Issue N. 30 - 2026

ARGO

New Frontiers in Practical Risk Management

iason
ESSENTIAL SERVICES FOR
FINANCIAL INSTITUTIONS

iason

Argo magazine

Year 2026 - Issue Number 30

Published in January 2026
First published in October 2013

Last published issues are available online:
http://www.iasonltd.com

Front Cover: **Umberto Boccioni**, *Dinamismo di un corpo umano*, 1913.

Milan Headquarter:
Corso Europa, 15
20122 Milan
Italy

London Headquarter:
58-60, Kensington Church Street
W8 4DB London
United Kingdom

Madrid Headquarter:
Calle Miguel Ángel, 16
28010 Madrid
Spain

Contact Information:
info@iasonltd.eu
www.iasonltd.com

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

Credit Risk

Dear Readers,

Welcome to the latest edition of Argo Magazine, where we explore the innovative approaches and emerging trends that are transforming the financial landscape.

This issue brings together a selection of articles that spanning key topics in credit risk management and financial innovation. They offer theoretical insights and practical frameworks for implementation.

*Opening with the Innovation section, the first article, "**iason-Nuant Optimal Cross-Pool Liquidity Provision Strategy**" by Jacopo Mancin and Antonio Scarinci, tackles the challenge of optimal automated market making in decentralized crypto-asset liquidity pools. Using advanced mathematical modelling and portfolio allocation techniques, the authors propose a multi-pool strategy that balances diversification with independent execution, mitigating impermanent losses while optimizing rebalancing costs.*

*The section concludes with "**Market Scenario Generation with GenAI**" by Michele Bonollo, Antonio Menegon, Giuseppe Crupi and Caterina Papetti. This forward-looking paper explores how generative artificial intelligence can revolutionize scenario generation for stress testing and risk management. By leveraging advanced models such as GANs, VAEs, and Transformer-based architectures, the authors outline a roadmap for deploying next-generation tools that enhance realism and interpretability in market simulations.*

*Closing this edition, the Credit Risk section features the article **"A Common Collateral Pooling Arrangement for Corporate Lending"** by Antonio Castagna, proposes a mutual collateral scheme designed to improve the credit quality of corporate loan portfolios. By analysing the extent of the credit protection provided and outlining compensation mechanisms based on actuarial and financial fairness principles, this work offers a compelling perspective on collaborative risk mitigation strategies.*

We hope that this edition will offer valuable insights into the technological and credit risk dimensions of modern finance, inspiring new ideas and encouraging informed discussions. As always, we invite you to engage with these topics and share your perspectives with us.

Enjoy your read!

Antonio Castagna
Luca Olivo
Giulia Perfetti

Just in Time

iason Notes

USA vs EU: Use and Oversight of AI in Financial Services

AI is transforming the financial sector in the US and the EU by increasing efficiency, reducing costs, and improving credit decisions, fraud detection, and customer services. The use of alternative data is expanding financial inclusion.

Regulators are adopting different approaches: a flexible, reactive model in the US; and a preventive, rigorous framework in the EU through the AI Act and harmonized standards.

read more

Date January 2026

AI Agents: An Introduction to Agentic Systems, Market Impact, and Future Risks

The rise of Agentic AI marks a significant step forward moment in the current AI revolution, marked by accelerating technological breakthroughs and massive capital deployment. Unlike previous AI developments, Agentic AI holds the promise of significant practical utility and adaptability, already demonstrating early capabilities as a powerful automation tool, which presents unprecedented opportunities alongside new complexities.

read more

Date November 2025

ICT Risk: Focus on Thread-Led Penetration Testing (TLPT)

Threat-Led Penetration Testing (TLPT) represents a cornerstone of the European Union's strategy to enhance the cyber resilience of financial institutions. Mandated by the Digital Operational Resilience Act (DORA) and operationalized through the TIBER - EU framework, TLPT simulates sophisticated cyberattacks based on real -world threat intelligence to assess an entity's ability to detect, respond to, and recover from advanced threats.

read more

Date August 2025

Artificial Intelligence: Financial Industry Market Overview

Starting 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

ECB: SREP 2025 and SSM Priorities 2026-2028

The 2025 Supervisory Review and Evaluation Process (SREP) shows that the banks supervised by the ECB have continued to exhibit strong capital positions, solid leverage ratios, and comfortable liquidity buffers. Profitability remains robust, with return on equity reaching 10.1%, up from 9.5% in Q4 2024. Asset quality remains sound overall, though non-performing loan (NPL) ratios are higher in the non-residential real estate and SME segments. Key areas of supervisory concern include elevated leveraged finance and insufficient provisioning for non-performing exposures.

read more

Date January 2026

Credit Risk Meets Large Language Models

The study addresses the issue of information asymmetry in peer-to-peer lending, where lenders often lack sufficient data to assess borrower credit. By applying BERT (Bidirectional Encoder Representations from Transformers), a Large Language Model, the authors generate a credit risk score based on the borrowers' description of the credit request. Once integrated into a XGBoost model with traditional inputs, this score improves predictive accuracy and AUC (Area Under Curve), demonstrating the value of combining textual insights with standard credit data.

read more

Date September 2025

Counterparty Credit Risk Exploratory Scenario Exercise

The ECB conducted an exploratory scenario exercise to assess counterparty credit risk (CCR) among 15 Euro-area banks selected based on quantitative criteria derived from 2023 EBA stress test. Participants were mainly banks with high CCR exposure towards non-bank financial institutions (NBFIs).

The exercise focused on vulnerabilities linked to exposures to NBFIs and evaluated banks' stress-testing capabilities under multiple adverse scenarios.

read more

Date September 2025

Guidelines for Integrating ESG Risks into Stress Testing

The Joint Guidelines developed by the European Supervisory Authorities (ESAs) aim to ensure the consistent integration of environmental, social, and governance (ESG) risks into supervisory stress testing by competent authorities.
These guidelines provide a structured approach for incorporating ESG risks either within existing frameworks or through complementary assessments. They emphasize the importance of clear methodological principles, adequate resource allocation, robust data infrastructure, and effective governance.

read more

Date August 2025

EBA Guidelines on Environmental Scenario Analysis: Usage and Limitations

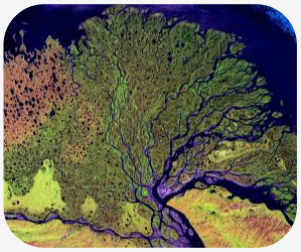

The EBA Guidelines on environmental scenario analysis (EBA/GL/2025/04) set out supervisory expectations for financial institutions to systematically conduct environmental scenario analysis, strengthening forward-looking risk assessment and management of environmental, especially climate-related, risks. They require institutions to integrate such scenario analyses into existing stress-testing frameworks to assess short-term impacts on capital and liquidity and to evaluate medium- and long-term business model resilience under plausible future conditions.

read more

Date January 2026

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!

iason ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

iason-Nuant Optimal Cross-Pool Liquidity Provision Strategy

About the Authors

Jacopo Mancin:
Senior Manager
He holds a degree in mathematics, a master degree in mathematical finance at the Paris-Est University and a PHD in financial mathematics at the LMU Munich. He worked as a front office quant in major international banks. He has also collaborated with crypto startups to develop risk management tools for digital assets and DeFi strategies. He gained practical experience in ML and AI working on topics such as fraud detection, time series forecasting and AI agents architectures.

Antonio Scarinci:
Financial Engineer
He holds a Bsc in Economics and Finance from Bocconi University and a Msc in Mathematical Engineering - Major in Quantitative Finance, earned after taking integrative courses in mathematics and engineering from the Bsc, from Politecnico di Milano. He worked as a bond arbitrage trader in an HFT firm and equity proprietary trader in the Equity Portfolio Management desk of an important Italian banking institution. He is currently a iason Financial Engineering consultant for the Financial Engineering - Rates, Credit Inflation Desk of one of the largest Italian banks.

In this work, we tackle the problem of optimal Automated Market Making in Decentralized Crypto-Asset Liquidity Pools. We reviewed the approaches defined in the literature for defining an optimal strategy to be executed on a single liquidity pool, taking into account both rebalancing costs and impermanent losses. Hence, supposing the spot exchange rates to be driven by an elliptically distributed noise and exploiting their spectral properties, we derive an optimal multi-pool portfolio allocation method able to achieve two goals at the same time: diversify effectively the exposure of the agent's wealth and allow for the execution of the aforementioned strategy independently in each liquidity pool.

Decentralized Finance (DeFi) has emerged as one of the most significant technological innovations of recent years. In this ecosystem, a wide range of crypto-assets has been introduced, together with cross-cryptocurrency exchange protocols that algorithmically regulate the interaction between supply and demand across different digital assets. Among these, Decentralized Exchange (DEX) protocols – such as Uniswap – rely on the presence of Liquidity Providers (LPs), who function as market makers. LPs supply liquidity by depositing pairs of cryptocurrencies into decentralized trading venues known as liquidity pools, adhering to a predefined set of protocol-specific rules.

Given the availability of multiple exchange rates and liquidity pools, an LP is faced with the problem of determining how to optimally allocate their capital across different pools. Thus, the liquidity provision decision can be framed as an optimal wealth allocation problem.

In this paper, we address this problem from two complementary perspectives. First, we analyze an optimal liquidity provision strategy tailored to cryptocurrency pairs traded under the Uniswap V3 protocol. Second, we consider liquidity provision strategies as investable assets and formulate a portfolio optimization problem at the LP level.

The structure of the paper is as follows. In Section "Optimal Liquidity Provision: Case of Single Pool", we review the optimal liquidity provision framework proposed in [3], which we adopt as a baseline model. In Section "Optimal Liquidity Provision: Case of Multiple Pools", we extend this framework to the setting in which an LP distributes capital across multiple liquidity pools. Section "Agnostic Risk Parity Portfolio" introduces the concept of Agnostic Risk Parity portfolios, which provide more diversified allocations and reduce sensitivity to estimation errors in the covariance matrix of expected returns. Finally, in Section "The Strategy", we combine these elements to construct an optimal multi-pool liquidity provision strategy. A numerical implementation will be explored in future work.

Constant Function Markets and Concentrated Liquidity

Constant Function Markets (CFM)

Consider two cryptocurrencies X and Y. A liquidity pool is a decentralized trading venue in which LPs lock some amounts of asset X and asset Y to make them available for trading activity to any willing agent, called liquidity taker (from now on LT), at an exchange rate Z. The economic rationale for an LP to provide its assets is the possibility to earn fees as compensation for the liquidity provision activity.

A Constant Function Market (CFM) is a liquidity pool in which:

$$f(q_X, q_Y) = \kappa^2.$$

For some function f monotonically increasing both in q_X and in q_Y. This entails that the total amount q_X and q_Y of asset X and asset Y respectively locked into the pool is

given by:

$$q_X = \varphi_\kappa(q_Y),$$

where φ_κ is a decreasing *level function* of the market function f at the *liquidity level* κ. Hence, whenever a LT wants to swap y units of asset Y, the amount x of asset X that he can obtain must verify the following relation:

$$f(q_X + (1-\tau)x, q_Y - y) = \kappa^2,$$

where τ represents the fee rate to be paid to LPs as compensation for their liquidity provision activity. As a result, the *execution exchange rate*, i.e. the exchange rate at which any liquidity taking trade is performed, is given by:

$$\tilde{Z}(y) = \frac{\varphi_\kappa(q_Y - y) - \varphi_\kappa(q_Y)}{(1-\tau)y}.$$

Similarly, whenever a LT wants to swap x units of asset X the amount y of asset Y that he can obtain satisfies:

$$f(q_X - x, q_Y + (1-\tau)y) = \kappa^2,$$

and the consequent execution exchange rate is given by:

$$\tilde{Z}(-y) = \frac{\varphi_\kappa(q_Y) - \varphi_\kappa(q_Y + (1-\tau)y)}{y}.$$

Hence, the spot exchange rate quoted on the liquidity pool is given by:

$$Z = \lim_{y \to 0} Z(y) = -\frac{1}{1-\tau}\varphi'_\kappa(q_Y).$$

On the other hand, liquidity provision should not impact the spot exchange rate. Hence, when liquidity provision is performed by LPs:

$$Z = \frac{q_X}{q_Y} = \frac{q_X + x}{q_Y + y}.$$

In the case of the UniswapV2 protocol, one of the most used Constant Function Market:

$$f(q_X, q_Y) \coloneqq q_X q_Y = \kappa^2, \qquad \varphi_\kappa(q_Y) = \frac{\kappa^2}{q_Y},$$
$$Z = \frac{q_X}{q_Y}.$$

CFMs with Concentrated Liquidity (CFMCL)

In CFMs, LPs are forced to offer liquidity at any possible exchange rate, determined by the dynamics of LT transactions. Constant Function Markets with Concentrated Liquidity, on the other hand, adds the possibility for LPs to select specific price boundaries to provide their liquidity within, whose bounds should be chosen among a set of discrete *price ticks* $\{Z^1, Z^2, ..., Z^n\}$. The smallest range $(Z^i, Z^{i+1}]$ is called *tick range*. Let the liquidity provider choose a price interval $(Z_l, Z_u]$. For any chosen range, the LP's asset amounts x and y are specified by key formulae, which determine what assets the LP holds depending on where the pool price currently sits relative to the chosen boundaries:

$$\begin{cases} x = 0 & \text{if } Z \le Z^\ell \\ x = \tilde{\kappa}\left(Z^{1/2} - (Z^\ell)^{1/2}\right) & \text{if } Z^\ell < Z \le Z^u \\ x = \tilde{\kappa}\left((Z^u)^{1/2} - (Z^\ell)^{1/2}\right) & \text{if } Z > Z^u \end{cases},$$

$$\begin{cases} y = \tilde{\kappa}\left((Z^\ell)^{-1/2} - (Z^u)^{-1/2}\right) & \text{if } Z \le Z^\ell \\ y = \tilde{\kappa}\left(Z^{-1/2} - (Z^u)^{-1/2}\right) & \text{if } Z^\ell < Z \le Z^u \\ y = 0 & \text{if } Z > Z^u \end{cases}.$$

This implies that:

- If $Z \le Z_l$ the LP only holds Y;
- If $Z > Z_u$ the LP only holds X.

