dependencies may emerge from common exposure scenarios across different lines of business; for example, a single catastrophic event like a hurricane can simultaneously impact P&C (property and casualty) and life insurance portfolios. The danger of oversimplifying these interdependencies is that a model may produce an overly optimistic picture of a firm's overall risk, even when each marginal component appears reasonably assessed. This issue becomes especially relevant when dependencies intensify during stressed market conditions, a phenomenon well-documented during financial crises, where risk factors that once appeared uncorrelated suddenly moved in the same direction. Therefore, capturing and accurately representing dependencies is crucial for a realistic evaluation of enterprise-wide risk.

Mathematical Representation of Dependency

Ideally, one would represent all dependencies via a comprehensive system of equations, capturing every causal and correlative link. However, such a system is not only infeasible to construct but also impossible to parametrise accurately with available data. Consequently, statistical modelling relies on more tractable, but often imperfect, tools to approximate these relationships.

Traditionally, dependency between two risks is quantified using the linear correlation coefficient, a single scalar statistic. While useful in certain contexts, linear correlation is inadequate for capturing the full range of dependency structures. The term dependency structure is preferred over mere correlation when relationships are non-linear or when dependencies vary across the distribution, particularly in the tails. For instance, two risks might show moderate average correlation but exhibit near-perfect dependency in extreme loss scenarios. This tail dependency is especially relevant in economic capital modelling, where extreme outcomes drive capital requirements.

Types of Dependency Measures

To better understand the nature of dependence between financial variables, it is essential to distinguish among different types of dependency measures. These measures can be broadly categorised into two families: those based on linear correlation and those that rely on rank-based correlations, which are capable of capturing monotonic but potentially non-linear relationships. Figure 3 illustrates this classification, highlighting Pearson correlation for linear associations, and Spearman's rho and Kendall's tau for rank-based associations. Each of these measures offers distinct insights into dependency structure, and the choice of measure should be guided by the nature of the underlying data and the objectives of the analysis. In the following subsections, as described by Shaw and Spivak [48], we discuss the main ones.

Linear Correlation Coefficient

The Pearson linear correlation coefficient remains the most widely recognised measure of dependency. Defined for

pairs of random variables X, Y with finite variance, it quantifies the degree of linear association:

$$\rho^{Pearson} = Cor(X, Y) = \frac{\mathbb{E}[(X - \mathbb{E}(X)) (Y - \mathbb{E}(Y))]}{\sqrt{Var(X) Var(Y)}}. \quad (1)$$

A value of +1 or -1 indicates a perfect increasing or decreasing linear relationship, respectively, while a value of 0 indicates no linear relationship.

However, linear correlation has several critical limitations:

- It detects only linear relationships and fails to capture non-linear dependencies;
- A zero correlation does not imply independence;
- It is sensitive to the marginal distributions of the variables;
- It is not invariant under non-linear transformations;
- It requires finite variance, making it unsuitable for heavy-tailed distributions, like Lévy or Pareto, common in financial modelling.

These deficiencies limit the usefulness of correlation in capturing complex, non-linear and tail-dependent behaviours often seen in real-world risk data.

Rank Correlation Measures

To address these limitations, rank-based correlation measures such as Spearman's rho and Kendall's tau have been developed. These statistics measure the strength of a monotonic relationship between variables based on the ranks rather than the raw values.

Spearman's rho

Spearman's rank correlation coefficient is the linear correlation between the ranked values of two variables:

$$\rho^{Spearman} = \rho(F_X(X), F_Y(Y)), \quad (2)$$

where F_X and F_Y are the cumulative distribution function (i.e. ranks) of X and Y .

Unlike Pearson's correlation, which captures only linear relationships, Spearman's coefficient measures the strength of monotonic associations by assessing how consistently one variable increases or decreases with the other. It is invariant under monotonic transformations and it does not need any assumption about the distribution, making it a flexible choice when data do not meet the hypothesis of linear correlation [50].

Kendall's tau

Kendall's tau, is another non-parametric measure that assesses the concordance between pairs of observations. Given a sample of n paired observations, define:

- C : number of concordant pairs;
- D : number of discordant pairs;
- $S = C - D$.

Then:

$$\tau = \frac{2S}{n(n-1)}. \quad (3)$$

It provides a more intuitive understanding of dependency, quantifying the probability that variables move in the same direction [39].

Both measures share the advantage of being distribution-free, they do not rely on any assumptions about the marginal distributions of the variables. This property makes them particularly useful in non-parametric settings and for calibrating copulas from empirical data.

Nevertheless, these rank-based measures are not without their own shortcomings, which partly explains their less




FIGURE 3: Different types of correlations

frequent use in industry practice. First, rank correlations are less interpretable in economic or financial terms: while a Pearson correlation can be readily understood as a linear link, a Spearman correlation refers to a monotonic trend but does not convey the magnitude of change. Second, these measures are less tractable analytically, especially in high-dimensional settings, making it harder to integrate them into standard statistical and actuarial models.

Moreover, rank-based measures are often less efficient when the true relationship between variables is linear. In such cases, Pearson correlation provides more statistical power. Computationally, calculating Kendall's tau can also be intensive for large datasets, as it involves pairwise comparisons, making it impractical for some large-scale applications. In practice, while rank correlations are invaluable for understanding non-linear or ordinal relationships, their limited adoption stems from practical, interpretative, and computational challenges, rather than from lack of merit.

As risk management evolves, it becomes increasingly critical to move beyond simple correlation metrics and adopt more sophisticated dependency structures that reflect the real-world behaviour of risks, especially under stress scenarios.

Beyond Correlation: Introduction to Copulas

Copulas provide a sophisticated mathematical framework to describe the dependency structure between random variables independently from their marginal distributions [44]. In risk modelling, this aligns well with the typical two-step process: first, modelling each individual risk's marginal distribution and, second, specifying how these risks interact or co-move, especially under stress. The second step is where copulas offer distinct advantages over correlation matrices. Sklar's Theorem underpins the use of copulas by demonstrating that any multivariate joint distribution can be decomposed into its marginal distributions and a copula that binds them together [42]. This allows for the construction of joint distributions with specified marginal behaviours and tailored dependency characteristics, making copulas indispensable in the modelling of aggregate risk.

Copula Mathematics

A copula is a multivariate distribution function defined on the unit hypercube $[0, 1]^n$, which captures the dependence structure between random variables, independent of their marginal distributions. Specifically, an n -dimensional copula is a joint distribution function $C(u_1, \dots, u_n)$ of a random vector (U_1, \dots, U_n) , where each component U_k is uniformly distributed on $[0, 1]$. This means that for all $k = 1, \dots, n$, the marginal distribution satisfies:

$$P(U_k \leq u) = u \text{ for all } u \in [0, 1].$$

Definition 3.1 (Copula). A function $C : [0, 1]^n \rightarrow [0, 1]$ is called a copula if it satisfies the following properties:

- $C(u_1, \dots, u_n) = 0$ whenever at least one $u_i = 0$. This ensures the distribution function starts at zero;

- For every $i \in \{1, \dots, n\}$,

$$C(1, \dots, 1, u_i, 1, \dots, 1) = u_i,$$

where all arguments are 1 except the i -th coordinate;

- The function C is n -increasing.

For example, in the bivariate case ($n = 2$), for all $(a_1, a_2), (b_1, b_2) \in [0, 1]^2$ such that $a_k \leq b_k$ for $k = 1, 2$, we have

$$C(b_1, b_2) - C(a_1, b_2) - C(b_1, a_2) + C(a_1, a_2) \geq 0.$$

This condition ensures that C defines a valid joint cumulative distribution function.

Sklar's Theorem

Sklar's Theorem is a foundational result in the theory of copulas [49]. It provides a formal connection between multivariate distribution functions and their marginals.

Theorem 3.1 (Sklar's Theorem). Let $F(x_1, \dots, x_n)$ be the joint cumulative distribution function of a random vector (X_1, \dots, X_n) , with marginal $F_1(x_1), \dots, F_n(x_n)$. Then, there exists a copula $C : [0, 1]^n \rightarrow [0, 1]$ such that for every $(x_1, \dots, x_n) \in \mathbb{R}^n$, the following holds:

$$F(x_1, \dots, x_n) = C(F_1(x_1), \dots, F_n(x_n)).$$

Moreover, if the marginal distributions F_1, \dots, F_n are continuous, then the copula C is unique.

This theorem provides a mathematical framework that separates the modelling of the dependence structure (via the copula) from the modelling of the marginal distributions. That is, once the marginals are known or estimated, the copula aggregates them into a full joint distribution. Furthermore, Sklar's Theorem implies an important invariance property: if C is a copula for the random vector (X_1, \dots, X_n) , then for any set of strictly increasing transformations T_1, \dots, T_n , the same copula C describes the transformed vector:

$$(T_1(X_1), \dots, T_n(X_n)).$$

This invariance under strictly increasing transformations is analogous to properties of rank correlation measures like Kendall's tau and Spearman's rho, which are also invariant under monotonic transformations of the variables [17]. It makes copulas especially useful in modelling dependencies that are unaffected by scaling or marginal distributional changes.

Aggregation Methodologies

The aggregation of risks represents a crucial step in modern financial and actuarial modelling, bridging the gap between isolated marginal distributions and a consolidated assessment of total portfolio risk. This chapter explores the

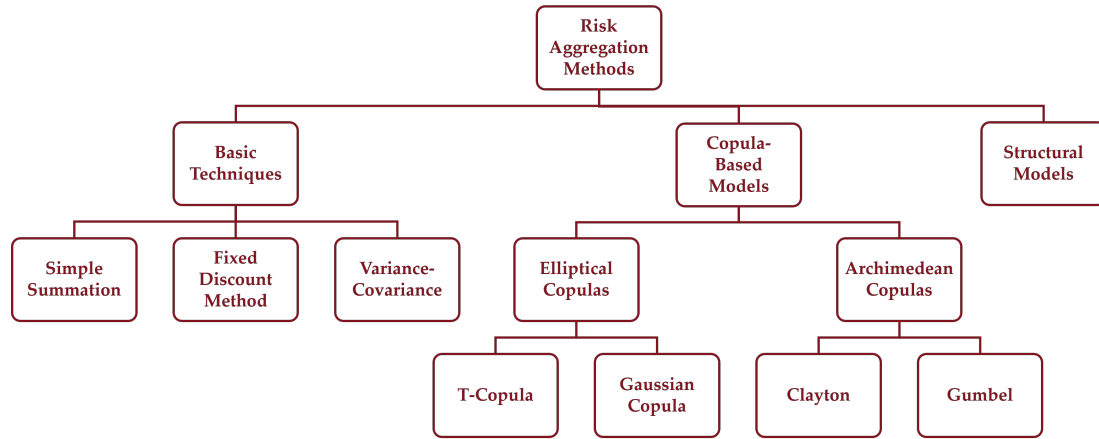


FIGURE 4: Overview of Risk Aggregation Methods

main methodologies used to aggregate risks across financial and insurance portfolios, with an emphasis on theoretical foundations, practical implementations and methodological rigour. A variety of methods have emerged, each with unique characteristics in terms of assumptions, computational complexity, regulatory acceptance and capacity to capture real-world dependencies. Figure 4 provides a high-level taxonomy of the principal methodologies employed in risk aggregation. These range from basic summation techniques to advanced structural models incorporating macroeconomic drivers. The classification illustrates how methodologies differ in complexity, modelling assumptions, and capacity to capture risk interdependencies.

Basic Aggregation Techniques and Limitations

Risk aggregation aims to provide a comprehensive understanding of overall portfolio exposure by combining multiple risk sources. The following discussion focuses on the insurance sector, particularly on the Solvency Capital Requirement (SCR), although similar reasoning can be applied to the banking sector.

The simplest approach is the summation of stand-alone capital requirements across risk types, which assumes perfect correlation (100%) and offers no diversification benefit. Formally, this is expressed as:

$$SCR_{agg} = \sum_i SCR_i. \quad (4)$$

Though conservative and easy to communicate, it significantly overstates capital needs and is seldom used in modern frameworks except for benchmarking. A marginal improvement is the use of a fixed diversification percentage, which adjusts the simple sum by a constant factor $k \in (0, 1)$:

$$SCR_{agg} = k \cdot \sum_i SCR_i. \quad (5)$$

However, this method is largely static, ignores changing inter-risk dynamics and lacks statistical grounding [48]. Both techniques are considered rudimentary and are primarily found in legacy systems or for conservative buffers. To overcome the limitations of such basic techniques, the variance-covariance matrix approach has become widely adopted. This methodology incorporates pairwise correlations between risk factors and aggregates capital using the

well-known square-root formula:

$$SCR = \sqrt{\sum_i \sum_j \rho_{ij} SCR_i SCR_j}, \quad (6)$$

where ρ_{ij} is the correlation coefficient between risks i and j , and SCR_i is the standalone capital for risk i . This formula assumes elliptical distributions and linear dependence. While efficient, its effectiveness depends on the correlation matrix, that is symmetric and positive semi-definite. It is widely used in regulatory frameworks, notably in the Solvency II standard formula, which provides standardised correlation parameters across risk modules. Otherwise, in partial or full internal models, the $\rho_{i,j}$ can be estimated using statistical metrics as described in the previous chapter. However, the true dependence structure in extreme events may deviate from these assumptions, prompting concerns over model robustness.

Copula-Based Models and Scenario-Based Aggregation

Copula-based modelling offers a more flexible framework by decoupling the marginal behaviour of individual risks from the dependence structure that binds them.

The Gaussian copula is derived from the multivariate normal distribution. It is defined as:

$$C^{\text{Gauss}}(u_1, \dots, u_d) = \Phi\left(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d)\right), \quad (7)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution function, and Φ denotes the joint CDF of a multivariate normal distribution.

The Gaussian copula is popular due to its mathematical tractability and ease of simulation. However, a significant limitation lies in its lack of tail dependence. This implies that even extreme events in individual risks are not likely to occur simultaneously, an unrealistic assumption in stress scenarios or market crashes. As such, while the Gaussian copula works well under normal conditions, it tends to underestimate joint risk during crises.

The (Student's) t copula offers a remedy to the Gaussian copula's limitations by incorporating tail dependence. It is constructed analogously from the multivariate t -distribution and is defined as:

$$C^t(u_1, \dots, u_d) = t_v\left(t_v^{-1}(u_1), \dots, t_v^{-1}(u_d)\right), \quad (8)$$




FIGURE 5: Different copulas

where t_v^{-1} is the quantile function of the univariate t-distribution with degrees of freedom ν and t_ν is the joint CDF of the multivariate t-distribution. This introduces a degree of freedom parameter ν , which governs the extent of tail dependence: lower values of ν produce stronger tail dependence, meaning that extreme co-movements become more likely. In copula-based modelling, as said in *Theorem 3.1*, the marginals and the copula are specified separately. The choice of marginal distributions is crucial, as it captures the individual behaviour of each risk. Common choices include the normal distribution for simplicity, or heavy-tailed distributions such as the Student's t or generalised Pareto when modeling financial or insurance data with extreme values. Importantly, the same copula can be combined with different marginals, allowing practitioners to tailor the model to the marginal properties observed in the data, independently of the dependence structure.

The t copula is elliptical and symmetric, implying that both upper and lower tails are treated equally. While this is an improvement over the Gaussian case, it can be limiting in applications focused on one tail only: in such cases, the symmetry of the t copula may lead to a misrepresentation of tail dependence, as its calibration is influenced equally by both tails. This can result in an inaccurate characterization of extreme co-movements in the tail of interest.

Thus, the t copula allows for a more realistic modelling of joint extremes, though it assumes uniform tail behaviour across all risk pairs.

While the Gaussian and t copulas are both elliptical and derived from multivariate distributions, Archimedean copulas offer a fundamentally different approach. They are typically simpler in form, highly flexible and particularly useful in modelling asymmetric dependency structures, including one-sided tail dependence.

Two prominent members of this family are the Clayton and

Gumbel copulas. The Clayton copula is defined as:

$$C_{\theta}^{\text{Clayton}}(u, v) = \left(u^{-\theta} + v^{-\theta} - 1 \right)^{-1/\theta}, \theta > 0. \quad (9)$$

It exhibits lower tail dependence but no upper tail dependence. This makes the Clayton copula ideal for applications where simultaneous extreme losses (e.g., defaults) are of concern. In contrast, the Gumbel copula focuses on upper tail dependence, and is given by:

$$C_{\theta}^{\text{Gumbel}}(u, v) = \exp \left[- \left((-\ln u)^{\theta} + (-\ln v)^{\theta} \right)^{1/\theta} \right], \theta \geq 1. \quad (10)$$

This copula is particularly well-suited to modelling the co-occurrence of extreme gains or large claim events, as common in catastrophe insurance or systemic market rallies.

In selecting an appropriate copula for stress testing or risk modelling, no single choice universally outperforms the others. As shown in Figure 5, the Gaussian copula does not present clear evidence of tail dependence, the t copula, instead, captures symmetric tail dependence and Archimedean copulas like Clayton and Gumbel allow for asymmetric extremes. Although the Gaussian copula remains dominant due to its simplicity, tractability, and widespread adoption. A key advantage of elliptical copulas lies in their simulation simplicity, particularly in higher dimensions.

The sampling procedure is straightforward relying on transformations from multivariate distributions. This simplicity is especially attractive in stress testing frameworks, where large-scale scenario generation is often required. The pseudo-algorithm of a typical simulation process is outlined in Algorithm 1.

In contrast, simulation from Archimedean copulas, especially in more than two dimensions, is often more involved. While conditional methods or generator inversion techniques exist, they require sampling from non-standard dis-

Algorithm 1 Simulation from an elliptical copula (Gaussian or t)

- 1: Choose an elliptical copula with parameters:
 - correlation matrix Σ
 - degrees of freedom ν (for t copula)
 - 2: Specify the marginal distributions F_1, \dots, F_d
 - 3: Generate a random vector $\mathbf{z} = (z_1, \dots, z_d) \sim N_d(0, \Sigma)$ or $t_{\nu, \Sigma}$
 - 4: **for** each component $i = 1, \dots, d$ **do**
 - 5: Compute $u_i = \Phi(z_i)$ (or $u_i = t_{\nu}(z_i)$)
 - 6: **end for**
 - 7: The resulting vector $\mathbf{u} = (u_1, \dots, u_d)$ follows the copula distribution
 - 8: Transform using marginals: $x_i = F_i^{-1}(u_i)$
-

tributions and are less efficient or less scalable in higher-dimensional settings.

This difference in computational tractability further supports the use of elliptical copulas in large-scale stress testing, even if they may be less expressive in capturing asymmetric tail behaviour.

Moreover, as highlighted by Koziol et al. (2015) [40], the Gaussian copula can generate severe stress scenarios when assuming extreme stress forecasts. This is because, although it lacks tail dependence in a global sense, it is an elliptical distribution, meaning that joint extreme events become more likely when analysis is focused on a narrow, stressed region of the tail. Consequently, with appropriate scenario design, the Gaussian copula remains a valid and effective choice in stress testing frameworks.

Structural Models and Scenario-Based Aggregation

Taking a more integrated approach, structural models simulate joint risk evolution by linking risk components through shared macroeconomic drivers. Consider an economic factor vector Z influencing each risk component via functions $X_i = f_i(Z)$.

To generate realistic economic scenarios for the factor Z , structural models rely on stochastic modelling frameworks. Common choices include Vector Autoregressive (VAR) models or Generalised Autoregressive Conditional Heteroskedasticity (GARCH) models. More generally, scenario generation can also be based on stochastic differential equations frameworks (such as the Heston stochastic volatility or Merton structural credit models). The joint distribution of total risk is then obtained via Monte Carlo simulations or analytical approximations. Letting $SCR_i(Z)$ represent risk under scenario Z , the total risk becomes:

$$SCR = \mathbb{E}_Z [g(SCR_1(Z), \dots, SCR_n(Z))], \quad (11)$$

where g is an aggregation function. Structural models offer realism and scenario analysis capability but come with high data and model governance demands. Their application aligns with the ORSA (Own Risk and Solvency Assessment) principle and supports forward-looking solvency analysis.

Comparative Insights and Practical Considerations

Bringing these methodologies into perspective, it becomes clear that each presents distinct advantages and limitations. The summation and fixed discount approaches, while easy to apply and communicate, are overly conservative and ignore real diversification effects. In contrast, the variance-covariance method offers a balance of practicality and statis-

tical grounding but remains sensitive to the assumption of linearity. Copula-based models enable nuanced dependency modelling and provide greater flexibility in tail-risk analysis yet require careful calibration and pose model selection challenges. Structural aggregation models stand out for their depth and coherence, integrating macroeconomic linkages directly into risk projections. However, their demand for data, expertise, and computational power may pose barriers to adoption.