These formulae dictate how the LP's holdings evolve within the range $(Z^l, Z^u]$[1] and what they receive upon exiting the pool. Liquidity depth κ within a tick range $(Z^l, Z^u]$ is constant unless more liquidity is deposited or withdrawn. When the marginal price crosses any of the tick

[1] A LP can quote several ranges at the same time.

boundaries, the pool must execute distinct trades using the appropriate depth for each tick. Aggregate liquidity depth in a given tick is determined by summing the depths of all LP positions in that range.

$$\kappa = \sum_i \tilde{\kappa}_i.$$

Let p the total amount of fees of the pool. The amount of fees earned by an LP is proportional to the ratio between the LP's share of liquidity $\tilde{\kappa}$ and the total liquidity of the pool κ:

$$\tilde{p} = \frac{\tilde{\kappa}}{\kappa}\, p \cdot \mathbb{1}_{Z_\ell < Z \leq Z_u}.$$

This implies that LPs who concentrate more liquidity in narrower ticks capture a greater share of fees in the range selected $(Z^l, Z^u]$, but this increases the concentration risk, i.e. the chance that the price leaves the tick range letting the LP liquidity switch from being *active* (i.e. available for LT transactions) to being *inactive* (i.e. unavailable for LT transactions).
UniswapV3 is an example of CFM with CL and will be the protocol in which the analysis of the liquidity provision strategies introduced in the next sections will be carried out.

Optimal Liquidity Provision: Case of Single Pool

Basic Assumptions

Two assets X and Y are assumed to be exchanged in a CFM with CL. Each LP must choose a range $(Z^l, Z^u]$ such that he is willing to buy or sell to any LT an amount x of the first or y second asset at an exchange rate within the bounds Z^l, Z^u of the range. Such bounds must be chosen by the LP in a grid $\{Z_i\}_{i=1}^N$.
The marginal exchange rate Z_t expressing the amount of asset X received per each unit of asset Y is assumed to follow the stochastic dynamics given by:

$$dZ_t = \mu_t Z_t dt + \sigma Z_t dW_t,$$

where:

- μ_t: Drift process with finite fourth moment;
- σ: Volatility coefficient, assumed to be constant;
- W_t: Standard Brownian motion.

The quantities x and y of assets X and Y allocated by an LP can be expressed in terms of the depth of the LP liquidity in the pool κ, i.e. the percentage of assets X and Y allocated by the LP over the whole amount allocated by all the LPs in the pool measured in a reference numeraire (for instance USDC):

$$x_t = \kappa \left((Z_t^u)^{1/2} - (Z_t^\ell)^{1/2} \right),$$

$$y_t = \kappa \left((Z_t^\ell)^{-1/2} - (Z_t^u)^{-1/2} \right).$$

Wealth Dynamics

In [3], the authors develop an optimal strategy for dynamically managing liquidity in a CFM. The goal is to maximize the LP's terminal wealth, which includes fees earned from trades and P&L from market-making. The strategy adjusts the liquidity range and skew based on the DEX spot.
For log-utility preferences, they derive an explicit solution that navigates the trade-off between collecting fees and managing impermanent loss. When volatility rises, the LP broaden the liquidity range to prevent the possibility of seeing her liquidity inactive. In extreme cases of high volatility, the LP may exit the pool entirely as liquidity provisioning can become unprofitable. On the other hand, increased fee potential due to higher trading activity encourages concentrating liquidity in a narrow band around the exchange rate, balanced against the risks of range concentration. When the marginal exchange rate exhibits stochastic drift, the strategy shifts liquidity placement to capture trading flows and capitalize on expected rate changes.
The LP's wealth, expressed in units of asset

X, can be decomposed additively:

$$\tilde{x}_t = \alpha_t + p_t + c_t,$$

where the processes α_t, p_t and c_t represent:

1. **Position Value** (α_t): total asset value of the LP in the liquidity pool;
2. **Fee Revenue** (p_t): income from fees earned by the LP and paid by LTs;
3. **Rebalancing Costs** (c_t): Costs incurred while adjusting liquidity ranges.

In the following we provide a brief description of each aforementioned component.

Position Value

The position value α_t evolves according to the following stochastic differential equation:

$$\begin{aligned} d\alpha_t &= \frac{\tilde{x}_t}{\delta_\ell + \delta_u}\left(-\frac{\sigma^2}{2}dt + \mu_t\delta_u dt + \sigma\delta_u dW_t\right) \\ &= d\mathrm{PL}_t + \frac{\tilde{x}_t}{\delta_\ell + \delta_u}\left(\mu_t\delta_u dt + \sigma\delta_u dW_t\right), \end{aligned}$$

where:

- $\tilde{x}_t$: the LP's wealth in the reference numeraire;
- δ_u and δ_ℓ: control parameters defining the upper and lower bounds of the liquidity range, such that the liquidity range is defined like:

$$\begin{cases} (Z_t^u)^{1/2} &= Z_t^{1/2}/(1-\delta_t^u/2), \\ (Z_t^l)^{1/2} &= Z_t^{1/2}(1-\delta_t^l/2). \end{cases} \quad (1)$$

The deterministic component:

$$d\mathrm{PL}_t = -\frac{\sigma^2}{2}\cdot\frac{\tilde{x}_t}{\delta_\ell + \delta_u}dt, \quad (2)$$

represents the **predictable loss**, i.e. losses due to arbitrages made available between the liquidity range quoted in the pool by the LP and other trading venues due to the latency between the time of variation in the exchange rate on other markets and the time of update of liquidity provision by the LP in the liquidity pool, known as **loss versus rebalancing** (LVR).
On the other hand, due to the observability of the drift component, the LP can skew the liquidity provision range upward or downward according to whether μ_t is positive or negative respectively in order to capture a larger share of fee revenues (see the next Section "Fee Revenue").
Hence, the LP faces a trade-off between a **drift-based strategic positioning** and a drag component scaling in time like the variance of the exchange rate caused by arbitrages.

Fee Revenue

The dynamics of p_t expresses how the fee generation process, adjusted for spread and concentration risk, accrues over time, hence is given by:

$$dp_t = \frac{4}{\delta_\ell + \delta_u}\pi_t\tilde{x}_t\,dt - \gamma\left(\frac{1}{\delta_\ell + \delta_u}\right)^2\tilde{x}_t\,dt, \quad (3)$$

where:

- $\delta_\ell + \delta_u = \delta_t$: the total spread of the liquidity position.
- π_t: the pool's fee rate, i.e. the rate of fees generated instantaneously by transactions executed in the pool.
- $\gamma > 0$: the concentration cost parameter, i.e an instantaneous (constant) friction coefficient penalizing the overall profitability of the LP activity according to the squared inverse of the liquidity range length.

As shown in (3), the fee revenue dynamics is made up of two distinct additive components:

- **Fee Revenue**:

$$\frac{4}{\delta_\ell + \delta_u}\pi_t\tilde{x}_t\,dt.$$

This term captures the proportional relationship between the LP's position size and the pool's profitability rate π_t. Narrower spreads (δ_t) yield higher fee income per trade.

- **Concentration Risk Penalty**:

$$-\gamma \left(\frac{1}{\delta_\ell + \delta_u} \right)^2 \tilde{x}_t \, dt.$$

This term accounts for the risk of the marginal rate Z_t exiting the LP's liquidity range, reducing the effectiveness of fee collection for smaller spreads.

Hence, similarly to the position value, the fee revenue captures a trade-off, namely the trade-off between profitability coming from the liquidity provision activity and concentration cost: the narrower the length of the liquidity range $\delta_t = \delta_t^l + \delta_t^u$, the higher the fee rate captured by the LP in the range $(\delta_t^l, \delta_t^u]$ but the higher cost-opportunity coming from quoting in a narrower range.

Rebalancing Cost

Finally, c_t, represents the rebalancing costs incurred by the LP as they adjust their liquidity position. Such costs are proportional to the holdings of asset Y that need adjustment and are given by:

$$dc_t = -\zeta \cdot \frac{\delta_u}{\delta_\ell + \delta_u} \tilde{x}_t \, dt, \tag{4}$$

where:

- ζ: a constant representing the proportional execution cost associated with rebalancing the liquidity position.

Optimal Liquidity Strategy

At this point, we introduce the position asymmetry function:

$$\rho_t = \frac{\delta_t^u}{\delta_t^l + \delta_t^u}, \tag{5}$$

which measures how much the position is skewed, i.e. how much the current spot Z_t is not equidistant from the boundaries $(Z_t^l, Z_t^u]$. In particular we have that:

- If $\rho = \frac{1}{2}$ the range is perfectly centered;
- If $\rho \to 0$ then $\delta_t^u \to 0$, hence $Z_t^u \to Z_t$: only *token1* is offered in the position $(Z_t^l, Z_t^u]$;
- If $\rho \to 1$ then $\delta_t^l \to 0$, hence $Z_t^l \to Z_t$: only *token0* is offered in the position $(Z_t^l, Z_t^u]$.

The author of [3] assume that:

$$\rho_t = \rho(\delta_t, \mu_t) = \frac{1}{2} + \frac{\mu_t}{\delta_t} = \frac{1}{2} + \frac{\mu_t}{\delta_t^l + \delta_t^u} \quad , \quad \forall t \in [0, T]. \tag{6}$$

In order to gain fees from the liquidity provision activity, the LP must keep the provided liquidity active, i.e. within a range enclosing the current spot price Z_t. Hence:

- If $\mu_t > 0$ the spot tends to move towards Z_t^u and the LP adjusts the liquidity range to the right;
- If $\mu_t < 0$ the spot tends to move towards Z_t^l and the LP adjusts the liquidity range to the left.