Overall, the selection of an aggregation methodology should align with the institution's risk profile, data environment and regulatory context. As regulatory expectations evolve, institutions increasingly adopt advanced frameworks validated through empirical testing and scenario analysis. A well-constructed aggregation strategy is essential for accurate capital assessment, product pricing and long-term financial resilience. Furthermore, collaboration between actuarial, risk management and regulatory functions enhances the transparency and reliability of aggregation outcomes, supporting both internal governance and external reporting obligations.

Tail Dependence Methodologies

Understanding and accurately capturing extreme co-movements between risk factors is a critical concern in modern financial and actuarial risk management. Standard correlation-based techniques, while useful in average-case scenarios, fail to describe the behaviour of risks in the extreme tails of distributions, precisely where systemic crises tend to manifest. This limitation can lead to the underestimation of joint losses, particularly in capital adequacy frameworks built on Value-at-Risk, Expected Shortfall, or Economic Capital (EC) [42].

To address this gap, the concept of tail dependence has emerged as a vital tool. Tail dependence quantifies the probability that multiple risks experience extreme outcomes simultaneously, offering deeper insights into contagion effects, systemic stress, and portfolio vulnerabilities.

In this chapter, we present a comprehensive treatment of tail dependence in the context of risk aggregation. We introduce both qualitative and quantitative measures of tail dependence, moving from intuitive visual diagnostics to formal asymptotic coefficients. We then conclude with advanced modelling frameworks, highlighting the role of copula models for simulating joint tail behaviour and Extreme Value Theory (EVT) for extrapolating beyond observed quantiles.

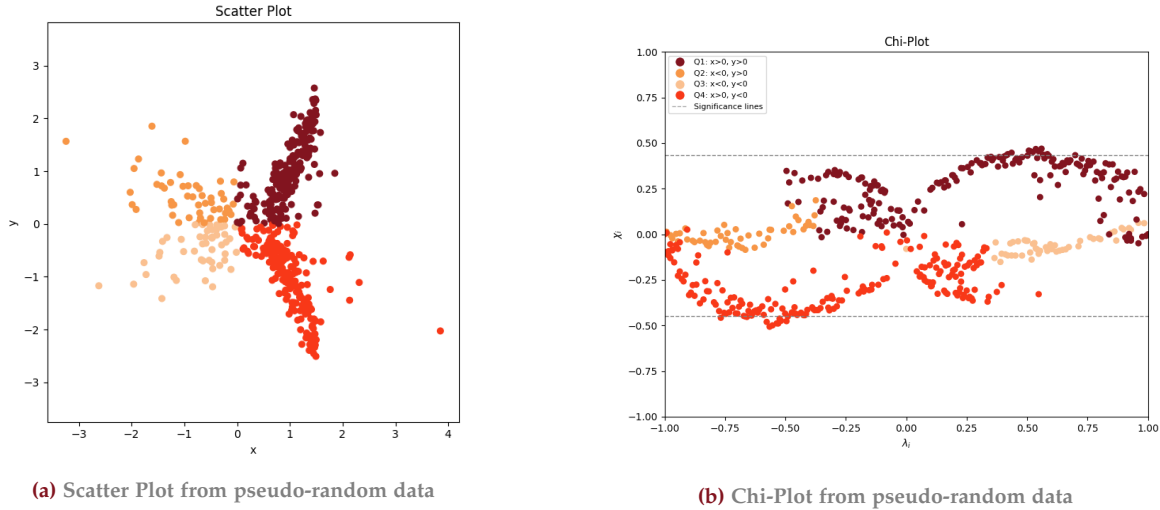


FIGURE 6

Qualitative Measures of Tail Dependence

Qualitative techniques provide intuitive and preliminary insights into the presence of tail dependence between risk factors. Although they do not yield formal numerical measures, these methods are essential for initial exploration of non-linear and extreme co-movement patterns that may not be visible through standard correlation analysis, providing the necessary intuition to motivate the use of formal quantitative tools [18].

Rolling Correlation

Rolling correlation is used to assess whether the correlation between two risk factors remains stable over time or varies during specific periods, such as market stress. For two time series X and Y , the rolling correlation over a window of size n at time t is defined by:

$$\rho_t^{(n)} = \frac{\sum_{i=0}^{n-1} (X_{t-i} - \bar{X}_t)(Y_{t-i} - \bar{Y}_t)}{\sqrt{\sum_{i=0}^{n-1} (X_{t-i} - \bar{X}_t)^2} \sqrt{\sum_{i=0}^{n-1} (Y_{t-i} - \bar{Y}_t)^2}}. \quad (12)$$

This moving estimate helps reveal temporal shifts in dependency structures, particularly relevant in financial systems where co-movements intensify under stress. However, rolling correlation primarily reflects average behaviour over short intervals and does not isolate dependence in the tails [42].

Scatter Plots

Scatter plots visually represent the joint distribution of two risk factors, as shown in Figure 6a. Under a Gaussian dependence structure, the observations tend to form elliptical contours, reflecting weak tail interaction. In contrast, tail-dependent distributions exhibit clustering in the joint corners, such as the upper-right or lower-left quadrants [18]. This concentration suggests that extreme outcomes in both variables are likely to occur simultaneously.

While scatter plots are easy to generate and interpret, they do not offer a formal measure of dependence and are sensitive to sample size. Moreover, they are limited to bivariate analysis and may not distinguish between strong correlation and true tail dependence.

Chi-Plots

The Chi-plot provides a more advanced graphical tool to investigate non-linear dependence, as shown in Figure 6b. For a sample of observation pairs (x_i, y_i) , the Chi-plot transforms ranks into coordinates (χ_i, λ_i) using empirical distribution functions [31]:

$$\begin{aligned} \chi_i &= \frac{H_i - F_i G_i}{\sqrt{F_i(1-F_i)G_i(1-G_i)}}, \\ \lambda_i &= 4S_i \cdot \max\left((F_i - 0.5)^2, (G_i - 0.5)^2\right). \end{aligned} \quad (13)$$

Here, H_i denotes the proportion of pairs (x, y) such that $x \leq x_i$ and $y \leq y_i$, and F_i, G_i are empirical cumulative probabilities of the data. S_i represents the sign of the quadrant (positive for the first and third, negative for the second and fourth). Large absolute values of χ_i indicate strong local dependence between the variables (in the example presented in Figure 6b, we considered values outside the 95% probability region to be significant, following the approach suggested in Fisher et al. (1985)[31]), while high values of λ_i emphasise observations that lie in the tails of the marginal distributions. This combination allows the Chi-plot to effectively detect tail-dependent structures.

Although Chi-plots reveal tail structure more clearly than scatter plots, they remain visual tools without a scalar summary statistic and are generally limited to bivariate cases.

Quantitative Measures of Tail Dependence

While qualitative techniques provide useful preliminary insight, they are limited by their subjectivity and lack of formal metrics. In contrast, quantitative measures offer a precise and replicable framework for assessing tail dependence are essential for building robust risk aggregation and capital allocation models. These measures aim to quantify the strength of joint extremes either through empirical probabilities or asymptotic coefficients derived from copulas.

Tail Dependence Coefficients

The most widely used theoretical measures of tail dependence are the tail dependence coefficients. They provide a formal way to quantify the asymptotic probability that one variable exceeds a high threshold given that the other does

as well.

For a bivariate random vector (X, Y) with continuous marginal distribution functions F_X and F_Y , the upper and lower tail dependence coefficients are defined as:

$$\lambda_U = \lim_{u \rightarrow 1^-} \mathbb{P}(Y > F_Y^{-1}(u) \mid X > F_X^{-1}(u)), \quad (14)$$

$$\lambda_L = \lim_{u \rightarrow 0^+} \mathbb{P}(Y < F_Y^{-1}(u) \mid X < F_X^{-1}(u)). \quad (15)$$

These limits capture the strength of co-movements in the tails of the distribution. If $\lambda_U > 0$, the X and Y are said to be asymptotically dependent in the upper tail; otherwise, they are asymptotically independent.

Tail dependence coefficients (TDCs) can also be expressed using the copula $C(u, v)$ of the joint distribution [49] [44]:

$$\lambda_U = \lim_{u \rightarrow 1^-} \frac{1 - 2u + C(u, u)}{1 - u}, \quad (16)$$

$$\lambda_L = \lim_{u \rightarrow 0^+} \frac{C(u, u)}{u}. \quad (17)$$

These coefficients are particularly useful for assessing dependence in risk aggregation contexts, as they isolate extreme co-occurrence from the overall dependence structure.

Parametric Estimators of Tail Dependence

Parametric estimators of tail dependence refer to analytical expressions for tail dependence coefficients derived under the assumption that the joint distribution of risk factors follows a known parametric family, such as the multivariate Student-t distribution or specific copula functions like Gumbel or Clayton.

These estimators do not constitute full dependence models themselves, but rather quantify tail dependence implied by specific parametric structures. While copulas play a central role in modelling, their tail properties can also be used through plug-in methods to estimate tail dependence coefficients based on fitted parameters.

Elliptical Distribution-Based Estimators

A foundational example of parametric estimation arises in elliptical distributions, particularly the bivariate Student-t distribution. Embrechts et al. (2002) [18] derive an expression for the upper tail dependence coefficient assuming that the pair (X, Y) follows a bivariate t-distribution with correlation ρ and degrees of freedom ν . The resulting tail dependence coefficient is given by:

$$\lambda_U = 2t_{\nu+1} \left(-\sqrt{\frac{(\nu+1)(1-\rho)}{1+\rho}} \right), \quad (18)$$

where $t_{\nu+1}(\cdot)$ is the cumulative function of the univariate Student-t distribution with $\nu + 1$ degrees of freedom. This formula shows that tail dependence increases as either the correlation ρ increases or the degrees of freedom ν decrease. As $\nu \rightarrow \infty$, the Student-t distribution converges to the Gaussian distribution and the tail dependence coefficient tends to zero. This is consistent with the known property that Gaussian distributions imply asymptotic independence, even in the presence of high correlation.

Copula-Based Plug-in Estimators

Another widely used approach involves calculating tail dependence from the closed-form expressions available for known copula families. As Table 2 summarises, these estimators use the estimated copula parameters as inputs to derive the implied tail dependence. For example, the Gumbel copula, which is suited to modelling upper tail

dependence, yields the coefficient:

$$\lambda_U = 2 - 2^{\frac{1}{\theta}}, \quad (19)$$

where $\theta \geq 1$ is the copula's dependence parameter. As θ increases, the upper tail dependence increases, approaching one in the limit. The Clayton copula, in contrast, is characterised by lower tail dependence. Its coefficient is given by:

$$\lambda_L = 2^{-\frac{1}{\theta}}, \quad (20)$$

where $\theta \geq 0$. In this case, lower values of θ correspond to weaker tail dependence. These expressions provide straightforward tools for quantifying the degree of dependence in specific parts of the joint distribution, which is particularly useful when the directionality of tail risk (e.g., lower vs upper) is known a priori, such as in credit risk or catastrophe insurance.

The Student-t copula, derived from the multivariate Student-t distribution, exhibits symmetric tail dependence in both the upper and lower tails. Its tail dependence coefficient is identical to that of the elliptical Student-t distribution discussed earlier. This makes the Student-t copula particularly attractive in financial risk modelling, where both large joint losses and gains are possible and need to be captured symmetrically.

In contrast, some copulas imply no tail dependence at all. The Gaussian copula, while widely used for its analytical tractability, always yields tail dependence coefficients of $\lambda_U = \lambda_L = 0$, regardless of the strength of correlation. This characteristic limits its applicability in scenarios involving systemic risk or stress testing, where co-extremes are likely [38] [44].

Plug-in Estimation Procedure

The parametric estimation process is typically carried out in two steps. First, the practitioner fits the chosen copula to data using methods such as maximum likelihood estimation or inversion of dependence measures like Kendall's tau. Second, the estimated parameter $\hat{\theta}$ is substituted into the copula's closed-form expression for λ_U or λ_L . This plug-in strategy provides computational efficiency and avoids the threshold selection problems inherent in non-parametric estimation.

However, it is important to recognise that parametric estimators are highly sensitive to model specification. Misidentifying the copula family or assuming elliptical dependence where none exists can lead to substantial under- or overestimation of tail risk. For this reason, parametric methods should be accompanied by goodness-of-fit tests, backtesting procedures, and, when necessary, supplemented with non-parametric or semi-parametric alternatives for robustness.

Non-Parametric Estimation of Tail Dependence

Non-parametric estimators of tail dependence provide a flexible, data-driven approach to measuring the strength of joint extremes without imposing strong assumptions about the underlying distribution or copula structure. These estimators are particularly useful when the true dependence structure is unknown or complex, or when reliability is a priority, such as in regulatory contexts. Unlike parametric approaches, which rely on known functional forms and closed-form expressions, non-parametric methods infer tail dependence directly from empirical data using ranks, thresholds, and empirical copulas [37].

Empirical Copula-Based Estimator

One of the earliest and most intuitive non-parametric estimators of tail dependence is based on the empirical copula proposed by Joe et al. (1992) [37]. This estimator counts

Copula	Formula	λ_U	λ_L
Gaussian	$c^{\text{Gauss}}(u_1, \dots, u_d) = \Phi(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_d))$	0	0
Student-t	$c^t(u_1, \dots, u_d) = t_v(t_v^{-1}(u_1), \dots, t_v^{-1}(u_d))$	$2t_{v+1}\left(-\sqrt{\frac{(v+1)(1-\rho)}{1+\rho}}\right)$	$2t_{v+1}\left(-\sqrt{\frac{(v+1)(1-\rho)}{1+\rho}}\right)$
Gumbel	$c_{\theta}^{\text{Gumbel}}(u, v) = \exp\left[-\left((-\ln u)^{\theta} + (-\ln v)^{\theta}\right)^{1/\theta}\right], \quad \theta \geq 1$	$2 - 2^{1/\theta}$	0
Clayton	$c_{\theta}^{\text{Clayton}}(u, v) = (u^{-\theta} + v^{-\theta} - 1)^{-1/\theta}, \quad \theta > 0$	0	$2^{-1/\theta}$

TABLE 2: Copulas and TDCs

Copula	Analytical λ_U	Analytical λ_L	$\hat{\lambda}_U^{JOE}$	$\hat{\lambda}_U^{SS}$	$\hat{\lambda}_L^{SS}$	$\hat{\lambda}_L^{CG}$
Gaussian	0	0	0.065	0.065	0.053	0.053
Student-t	0.391	0.391	0.384	0.384	0.382	0.382
Gumbel	0.586	0	0.574	0.574	0.055	0.055
Clayton	0	0.707	0.003	0.003	0.716	0.716

TABLE 3: Non-Parametric Estimators Results

the frequency of joint exceedances in the upper tail and adjusts for the finite sample size. The upper tail dependence coefficient is approximated as:

$$\hat{\lambda}_U^{JOE} = 2 - \frac{1 - \hat{C}_n\left(1 - \frac{k}{n}, 1 - \frac{k}{n}\right)}{1 - \frac{k}{n}}, \quad (21)$$

where $\hat{C}_n(u, v)$ denotes the empirical copula evaluated at empirical quantiles, n is the sample size, and k is the number of exceedances considered (typically with $k \ll n$).

$$\hat{C}_n\left(\frac{i}{n}, \frac{j}{n}\right) = \frac{\#\{(x, y) | x \leq x_{(i)} \text{ and } y \leq y_{(j)}\}}{n}. \quad (22)$$

This estimator provides a simple way to assess whether extreme co-movements occur more frequently than under independence. While conceptually straightforward and widely applicable, the choice of k significantly affects the stability of the estimate, with smaller values improving extremality at the expense of higher sampling variability.

Schmidt-Stadt Müller Estimator

To address the trade-off between parametric precision and non-parametric robustness, Schmidt and Stadtmüller (2006) [47] proposed a semi-parametric estimator that combines rank-based empirical counting with asymptotic principles. The estimators are defined as:

$$\hat{\lambda}_U^{SS} = \frac{1}{k} \sum_{i=1}^n \mathbb{I}\left(R_i^{(1)} > n - k, R_i^{(2)} > n - k\right), \quad (23)$$

$$\hat{\lambda}_L^{SS} = \frac{1}{k} \sum_{i=1}^n \mathbb{I}\left(R_i^{(1)} \leq k, R_i^{(2)} \leq k\right), \quad (24)$$

where $R_i^{(1)}$ and $R_i^{(2)}$ are the marginal ranks of the i -th observation, and $k \in \{1, \dots, n\}$. The indicator function $\mathbb{I}(\cdot)$ counts how many pairs of observations fall into the joint upper tail. This estimator has gained popularity in

financial and insurance applications due to its conservative nature and ease of computation. It is especially well-suited to regulatory settings under Solvency II or ICAAP, where tail clustering must be captured without strong structural assumptions. It also forms the basis for several backtesting and model validation procedures in operational and market risk.

Caillault-Guegan Estimator

Beyond classical threshold-based estimators, other non-parametric techniques have emerged using the concept of empirical copulas. A notable contribution in this direction is the Caillault-Guegan estimator [11], which defines a tail dependence estimator using:

$$\hat{\lambda}_{L,n}^{CG}\left(\frac{i}{n}\right) = \frac{\hat{C}\left(\frac{i}{n}, \frac{i}{n}\right)}{\left(\frac{i}{n}\right)}, \quad (25)$$

where \hat{C}_n is the empirical copula. A 'plateau' detection algorithm identifies a stable zone over which estimates can be averaged:

$$\hat{\lambda}_{L,n}^{CG} = \frac{1}{p^*} \sum_{k=1}^{p^*} \hat{\lambda}_{L,n}\left(\frac{k}{n}\right). \quad (26)$$

This estimator does not require assumptions about the marginal distributions and works well under weak structural assumptions. However, its reliability depends on carefully determining the smoothing range and identifying the stable zone correctly.

Implementation of Non-Parametric Estimators

All non-parametric estimators face common challenges, notably the sensitivity to the threshold k , the scarcity of tail observations, and the high variance of estimates in small samples. Careful selection of k , combined with bootstrap




FIGURE 7: Empirical Joint Quantile Exceedance Probability (JQEP) on pseudo-random data

procedures or smoothing techniques, is often required to ensure stability and statistical reliability. Moreover, in high-dimensional contexts, the "curse of dimensionality" can affect the efficacy of non-parametric tail estimation, necessitating the use of pairwise or vine decomposition techniques to manage complexity.

To demonstrate the performance of the estimators in a controlled setting, we implemented an illustrative simulation study based on synthetic data generated from copula models. This framework allowed us to validate the non-parametric estimates against known analytical benchmarks. The threshold parameter k was selected as a simple yet representative choice that aligns with the theoretical considerations discussed by Schmidt and Stadtmüller (2006) [47], specifically $k = \lfloor \sqrt{n} \rfloor$. The corresponding results are presented in Table 3.

Despite these limitations, non-parametric estimators remain a cornerstone of practical tail risk assessment. Their model-free nature and interpretability make them valuable for initial diagnostics, model validation, and scenarios where robustness is prioritised over analytical convenience.

Joint Exceedance-Based Measures of Tail Dependence based

Joint exceedance-based measures assess tail dependence by directly examining the frequency of extreme co-movements between risk factors. These methods offer intuitive, empirical tools for analysing co-exceedances, which makes them especially useful in practical contexts such as systemic risk analysis, stress testing, and portfolio risk aggregation. By focusing on the likelihood that both variables exceed certain thresholds, joint exceedance measures can reveal hidden dependence in the tail that standard correlation methods fail to capture.

Joint Exceedance Probability (JEP)

The Joint Exceedance Probability (JEP) captures the likelihood that two variables exceed or fall below a given threshold simultaneously. For transformed uniform variables $U = F_X(X)$ and $V = F_Y(Y)$, the upper and lower tail JEPs are expressed as:

$$\begin{aligned} RJEP(z) &= \mathbb{P}(U > z, V > z), \\ LJEP(z) &= \mathbb{P}(U < z, V < z). \end{aligned} \quad (27)$$

In the case of independence, these probabilities reduce to $(1 - z)^2$ and z^2 , respectively. Any upward deviation from

these values is a sign of tail dependence [18]. A more risk-focused variant uses the quantile values of X and Y directly, as in:

$$JEP(\alpha) = \mathbb{P}(X > u_\alpha, Y > v_\alpha), \quad (28)$$

where $u_\alpha = VaR_X(\alpha)$ and $v_\alpha = VaR_Y(\alpha)$.