The LP wants to optimize the final expected log-utility of its capital locked into the pool. It can be proved that, assuming expression (6) for the asymmetry function, the final expected log-utility problem can be formulated with just one control variable, given by the liquidity provision range length δ_t. Moreover, we need to assume that the asymmetry function is square-integrable and this entails that δ_s must be such that:

$$\mathcal{A}_t = \Big\{ \delta_s : [t, T] \to \mathbb{R} : \quad \delta_s \; is \; \mathbb{F}_t - adapted, \quad \int_t^T \delta_s^{-4} ds < \infty \quad \mathbb{P} - a.s. \Big\}.$$

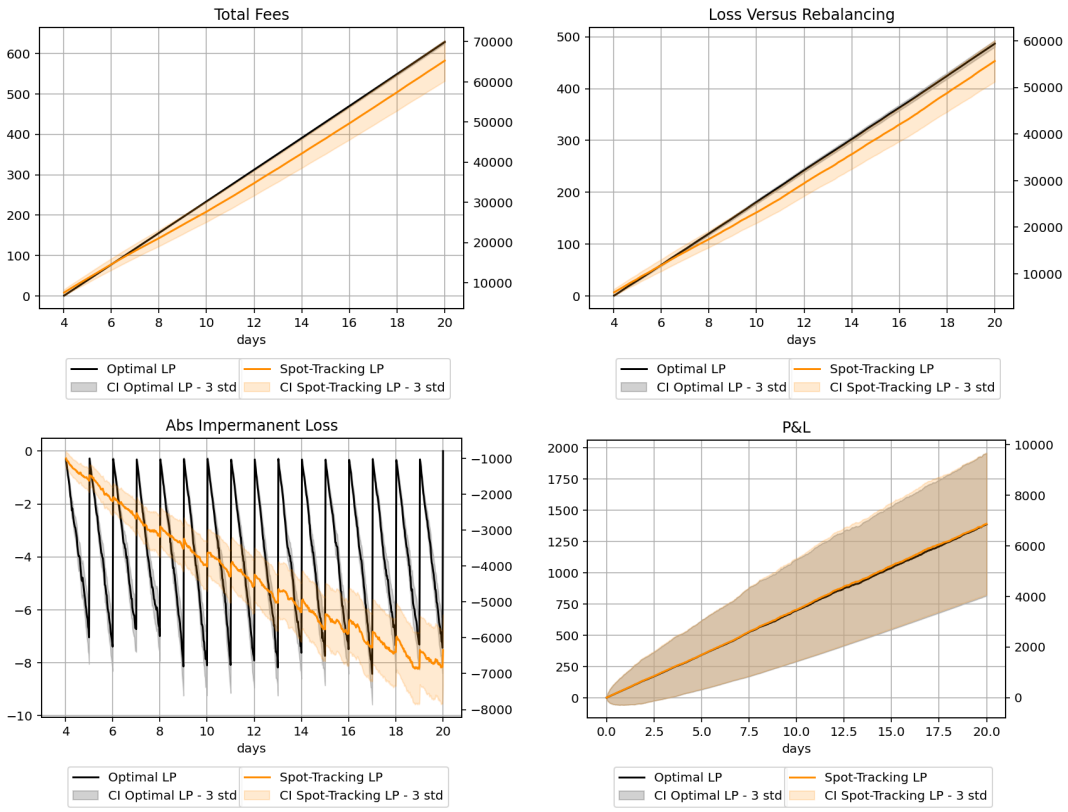

FIGURE 1: Total Fees (upper-left), losses-vs-rebalancing (LVR, upper-rigth), absolute impermanent losses (bottom-left) and total P&L (bottom-right).

Hence, under the square-integrability assumption of the asymmetry function and assuming the following **profitability condition**:

$$\pi_t \geq \eta_t = \frac{\sigma^2}{8} - \frac{\mu_t}{4}\left(\mu_t - \frac{\sigma^2}{2}\right),$$

the stochastic dynamic control problem can be formulated as:

$$\sup_{\delta \in \mathcal{A}_t} u^{\delta}(t, \tilde{x}, z, \pi, \mu) = \mathbb{E}_{t,\tilde{x},z,\pi,\mu}\left[\log(x_T^{\delta})\right]. \tag{7}$$

Under the previous assumptions, the resulting Hamilton-Jacobi-Bellman equations have a unique solution **optimal spread** given by:

$$\delta^* = \frac{2\gamma + \mu^2\sigma^2}{4\pi - \frac{\sigma^2}{2} + \mu\left(\mu - \frac{\sigma^2}{2}\right)}. \tag{8}$$

Once the trading frequency has been decided, the strategy consists in rebalancing the liquidity around the bounds defined by (1), with δ_u and δ_l defined by:

$$\begin{aligned} \delta_l &= \delta^*/2 - \mu_t, \\ \delta_u &= \delta^*/2 + \mu_t. \end{aligned} \tag{9}$$

It can be seen that the optimal range is strictly increasing in the volatility of the spot rate, so that the agent widens her liquidity bounds in highly volatile scenarios, to avoid the possibility of being left with inactive liquidity between the rebalancing times.
A higher fee rate, on the other hand, encourages the agent to narrow her position, in order to capture a larger portion of the trading fees.

Simulation Results

We assessed the performance of the aforementioned strategy over a set of 1000 simulations of the spot exchange rate dynamics. We simulated a WETH-USDT pool assuming:

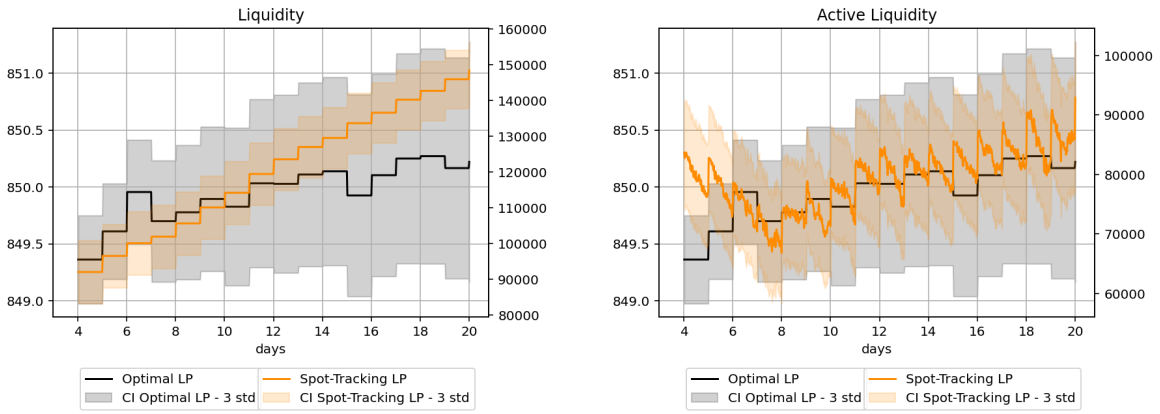

FIGURE 2: Liquidity (left) vs active liquidity (right).

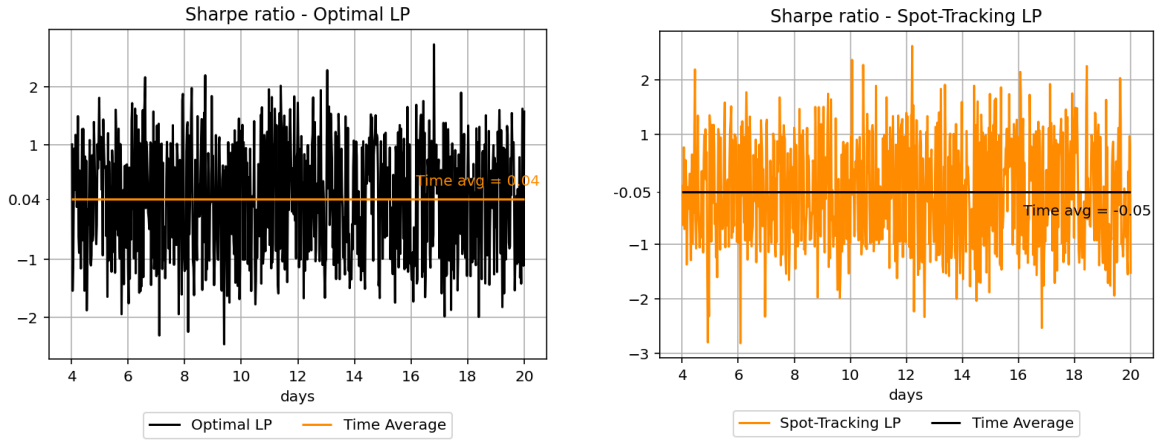

FIGURE 3: Simulated Sharpe Ratio of the Optimal LP (left) vs simulated Sharpe Ratio of the Spot-Tracking LP (right).

- A realistic level of volume/day ($\sim$ 20 MIO);
- Gas fees at 25 USDT/transaction;
- 1 block added every $\sim$ 14 seconds;
- Arbitrageurs taking arbitrage opportunities across the simulated WETH/USDT DEX and the corresponding CEX.

The CEX spot exchange drift and the volatility coefficients are set to $\mu = 50\%$ and $\sigma = 50\%$ (1:1 Sharpe Ratio of the Buy-and-Hold strategy assuming no risk-free rate). We tested the optimal AMM strategy outlined versus a simple Spot-Tracking strategy, keeping a liquidity provision range centered a +/-5% of the current spot price and adjusting it whenever the DEX spot price go out of such the corresponding boundaries. Both strategies were implemented with daily monitoring (i.e. at most 1 adjustment/day). Finally, we assumed the fee rate to be a martingale. Analyzing fig.(1), we notice that the Optimal Strategy is much more reactive than the Spot-Tracking one, being dependent also on the estimated fee rate. In facts, the absolute impermanent loss, defined as the difference between the capital value locked in the pool at time t and the capital value locked in the pool at the time of the last rebalancing $t - \Delta t$, where Δt is the rebalancing frequency (1 day in this case), is expected to be close to 0 at the rebalancing time whereas the spot-tracking case shows a clear negative drift, sign of a much lower probability of adjusting the liquidity at the rebalancing time. Moreover, analyzing the expected liquid-

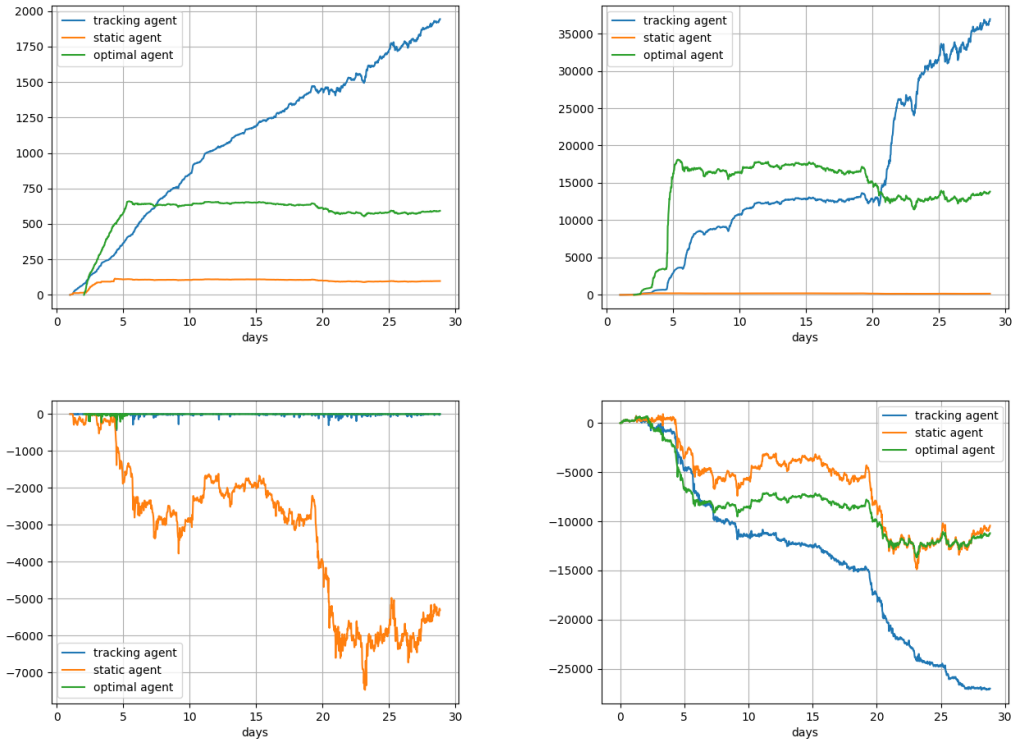

FIGURE 4: Total Fees (upper-left), losses-vs-rebalancing (LVR, upper-right), absolute impermanent losses (bottom-lefy) and total P&L (bottom-right) of the backtested strategies.

ity and active liquidity (i.e. the portion of liquidity provided in a range including the spot, hence available for swap at time t, fig.(2)), it is clear why total fees and losses vs rebalancing are 2 order of magnitude lower for the Optimal LP vs the Spot-Tracking LP: liquidity and active liquidity show the same ratio approximately.
Finally, what said about the absolute impermanent losses is confirmed: the Optimal Strategy keeps all its expected liquidity active through the whole period, whereas the Spot-Tracking strategy has active just a portion of the provided liquidity.
Analyzing the cumulated expected PnL (in base 100), the two strategies look equivalent: both the expected PnL and the confidence interval bounds almost overlap. Also the Sharpe Ratios look similar: noisy in both cases and averaging 0 in time.
Finally, we stress the fact that the overall profitability is sensitive to the assumptions made about the data generating process of the daily transactions, which determinate volumes, the level of gas fees and the initial amount of capital allocated to the pool: in other simulations with lower volumes and lower initial capital profitability decreases (not reported here).