Tail Concentration Function

To refine the view further, Joe (1997) [38] introduced the Tail Concentration Function, which considers the conditional probability of joint exceedance:

$$R(z) = \frac{\mathbb{P}(U > z, V > z)}{1 - z}, \quad L(z) = \frac{\mathbb{P}(U < z, V < z)}{z}. \quad (29)$$

These are interpreted as $\mathbb{P}(U > z | V > z)$ and $\mathbb{P}(U < z | V < z)$, respectively. In the presence of tail dependence, these ratios will exceed the values implied by independence. In fact, the upper tail dependence coefficient can be retrieved as the limit $\lambda_U = \lim_{z \rightarrow 1} R(z)$, reinforcing the connection between this conditional probability and asymptotic dependence structure.

Joint Quantile Exceedance Probability (JQEP)

To address situations where different quantile levels must be used for each risk factor, the Joint Quantile Exceedance Probability (JQEP) generalises JEP. Its empirical form is:

$$\begin{aligned} JQEP_{\text{empirical, lower}}(p_x, p_y) &= \\ &= \mathbb{P}(F_X(X) < p_x, F_Y(Y) < p_y), \\ JQEP_{\text{empirical, upper}}(p'_x, p'_y) &= \\ &= \mathbb{P}(F_X(X) > p'_x, F_Y(Y) > p'_y). \end{aligned} \quad (30)$$

Theoretical counterparts, based on copula models, are:

$$JQEP_{\text{theoretical, lower}} = \int \int_{(0,0)}^{(p_x, p_y)} C(u, v; \rho) du dv, \quad (31)$$

$$JQEP_{\text{theoretical, upper}} = \int \int_{(p'_x, p'_y)}^{(1,1)} C(u, v; \rho) du dv.$$

Here, the comparison between empirical and theoretical JQEP values helps assess the adequacy of the assumed dependence structure, particularly under a Gaussian or Student-t copula [42]. To further support the empirical

interpretation of JQEP behaviour, Figure 7 displays a simulated chart of the empirical upper JQEP values plotted against increasing quantile thresholds. As expected, JQEP values drop sharply as the quantile threshold increases, consistent with the theoretical tail behaviour predicted by the copula model.

Conditional Quantile Exceedance Probability (CQEP)

A conditional version of this idea leads to the Conditional Quantile Exceedance Probability (CQEP), which expresses the probability that one variable exceeds its quantile given that the other does. For the upper tail:

$$CQEP_{upper}(q) = \frac{\mathbb{P}(F_X(X) > q, F_Y(Y) > q)}{\mathbb{P}(F_Y(Y) > q)}. \quad (32)$$

Each of the quantitative methods presented provides a distinct lens on tail dependence. JEP and JQEP offer intuitive metrics based on observed co-exceedances. The Tail Concentration Function formalises these into conditional terms, while CQEP introduces directional insights. Finally, tail dependence coefficients offer asymptotic precision and theoretical elegance. These tools lay the foundation for estimation procedures and dependence modelling, particularly in the context of copulas.

Strengths and Limitations

Joint exceedance-based measures are highly intuitive and flexible, providing empirical tools to assess extreme co-movements in risk factors. These measures are particularly suited for backtesting and risk diagnostics, especially in cases where directional or asymmetric tail dependence is of interest. However, they do have limitations: their estimates can be noisy due to the scarcity of data in the tails, and they may not converge to the true tail dependence coefficients in small samples. To mitigate these issues, smoothing techniques and bootstrap methods are often employed, and high-dimensional dependence requires additional techniques like pair-copula constructions or vine copulas. Despite these limitations, joint exceedance measures provide a valuable and practical set of tools for identifying tail dependence, especially when no explicit parametric model is available or desired. They complement more formal asymptotic measures and serve as a critical diagnostic in risk aggregation and capital modelling.

Copula Models as Generative Tools for Joint Tail Simulation

While parametric estimators based on copula functions offer closed-form expressions for tail dependence coefficients, copulas play a more fundamental role as generative models for simulating joint distributions. This modelling capability arises from Sklar's Theorem [49], as presented in chapter "Sklar's Theorem", which separates the marginal distributions from the dependence structure, allowing for flexible construction of multivariate models.

This approach enables the simulation of risk vectors that conform to both the empirical marginals and a chosen dependence structure, making it particularly valuable for modelling joint loss scenarios in risk aggregation, stress testing, and solvency assessment.

For high-dimensional settings, the copula modelling approach can be extended using vine copulas, which allow for the construction of multivariate copulas from a cascade of bivariate building blocks. This method enables scalable yet flexible modelling of complex dependence structures, accommodating tail asymmetries and conditional relationships [1]. Vine copulas are especially relevant in insurance portfolios, operational risk modelling, and multi-line prod-

uct risk management, where pairwise tail behaviour plays a critical role.

Model calibration typically involves estimating the marginals independently, potentially using Extreme Value Theory (EVT) or generalised Pareto distributions, and then fitting the copula parameters using maximum likelihood methods, inference functions for margins (IFM), or inversion of rank-based measures such as Kendall's tau. Once calibrated, the model can be used to simulate multivariate loss distributions, compute joint Value-at-Risk, Expected Shortfall, or assess diversification benefits under regulatory frameworks such as Solvency II and Basel III [42][18].

Unlike tail dependence coefficients, which provide summary measures of joint extremes, copula models offer a full probabilistic representation of the dependence structure. They allow practitioners not only to quantify, but also to generate, realistic tail-dependent scenarios, a capability essential for solid capital estimation and systemic risk analysis.

EVT-Based Models and Extreme Quantile Analysis (EQA)

While copula-based models offer a versatile generative framework for simulating dependent risks, they remain limited by their reliance on a specific dependence structure and finite-sample calibrations. In contrast, EVT provides a complementary, asymptotic framework for modelling the behaviour of tail events without requiring a predefined copula family. EVT-based methods focus directly on the statistical behaviour of rare and extreme outcomes, making them particularly valuable for risk aggregation in sparse or highly volatile regimes.

EVT approaches model the tails of distributions either through the block maxima method, which relies on the Generalised Extreme Value (GEV) distribution, or the more flexible Peaks Over Threshold (POT) method, which models exceedances above a high threshold using the Generalised Pareto Distribution (GPD).

In the POT framework, the conditional distribution of excess losses $X - u$, given $X > u$, is approximated as:

$$\mathbb{P}(X > x \mid X > u) \approx \left(1 + \xi \frac{x - u}{\beta}\right)^{-1/\xi}, \quad \text{for } x > u, \quad (33)$$

where ξ is the shape parameter and β the scale parameter. This provides a flexible tool for modelling heavy-tailed behaviour in a univariate or marginal setting.

To extend this framework to multivariate and joint risk analysis, EVT can be integrated with copula techniques or applied through multivariate EVT constructions, such as multivariate threshold exceedances or angular measure models [32]. These allow the capture of extremal dependence beyond the limitations of traditional tail dependence coefficients. For example, even in cases where the upper tail dependence coefficient $\lambda_U = 0$, EQA may still uncover joint explosive behaviour in the far tails. This is particularly relevant for systemic risk modelling or operational loss estimation in banking and insurance [14] [42].

As a risk aggregation perspective, EVT-based methods may still detect joint clustering of extremes, a phenomenon especially relevant in systemic risk modelling, market crashes, or operational event losses [14].

EQA is a particularly powerful application of EVT in practice. EQA is designed to quantify loss behaviour beyond extreme quantile thresholds, typically beyond the 99.5th or 99.9th percentile, and is increasingly used in both Basel and Solvency II contexts to support conservative capital estimation. It relies on extrapolation methods that go beyond observed data, estimating the tail distribution in regions

of data sparsity, where traditional empirical measures are insufficient.

In the multivariate context, EQA employs techniques such as conditional multivariate tail modelling, extremal coefficients, or hybrid EVT-copula approaches to better understand the shape and concentration of tail risk. These models enable more realistic joint loss distribution estimation and are particularly effective in capturing non-linear dependence structures during stress conditions.

However, the application of EVT and EQA is not without challenges. Threshold selection is a critical issue, as poorly chosen thresholds can lead to instability and high variance in parameter estimates. Diagnostic tools such as mean residual life plots and stability plots are commonly used to assess the suitability of threshold choices. Moreover, EVT estimators tend to be sensitive to sample size, and estimation uncertainty must be addressed, especially in risk contexts with limited tail observations [15] [32].

In summary, Extreme Quantile Analysis strengthens the risk quantification toolkit by allowing practitioners to capture asymptotic tail behaviour in both univariate and multivariate settings. Its integration with copula modelling or as a standalone approach provides a critical lens for understanding systemic and co-extreme risks, especially in regimes where traditional dependence measures fail to capture explosive joint behaviour.

Applications of Tail Dependence in Practice

The relevance of tail dependence modelling extends beyond theoretical interest; it plays a central role in practical risk management and regulatory decision-making. This subsection highlights how tail dependence concepts are effectively applied in the insurance, banking, and pension sectors, as well as in supervisory frameworks.

Insurance and Reinsurance

In the insurance domain, understanding joint tail behaviour is critical when aggregating diverse risk types such as mortality, longevity, and catastrophe risk. The European Solvency II directive [26] emphasises the importance of modelling tail dependencies, particularly when aggregating risks under the Solvency Capital Requirement (SCR). Ignoring these dependencies can result in substantial underestimation of capital requirements, especially when risks are heavy-tailed and not linearly correlated. Copula-based approaches and extreme quantile techniques have been increasingly used to assess concentration effects and ensure sufficient capital buffers.

Banking Sector

In banking, tail dependence plays a pivotal role in stress testing and capital planning. Traditional correlation-based models often underestimate the likelihood of joint extreme losses across asset classes or sectors. Tail-dependent models, such as those involving the Student- t copula or empirical tail copulas, have been applied to assess contagion risk and systemic vulnerabilities. For example, during the global financial crisis, significant tail co-movements among credit instruments (e.g., reflected in iTraxx Crossover spreads) highlighted the necessity of using tail-aware models. Such methodologies inform the calculation of metrics like Value-at-Risk, Expected Shortfall, and CoVaR in portfolio- and institution-level risk aggregation [42] [38].

Pension Funds and Long-Term Investment Strategies

In pension fund management, understanding the tail behaviour between liabilities (e.g., longevity risk) and assets (e.g., equity risk) is crucial for long-term solvency planning. Tail dependence modelling supports the calibration of resilience scenarios, especially under shocks such as simultaneous market downturns and increases in life expectancy, which are better captured through dependence structures that extend beyond linear correlation [45].

Regulatory Guidelines and Supervisory Practice

Supervisory bodies, including the Bank for International Settlements (BIS), EIOPA, and the International Association of Insurance Supervisors (IAIS), explicitly caution against reliance on linear correlation for aggregating risks under stress [25] [33]. They recommend the use of models that incorporate tail dependence, particularly in the context of Pillar II evaluations, group supervision, and macroprudential surveillance. In operational risk and internal model validation, regulators increasingly request evidence that risk aggregation captures joint tail behaviour using credible, empirically supported methods.

The integration of tail dependence in real-world risk aggregation frameworks enables more resilient capital planning, accurate solvency estimation, and better preparation for systemic shocks. Whether through copulas, extreme quantile methods, or hybrid EVT frameworks, these tools help financial institutions move beyond naive assumptions of independence or normality, aligning modelling practices with the complex realities of joint risk behaviour.

Impact Analysis of Tail Dependence

As introduced in chapter “Risk Aggregation: Regulatory Evolution”, capital requirements are a foundational pillar of financial regulation, designed to ensure the solvency and resilience of institutions in the face of adverse events. Both banking and insurance sectors have long operated under risk-based capital frameworks, which determine how much capital an institution must hold against its aggregate risks. Central to this process is the modelling of dependencies between different risk types. Traditionally, regulatory and internal models have relied on correlation-based methods for aggregating risk exposures. However, as recent literature and empirical studies demonstrate, the choice of dependence structure has a profound impact on capital calculations, risk governance, and ultimately on the stability of financial institutions.

This chapter explores the critical role that dependency structures play in shaping diversification benefits and capital requirements, with a particular focus on tail dependence. We draw on quantitative tools, empirical findings, and conceptual frameworks to examine the material differences that emerge when correlation-based models are replaced with more advanced, tail-sensitive alternatives such as copula models.

The Limits of Correlation-Based Dependence Structures

Under both Solvency II and Basel III, risk aggregation is typically performed using linear correlation matrices. These approaches are popular due to their mathematical simplicity and computational efficiency. However, the use of Pearson

correlations inherently assumes linearity and elliptical distributions, which fails to capture the behaviour of extreme co-movements between risks, particularly during financial crises.

The CEIOPS in its *Calibration of the Solvency II standard formula: Technical Paper* [16] revealed that using a correlation matrix led to a market risk SCR of 61.9 (normalised units), whereas incorporating tail dependence raised the requirement to 82.5, showing a 33% underestimation when tail dependence was ignored.

This underestimation becomes particularly dangerous in times of systemic crises, when multiple risk factors become highly interdependent. Furthermore, correlation structures often produce misleading assessments of diversification benefits. During benign periods, linear correlations suggest substantial diversification, which may not materialise under stress. This mismatch between modelled and actual behaviour has prompted increasing scrutiny from regulators and internal model validation teams.

The Role of Tail Dependence and Copula Models

To address these limitations, modern risk aggregation has shifted toward the use of copula models, as presented in chapter “Measures of Dependency”, which allow for more flexible and realistic dependence structures. Copulas separate the marginal behaviour of individual risks from their joint dependence, enabling a tailored representation of tail events.

For instance, Tang and Valdez (2005) [51] examined Australian general insurance data and found that capital requirements computed at the 99.5% confidence level varied from 92% to 101% of premiums depending on the copula used. Heavy-tailed copulas like the Student-*t*, which model tail dependence, led to the highest capital charges. The authors conclude that improper modelling of tail dependence can materially distort capital and diversification outcomes. In the insurance context, Mejdoub and Ben Arab (2018) [43] used a D-vine copula model on non-life portfolios and showed that ignoring tail dependence in aggregation can significantly distort Value-at-Risk (VaR) and Tail Value-at-Risk (TVaR), leading to under- or overestimation of required capital depending on the dependence structure.

These findings are not merely academic. They have concrete implications for solvency, capital planning, and regulatory compliance. Institutions using copula-based models benefit from a more realistic depiction of joint loss events, enhancing their resilience to systemic shocks.

Dependency Structures and Diversification Efficiency

One of the most critical aspects of capital modelling is the assessment of diversification benefits. While diversification is generally seen as a source of capital efficiency, its measurement is highly sensitive to how dependencies are modelled. Correlation-based models often overstate diversification effects, particularly in portfolios exposed to multiple interacting risks. This can lead to underestimation of capital needs during market turbulence, where dependencies between risk factors intensify.

Rosenberg and Schuermann (2006) [46] analyzed integrated credit and market risks using copulas and showed that neglecting dependency across asset classes can lead to capital overestimation by 30–40% when no diversification is assumed. Though the Student-*t* copula framework improved tail sensitivity, the empirical increase in capital was modest in their baseline application, emphasizing the role of dependence assumptions on diversification measurement.

Copula models, by contrast, can differentiate between normal and stressed conditions. They reveal that under tail

dependence, diversification deteriorates as risks begin to move together, particularly in the upper quantiles of the loss distribution. This is where capital adequacy is truly tested.

Analytical Framework for Dependency Impact

To operationalise the evaluation of dependency structures, internal models have developed diagnostic tools that provide deep insight into the sensitivity and concentration of capital requirements. These tools are particularly useful for understanding how different risk components and assumptions about their interrelations affect overall capital needs. While this approach can be applied across financial sectors, this section focuses specifically on the insurance context, following the approach developed by EIOPA in its *Study on Diversification in Internal Models* [25].

Risk Multiplier

The risk multiplier quantifies how a single undiversified risk component impacts total capital under a variance-covariance method:

$$\text{riskMultiplier}_j = \sum_i \rho_{ij} \cdot \frac{\text{undivSCR}_i}{\text{SCR}^{\text{vc}}}. \quad (34)$$

A high multiplier implies both a significant standalone risk and strong interactivity with other components. For example, if market risk shows a 50% multiplier, even moderate shifts in its value can materially affect the total SCR, guiding portfolio rebalancing and additional buffer provisioning.

Correlation Multiplier

This metric measures how sensitive the SCR is to changes in specific correlations:

$$\text{correlationMultiplier}_{ij} = \frac{\text{undivSCR}_i}{\text{SCR}^{\text{vc}}} \cdot \frac{\text{undivSCR}_j}{\text{SCR}^{\text{vc}}}. \quad (35)$$

This metric is crucial for assessing model risk. Internal models often rely on subjective or expert-based assumptions for correlation matrices, especially where historical data is sparse or unreliable. A high correlation multiplier signals that even a small misspecification in correlation can have a disproportionately large impact on capital. In practice, this helps model owners and validators to identify where careful justification and stress testing is needed.

For example, if the correlation between market and non-life risks has a correlation multiplier of 25%, then an optimistic assumption could understate capital significantly. This metric thus anchors the model to reality-checks and encourages prudent assumptions.

Diversification Benefit Decomposition

Diversification is often presented as a total benefit, but this metric dissects it risk by risk, providing a nuanced view of how each contributes to the total capital relief.

Key metrics involved are:

- **Undiversified weight:** Share of the undiversified SCR each risk contributes.
- **Diversified weight:** Share of the diversified SCR.
- **Individual diversification benefit:** How much the capital for that risk is reduced due to diversification.
- **Diversification benefit weight:** Each risk's share of the total benefit.

$$\text{divBenefit}_i = 1 - \frac{\text{divSCR}_i}{\text{undivSCR}_i}, \quad (36)$$

$$\text{divBenefitWeight}_i = \frac{\text{undivSCR}_i - \text{divSCR}_i}{\sum_j (\text{undivSCR}_j - \text{divSCR}_j)}. \quad (37)$$

This disaggregation allows to highlight which risks generate most of the diversification benefit and which contribute least. A risk with a high undiversified weight but a low diversification benefit may be a concern, particularly if it dominates the capital structure. Conversely, a high diversification benefit may validate the strategic value of holding that risk in the portfolio.

Quantile-Based and Tail Analysis

In chapter “Tail Dependence Methodologies”, we introduced Extreme Quantile Analysis (EQA) as a powerful extension of traditional tail risk measurement, grounded in Extreme Value Theory (EVT). Here, we apply it alongside Landing Quantile Analysis to deepen our understanding of diversification under extreme stress.

Extreme Quantile Analysis applies Extreme Value Theory (EVT) to assess whether internal models realistically capture the far right tail of the loss distribution, typically beyond the 99.5th percentile. This is essential for identifying underestimation of rare, high-impact events like pandemics or financial crises.

Landing Quantile Analysis evaluates where the diversified SCR lies within the empirical loss distribution of a given risk. A high landing quantile suggests poor diversification, possibly due to persistent tail dependence, whereas a lower quantile implies that diversification is effectively reducing capital needs.

When performed across all risk types, landing quantile analysis provides a map of diversification efficiency, helping identify risks that behave poorly under aggregation or remain problematic in stress scenarios.

Dependency Structures and Capital Composition

Another layer of analysis tracks how risk composition changes across quantiles of the loss distribution. At the median or lower quantiles, risk contributions may appear balanced, but at higher quantiles, such as the 99.9th percentile, certain risks may dominate. This shift highlights how apparent diversification can collapse under stress. Smoothing techniques such as Gaussian kernels or Harrell-Davis estimators help visualise this dynamic distribution of risk, offering a more accurate picture of portfolio behaviour under adverse scenarios.

Measuring Concentration and Diversification Limits

Even the best dependency modelling cannot overcome structural concentration. The Gini coefficient offers a quantitative measure of how evenly risks contribute to the portfolio:

$$G = \frac{\sum_i \sum_j |x_i - x_j|}{2n \sum_i x_i}, \quad \text{where } x_i = \frac{\text{undivSCR}_i}{\sum_j \text{undivSCR}_j}. \quad (38)$$

A coefficient near zero implies broad diversification, while a value closer to 1 indicates that the portfolio is dominated by one or two risks, reducing the potential for effective capital relief.

This metric provides a realistic ceiling to diversification: even with ideal dependency assumptions, a highly concentrated portfolio cannot achieve strong capital efficiency.

To complement this, the **Diversification Score** compares actual diversification to the theoretical maximum under full independence:

$$\text{divScore} = \frac{\sum_i \text{undivSCR}_i - \text{divSCR}_{\text{total}}}{\sum_i \text{undivSCR}_i - \sqrt{\sum_i \text{undivSCR}_i^2}}. \quad (39)$$

A high score suggests that the portfolio is well-diversified relative to its potential, given the existing risk profile. A low score, on the other hand, may signal that correlation assumptions are overly conservative or that the model is failing to recognise legitimate diversification effects. For peer comparisons, this score helps benchmark institutions with similar diversified profiles but differing dependency structures. It provides a clean, normalised view of how effectively a firm is leveraging diversification.