Backtest and Live Results

We performed backtests on historical spot exchange rate. In this case we tested the Optimal LP versus both the Spot-tracking LP and the Static LP (i.e. a LP that simply put liquidity at the start period and doesn't adjust the liquidity range).
We estimated the drift parameter by a 20-day simple moving average and the volatility parameter by 20-day rolling standard deviation of the log-returns of the spot. The evidence (see fig. (4)) follows:

- LVR overcome Total fees;
- Active liquidity is heavily clustered: after the cluster at the start of the backtest period, the optimal strategy according to (8) is not to take part into

liquidity provisioning activities since the profitability condition is not met;

- The total PnL is negative and follows a path similar to the spot: this is explained by the fact that for most of the time the liquidity is inactive and so the PnL is mostly driven by the capital loss due to the spot decline and not by fees, which is actually the economic reason to provide liquidity in a DEX.

Similar results were obtained in other time windows.
Finally, as far as the live test is concerned, our partner Nuant A.G. deployed a smart contract tracking faithfully the performance of the mono-pool strategy starting from October 1st on the UniswapV4 blockchain. The strategy is not profitable so far (see fig. 5).

Optimal Liquidity Provision: Case of Multiple Pools

Multiple Spot Exchange Rate Dynamics

Let's extend the analysis to the case of multiple exchange rates available.
Let W_t a d-dimensional standard Wiener process and $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t\geq 0})$ the filtered space such that $\{\mathcal{F}_t\}_{t\geq 0}$ is the natural filtration of W_t. Let Z_t be a d-dimensional Ito process representing the exchange rate of a set of n tokens, hereafter $token_i$, $\forall i \in \{1, 2, ..., n\}$ with respect to another reference token, hereafter referenced as $token_0$ and let:

$$dZ_t = \mu I Z_t dt + \sigma I Z_t dW_t, \tag{10}$$

where $\mu \in \mathbb{R}^n$ is the vector of real coefficients, $\sigma \in \mathbb{M}_{n\times n}(\mathbb{R})$ is a matrix such that $\Sigma := \sigma\sigma^T$ is a positive definite matrix and $I \in \mathbb{M}_{n\times n}$ the $n \times n$ identity matrix, be the stochastic differential equation describing the dynamics of X_t. This entails the following dynamics of log-returns $\{X_t\}_{t\geq 0}$:

$$dX_t = \left(\mu - \frac{1}{2}\psi\right)dt + \sigma dW_t, \tag{11}$$

where:

$$\psi := diag(\Sigma), \tag{12}$$

given $diag : \mathbb{M}_{\mathbb{R}(n\times n)} \to \mathbb{R}^n$ the operator which yields the main diagonal of a matrix. By integrating eq. (11) in an interval $[0, t]$ we have that:

$$X_t = X_0 + \left(\mu - \frac{1}{2}\psi\right)t + \sigma W_t. \tag{13}$$

Hence:

$$\begin{aligned} Var(X_t) =& \sigma Var(W_t)\sigma^T = \sigma(\mathcal{I}t)\sigma^T \\ =& \sigma\sigma^T t = \Sigma t, \end{aligned} \tag{14}$$

where $\mathcal{I} \in \mathbb{M}_{n\times n}(\mathbb{R})$ is the $n \times n$ identity matrix.

Reformulation in the Principal Component Space...

Since Σ is a positive definite matrix by construction there exist $\Lambda, V \in \mathbb{M}_{n\times n}(\mathbb{R})$ matrix of eigenvalues and eigenvectors of Σ such that:

$$\Sigma = V\Lambda V^T = V\sqrt{\Lambda}\sqrt{\Lambda}V^T. \tag{15}$$

Recalling that, by definition, $\Sigma = \sigma\sigma^T$, it follows that:

$$\sigma = V\sqrt{\Lambda}. \tag{16}$$

Moreover, since $\sqrt{\Lambda} = \left(\sqrt{\Lambda}\right)^T$:

$$\sigma = V\sqrt{\Lambda} = \left(\sqrt{\Lambda}\right)^T V^T = \sqrt{\Lambda}V^T. \tag{17}$$

Hence, we can apply a suitable rotation to the original process of log-returns X_t and obtain a process $\tilde{X}_t$ such that:

$$\begin{aligned} \tilde{X}_t &= V^T X_t \\ &= V^T\left(\mu - \frac{1}{2}\psi\right)dt + V^T V\sqrt{\Lambda}dW_t \\ &= \left(\tilde{\mu} - \frac{1}{2}\tilde{\psi}\right)dt + \sqrt{\Lambda}dW_t, \end{aligned} \tag{18}$$

where:

$$\tilde{\mu} := V^T\mu, \tag{19}$$

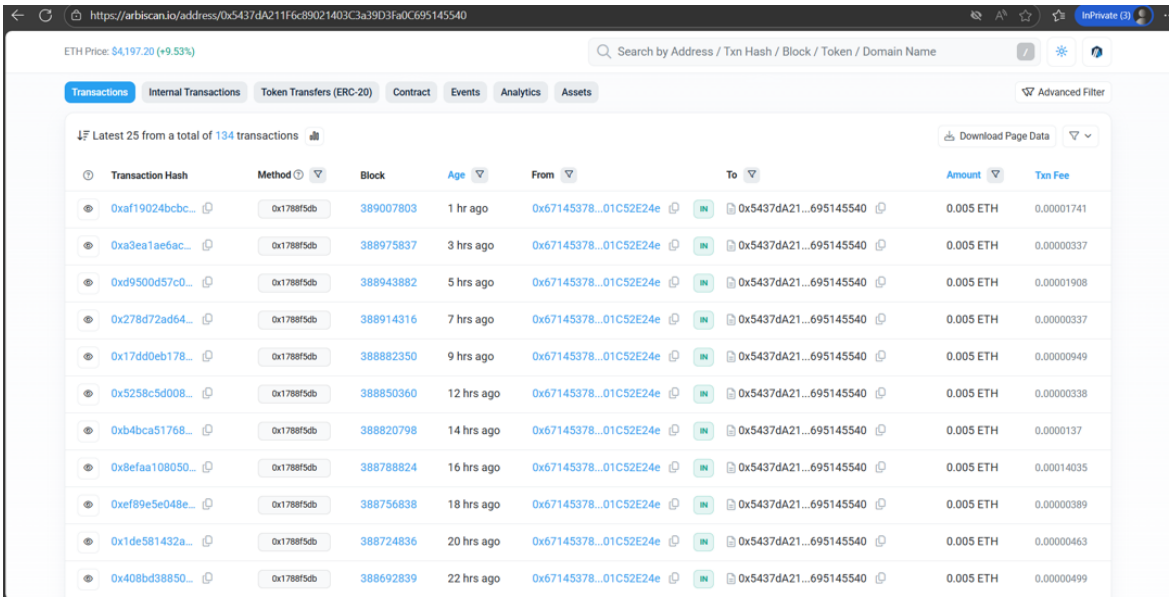

***FIGURE 5**: Transactions performed by the Smart Contract deployed by Nuant A.G. on UniswapV4 tracking the optimal strategy.*

and:

$$\tilde{\psi} := V^T diag(\Sigma) = V^T diag(V \Lambda V^T). \quad (20)$$

This paves the way leading to the possibility to recover the 1-dimensional optimal liquidity-providing strategy independently for each exchange rate by simply rotating the original log-return process X_t to obtain the process $\hat{X}_t$ representing the log-returns in an underlying-risk-factor space.
This is due to the fact that the optimal strategy developed in [3] is determined by maximizing an expected value and, hence, it is agnostic of the precise path followed by the driving process (either the original Wiener process W_t or its rotation $\hat{W}_t = V^T W_t$). Let the value process of the agent's portfolio at time t:

$$\mathcal{V}(t, X_t) := \sum_{i=1}^{n} z_i(t) \tilde{x}_i(t), \quad (21)$$

where $\tilde{x}_i(t)$ is the value process of the agent's wealth locked into a fictious liquidity pool exchanging the i-th underlying risk factor vs $token_0$:

$$d\tilde{x}_i(t) = \tilde{x}_i(t) \left(\frac{1}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \cdot \left[\left(4\tilde{\pi}_{t,i} - \frac{\tilde{\psi}_i}{2} + \tilde{\mu}_i \tilde{\delta}^u_{t,i} \right) dt + \sqrt{\lambda_i} \tilde{\delta}^u_{t,i} dW_t \right] - \gamma_i \left(\frac{1}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right)^2 \tilde{x}_i(t)\, dt, \quad (22)$$

where $\tilde{\pi}_{t,i}$ and $\tilde{\gamma}_i$ represent the fee rate and the concentration cost parameter for the exchange rate of the i-th underlying risk factor vs $token_0$, analogous to those defined in [3] for the exchange rate of $token_i$ vs $token_0$, $\forall i \in \{1, 2, ..., n\}$: the formers will be estimated conditioned to the latters as it will be explained in the next paragraph.
Now, by Ito's lemma, it follows that the log-wealth-process locked in the LP for the i-th underlying risk factor is given by:

$$\begin{aligned} d \log \tilde{x}_{t,i} &= \frac{1}{\tilde{x}_{t,i}} d\tilde{x}_{t,i} - \frac{1}{2\tilde{x}^2_{t,i}} d\langle \tilde{x}_{t,i} \rangle^2_t \\ &= \left(\frac{1}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \left[4\tilde{\pi}_{t,i} - \frac{\tilde{\psi}_i}{2} + \left(\frac{\tilde{\delta}^u_{t,i}}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \frac{\tilde{\psi}_i}{2} \right] dt \\ &\quad - \left(\frac{1}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \left[\left(\tilde{\gamma}_i + \frac{\lambda_i}{2} (\tilde{\delta}^u_{t,i})^2 \right) \left(\frac{1}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \right] dt \\ &\quad + \left(\frac{\tilde{\delta}^u_{t,i}}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) \left[\left(\tilde{\mu}_i - \frac{1}{2} \tilde{\psi}_i \right) dt + \sqrt{\lambda_i}\, dW_t \right] \\ &= \tilde{c}_i(\tilde{\pi}_{t,i}) + \left(\frac{\tilde{\delta}^u_{t,i}}{\tilde{\delta}^l_{t,i} + \tilde{\delta}^u_{t,i}} \right) d\tilde{X}_t, \\ &\forall i \in \{1, 2, \ldots, n\}. \end{aligned}$$

Note that $\tilde{\delta}^u_{t,i}$ and $\tilde{\delta}^l_{t,i}$ are a function of the fee

rate $\tilde{\pi}_{t,i}$ (see eq. 27 in [3]) and term $\tilde{c}_i(\pi_{t,i})$ in turn depends on them, other than $\tilde{\mu}_i$, $\tilde{\psi}_i$ λ_i and $\tilde{\gamma}_i$: the latter are assumed to be constant, hence the notation $\tilde{c}_i(\tilde{\pi}_{t,i})$.
In vector notation, let $d\tilde{\mathcal{X}}_t := [d \log \tilde{x}_{t,i}]_{i=1}^n$:

$$d\tilde{\mathcal{X}}_t = \tilde{c}(\tilde{\pi}_t) + \tilde{\Delta}_t^u d\tilde{X}_t, \tag{23}$$

with $\tilde{\Delta}_t^u \in \mathbb{M}_{\mathbb{R}}(n \times n)$ diagonal matrix containing the sequence $\left\{\frac{\tilde{\delta}_{t,i}^u}{\tilde{\delta}_{t,i}^l + \tilde{\delta}_{t,i}^u}\right\}_{i=1}^n$ on the main diagonal, $\tilde{\pi}_t \in \mathbb{R}^n$ vector of the fee rates $\{\tilde{\pi}_{t,i}\}_{i=1}^n$ and $\tilde{c}(\tilde{\pi}_t) \in \mathbb{R}^n$ containing the sequence $\{\tilde{c}_i(\pi_{t,i})\}_{i=1}^n$.
Note that the fee rates $\tilde{\pi}_t$ and the concentration cost parameters $\tilde{\gamma}$ of the pools in the fundamental risk factor space are unobservable. However, as it will be made clear, we will assume that they depend exclusively on the fee rate π_t in the real asset space. Moreover, the strategy is specified in terms of the liquidity provision policy $(\tilde{\delta}_t^u, \tilde{\delta}_t^u)$, which should dictate the policy in the asset space.
Hence, define the conditional expectation with respect to the σ-algebra generated by the fee rate as estimator of the future expected returns $\tilde{\pi}_t$, that is:

$$\tilde{p} := \mathbb{E}[d\tilde{\mathcal{X}}_t | \tilde{\pi}_t]. \tag{24}$$

Now, let:

$$\mu_{\tilde{p}} := \mathbb{E}[\tilde{p}] = \left[\mathbb{E}[\tilde{c}(\tilde{\pi}_t)] + \tilde{\Delta}_t^u \left(\tilde{\mu} - \frac{1}{2}\tilde{\psi}\right)\right] dt, \tag{25}$$

and:

$$\Sigma_{\tilde{p}} := Var(\tilde{p}) = \left(\tilde{\Delta}_t^u\right)^2 \Lambda dt. \tag{26}$$