Final Considerations

This chapter has shown that the structure of dependency modelling is not a technical footnote but a defining factor in the accuracy and prudence of capital requirement calculations. Traditional correlation-based models offer convenience but often fail to capture the true nature of joint tail events, particularly under stress. Copula-based approaches, though more complex, provide a much richer and more realistic framework for risk aggregation, resulting in higher, but more appropriate, capital requirements.

By applying quantitative diagnostics such as risk and correlation multipliers, quantile-based tail analyses, and concentration metrics, institutions can rigorously evaluate their internal models and improve their risk governance. Ultimately, the choice of dependency structure determines whether capital models merely comply with regulation or genuinely safeguard financial stability. As systemic risks evolve and regulatory scrutiny deepens, the ability to model and manage tail dependencies will be increasingly critical in the financial industry's toolkit.

Conclusions

The progressive evolution of regulatory frameworks within the banking and insurance sectors underscores a growing awareness of the pivotal role that tail dependence plays in the accurate quantification of systemic risk. While traditional correlation-based models remain foundational, they often fail to capture the complex interdependencies that emerge during extreme market events. As a result, regulators are increasingly advocating for more sophisticated approaches that consider the full spectrum of dependence structures, particularly in the tails.

This paper highlights the need to move beyond conventional correlation by examining alternative dependence modelling techniques, such as copula-based methods and other non-linear association measures. These approaches offer a more realistic representation of how risks interact under stress. In particular, aggregation methodologies reveal how assumptions about dependence can materially influence capital adequacy assessments. In this light, the selection of appropriate tail dependence metrics and estimators is not a merely technical concern: it carries substantial regulatory and financial implications.

Looking forward, the future of prudential regulation will increasingly depend on the ability of both supervisory authorities and financial institutions to integrate methodological innovation into risk management practices. Furthermore, as financial systems grow more interconnected and exposed to tail risks, driven by global economic, climatic, and technological shifts, the ability to anticipate and mitigate the consequences of joint extreme events will become central to

regulatory design.

A deeper understanding of tail dependence can strengthen the resilience of individual institutions and safeguard the broader stability of the financial system. Achieving this will require a thoughtful integration of advanced analytical methods with transparent and practical frameworks, capable of adapting to the complexities of a shifting risk environment.



References

- [1] **Aas, K., Czado, C., Frigessi, A. and Bakken, H.** *Pair-copula constructions of multiple dependence*. Insurance: Mathematics and Economics, 44(2): 182-198, 2009.
- [2] **Actuarial Standards Board (ASB).** *Actuarial Standard of Practice No. 55: Capital Adequacy Assessment*. 2019.
- [3] **Allianz Commercial.** *The predictive power of ESG for insurance*. August 2018.
- [4] **Basel Committee on Banking Supervision** *Basel III: A global regulatory framework for more resilient banks and banking systems*. Basel: Bank for International Settlements, June 2011.
- [5] **Basel Committee on Banking Supervision** *Fundamental review of the trading book: A revised market risk framework*. Basel: Bank for International Settlements, October 2013.
- [6] **Basel Committee on Banking Supervision** *International Convergence of Capital Measurement and Capital Standards*. Basel: Bank for International Settlements, July 1988.
- [7] **Basel Committee on Banking Supervision** *International Convergence of Capital Measurement and Capital Standards: A Revised Framework - Comprehensive Version*. Basel: Bank for International Settlements, June 2006.
- [8] **Basel Committee on Banking Supervision** *Revisions to the Basel II market risk framework*. Basel: Bank for International Settlements, February 2011.
- [9] **Bernard, C. and Vanduffel, S.** *Risk Aggregation and Diversification*. Society of Actuaries, 2016.
- [10] **Bruneton, J.P.** *Copula-based Hierarchical Aggregation of Correlated Risks. The behaviour of the diversification benefit in Gaussian and Lognormal Trees*. arXiv, 2011.
- [11] **Caillault, C. and Guégan, D.** *Empirical estimation of extreme risk measures using conditional copula*. Finance Research Letters, 2(2): 110-123, 2005.
- [12] **Cambou, M. and Filipovic, D.** *Model Uncertainty And Scenario Aggregation*. Mathematical Finance, Vol. 27, Issue 2, May 2017.
- [13] **Capéraà, P., Fougères, A.-L. and Genest, C.** *A nonparametric estimation procedure for bivariate extreme value copulas*. Biometrika, 84(3): 567-577, 1997.
- [14] **Chavez-Demoulin, V., Embrechts, P. and Hofert, M.** *An extreme value approach for modelling operational risk losses depending on covariates*. Journal of Risk and Insurance, 83(3): 735-776, 2016.
- [15] **Coles, S.** *An Introduction to Statistical modelling of Extreme Values*. Springer Series in Statistics, 2001.
- [16] **Committee of European Insurance and Occupational Pensions Supervisors (CEIOPS).** *Calibration of the Solvency II standard formula: Technical Paper*. CEIOPS, 2010.
- [17] **Denuit, M., Dhaene, J., Goovaerts, M. and Kaas, R.** *Actuarial Theory for Dependent Risks: Measures, Orders and Models*. Wiley, 2005.
- [18] **Embrechts, P., McNeil, A. J. and Straumann, D.** *Correlation and dependence in risk management: properties and pitfalls*. In: Dempster,

- M.A.H. (ed.), *Risk Management: Value at Risk and Beyond*, Cambridge University Press, 2002.
- [19] **Embrechts, P., Wang, B. and Wang R.** *Aggregation-Robustness and Model Uncertainty of Regulatory Risk Measures*. 2014.
- [20] **European Banking Authority.** *Basel III Monitoring Report as of December 2022*. September 2023.
- [21] **European Banking Authority.** *Guidelines for common procedures and methodologies for the supervisory review and evaluation process (SREP) and supervisory stress testing*. October 2022.
- [22] **European Banking Authority.** *Final Report: Guidelines on ICAAP and ILAAP*. 2016.
- [23] **European Central Bank.** *Financial Stability Review - November 2024*. European Central Bank, 2024.
- [24] **European Insurance and Occupational Pensions Authority.** *Guidelines on own risk and solvency assessment*. EIOPA, 2015.
- [25] **European Insurance and Occupational Pensions Authority.** *Study on Diversification in Internal Models*. EIOPA, 2020.
- [26] **European Parliament.** *Commission Delegated Regulation (EU) 2015/35 of 10 October 2014*. Official Journal of the European Union, January 2015.
- [27] **European Parliament and Council of the European Union.** *Capital Requirements Directive IV (EU) 2013/36*. Official Journal of the European Union, June 2013.
- [28] **European Parliament and Council of the European Union.** *Capital Requirements Regulation (EU) No 2013/575*. Official Journal of the European Union, June 2013.
- [29] **European Parliament and the Council of the European Union.** *Directive 2002/13/EC of the European Parliament and of the Council of 5 March 2002 amending Council Directive 73/239/EEC as regards the solvency margin requirements for non-life insurance undertakings*. Official Journal of the European Union, March 2002.
- [30] **European Parliament and the Council of the European Union.** *Directive 2002/83/EC of the European Parliament and of the Council of 5 November 2002 concerning life assurance*. Official Journal of the European Union, December 2002.
- [31] **Fisher, N. I. and Switzer, P.** *Chi-plots for assessing dependence*. *Biometrika*, 72(2): 253-265, 1985.
- [32] **Haan, L. and Ferreira, A.** *Extreme Value Theory: An Introduction*. Springer Series in Operations Research and Financial Engineering, 2006.
- [33] **International Association of Insurance Supervisors.** *Insurance Core Principle 16: Enterprise Risk Management for Solvency Purposes*. 2019.
- [34] **International Monetary Fund.** *Italy: Financial Sector Assessment Program-Technical Note*. IMF Country Report No. 20/233, August 2020.
- [35] **Istituto per la Vigilanza sulle Assicurazioni.** *Regolamento n. 20 del 3 Maggio 2016*. IVASS, May 2016.
- [36] **Istituto per la Vigilanza sulle Assicurazioni.** *Regolamento n. 32 del 9 Novembre 2016*. IVASS, November 2016.
- [37] **Joe, H., Li, X. and Nikoloulopoulos, A. K.** *Nonparametric estimation of multivariate tail probabilities and tail dependence coefficients*. *Journal of*

- Multivariate Analysis, 43(2): 215–238, 1992.
- [38] **Joe, H.** *Multivariate Models and Dependence Concepts*. Chapman & Hall, 1997.
- [39] **Kendall, M. G.** *A New Measure of Rank Correlation*. Biometrika, 30(1/2): 81–93, 1938.
- [40] **Koziol, P., Schell, C. and Eckhardt, M.** *Credit risk stress testing and copulas - is the Gaussian copula better than its reputation?* Deutsche Bundesbank No 46/2015, 2015.
- [41] **Marri, F. and Moutanabbir, K.** *Risk Aggregation and Capital Allocation Using a New Generalised Archimedean Copula*. arXiv, 2021.
- [42] **McNeil, A. J., Frey, R. and Embrechts, P.** *Quantitative Risk Management: Concepts, Techniques and Tools*. Princeton University Press, 2015.
- [43] **Mejdoub, H. and Ben Arab, M.** *Dependence modelling in non-life insurance: a D-vine copula approach*. Borsa Istanbul Review, 18(1): 30–41, 2018.
- [44] **Nelsen, R. B.** *An Introduction to Copulas, 2nd Edition*. Springer Series in Statistics, 2006.
- [45] **Owadally, I.** *Tail risk in pension funds: an analysis using ARCH models and bilinear processes*. Review of Quantitative Finance and Accounting, 43(2): 301–331, 2014.
- [46] **Rosenberg, J. V. and Schuermann, T.** *A general approach to integrated risk management with skewed, fat-tailed risks*. Journal of Financial Economics, 79(3): 569–614, 2006.
- [47] **Schmidt, R. and Stadtmüller, U.** *Non-parametric estimation of tail dependence*. Scandinavian Journal of Statistics, 33(2): 307–335, 2006.
- [48] **Shaw, R. and Spivak, G.** *Correlations and Dependencies in Economic Capital Models*. Institute and Faculty of Actuaries, September 2009.
- [49] **Sklar, A.** *Fonctions de Répartition à n Dimensions et Leurs Marges*. Publications de l'Institut de Statistique de l'Université de Paris, 8: 229–231, 1959.
- [50] **Spearman, C.** *The Proof and Measurement of Association Between Two Things*. American Journal of Psychology, 15(1): 72–101, 1904.
- [51] **Tang, A. and Valdez, E. A.** *Economic capital and the aggregation of risks using copulas*. Insurance: Mathematics and Economics, 37(1): 25–38, 2005.
- [52] **United Nations Environment Programme Finance Initiative.** *Guidance on the integration of ESG risks into insurance underwriting*. June 2020.



**A Comparison of Advanced Methods for the
Quantile Estimation in the Risk Management
Field**

About the Authors



Michele Bonollo:

Chief of Risk Methodologies Officer

He holds a degree in mathematics, a master degree in mathematical finance at the Paris VII University and a PHD in statistics. He worked as an executive in both large and regional Italian banks. He has also collaborated with some consultant companies in the broad area of risk management, asset management, pricing models, software systems and regulatory compliance. Along with his professional activities, he developed applied research in the above fields, with more than 20 papers published in scientific journals and dozens of speeches in international conferences. He currently gives seminars and lessons in some top ranked Italian universities.



Leonardo Mastrototaro:

Quantitative Analyst

Financial engineer and theoretical physicist with a Ph.D. in Physics. He has developed a robust theoretical and practical understanding of various mathematical, physical, and computational applications from a scientific perspective, through both academic studies and professional experience.



A Comparison of Advanced Methods for the Quantile Estimation in the Risk Management Field

Michele Bonollo

Leonardo Mastrototaro

In the financial sector the regulation prescribes what risk measures have to be implemented to guarantee a safe banking system. Due to the forthcoming new regulation about the market risk (FRTB), the banks have recently developed Montecarlo models to estimate the Default Risk Charge (DRC), namely a 1-year VaR at 99.9% confidence level related to the losses coming from the default events in the trading book. Banks have put a lot of effort in the modeling step of the process, i.e. how to build the simulation algorithm for the process: what are the risk factors, how to define the default event, how to infer the correlations among the obligors and so on. Despite the extreme quantile estimation is a well-known problem in the statistical field, it has received less attention by the banks modelers. In our paper we review the context and the existing literature, hence we compare on real data the performance of some advanced quantile estimators for the DRC measure that could be used to challenge the classical empirical quantile. For small-medium size samples the results are encouraging.

In recent years, the risk management field received an increasing attention, due to the growth in the banking regulation (the so-called Basel framework) that asked for more risk measures to the banks, to capture any potential source of risk (losses). Several quantitative disciplines were exploited to build more safe risk models, taking the techniques mainly from probability, stochastic processes, mathematical finance. Most of the attention of both the academic and financial community was devoted to build accurate models for the losses distributions, to define proper risk measures (e.g. the *Value at Risk* quantile vs. the *Expected Shortfall*), to decompose them to get a breakdown, to develop robust data quality steps for the market data. An exhaustive description of the risk management models is given in the following [9]. Some details about the risk measures breakdown are available in [5] and [6]. A recent deep analysis of the desirable properties of the risk measures is provided by [1]. Surprisingly, relative few attention was received by the last step of the risk management process, i.e. the measure estimation based on the empirical (or simulated) data. Given the model, given the data, it may happen that the risk estimation is not reliable, because of the usual uncertainty embedded in the historical or simulated data. This issue is very general, as it must be faced for any model (parametric or nonparametric) and for all risk measures (VaR, ES, ComponentVaR, etc.). In the forthcoming 'Basel IV' regulation the banks must estimate a very extreme quantile, namely a 99.9% quantile with 1 year horizon related to the default losses in the trading book, known as DRC, Default Risk Charge. Because of the lack of analytical models to calculate or to approximate satisfactory this quantile, all the banks adopt the Montecarlo simulation approach. Then we aim to compare some basic quantile estimators vs some more advanced tools, in order to check their performances in the bias-variance dimensions vs their computational complexity.

Financial Context. The DRC Risk Measure

The Default Risk Charge is a regulatory measure designed to capture default risk within the trading portfolio, as required by Basel standards, particularly within the framework of the Fundamental Review of the Trading Book (FRTB) outlined in the BCBS 457[2] document published by the Basel Committee on Banking Supervision. This model is designed to quantify the risk of loss resulting from the failure of a counterparty or issuer of financial instruments, including equity, bond, and derivative exposures. It is worth to note that banks could have *short* positions on the issuer debt, such as CDS, where the defaults imply a profit, not a loss. Then any DRC model requires a careful preliminary data management process where the granular positions related to each obligor are netted, aggregated, etc. The Default Risk Charge (DRC) regulatory set-up is described by Chapter 7 ('The Default Risk Charge') of the BCBS 457 document[2]. In the updated Basel Framework, the requirements related to the DRC are contained in section MAR 33, specifically in paragraphs MAR 33.18 - 33.38. These documents establish the criteria for calculating default risk, specifying that:

- It must be calculated over a one-year horizon.
- It must reflect a 99.9% threshold. The ECB EGIM [4] prescribes that the DRC measure must be provided with a confidence interval of 95%.
- It must include all exposures sensitive to default risk within the trading portfolio, excluding those specifically defined as "non-material risks".

The model aims to compute a portfolio loss distribution representative of overall credit risk. The DRC value is typically defined as the loss at the regulatory percentile (e.g., 99.9%)

over a one-year horizon, with the loss defined as:

$$PnL = \sum_n EAD_n \cdot \mathbb{1}_{D_n} \cdot LGD_n, \quad (40)$$

where EAD_n is the exposure on the n -th instrument, D_n express the default status for the legal entity linked to that instrument and LGD_n is the loss given default.

The DRC simulation is implemented through a Monte Carlo approach, a probabilistic methodology that allows for a detailed modeling of uncertainty related to risk factors. Each iteration of the simulation includes:

- **Simulation of Credit Drivers:** Stochastic scenarios for credit risk factors are generated for each Legal Entity in the portfolio. This process simulates the defaults of individual Legal Entities.
- **Calculation of Recovery Rates:** For each scenario, the recovery rate associated with the simulated defaults of Legal Entities is estimated. This parameter is crucial for quantifying the recoverable loss amount.
- **Determination of PnL:** Using the simulated credit drivers and calculated recovery rates, the Profit and Loss (PnL) associated with each financial instrument in the portfolio is estimated, consolidating the results to obtain the overall loss distribution.

Further details and methodological applications can be found in Basel documentation, including BCBS 352 ("Standards: Minimum capital requirements for market risk") and the aforementioned BCBS 457[2], which provide detailed guidelines for the implementation and reporting of the DRC.

Credit Worthiness Index

The random variable X_i is the Credit Worthiness Index (CWI) and reflects the credit quality of legal entity i . It is simulated for each legal entity considered through a stochastic process defined as:

$$\Delta X_i = \frac{\sum_j \beta_{ij} \Delta Z_j}{\rho_i} + \frac{\sigma_i \Delta W_i}{\rho_i}, \quad (41)$$

where i refers to the legal entity, j refers to the credit drivers associated with the legal entity, Z_j is a multidimensional stochastic process of credit driver returns with mean 0 and covariance matrix C , W_i is a Gaussian variable $\mathcal{N}(0,1)$ independent from Z_j , β_{ij} are the credit drivers' weights, ρ is the total variance matrix of the Credit Worthiness process and ν and σ are the volatility matrix of the systemic component and idiosyncratic component, respectively.

Specifically, the processes $X_i = \sum \Delta X_i \Delta t$ simulate various risk drivers as a multivariate Gaussian distribution using the historical correlation matrix of the drivers, while the idiosyncratic component is simulated as a standard Gaussian distribution $\mathcal{N}(0,1)$ scaled by the volatility σ_i of the reference issuer. The CWI processes are used to determine the default of a given obligor. This occurs when X_i falls below a given threshold T_i , related to their rankings:

$$D_n = \{X_i \leq \Phi^{-1}(\text{PD}_i)\}, \quad (42)$$

with Φ the cumulative distribution of a standard Gaussian and PD_i the probability of default of the i -th obligor.

Recovery Rates

Once defaults have been determined in each scenario, the recovery rate associated with each legal entity is computed.

The model currently implemented assumes that the recovery rate process for obligor i , R_i , is identical to the CWI process:

$$R_i = \frac{X_i}{\rho_i}. \quad (43)$$

Finally, the final recovery rate value is obtained using the inverse of the Beta distribution:

$$r_i = \beta_i^{-1}(F_i), \quad (44)$$

where F_i is a variable defined in $[0,1]$. The use of a β function is justified because the output is a variable in the range $[0,1]$. The parameters α and β of the Beta function are derived from input data and vary for each obligor. It is defined as:

$$F_i = \left(\frac{\Phi(\mathbf{R}_i / \sqrt{\Delta t})}{\Phi(\zeta_i / \sqrt{\Delta t})} \right), \quad (45)$$

where Φ is the cumulative Gaussian distribution, and ζ is the threshold determining the default of obligor i after time Δt . This is possible since the argument of the cumulative function corresponds to the CWI value, and the re-scaling ensures maximum recovery rate values when R_i equals the default threshold.

PnL Calculation

PnL is computed by aggregating the individual PnL for each instrument. The PnL for each instrument j is defined as:

$$PnL_j^S = (MtM_j - VaD_j^S) \times \mathbb{1}_{D_i}, \quad (46)$$

where MtM_j is the market value of the instrument, and VaD_j^S is the simulated value of instrument j and $\mathbb{1}_{D_i}$ takes into account the default of the legal entity i corresponding to instrument j . The actual calculation of VaD^S varies depending on the instrument and on the simulated recovery rate R_i^S . For example:

1. $VaD_j^S = MtM_j \times R_i^S$ if the instrument is a Bond;
2. $VaD_j^S = MtM_j \times (1 - R_i^S)$ if the instrument is a Credit Default Swap.

Once the PnL for each instrument j is computed, it is summed across all the instruments. This produces a PnL value for each scenario simulated and thus a vector with a length equal to the number of scenarios, whose 99.9th quantile represents the DRC value for the portfolio.