Since $\tilde{\delta}_t^l$ and $\tilde{\delta}_t^u$ are also a function of the vector of the fee rates $\tilde{\pi}_t$ (see eq. (9)), we have that[2]:

$$\tilde{p} \sim \mathcal{N}(\mu_{\tilde{p}}, \Sigma_{\tilde{p}}). \tag{27}$$

...and Back into the Asset Space

The linearity of both the conditional expectation and transformation between the underlying risk factor space and the asset space causes the conditional expectation of the log-wealth process in the asset space to be straightforward:

$$p := \mathbb{E}[d\mathcal{X}_t | \tilde{\pi}_t] = \mathbb{E}[V^T d\tilde{\mathcal{X}}_t | \tilde{\pi}_t]. \tag{28}$$

Hence, it's mean and covariance matrix are respectively:

$$\mu_p := \mathbb{E}[p] = \\ = \left[\mathbb{E}[V^T \tilde{c}(\tilde{\pi}_t)] + \tilde{\Delta}_t^u \left(\mu - \frac{1}{2}\psi\right)\right] dt,$$

and:

$$\Sigma_p := Var(p) = (\tilde{\Delta}_t^u)^2 \Sigma dt. \tag{29}$$

We stress that the latter passage is of paramount importance, since the allocation can be made just selecting weights expressed as percentage of the total wealth available to the LP in the numéraire asset (i.e. in the asset space, other than in the underlying risk factor space): this will be clear in the next section "Agnostic Risk Parity Portfolio".
At this point, we need to estimate the unobservable parameters $\tilde{\pi}_t$ and $\tilde{\gamma}$ and the optimal policy in the asset space. As said before, we assume that the fee rate in the fictious pool of the i-th risk factor $\tilde{\pi}_{i,t}$ depends just on the fee rate vector π, hence the sigma-algebra generated by $\tilde{\pi}$ is equal to the σ-algebra generated by π. Moreover:

$$p = V^T \tilde{p}. \tag{30}$$

Hence, it must hold that:

$$\begin{cases} \mathbb{E}[p|\tilde{\pi}_t] = \mathbb{E}[V^T \tilde{p}|\tilde{\pi}_t] \\ \qquad\qquad = \mathbb{E}[V^T \tilde{p}|\pi_t] = \mathbb{E}[p|\pi_t], \\ Var[p|\tilde{\pi}_t] = Var[V^T \tilde{p}|\tilde{\pi}_t] = Var[V^T \tilde{p}|\pi_t] \\ \qquad\qquad = Var[p|\pi_t]. \end{cases} \tag{31}$$

[2] We assumed a Gaussian driving process. The same would happen in case of any elliptically distributed driving process.

In our implementation of the optimal strategy described in [3], the fee rate $\tilde{\pi}_{t,i}$ is estimated pathwise. Hence, a reasonable way to estimate the unobservable parameters and establish the dependence of the optimal allocation policy in the asset space (δ_t^l, δ_t^u) from the optimal allocation policy in the fundamental risk factor space $(\tilde{\delta}_t^l, \tilde{\delta}_t^u)$ is to find out the expression of the right hand-side of eqq. (31) and solve the non-linear system for $(\tilde{\pi}_t, \tilde{\gamma}, \delta)$. Analogously to the 1-dimensional case, it shall be noted that the wealth process in the asset space is given by:

$$\begin{aligned} dx_i(t) = x_i(t) &\left(\frac{1}{\delta_{t,i}^l + \delta_{t,i}^u}\right) \\ &\cdot \left[\left(4\pi_{t,i} - \frac{\psi_i}{2} + \mu_i \delta_{t,i}^u\right) dt + \delta_{t,i}^u \sigma^T dW_t\right] \\ &- \gamma_i \left(\frac{1}{\delta_{t,i}^l + \delta_{t,i}^u}\right)^2 x_i(t) dt, \end{aligned}$$

which implies that the log-wealth process is given by:

$$\begin{aligned} d\log x_{t,i} &= \frac{1}{x_{t,i}} dx_{t,i} - \frac{1}{2x_{t,i}^2} d < x_{t,i} >_t^2 \\ &= \left(\frac{1}{\delta_{t,i}^l + \delta_{t,i}^u}\right)\left[4\pi_{t,i} - \frac{\psi_i}{2} + \left(\frac{\delta_{t,i}^u}{\delta_{t,i}^l + \delta_{t,i}^u}\right)\frac{\psi_i}{2}\right] dt \\ &- \left(\frac{1}{\delta_{t,i}^l + \delta_{t,i}^u}\right)\left[\left(\gamma_i + \frac{\psi_i}{2}(\delta_{t,i}^u)^2\right)\left(\frac{1}{\delta_{t,i}^l + \delta_{t,i}^u}\right)\right] dt \\ &+ \left(\frac{\delta_{t,i}^u}{\delta_{t,i}^l + \delta_{t,i}^u}\right)\left[\left(\mu_i - \frac{1}{2}\psi_i\right) dt + \sigma_i^T dW_t\right] \\ &= c_i(\pi_{t,i}) + \left(\frac{\delta_{t,i}^u}{\delta_{t,i}^l + \delta_{t,i}^u}\right) dX_t \end{aligned}$$

$\forall i \in \{1, 2, ...n\}$.

In vector form:

$$d\mathcal{X}_t = c(\pi_t) + \Delta_t^u dX_t. \tag{32}$$

Hence:

$$\mathbb{E}[p|\pi_t] = c(\pi_t) + \Delta_t^u \left(\mu - \frac{1}{2}\psi\right) dt, \tag{33}$$

$$Var[p|\pi_t] = (\Delta_t^u)^2 \Sigma dt. \tag{34}$$

Remembering eq. [8]

$$\tilde{\delta}_i^* = \frac{2\tilde{\gamma}_i + \mu_i^2 \tilde{\psi}_i^2}{4\tilde{\pi} - \frac{\tilde{\psi}_i^2}{2} + \tilde{\mu}_i\left(\tilde{\mu}_i - \frac{\tilde{\psi}_i^2}{2}\right)}, \tag{35}$$

and eqq. [9]:

$$\begin{aligned} \tilde{\delta}_i^u &= \frac{\tilde{\delta}_i^*}{2} + \tilde{\mu}_i, \\ \delta_i^l &= \frac{\delta_i^*}{2} - \mu_i. \end{aligned} \tag{36}$$

Plugging eqq. (35) and (36) into the left hand-side of eq. (31) and eqq. (9), (33) and (34) into the right hand-side of eq. (31), we get a system on n equations depending on $(\tilde{\pi}, \tilde{\gamma}, \delta)$. Hence, the optimal policy can be found by solving a system of non-linear equations imposing the following constraints, which guarantee the profitability of the solution in the asset space:

$$\begin{aligned} \delta_i \in (0, 4) \quad \forall i \in \{1, 2, ..., n\}, \\ \tilde{\gamma}_i > 0 \quad \forall i \in \{1, 2, ..., n\}. \end{aligned} \tag{37}$$

Finally, in order to choose a set $\{z_i(t)\}_{i=1}^n$ of optimal weights at time[3] t, the Agnostic Risk-Parity Portfolio can be adopted.

Agnostic Risk Parity Portfolio

At this point, we aim at treating each implementation of the optimal LP strategy discussed in sec. "Optimal Liquidity Provision: Case of Single Pool" as a synthetic asset in a portfolio optimization framework. In particular, we want to adjust the overall portfolio exposure of the LP in a way such that:

- The Optimal Liquidity Provision strategy is performed on each pool;
- The execution is coherent with a notion of optimal portfolio allocation to be chosen.

We evaluated a certain number of different methods (see Annex for a brief review). Typically, in any portfolio allocation problem, an investor aims at minimizing its risk

[3] Note that, given the current setting (constant μ and σ) such weights would be time-independent.

under some constraint on expected returns, implying the need for the solution of a Markowitz(-like) optimal portfolio allocation problem. Some issues eventually arise in such cases:

- Usually, risk is measured by the mean of the standard deviation of the portfolio, depending on the covariance matrix. In such cases, the risk minimization approach leads to a concentration risk, arising from the fact that weights tend to be higher along the eigenvectors of the covariance matrix whose associated eigenvalues are lower.
- Parameter estimation uncertainty, arising from the intrinsic non-stationarity of financial time-series (at least unconditionally from the past state). It makes statistical estimators of expected returns and covariance matrices very uncertain.

The combination of both of the above issues often leads to poor out-of-sample performance of such portfolios.
To cope with the excessive concentration issue, Risk-Budgeting-constraint on asset holdings other that "underlying risk factor" holding have been proposed in literature (see Annex): if the first one actually doesn't diversify risk properly since it doesn't take into account asset correlations in the constraints, the second one usually leads to risk factors whose statistics differ a lot from the statistics of the individual assets, consequently implying portfolio allocations lacking of "financial soundness". Moreover, the difficulty in obtaining reliable estimations of the relevant parameters (i.e. expected values and covariance matrices) reverberate in the estimation of eigenvalues and eigenvectors, above all for smaller eigenvalues.
Those issues are efficiently dealt with by the Agnostic Risk Parity approach [1], briefly described hereafter.
The target is to obtain a method which:

1. Minimizes a certain distance between the risk factors in the "asset space" and the risk factors in the "spectral space";
2. Minimizes the bias due to uncertainty in parameter estimation, due to intrinsic non-stationarity of asset price and, consequently, return time-series.

Hence, let $r \in \mathbb{R}^n$ the standardized (i.e. demeaned and rescaled to unit variance) returns of each strategy over a time interval Δt and let $\Sigma^{std} \in \mathbb{M}_{\mathbb{R}}(n \times n)$ be their covariance matrix. Note that, since the returns are standardized, such matrix is actually a correlation matrix.
Moreover, let the vector of a (possibly biased) estimator of standardized expected returns $p \in \mathbb{R}^n$ arbitrarily chosen and let $C \in \mathbb{M}_{\mathbb{R}}(n \times n)$ an estimation of their covariance matrix, $\{\gamma_i\}_{i=1}^n \subset \mathbb{R}$ and $\{u\}_{i=1}^n \in \mathbb{R}^n$ being the related sequences of eigenvectors and eigenvalues, respectively.
As far as the first point is concerned, we aim at a representation of true returns and expected returns $\hat{r}$ and $\hat{p}$ respectively such that their Mahalanobis distance from their "asset space" counterpart r and p respectively is minimized, that is:

$$\underset{M \in \mathbb{M}_{\mathbb{R}}(n \times n) \,:\, MM^T = I}{argmin} \quad \mathbb{E}\big[(\hat{a}(M) - r)^T K^{-1} (\hat{a}(M) - r)\big] = I, \tag{38}$$

with $\hat{a}(M) = M\Sigma_{std}^{-1} a$, $\forall a \in \{r, p\}$, $\forall K \in \{\Sigma^{std}, C\}$.
In [1] (Annex), it is shown that the solution to the aforementioned problem is given by the identity matrix $I \in \mathbb{R}(n \times n)$. Accordingly, the following rotations of r and p are performed:

$$\hat{r} = \Sigma_{std}^{-1/2} r, \tag{39}$$

where:

$$\Sigma_{std}^{-1/2} := \sum_{i=1}^{n} \frac{1}{\sqrt{\lambda_i^{std}}} v_i^{std} (v_i^{std})^T. \tag{40}$$

$\{\lambda_i^{std}\}_{i=1}^n$ with $\{v_i^{std}\}_{i=1}^n \subset \mathbb{R}^n$ components of the matrix $V^{std} \in \mathbb{M}_{\mathbb{R}}(n \times n)$ of the (orthogonal) eigenvectors of Σ^{std}.

Analogously, for vector p:

$$\hat{p} = C^{-1/2} p, \tag{41}$$

$$C^{-1/2} := \sum_{i=1}^{n} \frac{1}{\sqrt{\gamma_i}} h_i h_i^T, \tag{42}$$

with $\{h_i\}_{i=1}^n \subset \mathbb{R}^n$ components of the matrix $H \in \mathbb{M}_{\mathbb{R}}(n \times n)$ of the (orthogonal) eigenvectors of C.
Now, define the portfolio:

$$g := \omega < \hat{r}, \hat{p} >= \omega \hat{r}^T \hat{p} =< \hat{\pi}, \hat{r} >, \tag{43}$$

where $\omega \in \mathbb{R}^+$ is a suitable normalization factor and $\hat{\pi} := \omega \hat{p}$ are the portfolio weights.
Now, let $g_i := \pi_i \hat{r}_i = \omega \hat{p}_i \hat{r}_i$. This portfolio, depending on the expected returns of the underlying risk factors $\hat{r}$, has the following desirable features:

- $\mathbb{E}[g_i^2] = 1 \quad \forall i, j \in \{1, 2, ..., n\} : i \neq j \quad \forall i, j \in \{1, 2, ..., n\} : i = j$ (unitary risk associated to each underlying risk factor);
- $\mathbb{E}[g_i g_j] = 0 \quad \forall i, j \in \{1, 2, ..., n\} : i \neq j$ (uncorrelation);
- $g_R :=< \mathcal{R}(r), \mathcal{R}(p) >= r^T R^T R p = r^T p$ (rotation invariance).

Where $R \in \mathbb{M}_{\mathbb{R}}(n \times n)$ is the rotation matrix representing the rotation operator $\mathcal{R}$.
In particular, unitary risk associated to each risk factor and null correlation define the scale invariance property, which is important in order to let all the assets be comparable, no matter the scale of each one. Moreover, the rotation invariance implies that rotation leads to the same portfolio expected return, no matter.
By substituting eq. (41) and (39) into eq. (43) we obtain:

$$\begin{aligned} g &= \omega < C^{-1/2} p, \Sigma_{std}^{-1/2} r > \\ &= \omega p^T \left(C^{-1/2}\right)^T \Sigma_{std}^{-1/2} r \\ &= \pi^T r =< \pi, r >, \end{aligned}$$

where:

$$\pi := \omega p^T \left(C^{-1/2}\right)^T \Sigma_{std}^{-1/2} = \omega (\Sigma_{std}^{-1/2})^T C^{-1/2} p. \tag{44}$$

Such portfolio is defined *Eigenrisk-Parity Portfolio*. In facts, let the second order moment of π:

$$\begin{aligned} &Var[\pi \cdot \sqrt{\gamma_i} v_i^{std}] = \mathbb{E}[(\pi \cdot v_i^{std})^2] \gamma_i \\ &= \omega^2 \left(\frac{1}{\sqrt{\gamma_i}}\right)^2 \mathbb{E}[(v_i^{std} \cdot C^{-1/2} p)^2] \gamma_i \\ &= \omega^2 \\ &\forall i \in \{1, 2, ..., n\}. \end{aligned}$$

That is, the variability is the same for all directions (eigenvectors) in which risks spread out. Moreover, by avoiding to target minimum variance, it effectively tackles the issue of over-allocation to small eigenmodes, whose estimation is usually quite unstable. It rather aims at an "Occam razor" allocation, i.e. relying on a minimum information set and (implicit or explicit) assumptions about the data generating process of both realized asset returns and expected asset returns.
In particular, it is referred to as *Agnostic Risk-Parity portfolio* in case no particular relationship is supposed among expected return estimators, i.e. in case $C := I$. Hence, the Agnostic Risk-Parity portfolio is given by:

$$\pi_{ARP} := \omega \Sigma_{std}^{-1/2} p. \tag{45}$$

Instead, in case $C := \Sigma_{std}$ and $\omega = \frac{1}{\mathbf{1}^T \Sigma_{std}^{-1} w}$ we recover the Markowitz optimal portfolio (48).
In order obtain an unbiased estimation of the covariance matrix, the authors of [1] suggest to use the Rotation Invariant Estimator introduced in [2]: as the name suggests, it is invariant under (random) variations of the eigen-directions of the rotation represented by the eigenvector matrix of the covariance matrix.

The Strategy

Hence, a possible algorithm follows:

- At time $t_0 = 0$:
  1. Estimate the vector of means and the covariance matrix. For an efficient estimator of the covariance matrix one can resort to the Rotational Invariant Estimator (see [[2]]).
  2. Recast the optimal allocation problem by representing the log-returns of the spot exchange rates X_t in the "underlying risk factor" space as done in eq. (18) and obtain the wealth process in eq. (22).
  3. Select a rebalancing frequency Δt and, hence, a time-grid $\{t_i\}_{i=1}^{n}$ such that $t_k = k\Delta t \quad \forall k \in \{1, 2, ..., n\}$.
- At each time-step t_n:
  1. Find out the optimal strategy as prescribed by [3] for each (independent) underlying risk factor.
  2. Solve the non-linear system of eq. (31) and find out the optimal policy δ^* to be executed in each real pool.
  3. Update the vector of expected returns and rebalance the exposures to each real liquidity pool according to the Agnostic Risk Parity portfolio weights.
  4. Execute the optimal liquidity provision strategy in each real liquidity pool.

Conclusions

In this paper we dealt with optimal liquidity provision strategies in crypto-currency Decentralized Exchanges (DEXs) working on the UniswapV3 protocol. We outlined the optimal liquidity strategy proposed in [3] and we assessed its performance both in a simulated WETH-USDC pool and in a backtesting framework. We found out that the strategy did not perform better than some simpler liquidity provision strategy, namely the spot-tracking LP strategy and, in the backtest case, the static LP strategy. Hence, we went on proposing an extension to a multi-pool setting: the liquidity provider wants to allocate its capital optimally across more than one exchange. Exploiting the spectral properties of the Wiener process, we recovered the possibility to use the optimal single-pool strategy according some optimality conditions about the overall capital allocation. The implementation and assessment of the performance will be carried out in a future work.

References

[1] **Benichou, R. et al.** *Agnostic Risk Parity: Taming Known and Unknown-Unknowns*. SSRN Electronic Journal, Elsevier, 2017.

[2] **Bun, J., Bouchaud, J.P. and Potters, M.** *Cleaning Correlation Matrices*. Risk, March 2016.

[3] **Cartea, Á., Drissi, F. and Monga, M.** *Decentralized Finance and Automated Market Making: Predictable Loss and Optimal Liquidity Provision*. SIAM Journal on Financial Mathematics, Vol. 15, N. 3, 931-959, 2024.

[4] **Deguest, R., Martellini, L. and Meucci, A.** *Risk Parity and Beyond — From Asset Allocation to Risk Allocation Decisions*. The Journal of Portfolio Management, Institutional Investor Inc., Vol. 48, pp. 108-135, 2022.

[5] **Kind, C.** *Risk-Based Allocation of Principal Portfolios*. April 2013.

[6] **Markowitz, H.** *Portfolio Selection*. Journal of Finance, Vol. 7, N. 1, pp. 77-91, March 1952.

[7] **Meucci, A., Santangelo, A. and Deguest, R.** *Risk Budgeting and Diversification Based on Optimized Uncorrelated Factors*. Risk, Vol. 11, pp. 70-75, 2015.

[8] **Thierry R.** *Introduction to Risk Parity and Budgeting*. Chapman and Hall/CRC, 2014.

Annex

Mean-Variance Portfolio Optimization

The most famous portfolio optimization framework was proposed, in 1952, by Harry Markowitz [6], earning him the Nobel Price in Economics in 1990: the so-called Mean-Variance Optimal Portfolio is supposed to be the one that minimizes the overall portfolio variance under a minimum expected return requirement $\mu^* \in \mathbb{R}$ and the so-called full-invested constraint, requiring that all the capital is invested in the available universe of assets. In formulas:

$$\begin{array}{rll} \min\limits_{w} & Var_{ptf}(w) := w^T \Sigma w & \\ & \text{s.t.} & \\ & w^T \mu & \geq \quad \mu^* \\ & w^T \mathbf{1} & = \quad 1. \end{array} \tag{46}$$

Equivalently, the problem can be reformulated in terms of maximum Sharpe Ratio (SR) under full-invested constraint:

$$\begin{array}{rll} \max\limits_{w} & SR_{ptf}(w) := \frac{w^T \mu}{\sqrt{w^T \Sigma w}} & \\ & \text{s.t.} & \\ & w^T \mathbf{1} & = \quad 1. \end{array} \tag{47}$$

The solution exists in closed form and is given by:

$$w^*_{\text{MVO}} = \frac{\Sigma^{-1}\mu}{\mathbf{1}^T \Sigma^{-1} \mu}. \tag{48}$$

In case of absence of minimum expected return constraint in (46), the Markowitz optimal portfolio is referred to as the Global Minimum Variance (GMV) portfolio. The solution is given by:

$$w^*_{\text{GMV}} = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}^T \Sigma^{-1} \mathbf{1}}. \tag{49}$$

Portfolios of Markowitz-type (either MSR or GMV) resort to a principle of minimum variance. Hence, they lead to portfolios which are very concentrated in the lowest variance asset. This is even more evident if we recast the problem in the underlying risk factor space, that is, formulating the problem with respect to the principal components of the covariance matrix other than the physical asset space. By substituting eq. (15) into the objective function in eq. (46) and defining $w := Vx$ we obtain:

$$\begin{array}{rll} \min\limits_{w} & Var_{ptf}(w) := w^T \Lambda w & \\ & \text{s.t.} & \\ & w^T \mu & \geq \quad \mu^* \\ & w^T \mathbf{1} & = \quad 1. \end{array} \tag{50}$$

In order to attain the minimum possible variance, it is clear that the weights will be higher for the principal component with lower eigenvalues, exposing the portfolio to a material concentration risk. Indeed, this would be optimal under the assumption that the investor is able to estimate accurately the covariance matrix Σ and the expected return vector μ. However, this is usually very difficult when dealing with returns of financial assets, often suffering from lack of stationarity or even ergodicity. Hence, the lack of reliable estimators

for covariance matrices and expected returns causes the over-allocation to a lower-variance asset a drawback of Markowitz-type optimal portfolios.
In order to deal with such drawbacks, a certain number of approaches have been proposed in literature: most of them, aim at sorting out the concentration problem by increasing the level of diversification of optimal portfolios by imposing so-called risk budget constraints and/or exploiting rotation invariance properties of elliptical probability distributions (e.g. recasting the problem in some equivalent space characterized by the spectral properties of estimated covariance matrices, from now on called underlying risk factor space).
However, the latter fixes miss to deal with the main issue of Markowitz-type optimal portfolios, that is, the intrinsic uncertainty in the estimation of expected returns and covariances. A possible solution to this issue, along with some other drawbacks arising in the latter alternative approaches, is proposed in [1] which will be described in the last section (sec. "Max-Entropy").

Risk-Budgeting and Risk-Parity

A suitable approach to increase the diversification of a Markowitz-type optimal portfolio is the risk-budgeting approach, which consists in introducing a set of nonlinear constraints to the Markowitz optimal portfolio such that the overall quote of variance of the portfolio that should be allocated to a single asset in eq. (46) is fixed:

$$\begin{array}{rrcl} \min\limits_{w} & Var_{ptf}(w) := w^T\Sigma w & & \\ & \text{s.t.} & & \\ & w^T\mu & \geq & \mu^* \\ & \frac{\Sigma w}{w^T\Sigma w} \times w & = & b \\ & w^T\mathbf{1} & = & 1, \end{array} \tag{51}$$

where the entries of $b \in [0,1]^n$ represent the maximum portfolio variance fraction allowed to the i-th un derlying risk factor (see [8], page 80 for the proof). In particular, when $b := \mathbf{1} \cdot \frac{1}{n}$, the risk-budgeting constraint is referred to in a self-explicable way: *risk-parity* constraint. Such a problem is convex: although no closed-form solution is available, it is solvable by mean of any numerical algorithm.
In particular, removing the minimum expected return required constraint and setting the risk-parity constraint, lead to the following modification of the Global Minimum Variance optimal portfolio problem (49):

$$\begin{array}{rrcl} \min\limits_{w} & Var_{ptf}(w) := w^T\Sigma w & & \\ & \text{s.t.} & & \\ & \frac{\Sigma w}{w^T\Sigma w} \times w & = & \mathbf{1} \cdot \frac{1}{n} \\ & w^T\mathbf{1} & = & 1. \end{array} \tag{52}$$

However, the optimal solutions to the problem just introduced are specified in the asset domain, that is, the optimal weights are set such that they satisfy a constraint expressed in terms of variance of the single asset, without taking into account that correlations among asset make drawdowns, due to a shock in the returns of one asset, propagate also to other assets, hence not taking into account the underlying risk factors potentially driving more than asset.
In order to cope with this issue, the same problem can be specified in the underlying risk factor domain, analogously to what was done in (52) by imposing risk-parity constraints

with respect to underlying risk factors:

$$\begin{aligned} \min_{w} \quad & Var_{ptf}(\hat{w}) := \hat{w}^T \Lambda \hat{w} \\ & \text{s.t.} \\ & \frac{\Lambda \hat{w}}{\hat{w}^T \Lambda \hat{w}} \times \hat{w} = \mathbf{1} \cdot \frac{1}{n} \\ & \hat{w}^T \mathbf{1} = 1. \end{aligned} \tag{53}$$