The Quantile Estimation

Review of Classical Results

The empirical quantile is used very often as a "plug and play" tool to estimate from the data the unknown true quantile, but at the end it is just one of the many approaches for the estimation purpose, exactly as the arithmetic median is an alternative to the classical arithmetic average to estimate an expected value. Let us introduce some simple notation. Let be X the random variable source of our data and suppose to have an i.i.d sample drawn from this distribution. We indicate with $X_{(n)}$ the n -th order statistics, i.e. the n -th value after sorting by ascending order all the outcomes. Equipped with this notation, the distribution F_n of the n -th order statistics is given by:

$$F_n(x) = \Pr\{X_{(n)} \leq x\} = \Pr\{\text{all } X_i \leq x\} = F^n(x), \quad (47)$$

while the empirical quantile $Q_n(\alpha)$ (supposing that the positive values represent the profits, the negative values the

	Empirical	RF	TF	HD	EP	Empirical DRC
10⁴ Scenarios						
1st case	32.3%	16.7%	15.4%	14.8%	15.5%	311,050,198
2nd case	42.6%	36.4%	34.1%	29.1%	31.3%	97,262,380
3rd case	73.2%	36,2%	45,0%	41.2%	39.2%	189,501,850
2x10⁵ Scenarios						
1st case	6,66%	6,08%	5,96%	5,61%	5,98%	334,983,354
2nd case	23.4%	6.51%	6.37%	6.84%	6.34%	113,399,158
3rd case	35.5%	14.7%	14.3%	14.7%	14.2%	187,788,177
16x10⁶ Scenarios						
1st case	1.57%	1.44%	1.40%	1.37%	1.41%	329,265,406
2nd case	1.55%	1.38%	1.41%	1.68%	1.41%	112,998,640
3rd case	1.28%	1.16%	1.17%	1.30%	1.17%	186,601,892

TABLE 4: Uncertainty for 3 different portfolio, with 10^4 , 2×10^5 and 16×10^6 scenarios. In the table RF, TF, HD and EP stand respectively for Rectangular filter, triangular filter, Harrell-Davis estimator, Epanechnikov estimator. In the last column we have reported the empirical DRC value.

losses) writes as below:

$$Q_n(\alpha | X_1, \dots, X_n) = X_{[(1-\alpha) \times n]}. \quad (48)$$

The integer part operator $[]$ is needed to take an actual outcome from the sample. There are some slightly different versions, according to less or more conservative (prudent) approaches. The most popular statistical tools (R, SAS, Matlab, Excel) also allow for different implementations of the quantile. A very relevant result for the order statistics is their asymptotic distribution. It can be shown that while the $\min()$ and the $\max()$ of the distribution never converge to the gaussian random variable, it happens for all the other order statistics. Namely we have the following result

$$E(Q_n(\alpha)) = Q(\alpha) - \frac{\alpha(1-\alpha)f'(Q(\alpha))}{2(n+2)f^3(Q(\alpha))} + O(1/n^2); \quad (49)$$

$$\text{Var}(Q_n(\alpha)) = \frac{\alpha(1-\alpha)}{(n+2)f^2(Q(\alpha))} + O(1/n^2). \quad (50)$$

A seminal reference in this field is the textbook by [8] and in the work by [3]. If one analyzes the uncertainty of the empirical quantile as the parameters (n, α) change, one easily finds that the more extreme is the confidence level α and smaller is the sample size n , then less accurate (high variance) is the quantile estimator. Several attempts have been made in the inferential statistics field to define better estimators in the usual bias-variance trade-off.

Advanced Estimators from Order Statistics

We select some alternative estimators of the quantile as competitors of the basic empirical estimation. Considering the definition for an L -estimator as:

$$Q_n = \sum_i w_i X_i. \quad (51)$$

The *rectangular* filter is defined by $w_i = 1/n$ in an interval $[\alpha - \epsilon, \alpha + \epsilon]$.

The *triangular* filter is defined in the same interval with w_i maximum at α and symmetrically decreasing to 0 at $\alpha \pm \epsilon$. *Harrell-Davis* estimator [7]:

$$w_i = I_{j/n}(a, b) - I_{(j-1)/n}(a, b),$$

with

$$a = \alpha(n+1) \quad b = (1-\alpha)(n+1) \quad I_x(a, b) = \phi(\beta(a, b)),$$

where α is the confidence level, n total number of scenarios, ϕ cumulative distribution function and β is the beta Euler function.

Finally we have the *Epanechnikov* estimator:

$$w_i = K_{j/n}(\alpha, h) - K_{(j-1)/n}(\alpha, h),$$

where

$$K_x(\alpha, h) = \begin{cases} 0 & \text{if } x \leq \alpha - h \\ \frac{1}{2} + \frac{3}{4} \frac{x-\alpha}{h} - \frac{1}{4} \left(\frac{x-\alpha}{h} \right)^3 & \text{if } \alpha - h < x < \alpha + h \\ 1 & \text{if } x \geq \alpha + h \end{cases}$$

$h = 0.0005$.

Comparison of the Different Estimators

We extracted the PnL results from the DRC described in the previous section and applied the estimators presented in the "Introduction". The errors were assessed using the jackknife resampling method. Our findings indicate that the results obtained from the different estimators are consistent, with error estimates that remain comparable across 16 million scenarios. Indeed, for a sample portfolio composed of some thousands of positions (bonds, derivatives), belonging to




FIGURE 8: Uncertainty for 3 different portfolio, with 10^4 scenarios. In the legend RF, TF, HD and EP stand respectively for Rectangular filter, triangular filter, Harrell-Davis estimator, Epanechnikov estimator.




FIGURE 9: Uncertainty for different scenarios. In the legend RF, TF, HD and EP stand respectively for Rectangular filter, triangular filter, Harrell-Davis estimator, Epanechnikov estimator.

about thousand obligors, we obtained the relative uncertainties expressed in Tab. 4. The uncertainty in the table represent the half width of a 95% confidence level of the quantile estimator. It is calculated by bootstrap approach for the more advanced tools, by the analytical results in Section "Review of Classical Results" for the empirical quantile. The same results are graphically expressed in Fig. 8 and 9.

The best result was obtained with the Harrell-Davis estimator[7], but the improvement is only 16% compared to the uncertainty of the empirical quantile. We analyzed these results by considering correlation values in the range $[0.9985, 0.9995]$. Specifically, we computed the correlation between the vector containing the elements at positions

$[0.9985, 0.999]$ and the vector $[0.9985 + \epsilon, 0.999 + \epsilon]$, with $0 < \epsilon < 0.0005$. We obtained a mean correlation of 0.999, which indicates that, in this range, the points are highly correlated. Consequently, there is no significant difference whether we consider the position 0.999 directly or a nearby range. In other words, the correlation between any order statistics and the next one is so high that the smoothing technique embedded in the L-Estimators does not provide a relevant benefit in reducing the variance. Furthermore, since this is an extreme point, the range cannot be symmetrically extended beyond $\epsilon = 0.001$. To further support our analysis, we tested different distributions, extracting datasets of size 10^5 , 10^6 , and 10^7 , and comparing the results. We consid-

ered both the normal distribution (not heavy-tailed) and the lognormal distribution (heavy-tailed), but we present here only the results for the lognormal distribution, as they are similar to those obtained with the normal distribution. We obtained the following absolute differences in percentage error:

- 10^5 points: $\sim 24\%$;
- 10^6 points: $\sim 10\%$;
- 10^7 points: $\sim 0.05\%$.


These values are computed as:

$$\frac{|\sigma(HD) - \sigma(\text{emp})|}{q}, \quad (52)$$

where q represents the empirical quantile value.

This result demonstrates that, beyond a certain dataset size, the different estimators yield the same uncertainty. Consequently, their adoption does not provide a significant advantage over the empirical quantile estimator.

Conclusions

We exploit several techniques to face a very hard statistical problem, the estimation of an extreme quantile required by the banking regulations. The benchmark was the empirical quantile, some competitors come from classical theory, such as the filters, other were proposed by the statistical literature. We run some exercises on a real world large portfolio. We found that for a relative small number of simulations n , such as $O(n) = 10^4, 10^5$, the Harrel-Davis and the Epanechnikov methods show a relevant improvement in the variance reduction goal. when we can perform 10^6 or more simulations, the uncertainty of the more advanced tools becomes very close to the basic empirical quantile. To summarize, a bank should properly combine its hardware and software resources (and constraints) with its accuracy target, in order to achieve an adequate and sustainable risk measurement process. 

References

- [1] **Acerbi, C. and Székely, B.** *Backtestability and the ridge backtest.* *Frontiers of Mathematical Finance*, Vol. 2, issue 4, pp. 497-521, 2023.
- [2] **Basel Committee on Banking Supervision.** *Minimum Capital Requirements for Market Risk.* BCBS 457, 2019.
- [3] **Chen, E. and Kelton, D.** *Two-Phases Quantile Estimation.* *Proceedings of the 2001 Winter Simulation Conference*, pp. 451-459 January, 2002.
- [4] **European Central Bank.** *ECB Guide to Internal Models.* ECB, October 2019.
- [5] **Garman, M.** *Taking VaR to pieces.* *Risk Magazine*, vol. 10, October 1997.
- [6] **Hallerbach, W.** *Decomposing Portfolio Value-at-Risk: A General Analysis.* *The Journal of Risk*, vol. 5, November 1999.
- [7] **Harrell, F. E. and Davis, C. E.** *A new distribution-free quantile estimator.* *Biometrika*, vol. 69, issue 3, pp. 635-640, 1982.
- [8] **Herbert, D. and Nagaraja, H. N.** *Order Statistics.* *Journal of the American Statistical Association*, vol. 67, April 2004.
- [9] **McNeil, Alexander and Frey, Rüdiger and Embrechts, P.** *Quantitative Risk Management: Concepts, Techniques, and Tools.* *Journal of The American Statistical Association - J AMER STATIST ASSN*, vol. 101, October 2005.



AI Risk Management Frameworks

About the Authors



Nicola Mazzone:

Senior Project Manager

He holds an MSc in Economics For Finance from the Catholic University of Milan and has furthered his expertise with postgraduate master's degrees in Business Analytics, Corporate Finance, and Innovation and Sustainability Management. Currently, he is actively engaged in Business and IT process implementation projects at one of Europe's most prominent banking institutions.



Sara Martucci:

Project Manager

She holds an MSc in Mathematical Engineering with a focus on Quantitative Finance from Politecnico di Milano. She has gained significant experience as a Project Manager, leading international projects and developing strong expertise in project governance and delivery. Currently, she is involved in engineering and automation projects within the CIB division of one of Europe's leading banking groups, where she contributes to the transformation of complex processes into scalable, reliable software solutions.





Margherita Ranieri:

Analyst

She holds an MSc in Management Engineering with a specialization in Finance from the Politecnico di Milano. She is currently working as a business analyst, actively engaged in Business and IT process implementation projects at one of Europe's leading banking institutions.



AI Risk Management Frameworks

Nicola Mazzoni

Sara Martucci

Margherita Ranieri

This study aims to analyze AI Risk Management Frameworks (AI RMFs), exploring their role in promoting the safe, accountable, and transparent adoption of AI technologies within economic systems. The first part of the research provides a broad overview of the evolution of the AI market and its growing impact on strategic and operational processes, with a particular focus on the financial sector. The second part discusses the unique risks posed by AI systems, while the third part explores the regulatory responses to manage AI unique risks, with a particular focus on the EU AI Act. Finally, the fourth part analyzes several major AI RMFs developed by international and regional institutions, examining their guiding principles, technical requirements, and governance mechanisms. The study ultimately identifies common principles shared across regulations, guidelines, and AI RMFs, highlighting the strategic relevance of integrating AI governance into corporate strategy.

ARTIFICIAL Intelligence (AI), and particularly Generative AI (GenAI), is rapidly transforming industries worldwide, reshaping business models, and becoming a cornerstone of modern digital transformation. Driven by significant capital and human investments and a growing focus on productivity, efficiency, and innovation, AI technologies are being adopted across a wide range of sectors. However, this growing integration of AI also brings a complex set of challenges. As AI systems increasingly influence critical decisions, their unique characteristics, such as autonomy, opacity, and capacity for scale, raise important concerns around transparency, accountability, ethical use, and regulatory compliance. These challenges are amplified in highly regulated sectors like finance, where AI adoption must be carefully aligned with principles of financial stability, consumer protection, and institutional trust. Despite these concerns, the industry outlook for AI remains overwhelmingly positive. Barriers to adoption are progressively declining, supported by technological advances, cost reductions, and expanding availability of AI infrastructure. The continued rise in private investment, particularly from global leaders like the United States and China, further underlines the strategic relevance of AI for economic competitiveness and innovation leadership. The widespread deployment of AI technologies has intensified the need for coherent, harmonized regulatory frameworks capable of addressing the societal, ethical, legal, and economic risks associated with AI. Traditional regulatory approaches are increasingly proving insufficient, as AI's cross-sectoral and cross-border implications require a level of coordination that transcends national boundaries. In response, international and supranational bodies such as the OECD and UNESCO have developed guidelines to assist governments in navigating these challenges. At the same time, regional and national authorities have begun drafting or implementing specific regulatory mechanisms to manage AI risks, aiming to prevent legal fragmentation and ensure consistent oversight. A leading example of this evolution is the European Union's AI Act, which represents the first cross-jurisdictional regulatory framework explicitly dedicated to AI. The Act adopts a risk-based approach that classifies AI systems according to their potential impact on safety, fundamental rights, and societal stability. This paper aims to

provide a concise overview of the evolution of the AI market and the corresponding regulatory responses to emerging challenges, with a particular focus on risk management frameworks. While offering context on the development of the AI market and the evolving regulatory landscape, it seeks to examine the methodological approaches developed to assess and manage AI-related risks. Specifically, the paper explores the key features and practical implementation of structured risk management frameworks to evaluate how institutions are addressing major AI-specific concerns such as bias, opacity, accountability gaps, and systemic instability.

AI: a Brief View

With its spread across various industries and daily activities, the term "AI" is becoming increasingly overused. While artificial intelligence is a branch of study that encompasses several types of models and algorithms, the term is often incorrectly used to indiscriminately refer to all potential AI applications that fall under the broader AI umbrella. Artificial Intelligence (AI) "is the science and engineering of making intelligent machines, especially intelligent computer programs. It is related to the similar task of using computers to understand human intelligence, but AI does not have to confine itself to methods that are biologically observable".^[15] The words of J. McCharty¹ describe a branch of computer science focused on studying and developing systems capable of performing complex tasks typically requiring or associated with human intelligence. These include learning from data, understanding natural language, recognizing patterns, solving problems, and making autonomous decisions. The field of study of AI encompasses several different typologies of approaches and models that can serve different purposes:

- **"Machine Learning:** mathematical and statistical methods enabling machines to learn from data and improve with experience. It comprehends:
 - **Supervised Learning:** models that learn from input features and targets (training dataset)

¹J. McCarthy is considered one of the founders of Computer Science. His 1956 speech at Dartmouth University introduced for the first time the term Artificial Intelligence.




FIGURE 10: AI Methods

to generalize the model and make predictions on unseen data (e.g. identifying spam emails, image classification, etc.).

- **Unsupervised Learning:** models that work with unlabeled data, aiming to discover patterns and relationships within datasets that lack predefined target labels (e.g. clustering, anomaly detection).
- **Reinforcement Learning:** models that rely on agents that learn through their interaction with the environment, thanks to a reward system that assesses the quality of the agent's actions."[19]
- **Deep Learning:** a sub-branch of AI that refers to those machine learning models whose structure is based on the study of multiple successive layers, allowing algorithms to learn data representations at multiple levels of abstraction.
- **Generative AI:** GenAI is a subset of deep learning models focused on creating content. It is capable of generating diverse types of data, such as text, images, video, audio, and code, by learning patterns from large datasets and using that knowledge to produce new and original outputs. GenAI models are typically based on neural networks that are trained on massive amounts of data to understand the statistical relationships between elements such as words, pixels, or sounds. Once trained, these models can generate highly realistic and complex content that mimics human creativity, supports human labor by automating repetitive or time-consuming tasks, and serves as a powerful tool for enhancing productivity across multiple domains. They achieve this by predicting the next most likely element (e.g. word in a sentence) based on the patterns learned during training.
- **Large Language Models:** LLMs are a subset of GenAI models designed to understand and produce human-like language. Trained on vast text datasets, they aim to generate coherent and contextually relevant text by using statistical and probabilistic models to predict the next word or token in a sequence.
- **Agents:** agents are autonomous systems powered by Generative AI that can make decisions, complete tasks, and learn from experience. They can interact

with their environment or users, adapt to changing conditions, and automate complex workflows across various domains such as customer service, robotics, and data analysis.

AI Market Evolution

Since 2022, with the first public release of ChatGPT, AI has undergone a profound shift, transitioning from a specialized domain to a central focus of public discourse and strategic planning across several industries. The rapid adoption of these technologies (in particular GenAI) has led to widespread interest among both businesses and policymakers, driven by AI's demonstrably transformative impact on productivity and economic development. Enterprises across diverse sectors, recognizing the transformative potential of AI, have seized the momentum to increasingly integrate it into both their operational frameworks and strategic planning. This burgeoning interest is mirrored by a significant increase in global AI investment over the past decade, now amounting to hundreds of billions of dollars worldwide. Investments in the industry (ref. Figure 11) have seen a substantial increase: from 2023 to 2024, investments in AI grew by approximately 26%, rising from 201\$ billion to 252\$ billion. Notably, over the past eleven years, total investment in AI surged from around 14\$ billion in 2013 to 252\$ billion in 2024, with a historical peak of 360\$ billion reached in 2021, prior to the public release of ChatGPT. [18]

In the same period, the value of private investment (defined as the share of total investment excluding M&A, public offerings, and minority stakes) has also seen a substantial increase over the past year (ref. Figure 12), rising from 104\$ billion to slightly over 150\$ billion, marking a growth of approximately +45%. Looking at the period since 2013, private investment, despite a temporary slowdown in 2022 and 2023, has followed an overall upward trend, with its value more than tenfold over the past eleven years, growing from 13,34\$ billion in 2013 to 150\$ billion in 2024. [18] Taking a deeper look at who is driving most of the private investment in AI (ref. Figure 13), it's not surprising that over the last eleven years, the United States and China have led the field, accounting for a combined 81% (590\$ billion) of the total aggregate investment among the top 15 countries investing in AI. [18]




FIGURE 11: Global Corporate Investment in AI [18]




FIGURE 12: Global Private Investment in AI [18]

In the last 11 years, only three EU countries, Germany, France, and Sweden, have consistently carried out investments that placed them among the top providers of private AI investment. However, a closer look at the past three years reveals interesting insights. While the United States and China continue to lead AI investments, other European countries such as Italy (0,86\$ billion), Austria (1,15\$ billion), and the Netherlands (1,09\$ billion) have significantly increased their share of private AI investment entering among the global top 15 countries for private AI investment in 2024 (ref. Figure 14). In the same period, Singapore and Australia, which were among the top 15 investors in 2022 and 2023, fell out of the top 15 in 2024, although they still rank among the top investors in terms of aggregate investment over the past 11 years. [18] [17] [16]

The data [18] confirms that AI is increasingly becoming a key growth market, as evidenced by the ever-growing flow of investment capital. Market outlooks project a compound annual growth rate (CAGR) of over 34% for the next five years, with the total market value expected to exceed 3.000\$

billion[21]. The growing importance of AI across industries is further demonstrated by rising adoption rates. In the EU (ref. Figure 15) [28], between 2021 and 2024, AI adoption has shown a consistent upward trend across various sectors, with an average increase of approximately 42% between 2021 and 2024. Notably, industries traditionally characterized by technology-intensive operations, such as IT (+96%) and pharmaceuticals (+72%), have experienced a steep increase in AI integration. However, it is particularly significant that sectors like administrative services (+103%), business services (+106%), and wholesale and retail commerce (+117%) have shown an even more substantial rise in AI adoption. [28]

It is clear that AI has expanded through both private and professional adoption. However, the emergence of ChatGPT has placed a strong spotlight on Generative AI. GenAI has gathered substantial and growing attention in recent years, primarily because of its accelerated adoption across an increasingly diverse array of industries. In a relatively short timeframe, GenAI technologies have exhibited a po-




FIGURE 13: 2013-2024 AI Private Investment by Country [18]

tential capacity to redefine operational workflows, optimize resource allocation, and support the automation of complex tasks. This rapid diffusion underscores GenAI's transformative potential, not only in terms of productivity gains and cost efficiency but also in enabling entirely new modes of value creation. Looking ahead, the continued integration of GenAI into organizational structures is anticipated to drive fundamental shifts in business and institutional models, fostering enhanced innovation, scalability, and adaptability across both the private and public sectors. To better understand the extent of GenAI rapid breakthrough, it is useful to compare its adoption rate with that of other major disruptive technologies (ref. Figure 16)[2]. Within just two years of becoming widely available, GenAI has reached a workplace adoption rate close to 40% (U.S. data). To better understand the magnitude of this data, it is useful to consider that the internet required approximately five years to reach a comparable level of adoption, while personal computers (PCs) took nearly twelve years to achieve the same diffusion.[2] Following the broader trend in Artificial Intelligence, Generative AI has attracted substantial private investment in recent years (ref. Figure 17)[18]. Notably, the growth in GenAI-related investments has significantly outpaced the average trend observed across the AI sector. Since 2019, private investment in GenAI has increased by 41 times within just five years. From accounting for a modest 0,8% of total private AI investment in 2019, equivalent to 0,8\$ billion, GenAI has rapidly expanded its share to nearly 14% in 2024, reaching a total value of almost 34\$ billion.[18]

AI Adoption in Financial Services

The financial services industry has historically served as a frontrunner in the adoption of technological innovations, leveraging them to enhance operational efficiency and maximize returns. In recent years, the integration of various forms of AI has expanded rapidly[12], with financial institutions deploying AI across a wide range of operational and business processes, from generic administrative tasks to highly specialized, industry-specific functions (ref. Figure 18) [12].