Max-Entropy

The max-entropy approach tackles the problem of risk concentration differently from the risk-budgeting approach: the diversification requirement is expressed by looking for the set of weights of the different underlying risk factors that maximize an entropic measure.
First of all, we introduce the following entropic measures, known as Rényi entropies, to be used as measures of portfolio diversification:

$$\begin{gathered} \mathrm{E}_{\alpha}(w) = ||w||_{\alpha}^{\frac{\alpha}{1-\alpha}} = \left(\sum_{k=1}^{N} w_k^{\alpha} \right)^{\frac{1}{1-\alpha}}, \\ \alpha \geq 0, \quad \alpha \neq 1. \end{gathered} \tag{54}$$

In particular, for $\alpha \to 1$, we recover the Shannon Entropy:

$$\mathrm{E}_1(w) = \exp\left(-\sum_{k=1}^{N} w_k \ln(w_k) \right), \tag{55}$$

whereas, for $\alpha = 2$:

$$\mathrm{E}_2(w) = \frac{1}{\sum_{i=1}^{n} w_k^2}. \tag{56}$$

Such measures vary from 1 to n depending on whether the portfolio is fully invested in one asset (i.e. $\exists k \in \{1, 2, ..., n\} : w_k = 1;\ w_j = 0 \quad \forall j \in \{1, 2, ..., n\} : j \neq k$) or in all assets in the same measure (i.e. $w_k = \frac{1}{n}\ \forall k \in \{1, 2, ..., n\}$).
Then, let:

$$p(\hat{w}) := \Lambda \hat{w} \times \hat{w}. \tag{57}$$

Then, the optimum portfolio weights $\hat{w} \in [0, 1]^n$ are the weights solving the following maximization problem:

$$\begin{aligned} \max_{\hat{w}} \quad & E_{\alpha}(\hat{w}) \\ & \text{s.t.} \\ & \hat{w}^T \mathbf{1} = 1. \end{aligned} \tag{58}$$

It can be shown that the max-entropy approach for $\alpha = 2$ is equivalent to impose a shrinkage to the covariance matrix in the Risk-Parity Global Minimum Variance optimal portfolio problem (see Proposition 4 of [4]).

Market Scenario Generation with GenAI

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.

Antonio Menegon:
Chief Innovation Office
He holds bachelor and master degrees in mathematics, with a specialization in quantitative finance, complemented by post-degree studies in artificial intelligence and machine learning. He began his career as a risk quant and risk analyst before moving into leadership role advising major banking groups. He guided teams across risk management, front office, and ICT departments, driving model development, risk assessments and monitoring, digital transformation, governance, and AI initiatives. He leads the Group's Innovation Center, focusing on the application of technology and artificial intelligence in financial services.

Giuseppe Crupi:
Quant Analyst
He graduated in Data Physics from the Università Statale di Milano. His thesis, "Development of a Machine Learning framework for Anomaly Detection in Financial Time Series", demonstrates how machine learning and finance can coexist to drive future advancements. In 2023, he joined iason working on projects focused on Market Risk with one of the major Italian banks.

Caterina Papetti:
Quant Analyst
She holds a M.Sc. in Mathematical Engineering, Quantitative Finance, at Politecnico di Milano. She joined iason in 2024 and, as a Quantitative Analyst, she is currently working on various strategic projects related to Market and Credit Risk modelling at major Italian banks.

Scenario generation is fundamental in financial risk management, enabling institutions to explore possible market evolutions, assess portfolio resilience, and design stress test exercises. Traditional approaches, such as historical simulation and Monte Carlo methods, are the go-to solutions, with practitioners and modellers who found ways to manage the sometimes restrictive assumptions and limitation which used to affect these approaches (e.g., heavy tails, volatility clustering, non-linear cross-asset dependencies, etc.).
Innovations in Machine Learning (ML), especially in the rapidly advancing field of Generative Artificial Intelligence (GenAI), are introducing promising alternatives for market data augmentation and synthetic scenario creation. Deep generative models, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and Gaussian Mixture Models (GMMs), can learn non-linear dependencies and high-dimensional probability distributions directly from empirical data, promising to move beyond traditional resampling methods.
We evaluated academia and industry literature, including the emerging research adapting Transformer-like architectures (the basis of Large Language Models - LLMs) for integrating unstructured data (e.g., news or sentiment) into conditional factor generation. By integrating insights and critically comparing leading methodologies (VAE, GAN, GMM, and Transformers) across core metrics (e.g., interpretability, stability, data efficiency), this article wants to provide a practical overview to assess whether AI (and in particular GenAI) can be really production-ready for financial institutions in complement traditional approaches in market scenario generation and data augmentation.

Scenario generation is a crucial, broad-spectrum component of risk management and strategic decision-making across the financial sector. Depending on the institution — be it a bank, an insurance firm, or an asset manager — and the specific use case, scenario modeling underpins a diverse range of critical functions. Applications span, for example, Market, Credit, and Solvency VaR, future exposure prediction in Counterparty Credit Risk (CCR), and forward-looking simulations of possible events for portfolio optimization. Independently by the use case, it is always important for financial institutions to rely on trustable and sound data. In this regards, generating plausible future evolutions of key risk factors (e.g., interest rates, stock prices, credit spreads, etc.) to assess portfolio resilience and inform capital allocation decisions is the first goal to be achieved. Secondly, an equally important objective is to have sound data *augmentation techiques* when needed. This includes reconstructing complete historical datasets by filling missing data points (due to non-quotations, trading suspensions, or technical errors) and generating synthetic-but-plausible historical paths. This augmentation is especially critical for stabilizing risk metrics, like long-horizon Historical VaR, where native time series scarcity compromises reliability. The foundational "classic methods" to achieve these goals are based on **historical simulation** and **Monte Carlo** techniques. Widely used, these methods have been long fine-tuned and improved in the years to handle potentially critical reliance on strong parametric assumptions (e.g., *Gaussianity*, *stationarity*). Among the complexities to be handled, we acknowledge critical market stylized facts such as *heavy tails*, *volatility clustering*, and complex, *non-linear cross-asset dependencies*. These methodological challenges highlighted the need for more adaptive and data-driven frameworks. In the subsequent sections, we introduce Generative AI-based approaches as a complementary modern solution for both scenario generation and data augmentation. These methods, leveraging on architectures as Generative Adversial Networks (GANs), Variational AutoEncoders (VAEs), and Gaussian Mixture Models (GMMs), learn market distributions directly from empirical data, trying to over-

come some of the constraints of traditional techniques. We will present the methodologies currently proposed in the literature, discussing whether these AI-driven simulations can offer realistic, diverse, adaptive, and financially coherent solutions for risk management.

Classic Methods and Their Limits

Traditional approaches to scenario generation rely on well-established statistical techniques, primarily categorized into historical and model-driven (Monte Carlo) simulations. Foundational, they have been revised by both the academia and the industry in the past decades not to rely solely on too simplified assumptions and not to fall short to properly fit the modern, complex financial markets.

Historical Simulation

The **historical simulation** generates scenarios by directly resampling from historical time series, often adjusting for recent volatility using techniques like rescaling or normalization. This approach is widely favored (e.g., *classic Market Risk VaR* used in the banking sector) due to its methodological simplicity and its reliance on realized data, sidestepping the need for explicit parametric modeling of risk factor distributions. However, one limitation of the standard historical simulation is its implicit assumption of identically distributed returns, which may fail to capture time-varying volatility and clustering effects observed in financial data. To address this, **Filtered Historical Simulation** (FHS) has been proposed as an enhancement by Barone-Adesi and Giannopoulos (2001) [5]. In FHS, returns are first filtered through a volatility model (typically a *GARCH-type* process), and the standardized residuals are then resampled to generate new scenarios. The simulated returns are reconstructed by reapplying the estimated conditional volatility, allowing the model to account for recent changes in market dynamics. Closely related is the work of Bonollo et al. [7], who develop an enhanced risk management framework based on the FHS model. Their approach extends *FHS-VaR* to produce scenario-based P&L distributions over a one-year horizon, maintaining consistency with underlying risk-factor dynamics and annual realized volatilities via a bootstrap procedure. By linking realized and conditional volatility, the model achieves proportional control over simulated returns, enhancing responsiveness to market shocks while improving backtesting performance and capital efficiency under stress.

Model-Driven (Monte Carlo) Simulation

Monte Carlo simulation remains a foundational quantitative tool, offering significant flexibility and robustness by simulating a large number of potential market paths. This is achieved by randomly sampling from explicit stochastic processes that model the time-evolution of financial risk factors. Once a model (e.g., *Black-Scholes*, *Vasicek*, *GARCH*, etc.) is specified, Monte Carlo is adept at accommodating complex dependencies, non-linear pay-offs, and path-dependent derivative valuations. This is the classic approach used in use case as, for example, *Counterparty Credit Risk metrics* and *Solvency II Internal Model's SCR*.

Systemic Limitations of Classical Modeling

While traditional models are indubitably still the industry best practice, their practical utility can be compromised by their dependence on simplifying assumptions that clash with empirical market data. Among these, we can acknowledge difficulties that have been long studied as:

- **Stylized Facts and Distributional Shape**: in most of their base settings, classical methods impose assumptions as Gaussianity (normal returns) and stationarity. This can lead to a

FIGURE 6: Summary of the main limitations of traditional models in scenario generation.

systemic failure ([39]) to capture critical market characteristics that define the empirical distribution shape, including heavy tails and skewness (underestimating extreme events), serial and long-memory effects (autocorrelation), reliance on the i.i.d. bootstrap assumption which is often fundamentally violated in real data. Classical approaches addressing these limitations include the use of heavy-tailed distributions, such as the *Student-t distribution* [31], which better capture the fat tails and excess kurtosis typically observed in financial returns. Another widely adopted solution involves copula models, which decouple marginal distributions from their dependence structure, enabling a flexible modeling of nonlinear correlations and tail dependencies across risk factors [11, 22].

- **Rare Event Modeling**: due to parametrization choices and calibration constraints, rare events (e.g., crises, sudden market shocks) could be poorely represented. The probabilities of such extreme movements can be underestimated, leading to overconfident risk estimates and fragile stress tests. Among the methods used to overcome such a problem, one of the famously proposed solution entails the Extreme Value Theory (EVT) as discussed by Embrechts, Klüppelberg, and Mikosch (1997) [14] or Longin (2000) [29].

- **Dependency and Scalability Challenges**: models struggle with nonlinear and time-varying cross-asset dependencies, resulting in oversimplified correlation structures. This issue is compounded by the curse of dimensionality, where modeling high-dimensional, multi-factor dependencies becomes computationally expensive and statistically more complex. A method highly discussed in literature

and used in practice, from the banking to the insurance secotr, involves the use of copulas, as it has been proposed in works as Embrechts, McNeil and Straumann (2002) [15] and Demarta, McNeil (2005) [11].

- **Data Integration and Out-of-Sample Performance**: classical statistical approaches are traditionally limited to structured, tabular data, and mostly unequipped to incorporate unstructured inputs like financial news or sentiment. Indeed, integrating financial news into market scenario generation without AI primarily involves leveraging traditional statistical, econometric, and rule-based methodologies. These approaches focus on quantifying the impact of news content and incorporating it into models through historical correlations, sentiment indices, or regression-based techniques [2]. The fundamental challenge lies in systematically transforming unstructured textual information into structured data that can feed into quantitative financial models. Furthermore, calibrating models on historical data and still be able to generalize to new market regimes has always been a genuinely hard task.