This trend holds true also in the adoption of GenAI, which has experienced a rapid and widespread acceleration within the financial services sector. Over the past two years (ref. Figure 19), the adoption rate of GenAI has increased by

more than 300%[20], reflecting its integration across a wide range of business functions. This substantial growth is not confined to general-purpose applications; rather, it encompasses both support activities and critical operational areas. Data[25] confirms this broad diffusion and highlights a growing reliance on GenAI for high-value tasks such as software engineering, code generation, and the automation of complex processes. This development points to an increasing awareness among financial institutions of GenAI's potential to foster innovation, enhance operational performance, and support strategic differentiation in a highly competitive market environment. Crucially, two of the three fastest-growing use cases, pricing and risk management, which increased by +167% between 2023 and 2024, and trading and portfolio optimization, which increased by +153% between 2023 and 2024, are core, industry-specific functions. This emphasizes the strategic relevance of GenAI in the financial sector and signals a transition from exploratory adoption to its institutionalization as a tool for sustained competitive advantage.[25]

Looking ahead, the industry outlook confirms a sustained commitment to increasing investment in AI and GenAI technologies. These technologies are increasingly regarded as strategic levers for business development, supported by the expansion of an AI-specialized workforce and the growing reliance on third-party partners to improve process efficiency and accelerate solution development. While overall investment levels are expected to grow, emerging evidence suggests a shift in budget allocation priorities (ref. Figure 20)[25]. Specifically, there appears to be a reduction in funding directed toward exploratory research into novel AI applications. Instead, strategic emphasis is being placed on talent acquisition, the reinforcement of collaborations with external service providers, and the enhancement of technological infrastructure required to develop, deploy, and maintain advanced AI systems at scale.[25]




FIGURE 14: 2022, 2023, 2024 AI Private Investment by Country Comparison [18] [17] [16]

Risks and the Increasing Need for Regulatory and Risk Management Frameworks

As shown in previous chapters, AI, and particularly GenAI, are ushering a profound transformation in how organizations across several industries operate, compete, and deliver value. From automating decision-making processes to enabling scalable customer engagement, AI is becoming a foundational component of enterprise strategy across nearly every industry. As adoption accelerates, organizations are not only integrating AI into discrete functions but are increasingly embedding it into the very structure of their business models. However, the growing reliance on AI technologies also introduces a range of complex and potentially high-impact risks. In this context, the development of clear regulatory frameworks and the implementation of robust AI Risk Management Frameworks (AI RMFs) has become essential. These efforts are not only necessary to mitigate potential harm but also to support the long-term sustainability, resilience, and trustworthiness of economic systems across all sectors. AI systems, while powerful, are not inherently neutral. They are built, trained, and operated by humans, and as such, they are susceptible to the full spectrum of human error, bias, and oversight. The risks arising from the deployment of AI technologies are multifaceted and increasingly material, comprising:

- **Inaccuracy and Hallucinations:** generative models may produce outputs that appear plausible but are factually incorrect or misleading, posing significant risks in contexts requiring accuracy and reliability;
- **Algorithmic Bias and Discrimination:** AI systems trained on biased data can replicate or even exacerbate social inequalities, affecting decisions related to employment, lending, healthcare, or law enforcement;

- **Data Privacy and Security:** AI often relies on sensitive personal or proprietary data, which, if mismanaged, can result in regulatory violations or large-scale data breaches;
- **Intellectual property concerns:** the use of AI tools to generate content based on vast datasets raises legal and ethical issues related to ownership, copyright, and content originality;
- **Lack of explainability:** the decision-making processes of many AI systems are opaque, limiting transparency and accountability;
- **Cybersecurity Vulnerabilities:** AI systems can be exploited through adversarial attacks or model manipulation, increasing the surface area for cyber threats;
- **Reputational Damage and Stakeholder Mistrust:** public failures or misuse of AI can erode trust in an organization, affecting customer loyalty, investor confidence, and employee morale.

Without proactive governance, these risks can lead to severe consequences, including regulatory sanctions, operational disruption, financial losses, and societal harm. To mitigate these risks and fully harness the transformative potential of artificial intelligence, organizations must transition from reactive risk responses to a proactive and structured governance strategy. This requires the implementation of cross-functional, enterprise-wide systems that embed responsible AI principles across every stage of the AI lifecycle, from initial design and data sourcing to model development, deployment, post-deployment monitoring, and eventual system decommissioning or retirement. This is where AI Risk Management Frameworks (AI RMFs) become indispensable. These frameworks provide a systematic and repeatable structure for integrating risk identification, mitigation, and oversight into an organization's broader technology and governance ecosystems. When properly implemented, AI RMFs enable organizations to:

- Identify and assess AI-specific risks early in the development cycle, ensuring that potential harms re-




FIGURE 15: EU Cross-Industry AI Adoption Between 2021-2024 [28]




FIGURE 16: Technologies Adoption Rate, US [21], Data [2]

lated to bias, inaccuracy, misuse, or non-compliance are detected before models are deployed at scale;

- Implement technical, ethical, and procedural safeguards that are tailored to the organization's risk tolerance and legal obligations, such as differential privacy measures, adversarial robustness, fairness constraints, or explainability features;
- Ensure alignment with applicable laws, standards, and ethical guidelines, including emerging national and international AI regulations, industry codes of conduct, and human rights frameworks;
- Establish formal accountability structures by clearly defining ownership and responsibilities for AI oversight across business units, data science teams, legal departments, and executive leadership, including escalation procedures for adverse events;
- Continuously monitor AI systems post-deployment through performance metrics, fairness audits, real-time alerts, and incident reporting mechanisms to

ensure that the system continues to operate in a safe, effective, and equitable manner throughout its lifecycle.

Guidelines issued by supranational bodies could play a pivotal role in supporting the establishment of robust and coherent regulatory frameworks for AI. These regulatory frameworks, grounded in shared principles of safety, transparency, ethics, and accountability, could serve as essential references for national legislators and supervisory authorities, thereby potentially ensuring a harmonized and comprehensive approach to AI risk management at the global level. In turn, these regulatory frameworks could provide the foundation for the development and implementation of AI RMFs within organizations. On the other hand, well-designed AI RMF could enable organizations to translate regulatory guidelines into concrete operational practices, facilitating the identification, assessment, and mitigation of risks associated with AI deployment. When designed strategically, they become powerful enablers of innovation. By embedding trust, transparency, and ethical integrity into




FIGURE 17: Total Private Investment in GenAI [18]




FIGURE 18: Total FI AI Application Across Business Functions[12]

AI development, these frameworks provide organizations with the confidence and legitimacy needed to scale AI adoption responsibly. They help streamline decision-making, eliminate uncertainty, and foster internal alignment, allowing cross-disciplinary teams to collaborate more effectively and accelerate delivery cycles without compromising risk standards. Furthermore, by demonstrating a proactive commitment to ethical and responsible AI, organizations can strengthen stakeholder trust, enhance brand reputation, and distinguish themselves in increasingly competitive and scrutinized markets. This trust becomes particularly valuable in customer-facing industries or regulated sectors, where public perception and compliance readiness are essential to long-term success.

AI Risks: A View on the Financial Industry

As shown in the previous chapter, financial institutions have already started integrating AI and GenAI into both operational processes and business-specific tasks, and they plan to continue investing in their development and strategic

integration in the coming years. However, AI, and GenAI in particular, raise several concerns regarding potential risks for industry players. In 2025, both IOSCO[12] and the Japan FSA[13] conducted comprehensive analyses exploring the current use and outlook of AI in the financial sector through surveys of various financial entities. The results highlighted that, despite the significant potential benefits for business innovation and efficiency, it is essential to thoroughly assess and manage the risks associated with these emerging technologies. These risks and challenges can be summarized as follows:

- **Lack of Explainability:** many AI systems operate as "black boxes," making it difficult for institutions to interpret or justify decisions, posing challenges for transparency and accountability;
- **Data Privacy and Quality:** AI's performance heavily depends on the quality, relevance, and accuracy of data. Poor data management can result in flawed outputs, while over-reliance on personal data raises privacy and ethical concerns;
- **Cybersecurity and Operational Risk:** AI systems




FIGURE 19: Top Generative AI Use Cases in Financial Services Industry [25]

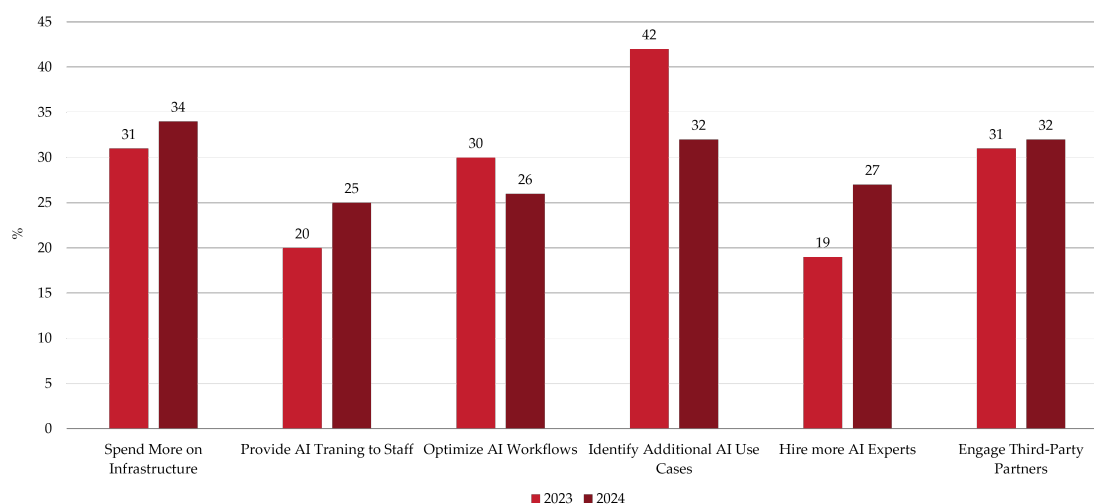


FIGURE 20: Financial Services Industry Investment Plans for 1 Year [25]

may introduce new vulnerabilities, including susceptibility to adversarial attacks, model drift, and technical failures that disrupt financial operations or cause regulatory breaches;

- **Third-Party Risk:** many institutions rely on external vendors for AI tools, raising concerns about oversight, vendor lock-in, and exposure to unregulated service providers;
- **Difficulties in Model Governance:** AI introduces unique model risks due to its broad applicability, non-deterministic behavior, and reliance on externally managed foundation models. These characteristics pose several challenges to traditional risk management frameworks, potentially resulting in an insufficient ability to fully assess, control, and explain the risks associated with generative AI systems.

Other than those mentioned above, the Japan FSA report[13] underlined some new risks and concerns arising from the adoption of GenAI models:

- **Hallucination:** hallucination risks raise several concerns due to the potential for misleading outputs

and information when generative AI models are used in business applications or decision-making processes. It is therefore essential to implement systems that maintain a "human-in-the-loop" to ensure robust oversight and avoid potential critical damages.

- **Financial Crime:** criminal methods are becoming increasingly sophisticated due to the adoption of AI and GenAI, particularly amplifying the potential risks to financial institutions and their customers. For example, the advent of generative AI further increases risks by automating the production of highly believable written text, audio, and images.
- **Systemic Risk:** the growing integration of GenAI into business functions and decision-making processes may increase the risk of highly correlated behaviors of market players, as similar AI-generated signals lead to uniform decisions. This convergence could foster herd behavior, amplifying market volatility and significantly increasing systemic risk.

Most jurisdictions have not adopted AI-specific regulations for the financial sector. Instead, they apply existing

technology-neutral regulatory frameworks, which already cover key areas such as risk management, governance, cybersecurity, data protection, and consumer protection. As a result, many of the cross-sectoral themes relevant to AI are broadly addressed under current financial regulations, making the need for dedicated AI-specific financial rules debatable. As reported by BIS[4] his regulatory stance likely explains why financial authorities are not planning new AI-specific rules in the near term, while actively evaluating whether additional measures are needed to address AI-specific risks in the currently developed frameworks[27].

AI Regulatory Framework

As illustrated in previous sections, AI adoption has accelerated significantly, increasing the urgency for robust oversight frameworks capable of addressing its complex challenges and potential risks. In response, and in a context marked by fragmented national policies and limited institutional expertise, supranational bodies have begun implementing targeted frameworks and proposing guidelines aimed at fostering openness in AI practices, strengthening governance responsibilities, and aligning supervisory approaches across jurisdictions. Within these efforts [21] led by supranational bodies, the first globally recognized initiative to establish a normative framework for AI governance came from the OECD, which adopted the OECD Principles on Artificial Intelligence in May 2019 (updated in 2024) [29]. Endorsed by over 40 countries, including all OECD members and several non-member states, these principles are intended to promote the responsible stewardship of trustworthy AI. The OECD Recommendation on AI seeks to advance the development and use of AI that is trustworthy and human-centric, to support responsible innovation, to safeguard human rights and democratic values, to foster international cooperation, and to guide public policy through a shared global framework. The Recommendation is structured around five key principles for responsible AI and five corresponding policy recommendations for national and international action. The five value-based principles for the responsible development and use of AI are [29]:

1. **Inclusive Growth, Sustainable Development, and Well-Being:** AI should benefit people and the planet by driving inclusive economic growth and sustainable development;
2. **Human-Centered Values and Fairness:** AI systems should be designed in a way that respects human rights and democratic values, including privacy, liberty, and equality;
3. **Transparency and Explainability:** the functioning of AI systems should be transparent to users and regulators, and decisions should be explainable where possible;
4. **Robustness, Security, and Safety:** AI systems must be technically robust and secure, and should function appropriately throughout their lifecycle;
5. **Accountability:** organizations and individuals developing, deploying, or operating AI systems should be accountable for their proper functioning.

In addition to these principles, the OECD outlines five strategic recommendations for policymakers [29]:

1. Promote investment in AI research and development;
2. Foster a digital ecosystem for AI;

3. Ensure a policy environment that promotes trustworthy AI;
4. Equip people with the necessary skills to interact with and benefit from AI;
5. Encourage international cooperation to ensure the global alignment of AI governance.

These principles and recommendations serve as a foundational reference for the development of numerous national and international AI policy frameworks that followed.

EU AI Act

Inspired by the OECD principles [29], the European Union has been pioneering in the definition of an AI regulatory framework. The AI Act², which represents the first cross-jurisdictional regulatory framework focused on artificial intelligence, establishing a harmonized set of rules for development, market introduction, deployment, and use of AI across the EU, entered into force in August 2024. It introduces a regulatory roadmap that outlines the gradual application of several rules, which will come fully into effect by August 2026. Key innovations of the AI Act include:

- Binding rules for high-risk applications (e.g. biometric identification, credit scoring, recruitment tools);
- Strict requirements for data quality, documentation, human oversight, and cybersecurity;
- Creation of national supervisory authorities and a European AI Office;
- Enforcement mechanisms with fines up to 7% of global turnover[19].

An Harmonized Regulatory Framework

The AI Act represents a significant regulatory milestone within the EU's digital policy agenda. It is part of the broader initiative Europe Fit for the Digital Age, which, among other objectives, seeks to position the European Union as a global leader in the development of trustworthy and human-centric AI. Unlike voluntary guidelines or fragmented national regulations, the AI Act establishes uniform, directly applicable rules across all Member States. Its legal basis derives primarily from Article 114 of the Treaty on the Functioning of the European Union (TFEU)[9], which enables the EU to adopt measures for the approximation of laws to ensure the functioning of the internal market. However, the AI Act goes beyond market concerns by explicitly incorporating the protection of fundamental rights, as enshrined in the Charter of Fundamental Rights of the European Union[5]. In doing so, the Act reflects the EU's dual objective: fostering innovation in AI while ensuring that such innovation does not come at the expense of safety, human dignity, and democratic values. Before describing regulatory measures, it is important to highlight a fundamental challenge: the lack of a universally accepted definition of Artificial Intelligence. As emphasized in the OECD 2024 report "Regulatory Approaches to Artificial Intelligence in Finance"[27], the absence of a shared understanding of what constitutes AI complicates regulation and supervision. The European Union has adopted a formal definition of AI, while many countries rely on non-binding, non-prescriptive definitions or lack a standardized description altogether. Against this backdrop, the AI Act provides a formal and legally binding definition of Artificial Intelligence, aiming to bring clarity and legal certainty to stakeholders operating across the EU. The definition adopted by the Act is closely aligned with the approach developed by the OECD in its

²For an extensive analysis, we suggest[5].

2023 revision, reflecting international efforts towards regulatory convergence. According to Article 3[9] of the AI Act, an AI system is defined as: "A machine-based system that is designed to operate with varying levels of autonomy and that may exhibit adaptiveness after deployment, and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments". [9] This definition does not aim to establish a fixed list of AI systems but rather provides a flexible, technology-neutral guideline to respond to the rapid technological evolution of the market. As defined in Recital 12 of the AI Act, this definition should not be applied mechanically to every AI system; instead, each system must be assessed based on its specific characteristics. AI Act Recital 12[9], together with the "Guidelines on the definition of an artificial intelligence system"[8], specifies that a system qualifies as AI only if it meets certain criteria. AI systems must be able to infer outputs from the input data they receive, enabling them to generate predictions, recommendations, and decisions by using models and algorithms (e.g., machine learning approaches, logic- and knowledge-based approaches). Conventional, rule-based software or deterministic algorithms that do not perform such inference particularly if they lack autonomy or adaptiveness (e.g. systems for improving mathematical optimization, basic data processing systems, systems based on classical heuristics, simple prediction systems), fall outside the scope of this definition. By adopting this definition, the AI Act ensures that its provisions apply not only to fully autonomous AI systems but also to systems that operate with partial autonomy or human oversight, if their outputs have the potential to influence decision-making processes or produce effects in physical or digital environments. This definitional clarity represents a crucial step towards ensuring legal consistency, preventing regulatory loopholes, and enabling effective enforcement of the rules across the internal market. The AI Act applies to a broad range of actors involved in the lifecycle of AI systems, regardless of their geographical location, provided their activities have an impact within the EU. The regulation covers:

- **Providers:** any natural or legal person, public authority, agency, or other body that develops an AI system or has it developed and places it on the market or puts it into service under their name or trademark;
- **Users:** individuals or organizations deploying AI systems within the EU in the context of their professional activities;
- **Importers and Distributors:** entities responsible for ensuring compliance when AI systems from outside the EU are introduced into the European market;
- **Third-Country Providers:** AI system providers established outside the EU whose products or services are offered to users within the Union.

This extraterritorial scope mirrors similar approaches adopted in other EU regulations, such as the General Data Protection Regulation (GDPR), reinforcing the EU's ambition to influence global AI governance standards.

Risk-Based Classification of AI Systems

The AI Act introduces a tiered regulatory approach based on the potential risks AI systems pose to health, safety, and fundamental rights. This risk-based framework allows for proportional obligations, ensuring that regulatory intervention corresponds to the level of risk. The four levels of risk identified are:

1. Unacceptable Risk;

2. High-Risk;
3. Limited-Risk;
4. Lower-Risk.

AI Systems Prohibited Due to Unacceptable Risk

Certain AI applications are considered fundamentally incompatible with EU values, such as human dignity and democracy, and are therefore prohibited outright. These include:

- AI systems that use subliminal techniques to manipulate individuals' behavior in a manner that may cause harm are strictly prohibited. This restriction is designed to safeguard individual autonomy and ensure freedom of thought and decision-making.
- Exploitative AI that takes advantage of vulnerable groups, such as children, individuals with disabilities or groups defined by socio-economic status.
- The deployment of real-time biometric identification systems by law enforcement in publicly accessible areas is generally not allowed. However, narrowly defined exceptions apply, limited to situations where such technology is strictly required to locate victims of crimes such as abduction, human trafficking, or sexual exploitation; to prevent an imminent and serious threat to individuals' lives or physical safety, including terrorist threats; or to identify suspects involved in criminal offences, for the purposes of investigation, prosecution, or enforcing judicial decisions.
- AI-based social scoring by public authorities, particularly when it leads to discriminatory or unjustified treatment of individuals.
- In addition to the restrictions already mentioned, the AI Act also prohibits the use of AI systems for emotion recognition in sensitive contexts such as the workplace and educational institutions, unless strictly necessary for health or safety purposes. This measure is intended to safeguard individuals from intrusive evaluations that could result in discrimination or negative consequences based on inferred emotional states.

Since February 2025[19], the use of AI systems classified as posing an unacceptable risk has been officially prohibited under Article 5 of the AI Act [9]. This prohibition reflects the EU's emphasis on preserving human autonomy, dignity, and protection from undue surveillance or manipulation. To assist stakeholders in identifying such prohibited AI practices, the European Commission published the Guidelines on Prohibited AI Practices[7] in February 2025, providing further clarity on the types of applications deemed fundamentally incompatible with EU values. These aim to clarify the scope and concrete application of the prohibitions set out in the mentioned Article 5, ensuring consistent enforcement across the European Union. Specifically, the Guidelines provide:

- A set of practical indicators for determining whether an AI system uses subliminal techniques in a way that may significantly impair an individual's ability to make autonomous decisions;
- Clarifications on what constitutes exploitation of vulnerabilities, including illustrative scenarios involving minors, persons with disabilities, or socio-economically disadvantaged groups;
- Concrete examples of AI-based social scoring practices by public authorities that are likely to produce unjustified or disproportionate negative effects on individuals, thereby falling under the prohibition.

- Specific parameters for assessing the use of real-time remote biometric identification by law enforcement, including how to evaluate the existence of exceptional circumstances that justify its deployment, as well as procedural safeguards required under such circumstances;
- Criteria for identifying prohibited emotion recognition applications in sensitive environments, including how to distinguish between acceptable uses for health or safety reasons and practices that may result in invasive monitoring or discriminatory consequences.

High-Risk AI Systems and Compliance Requirements

Systems are classified as high-risk when their use may significantly affect the health, safety, or fundamental rights of individuals, as set out in AI Act Article 6 [9]. This category includes AI systems intended to serve as safety components of products subject to third-party conformity assessments, as well as AI systems performing profiling of natural persons. Moreover, AI systems are considered high-risk when their deployment, due to their intended purpose or the specific context of use, carries a tangible risk of harm or adverse impact on fundamental rights. Conversely, systems that merely perform supporting or preparatory tasks, without substantially influencing decision-making or replacing human judgment, are generally excluded from the high-risk classification. The list of high-risk AI use cases, defined primarily in Annex III of the AI Act, may be updated by the European Commission in light of technological developments or emerging risks. Regarding so, at the beginning of June 2025, the European Commission's AI Office has launched a 6 weeks targeted stakeholder consultation in order to "to collect input from stakeholders on practical examples of AI systems and issues to be clarified in the Commission's guidelines on the classification of high-risk AI systems and future guidelines on high-risk requirements and obligations, as well as responsibilities along the AI value chain". [10] To mitigate potential risks, high-risk AI systems must comply with a comprehensive set of obligations throughout their lifecycle. These include:

- **Risk Management and Mitigation:** providers are required to establish and maintain a comprehensive risk management system that applies throughout the entire lifecycle of the AI system. This system must ensure the continuous identification, assessment, and mitigation of risks, taking into account both the intended use and reasonably foreseeable misuse scenarios. Moreover, the risk management process must be subject to regular, systematic review and, where necessary, updated to reflect new information, evolving use cases, or emerging risks.
- **Data Governance and Quality:** AI providers are required to implement robust data governance and quality management processes covering all phases of the system's lifecycle. These processes must guarantee that datasets used for training, validation, and testing are sufficiently relevant, representative, and appropriately curated to minimize biases and inaccuracies. The objective is to ensure that AI systems operate in a fair, reliable, and non-discriminatory manner, fully aligned with the requirements of the regulatory framework.
- **Technical Documentation and Record-Keeping:** providers are required to prepare and maintain comprehensive, up-to-date technical documentation demonstrating the AI system's compliance with the

obligations set out in the AI Act. This documentation must cover all relevant aspects of the system, including its design, development processes, intended purpose, risk management measures, and performance evaluation. In addition, AI systems must be designed to enable systematic record-keeping of key operational events, including those that may have an impact on safety, fundamental rights, or that indicate substantial modifications to the system. These requirements ensure that competent authorities can effectively assess compliance and investigate potential incidents or risks.

- **Transparency and User Instructions:** to ensure transparency and promote responsible use, providers must supply clear, accurate, and accessible information to users regarding the AI system. This includes comprehensive instructions for use, a description of the system's capabilities and intended purpose, as well as its known limitations and potential risks. Users must also be informed about any conditions or constraints under which the system may or may not perform reliably. These requirements are essential to enable users to make informed decisions, use the system appropriately, and avoid unintended consequences.
- **Human Oversight:** systems must incorporate safeguards that allow effective human monitoring and intervention, including the possibility to override or disable automated operations.
- **Robustness, Accuracy, and Cybersecurity:** AI systems must meet high technical standards for reliability, accuracy, resilience against manipulation, and protection against cyberattacks.

Before being placed on the market, high-risk systems must undergo a conformity assessment, typically conducted internally, but requiring third-party certification by notified bodies for specific applications. After deployment, providers must establish post-market monitoring systems and report serious incidents or malfunctions to competent authorities.

Regulatory Approach to Limited and Minimal Risk AI Systems

In addition to the stringent requirements imposed on high-risk and prohibited AI applications, the AI Act introduces a differentiated regulatory regime for AI systems classified as posing limited or minimal risk. This approach reflects the EU's intention to strike a balance between safeguarding fundamental rights and promoting innovation, applying obligations proportionate to the potential risks associated with the use of AI.

Limited Risk AI Systems

AI systems falling under the category of limited risk are not subject to strict compliance requirements but must adhere to specific transparency obligations, as explicitly outlined in Article 52 of the AI Act. These obligations are designed to ensure that users are aware when they are interacting with an AI system, particularly in cases where the AI might influence their perceptions, decisions, or behavior without their full awareness. Examples of limited risk AI systems include:

- **AI Chatbots and Virtual Assistants:** users must be clearly informed that they are engaging with an AI-driven system rather than a human. This is intended to avoid confusion or deception in digital interactions.

- **AI-Generated Contents (e.g. deepfakes):** when an AI system produces synthetic audio, images, video, or text intended to resemble authentic content, it must be explicitly labeled as artificially generated or manipulated. This measure aims to prevent misinformation and protect individuals from deception.
- **Emotion Recognition and Biometric Categorization Systems (outside of high-risk contexts):** in these cases, transparency requirements apply to inform individuals about the use of such technologies. It is important to note, however, that when these systems are deployed in sensitive environments, such as workplaces, educational settings, or law enforcement, their risk classification may be elevated to high-risk, triggering more stringent obligations.

The transparency requirements for limited risk AI systems do not extend to imposing technical or organizational controls beyond the obligation to inform users. Nonetheless, these provisions play a crucial role in promoting trust and user awareness in AI interactions.

Minimal Risk AI Systems

Minimal, or lower risk, AI systems are those whose use is considered to entail negligible or no risk to fundamental rights, safety, or public interests. These applications are largely excluded from binding legal obligations under the AI Act. Examples include:

- Spam filters, which use AI to automatically detect and filter unsolicited communications;
- Recommendation algorithms in entertainment platforms, such as AI systems suggesting movies, music, or video games based on user preferences;
- AI-based functionalities in video games, including non-player character (NPC) behaviors or adaptive difficulty systems;
- Autocorrect or grammar suggestion tools integrated into word processors or messaging applications.

While these systems are not subject to specific mandatory requirements under the AI Act, the European Commission and other regulatory bodies encourage providers of minimal risk AI to voluntarily adhere to codes of conduct, industry best practices, and principles for trustworthy AI.

AI Risk Management Frameworks

The AI Act has served as an inspiration for regulatory frameworks on AI, laying the foundation for a technology and sector-neutral approach. It addresses one of the key issues raised by both the OECD[27] and BIS[4], the need for a shared cross-jurisdictional definition of AI systems and establishes a clear risk-based classification of AI systems into different risk categories. This growing regulatory focus on AI requires organizations to implement robust frameworks and practices, not only to ensure compliance but also to demonstrate AI governance maturity and maintain public trust. In this context, several supranational bodies have proposed and developed guidelines and frameworks to support organizations in establishing procedures and practices that span the entire organization and effectively manage AI-related risks. The successful deployment of AI systems requires not only technical integration but also substantial organizational adaptation. AI is not a standalone tool or isolated technology; it is a systemic capability that touches every facet of an organization's operations, decision-making,

and stakeholder engagement. As such, its governance cannot be relegated to individual departments or ad hoc teams. It must be institutionalized through deliberate structural realignment. Leading organizations are beginning to recognize that to realize the full benefits of AI, they must embed it deeply into their strategic architecture and align it with their governance, culture, and talent management systems. This transformation is unfolding through several critical changes:

- **Redesigning Workflows:** companies are revisiting existing processes to determine where AI can enhance speed, accuracy, or efficiency. This includes automating repetitive tasks, augmenting human judgment in complex decisions, and re-engineering customer service, supply chain, or analytics functions to be AI-enabled by default.
- **Appointing Executive Leadership:** AI governance is increasingly being elevated to the C-suite. Organizations are naming Chief AI Officers or designating senior executives with formal authority to oversee AI strategy, risk management, and ethical compliance. This leadership is critical for securing resources, aligning cross-functional teams, and ensuring that AI initiatives support the organization's mission and risk appetite.
- **Centralizing Governance Functions:** functions such as data governance, algorithmic accountability, and risk oversight are being consolidated into Centers of Excellence or transformation offices. These structures serve as custodians of best practices and ensure that AI initiatives across business units adhere to consistent standards.
- **Creating Cross-Functional AI Governance Bodies:** effective AI governance requires a blend of perspectives. Legal experts, data scientists, cybersecurity professionals, ethicists, and business leaders are increasingly collaborating through formal committees or steering groups to assess model performance, regulatory exposure, and ethical implications.
- **Developing Adoption Roadmaps:** rather than deploying AI in an uncoordinated fashion, organizations are crafting strategic roll-out plans that define where and how AI will be introduced. These roadmaps include phased adoption schedules, integration milestones, and mechanisms for evaluation and iteration.
- **Institutionalizing Role-Based Training Programs:** as AI reshapes the nature of work, employees at all levels must be equipped with the knowledge to understand and interact with AI systems responsibly. Training is being tailored by function, for example, developers on fairness auditing, marketing teams on content validation, and compliance officers on risk classification, ensuring that each stakeholder understands both the capabilities and limitations of the AI tools they use.

These structural changes do more than reduce risk; they signal a broader organizational evolution. They reflect a growing recognition that responsible innovation must be an organizational value, not just a technical feature. Companies that succeed in operationalizing these changes are building a foundation where AI is not only scalable but also trustworthy, explainable, and socially acceptable. While awareness of AI's opportunities and challenges is growing, the implementation of risk management best practices remains immature in many organizations. Despite the expansion of AI use cases, significant gaps persist in how

companies measure, monitor, and govern AI systems. Many companies still lack of:

1. Defined KPIs tailored to AI;
2. Formalized adoption and risk mitigation strategies;
3. Centralized governance structures for AI oversight;
4. Processes for reviewing AI-generated outputs for accuracy and ethical compliance.

To build maturity in AI governance, organizations must move from experimentation to institutionalization. This involves embedding a set of structured and repeatable best practices across the AI lifecycle. Key practices include:

- **Executive-Level Risk Ownership:** assigning responsibility for AI governance at the highest levels, such as the CEO, CRO, CIO or Board of Directors, ensures accountability and signals strategic importance. This helps align AI initiatives with the organization's broader risk framework.
- **Tracking Performance Through Measurable KPIs:** establishing indicators related to accuracy, fairness, interpretability, and real-world impact is essential for both transparency and optimization. These metrics should be updated regularly and tied to business objectives and ethical commitments.
- **Phased Deployment with Pilot Programs:** introducing AI incrementally allows organizations to test, refine, and scale technologies with greater control. Pilot programs help surface potential risks before full implementation and allow for stakeholder feedback to inform adjustments.
- **Comprehensive Risk Assessments:** evaluations must cover technical, legal, ethical, and societal dimensions. This includes impact assessments on privacy, discrimination, security vulnerabilities, explainability, and regulatory exposure, both before deployment and during continuous operation.
- **Real-Time Monitoring and Alert Systems:** AI models are dynamic; without ongoing surveillance, their behavior may drift or degrade. Monitoring mechanisms, such as automated alerts for anomalies or bias shifts, are essential to ensure consistent performance and safety.
- **Transparent Internal Communication:** employees must understand the purpose, limitations, and oversight protocols of AI systems. Internal transparency fosters alignment and ensures that AI is not perceived as opaque or arbitrary.
- **Ongoing Employee Education:** training programs should evolve alongside AI tools and regulations. This includes both technical training (e.g. model validation) and ethical training (e.g. human-in-the-loop decision-making).
- **Customer-Facing Trust Strategies:** public trust is essential. Organizations should implement mechanisms such as user disclosures, consent protocols, explainable interfaces, and opt-out functionalities to ensure end-users understand when and how AI is being applied.

By institutionalizing these practices, organizations ensure that AI systems remain aligned with organizational values and stakeholder expectations even as technologies and markets evolve.

OECD Framework for the Classification and Risk Management of AI Systems

The OECD Framework for the Classification of AI Systems[26], developed within the broader context of the OECD AI Principles [29], has become a foundational reference for organizations aiming to implement effective procedures and processes to address AI-related challenges. Notably, regulations such as the EU AI Act and AI RMFs like the U.S. National Institute of Standards and Technology (NIST) AI Risk Management Framework [23] have incorporated many of the practices and principles that align with the OECD's vision and objectives, offering a foundational instrument for both public and private sectors to evaluate and manage AI technologies with greater clarity and foresight. This framework serves a unique function: rather than prescribing technical or legal obligations, it provides a structured and policy-relevant lens for categorizing AI systems according to their core characteristics, contexts of application, and potential impacts. Rooted in the OECD AI Principles [29], which emphasize human-centered values, transparency, accountability, and robustness, this classification tool is designed to support the development of proportionate, evidence-based governance strategies. The OECD Framework provides a consistent method to understand and compare the risks, benefits, and operational realities of AI systems by offering a non-normative but comprehensive classification structure. It empowers policymakers and organizations to align AI deployments with public interest, mitigate potential harms, and guide innovation in a direction that supports democratic values, sustainable development, and economic resilience.

Core Content

The OECD Framework for the Classification of AI Systems is structured around five foundational dimensions (ref. Figure 21) [26]. These dimensions work synergistically to facilitate consistent classification and comparative analysis, serving as the backbone for public policy design, institutional accountability, and ethical oversight: [26]

- **People and Planet:** explores the interface between AI systems, human well-being, and environmental sustainability. It captures the roles of actors involved in the development, deployment, and use of AI technologies, such as providers, end-users, impacted communities, and vulnerable groups, and examines how the system aligns with democratic values, fundamental rights, and sustainable development goals. This dimension assesses the potential for system-wide harms, such as those arising from biased outcomes, power asymmetries, labor displacement, or environmental degradation. It considers whether the system is deployed in contexts of power imbalance, such as law enforcement or employment screening, where contestability and recourse mechanisms are essential. Additionally, it addresses the degree of user dependency and control, evaluating whether humans can opt out or meaningfully override automated decisions. It also touches on transparency of communication, including disclosures to users about the presence and functioning of AI systems, and whether redress mechanisms are available when harm occurs.
- **Economic Context:** analyzes the sectoral and institutional environment in which the AI system operates. It distinguishes between sectors with high criticality and public service functions, such as healthcare, education, finance, transportation, and justice, and




FIGURE 21: The Five Dimensions of OECD Framework [26]

- sectors with relatively lower systemic risk. The dimension also considers the economic role of the AI system: whether it supports productivity, cost reduction, strategic decision-making, or personalized services. Importantly, it accounts for the business model dependencies, such as monetization through user data or third-party AI licensing, and the potential for market concentration or vendor lock-in. The framework invites analysis of whether the system is central to mission-critical operations or infrastructure, and if its failure would result in disproportionate economic or societal disruption. Moreover, it incorporates the scalability and diffusion potential of AI systems, recognizing that models with high replicability across markets may pose amplified systemic risks or lead to widespread behavioral impacts.
- **Data and Input:** examines the lifecycle of the data used in AI systems, focusing not only on type and provenance, but also on the mechanisms of collection, annotation, transformation, and use. It distinguishes among first-party, third-party, synthetic, derived, and public data, emphasizing the importance of data integrity and representativeness in minimizing algorithmic bias. The framework places particular emphasis on whether sensitive data, such as biometric, financial, or health-related information, is involved, and whether appropriate safeguards are in place to ensure lawful and ethical processing. It also assesses the degree of automation in data collection, the use of sensor-driven input streams (such as from wearables or IoT devices), and the system's ability to generate or infer additional data. This dimension is key to evaluating compliance with privacy regulations, exposure to adversarial data attacks, and the traceability of input-output linkages. Furthermore, it supports evaluation of documentation practices, including dataset documentation standards (e.g. datasheets for datasets) and versioning policies.
 - **AI Model:** delves into the system's internal logic and technical properties. It characterizes the model type, statistical, symbolic, hybrid, and the learning paradigm used, such as supervised, unsupervised, reinforcement, or transfer learning. It also considers whether the system is trained once and then fixed, or whether it is dynamic and self-learning, which significantly affects its risk profile and auditability. The framework emphasizes the degree of explainability and interpretability, considering whether stakeholders can understand, challenge, or replicate the rationale behind model outputs. Models deployed in high-stakes decisions are expected to offer some level of intelligibility, either inherently or through post-hoc methods. Furthermore, the dimension evaluates model transparency, including the disclosure of architecture, parameters, training data, and de-

sign decisions, especially when the system is offered commercially or as open source. Attention is also paid to the model's robustness to perturbations, security vulnerabilities, and capacity to generalize beyond the training environment. Where relevant, the framework supports the assessment of model cards or system documentation, especially for foundation models that may be reused across multiple downstream applications.

- **Task and Output:** describes the purpose, behavior, and real-world implications of the AI system. It categorizes the types of tasks the system performs, such as generation, prediction, optimization, detection, personalization, or decision support, and distinguishes between systems designed for supportive use and those intended for fully autonomous execution. This dimension incorporates an assessment of the level of human oversight, including whether human intervention is active, passive, or entirely absent at the point of decision-making. It evaluates the consequences of erroneous outputs, particularly in critical domains such as medical diagnostics or autonomous driving, where output reliability has life-and-death implications. The framework also encourages consideration of feedback mechanisms, such as whether the system's outputs are monitored, logged, and corrected post-deployment. Another relevant factor is the contextual use of outputs, including whether the AI results feed into final decisions or are mediated by human judgment. Finally, this dimension accounts for whether the task supports core public interest functions, which may raise regulatory obligations or justify heightened scrutiny.

Each of these dimensions is not static; they are intentionally designed to be interoperable and adaptable across sectors, jurisdictions, and levels of AI maturity. Their application enables a granular and multidimensional characterization of AI systems, fostering clarity in governance, consistency in comparative analysis, and foresight in risk mitigation.

NIST Artificial Intelligence Risk Management Framework

In response to the challenges posed by AI usage proliferation, the U.S. National Institute of Standards and Technology (NIST) developed the AI RMF [23] to assist organizations in designing, developing, deploying, and using AI systems in ways that are trustworthy and responsible. Released in 2023, the NIST AI RMF is a voluntary, sector-agnostic, and use-case-neutral framework, structured to promote flexibility and broad applicability across diverse organizational contexts. This framework serves to help organizations "to better manage risks across the AI lifecycle, aiming to:

- The development of innovative approaches to address characteristics of trustworthiness, including ac-

- curacy, explainability and interpretability, reliability, privacy, robustness, safety, security, and mitigation of unintended and/or harmful bias, as well as of harmful uses;
- Consider and encompass principles such as transparency, fairness, and accountability during design, deployment, use, and evaluation of AI technologies and systems;
- Consider risks from unintentional, unanticipated, or harmful outcomes that arise from intended uses, secondary uses, and misuses of the AI."[6]

The NIST Framework is complemented by the NIST AI RMF Playbook [22], a practical guide that provides operational recommendations for implementing the framework within organizations. This document helps stakeholders understand how to translate the core principles of the AI RMF into actionable steps within different organizational contexts. The Playbook includes suggested actions, references, and related guidance to support the achievement of desired outcomes across the AI lifecycle. In 2024, NIST also released the Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile[24], a guideline specifically designed to address the unique risks associated with generative AI. This profile adapts the NIST AI RMF to the specific challenges posed by GenAI systems, offering tailored risk assessment and mitigation strategies for developers, deployers, and users of large-scale generative models.

Core Content

The framework is built around four interdependent core functions (ref. Figure 22):

- **Govern:** the Govern function cultivates and implements a culture of risk management within organizations involved in designing, developing, deploying, evaluating, or acquiring AI systems. It establishes clear processes, documentation, organizational frameworks, and procedures to achieve desired risk management outcomes while incorporating assessments of potential impacts. This function provides a structure that ensures all activities align with the organization's principles, policies, and strategic priorities. By connecting the technical aspects of AI with organizational values and principles, it supports individuals responsible for acquiring, training, deploying, and monitoring AI systems. Additionally, it addresses the full product lifecycle and associated processes, including the management of any issues that may arise. The playbook[22] distinguishes six main domains related to the Govern function, which everyone comprehends as a subset of practices that should be put in place:
 - **Govern 1:** policies, processes, procedures, and practices across the organization related to the mapping, measuring, and managing of AI risks are in place, transparent, and implemented effectively;
 - **Govern 2:** accountability structures are in place so that the appropriate teams and individuals are empowered, responsible, and trained for mapping, measuring, and managing AI risks;
 - **Govern 3:** workforce diversity, equity, inclusion, and accessibility processes are prioritized in the mapping, measuring, and managing of AI risks throughout the lifecycle;

- **Govern 4:** organizational teams should be committed to a culture that considers and communicates AI risk;
- **Govern 5:** processes are in place for robust engagement with relevant AI actors;
- **Govern 6:** policies and procedures are in place to address AI risks and benefits arising from third-party software and data, and other supply chain issues.
- **Map:** The Map function defines the context required to identify and assess risks associated with an AI system across its lifecycle. AI development involves multiple interdependent activities, often managed by different actors who may lack full visibility or control over the entire process. This fragmentation can result in unforeseen impacts, as early design decisions may influence how the system behaves and how it interacts with its deployment environment. Such complexity introduces uncertainty into risk management, which the Map function seeks to reduce by gathering contextual knowledge and identifying potential sources of negative risk. This information supports informed decision-making and lays the foundation for the Measure and Manage functions. The Map function also encourages inclusion of diverse internal perspectives and, where relevant, engagement with external stakeholders such as users, affected communities, or collaborators. The engagement of different levels of stakeholders could help organizations to better understand the context of use and recognize both potential benefits and foreseeable negative impacts. In the end, the Map function will provide sufficient insight to determine whether the development or deployment of an AI system is appropriate. If the decision is to proceed, organizations should continue to apply the Map function throughout the system's lifecycle as risks, capabilities, and contexts evolve. The playbook[22] distinguishes five main domains related to the Map function:
 - **Map 1:** context is established and understood;
 - **Map 2:** categorization of the AI system is performed;
 - **Map 3:** AI capabilities, targeted usage, goals, and expected benefits and costs compared with appropriate benchmarks are understood;
 - **Map 4:** risks and benefits are mapped for all components of the AI system, including third-party software and data;
 - **Map 5:** impacts to individuals, groups, communities, organizations, and society are characterized.
- **Measure:** The Measure function applies quantitative, qualitative, or mixed-method approaches to analyze, assess, benchmark, and monitor AI risks and their impacts. It builds on the contextual understanding gained through the Map function and provides critical input to the Manage function. AI systems should be continuously tested and evaluated not only for performance but also for their social impact, human-AI interactions, and alignment with trustworthy characteristics. Where trade-offs among these characteristics occur, measurement serves as a traceable foundation to support informed management decisions. The playbook[22] distinguishes four main domains related to the Measure function:
 - **Measure 1:** appropriate methods and metrics are identified and applied;

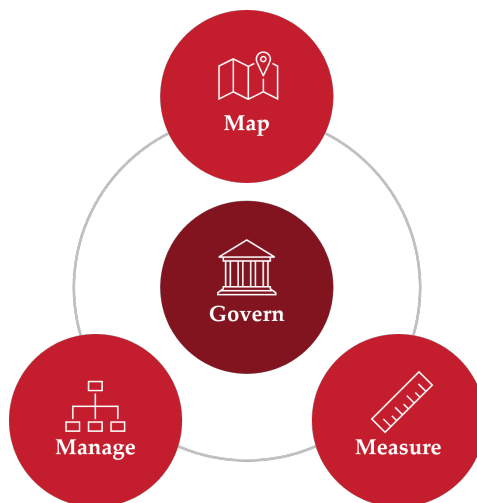


FIGURE 22: Four Interdependent Core Functions [23]

- **Measure 2:** AI systems are evaluated for trustworthy characteristics;
- **Measure 3:** mechanisms for tracking identified AI risks over time are in place;
- **Measure 4:** feedback about the efficacy of measurement is gathered and assessed.
- **Manage:** The Manage function enables organizations to prioritize, mitigate, accept, or avoid AI risks. It supports feedback loops, incident response planning, and risk communication strategies, allowing continuous improvement of AI systems. The playbook[22] distinguishes four main domains related to the Manage function:
 - **Manage 1:** AI risks based on assessments and other analytical output from the Map and Measure functions are prioritized, responded to, and managed;
 - **Manage 2:** strategies to maximize AI benefits and minimize negative impacts are planned, prepared, implemented, documented, and informed by input from relevant AI actors;
 - **Manage 3:** AI risks and benefits from third-party entities are managed;
 - **Manage 4:** risk treatments, including response and recovery, and communication plans for the identified and measured AI risks are documented and monitored regularly.

The AI RMF is designed to be both scalable and iterative, offering practical guidance for organizations ranging from early-stage AI developers to large, mature enterprises. One of its primary utilities lies in enhancing AI trustworthiness by embedding principles such as reliability, fairness, transparency, and privacy into the design and deployment processes. These safeguards not only improve the integrity of AI systems but also help reduce reputational and operational risks. Another key strength of the framework is its support for regulatory readiness; it aligns closely with global standards and emerging legal frameworks, including the OECD classification and the EU AI Act[19], making it a valuable foundation for organizations seeking to meet compliance requirements across jurisdictions. Furthermore, the AI RMF promotes meaningful cross-functional collaboration by encouraging active participation from legal, technical, governance, and policy stakeholders. Organizations that

adopt the AI RMF typically experience several significant improvements across their governance and operational practices. One of the key benefits is the enhanced clarity in defining internal roles and responsibilities, ensuring that each team involved in the AI lifecycle, whether in development, oversight, or deployment, understands its specific accountability in managing AI-related risks. The framework also promotes more robust documentation and traceability of risks, enabling organizations to track, audit, and respond to potential issues systematically over time. In addition, it strengthens processes for evaluating AI safety and preparing for incident response, helping institutions anticipate, detect, and mitigate adverse events more effectively. Perhaps most importantly, the adoption of the NIST AI RMF encourages a cultural shift toward ethical deployment, fostering greater alignment between the outcomes of AI systems and broader human values, societal expectations, and fundamental rights.

The Japanese AI Guidelines for Business

The AI Guidelines for Business[14], developed within the Japanese policy ecosystem and aligned with the national vision of Society 5.0[3], represent a forward-looking effort to foster responsible AI adoption across industrial sectors. The guidelines serve as a comprehensive, non-binding framework that encourages companies to voluntarily adopt risk-based governance principles, spanning the full life cycle of AI systems. Japan has adopted a goal-based, soft-law approach, whereby ethical, technical, and organizational recommendations support practical governance without imposing rigid constraints. The guidelines are informed by both international discussions (e.g. OECD and G7 Hiroshima Process) and domestic principles developed in earlier documents, such as the Social Principles for Human-Centric AI and the Governance Guidelines for Implementation of AI Principles ver. 1.1 [1].

Core Content

The AI Guidelines for Business are organized into five principal sections [14], forming a governance framework for entities involved in the development, provision, and use of artificial intelligence systems. The normative foundation of the document outlines a shared societal vision for AI aligned with Japan's strategic concept of Society 5.0,




FIGURE 23: Basic Philosophies [14]

anchored in three foundational philosophies: Human Dignity, Diversity and Inclusion, and Sustainability (ref. Figure 23). These principles emphasize the role of AI to support social advancement, individual autonomy, and inclusive development, while ensuring that its integration into society contributes to long-term well-being and equitable access to its benefits.

These principles encompass human-centricity, privacy protection, safety, fairness, transparency, accountability, and education. To facilitate their implementation, the guidelines introduce a range of concrete measures intended to mitigate societal risks. These include the prevention of manipulative system behaviors, attention to the informational impacts of filter bubbles and disinformation, the promotion of explainability through mechanisms that trace decision-making processes, and the adoption of documentation practices that support both auditability and external validation. Beyond these general principles, the framework addresses the governance of advanced AI systems, including generative and autonomous technologies. In alignment with the Hiroshima Process International Code of Conduct for Organizations Developing Advanced AI Systems[11], the guidelines recommend a lifecycle-based approach to AI oversight. This encompasses pre-deployment testing methods such as red-teaming, incident documentation protocols, post-deployment monitoring, and public disclosure of system capabilities, limitations, and intended uses. The guidelines also advocate for secure development practices, multistakeholder collaboration, and the use of content authentication technologies, such as watermarking, to detect and prevent malicious use of AI-generated outputs.

The core content of the framework is structured into 5 parts [14]:

- **Part 1 - Definitions:** introduces key definitions and conceptual distinctions, establishing a shared vocabulary that clarifies the scope and application of the framework. This groundwork enables a structured understanding of the differentiated responsibilities that follow in Parts 3 to 5, which delineate the roles of developers, providers, and users within the AI value chain.
- **Part 2 - Society to aim for with AI and matters each AI business actor works on:** it sets the normative foundation of the guidelines by articulating the societal vision for AI based on three core philosophies and by establishing a set of common guiding principles accompanied by operational recommendations to promote responsible AI development and use.
- **Part 3 - Matters related to AI Developers:** outlines the requirements for AI developers, defined as entities responsible for creating AI systems, including the design of algorithms, models, and training pipelines. Given their upstream position in the AI lifecycle, developers carry significant responsibility in shaping the behavior, reliability, and risks of AI systems. They are expected to ensure the quality, appropriateness, and legal compliance of training data; to implement bias detection and mitigation techniques during model development; and to apply privacy- and security-by-design principles from the earliest phases of system construction. Developers must document their design decisions, training procedures, data handling protocols, and evaluation methodologies to enable transparency and future auditability. Furthermore, they are encouraged to assess potential downstream impacts of the technologies they produce, especially in high-risk applications, and to engage in continuous research and collaboration to align their practices with emerging technical standards and evolving societal needs.
- **Part 4 - Matters related to AI Providers:** outlines the responsibilities of AI providers, who act as intermediaries between developers and end users by embedding AI models into applications, systems, or services and distributing them for practical use. Providers are tasked with ensuring that AI systems are properly configured, integrated, and validated for the intended use cases. They must test and verify the performance, accuracy, robustness, and resilience of AI systems under real-world conditions, and define clear operational parameters, including intended purpose, constraints, and potential risks. Providers are also required to develop comprehensive user documentation, including guidelines for appropriate use, explanation of system functionality, and information about known limitations or possible failure modes. In cases where retraining or updates are necessary, providers should establish maintenance protocols and communicate with developers and users to coordinate improvements. Their role extends to implementing incident-handling systems, facilitating post-deployment monitoring, and ensuring that end users receive adequate support and training. In addition, providers are expected to uphold transparency obligations by communicating essential information in a manner accessible and

appropriate to the technical capacities of the users.

- **Part 5 - Matters related to AI Business Users:** addresses AI business users, defined as organizations or entities that apply AI systems within their internal operations or customer-facing processes. As the closest actors to the end effects of AI deployment, business users are responsible for ensuring that AI tools are used in accordance with the provider's specifications and within the bounds of ethical, legal, and sectoral norms. Users must continuously monitor the behavior and outputs of deployed AI systems, identify irregularities or deviations, and report them through established feedback channels. They are also required to evaluate the potential impact of AI use on individuals, institutions, or society, particularly in contexts involving decisions about employment, credit, healthcare, law enforcement, or public administration. Business users must implement appropriate safeguards to prevent unintended harm, including mechanisms for human oversight, redress, and the protection of fundamental rights. Internally, they must ensure that relevant staff are properly trained in the use of AI systems and that operational procedures align with the governance principles established in the earlier parts of the guidelines. Where AI deployment intersects with public-facing services, users are also encouraged to engage with affected stakeholders, maintain transparency, and uphold accountability regarding how AI is used and governed within the organization.

Taken together, Parts 3 to 5 outline a role-specific governance architecture that distributes accountability and responsibility across the full AI lifecycle. A distinguishing feature of the guidelines is the promotion of an agile governance model, which encourages organizations to move beyond static, rule-based compliance toward dynamic, iterative oversight. This approach is grounded in continuous risk assessment, regular updates of governance protocols, and responsiveness to changes in technology, regulation, and stakeholder expectations. The guidelines emphasize that effective governance cannot rely on uniform rules alone but must be tailored to the roles and influence of each actor across the AI value chain. They advocate for embedding governance within broader business strategies and institutional cultures, promoting coordination, risk sensitivity, and proactive engagement with evolving AI-related challenges. This integrated approach aims to ensure long-term resilience, accountability, and alignment with the public interest in a rapidly evolving technological landscape.

Conclusions

The analysis has shown how the increasing and widespread adoption of AI across industries has accelerated the need to establish both regulatory frameworks and robust governance mechanisms capable of addressing emerging risks and promoting the responsible use of the technology. In recent years, several frameworks have been developed to manage and mitigate AI-specific risks. Despite differences in organizational structures and implementation approaches, these models converge around a common set of core principles, the same principles, inspired by the OECD and embedded in the EU AI Act, and can be summarized as follows:

- **Centrality of Ethics and Human Rights:** guidelines and regulatory frameworks are driven by the need to protect human rights, requiring high standards

of transparency, accountability, non-discrimination, privacy, robustness, and security.

- **AI Lifecycle as the Foundation of Governance:** effective AI risk management demands a holistic approach that spans the entire lifecycle of AI systems, from development to decommissioning. Frameworks emphasize the need for continuous risk monitoring and mitigation across all the life-cycle phases.
- **Human Oversight as a Safeguard:** the necessity of maintaining human-in-the-loop oversight over AI systems, especially in high-risk or decision-critical applications, is a crucial pillar across all frameworks. This requires the development of mechanisms that enable humans to oversee on AI systems, ensuring accountability, safety, and control throughout the system's lifecycle.
- **Risk Tiering and Proportionality:** a key principle is the assessment and classification of AI systems based on their risk level, ensuring that the intensity of controls is proportionate to the potential impact of the system.
- **Organizational Approaches:** the importance of embedding AI governance across organizational structures is underlined by all the frameworks. This requires the adoption of AI risk management not only as procedures and processes but as a corporate mindset involving different functions and teams to ensure effective oversight and alignment with enterprise goals.

To conclude, too often, risk management is perceived as a compliance exercise, a necessary but limiting set of controls. In the context of AI, this mindset is not only outdated but also strategically shortsighted. When designed and executed effectively, AI governance Frameworks become a powerful driver of competitive advantage. Robust governance empowers organizations to:

- **Accelerate Safe Innovation:** by establishing clear boundaries, escalation paths, and validation protocols, governance reduces uncertainty and enables faster experimentation and deployment.
- **Build Stakeholder Trust:** consumers, investors, regulators, and the public are increasingly concerned about the ethical use of AI. Demonstrating a credible governance framework enhances reputation, supports brand integrity, and attracts ethically conscious partners and clients.
- **Enable Internal Coherence:** a standardized approach to AI governance facilitates cross-departmental collaboration, minimizes duplication of efforts, and ensures consistency in how decisions are made and risks are managed.
- **Enhance Regulatory Readiness:** as AI regulation evolves across jurisdictions, proactive governance allows organizations to anticipate requirements, reduce compliance burdens, and respond swiftly to legal changes.
- **Foster Long-term Adaptability:** with technology evolving rapidly, static or informal practices are insufficient. A governance model that is scalable, flexible, and principle-based equips organizations to manage future use cases, risks, and opportunities more effectively.

References


- [1] **AI Strategy Council, Japan.** *Governance Guidelines for the Implementation of AI Principles ver. 1.1.* January 2022.
- [2] **Bick, A., Blandin, A. and Deming, D.** *The Rapid Adoption of Generative AI.* NBER Working Paper Series N. 32966, February 2025.
- [3] **Cabinet Office, Government of Japan.** *Social Principles of Human-Centric AI.* Official Publication (Japan), March 2019.
- [4] **Cristanto, J. C., Leuterio, C. B., Prenio, J. and Yong, J.** *Regulating AI in the Financial Sector: Recent Developments and Main Challenges.* BIS FIS Insight n. 63, December 2024.
- [5] **Esposito, D., Carrozino, P., Ghilardi, B. and Cecchin, M.** *Artificial Intelligence Act (AI Act).* Argo N. 26, August 2024.
- [6] **Esposito, D., Ghilardi, B. and Cecchin, M.** *AI Risk Management Framework: National Institute of Standards and Technology (NIST).* Just in Time, February 2023.
- [7] **European Commission.** *Commission Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689 (AI Act).* February 2025.
- [8] **European Commission.** *Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689 (AI Act).* February 2025.
- [9] **European Commission.** *Regulation (EU) 2024/1689 on Artificial Intelligence (AI Act).* Official Publication of the European Union, July 2024.
- [10] **European Commission.** *Targeted stakeholder consultation on classification of AI systems as high-risk.* June 2025.
- [11] **G7 Leaders.** *Hiroshima Process - International Guiding Principles for Organizations Developing Advanced AI Systems.* G7, October 2023.
- [12] **IOSCO.** *Artificial Intelligence in Capital Markets: Use Cases, Risks, and Challenges.* Report of the Board of IOSCO, March 2025.
- [13] **Japan FSA.** *Preliminary Discussion Points for Promoting the Sound Utilization of AI in the Financial Sector.* AI Discussion paper, March 2025.
- [14] **Japan Ministry of Internal Affairs and Communications, Ministry of Economy, Trade and Industry.** *AI Guidelines for Business Ver1.0.* Official Publication (Japan), April 2024.
- [15] **McCarthy, J.** *What is Artificial Intelligence?* Stanford Press, November 2007.
- [16] **Maslej et al.** *The AI Index 2023 Annual Report.* AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, April 2023.
- [17] **Maslej, N. et al.** *The AI Index 2024 Annual Report.* AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, April 2024.
- [18] **Maslej, N. et al.** *The AI Index 2025 Annual Report.* AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, April 2025.
- [19] **Mazzoni, N.** *AI Act: AI System Definition and Prohibited AI Practices.* Just in Time, April 2025.

- [20] **Mazzoni, N., Figuriello, J. and Campaniolo, G.** *Artificial Intelligence Financial Industry Market Overview*. Just in Time, July 2025.
- [21] **Mazzoni, N., Ranieri, M., Gentilavigna, F. and Bandini, L.** *AI: Regulation, Rise and Challenge*. Just in Time, April 2025.
- [22] **National Institute of Standards and Technology.** *AI RMF Playbook*. January 2023.
- [23] **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. January 2023.
- [24] **National Institute of Standards and Technology.** *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. July 2024.
- [25] **NVIDIA.** *State of AI in Financial Services: 2025 Trends*. NVIDIA Survey Report, February 2025.
- [26] **OECD.** *OECD Framework for the Classification of AI Systems*. OECD Digital Economy Papers No. 323, February 2022.
- [27] **OECD.** *Regulatory Approaches to Artificial Intelligence in Finance*. OECD Artificial Intelligence Papers, September 2024.
- [28] **Eurostat.** *Website*.
- [29] **OECD AI Principles.** *Website*.

Sitography

In the previous issue

Issue N. 28 - 2025



Issue N. 28 - 2025

INNOVATION

Atomic Settlement: Potential Implications of DLT-based Compressed Settlement Cycles

CBDC For Financial Inclusion: Evidences from Sand Dollar and e-Naira

ENVIRONMENTAL, SOCIAL AND GOVERNANCE RISKS

EBA Final Guidelines on the Management of ESG Risks

Last issues are available at www.iasonltd.com/research/argo/