The AI Paradigm: Data-Driven Modeling for Finance

The modeling challenges classical approaches have been facing led both the academia and the industry studying more data-driven and adaptive frameworks. **Advanced Machine Learning** (ML) and **Deep Learning** (DL) techniques started to proliferate within the financial sector research space, most often for business and managerial risk management applications.
Traditional AI applications, such as credit scoring or fraud detection, typically employ discriminative models, which learn mappings from inputs to outputs by estimating conditional probabilities $P(Y \mid X)$. For instance, a Random Forest classifier distinguishes between "likely to default" and "not likely to default," while models such as LSTMs (Long-Short Term Memories) are used for time-series forecasting. Generative AI, by contrast, aims to model the joint structure of the data itself (i.e., $P(X)$). Rather than predicting labels, it learns to represent and reproduce the statistical regularities of observed financial phenomena. This enables the synthesis of new, realistic data samples and supports applications such as scenario generation and data augmentation. Prominent architectures adapted for such use cases in finance include **Generative Adversarial Networks** (GANs) [19], **Variational Autoencoders** (VAEs) [27], and **Gaussian Mixture Models** (GMMs) [32].
Among these models, **Generative Adversarial Networks** (GANs) learn to produce realistic synthetic data through the competition between a generator and a discriminator [19]. Their conditional variant, cGAN, enables controlled scenario creation by conditioning on specific variables such as volatility regimes or macro factors [34]. Similarly, **Variational Autoencoders** (VAEs) generate data by encoding and decoding through a probabilistic latent space [27]; their conditional form (*CVAE*) and temporal extension (*TimeVAE*) allow for richer, time-dependent representations of financial series.
Alongside these neural models, **Gaussian Mixture Models** (GMMs) provide a classical probabilistic approach [37] and a particularly interpretable and statistically grounded framework. By expressing the data distribution as a weighted sum of Gaussian components, GMMs capture nonlinearities, skewness, and multi-modality in empirical financial return distributions. This flexibility makes them valuable for modeling heterogeneous market regimes or mixtures of volatility states. As discussed by Brigo and Mercurio (2006) [8], mixture models naturally extend classical parametric settings by combining multiple Gaussian

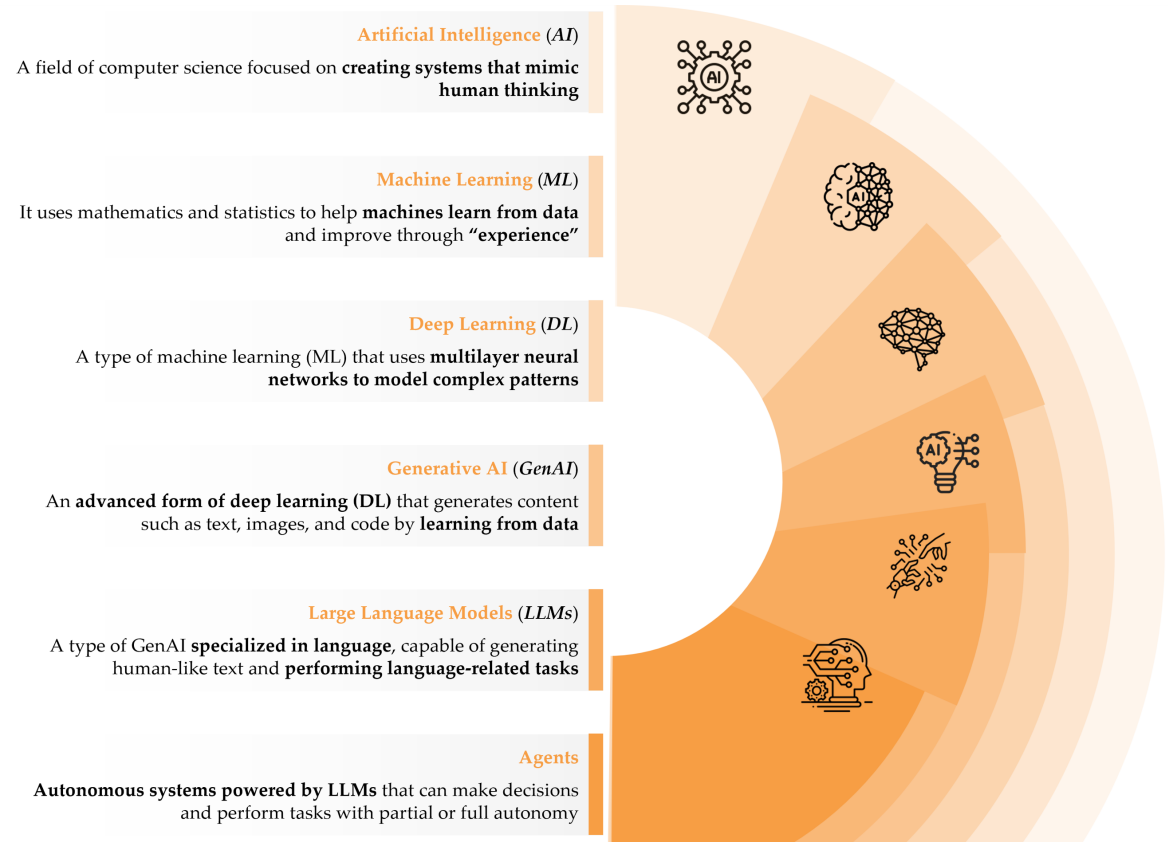

FIGURE 7: Overview of different modern AI methodologies.

densities to approximate complex empirical distributions while preserving analytical tractability. Their use in finance ranges from modeling implied volatility surfaces and interest rate dynamics to generating realistic risk-factor scenarios within a probabilistic framework.

To address the sequential nature of market data, **Recurrent Neural Networks** (RNNs) and their improved version, the **Long Short-Term Memory** (LSTM) network, capture temporal dependencies for forecasting applications [23]. More recently, **Transformers**, based on attention mechanisms [42] rather than recurrence, have emerged as efficient architectures to try to model both short- and long-range dependencies in time series, as well as mean to encode automatically unstructured data into scenario generation pipelines [12, 28, 3].

Together, these models form a methodological spectrum: while GANs, VAEs, and GMMs are mainly used for sampling realistic scenarios, LSTMs focus on forecasting temporal dynamics, and Transformers are better used to include unstructured data — a distinction central to the framework developed in this study.

Operating without restrictive parametric assumptions, AI (and particularly GenAI) models promise to emulate historical price dynamics, volatility clustering, and asset interdependencies with high fidelity. By capturing nonlinear relationships and cross-sectional structures, they are proposed to achieve a richer representation of stylized facts — such as heavy tails, skewness, and volatility clustering — while aiming at scaling effectively to high-dimensional settings.

Limits	Description	Classic Solutions	How GenAI addresses it
Stationarity	Models assume constant statistical properties over time.	• Differencing and Return Transformation [Tunnicliffe]; • GARCH and Stochastic Volatility Models [Bollerslev].	Captures evolving market regimes and structural breaks through dynamic, data-driven learning.
Gaussian Assumption	Returns are assumed to follow a normal distribution.	• Heavy-Tailed Distributions (Student-t)[Mandelbrot]; • Copula Models [Haugh]; • Extreme Value Theory (EVT)[Embrechts][Longin].	Learns non-Gaussian, heavy-tailed, and skewed distributions directly from data, improving tail risk representation.
Independence	Ignores autocorrelation and temporal dependencies in returns and volatility.	• ARCH and GARCH Models [Bollerslev]; • Multivariate GARCH (MGARCH, DCC-GARCH) [Engle].	Models temporal and cross-asset dependencies using recurrent or attention-based architectures.
Poor Representation of Rare Events	Underestimates the probability of crises or extreme scenarios.	• Extreme Value Theory (EVT) [Embrechts] [Longin]; • Copula Models for Tail Dependence [Haugh]; • Mixture Models (Gaussian Mixtures)[McLaugh].	Simulates rare and extreme events by learning from both typical and atypical historical patterns.
Markovianity	Assumes future states depend only on current ones, ignoring long-term memory effects.	• ARMA and ARIMA Models (Higher-Order Lags) [Tunnicliff]; • Long-Memory GARCH (FIGARCH, FIEGARCH)[Ballie].	Learns complex, path-dependent relationships.
Curse of Dimensionality	Estimating joint distributions becomes unstable in high-dimensional spaces.	• Principal Component Analysis (PCA) [Jolliffe]; • Copula Decomposition [Embrechts].	Learns high-dimensional dependencies efficiently using latent representations that preserve structure.

TABLE 1: Classical limitation across classical financial modeling approaches and most Generative AI methods.

Challenges of Market Scenario Generation

While both academia and industry are assessing whether Generative AI can be a complement (if not in some cases an alternative) to traditional financial modeling, its application to market scenario generation and data augmentation remains far from straightforward. As such models gain attention for risk management and stress testing use cases, it becomes essential to critically assess their reliability, consistency, and theoretical validity within financial contexts. Three broad areas of challenge emerge at the intersection of finance and AI.

First, there is the challenge of **intra-class heterogeneity**. Risk classes — such as equities, interest rates, etc. — are not monolithic. Each is composed of multiple distinct risk factor types governed by their own dynamics, time scales, and statistical properties. For example, within the single Equity class, a model must simultaneously capture the behavior of spot prices, dividend curves, and implied volatility surfaces. Designing neural architectures capable of jointly capturing these diverse internal components without loss of fidelity is a non-trivial task. Second is the modeling of **inter-class dependencies**. Financial markets are deeply interconnected systems where shocks prop-

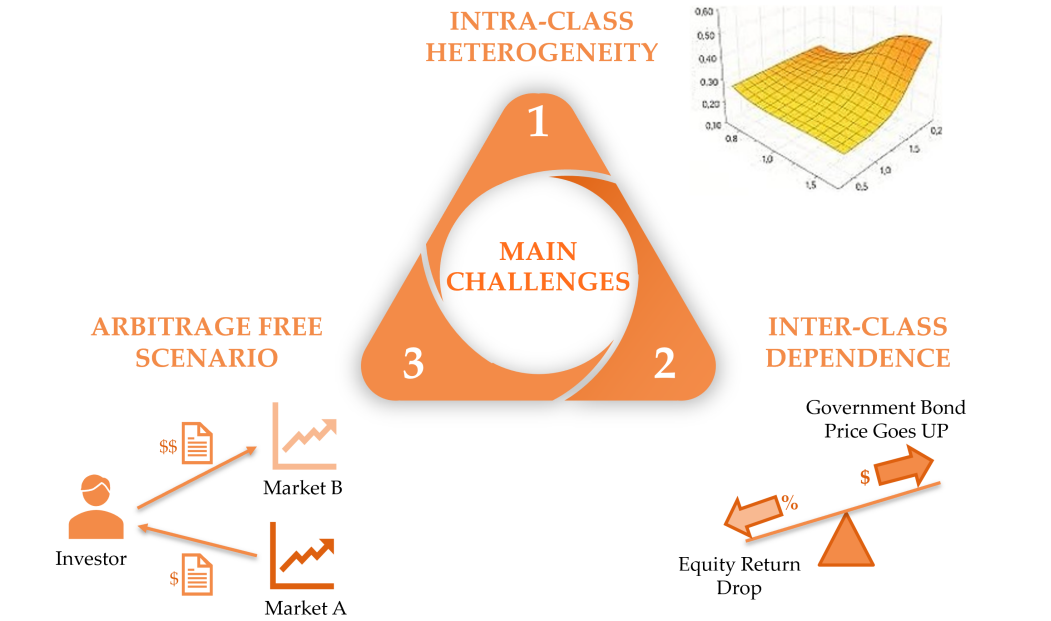

FIGURE 8: The figure illustrates the three key challenges in applying Generative AI to financial scenario generation. The first challenge is the modeling of multidimensional risk classes, such as spots, curves and volatility surfaces. The second challenge concerns inter-class dependencies, such as the example shown between government bond and equity returns during a 'flight-to-quality' period, when prices of the latter goes down, in favor of government bonds. The third and final challenge is the enforcement of financial consistency conditions — e.g., absence of arbitrage opportunities.

agate across assets through non-linear and time-varying channels. A shock to interest rates, for instance, can cascade into equity valuations and credit spreads. Generative models must therefore go beyond marginal accuracy (i.e., modeling each class in isolation) to reproduce the complex web of co-movements and contagion effects that define real-world market behavior — especially under stress conditions.

Finally, the enforcement of financial consistency, most notably the **absence of arbitrage**, represents a foundational constraint often violated by unconstrained neural networks. Ensuring that generated scenarios remain economically valid requires integrating domain knowledge, such as pricing relationships and no-arbitrage principles, directly into model design or training objectives.

In the following sections, we examine these challenges in greater depth. We begin with the problem of modeling multidimensional risk classes, proceed to the representation of inter-class dependencies, and conclude with the mechanisms required to guarantee arbitrage-free scenario generation.

Intra-Class Heterogeneity

Within each financial risk class lies a multidimensional structure that generative models must reproduce accurately. Interest rates, for instance, are represented through curves across maturities; credit spreads vary across tenors and counterparties; equity or FX markets are characterized by both price levels and implied volatilities. Each of these elements embodies a distinct but interrelated risk dimension, driven by different economic mechanisms and data characteristics. Capturing their joint dynamics within a coherent framework is one of the most demanding tasks in AI-based scenario generation.

The challenge arises from the coexistence of heterogeneous data types — scalar, vector, and surface representations — each with different statistical properties, time scales,

and sensitivities. Traditional neural architectures are typically optimized for homogeneous inputs, yet financial data within a single risk class can range from dense, high-frequency time series (e.g., spot prices) to sparse, structured surfaces (e.g., yield or volatility curves). Ensuring that models learn consistent patterns across these domains without distorting their internal relationships is non-trivial.

A key risk in this setting is the loss of internal coherence. For example, a generative model trained separately on yield curves and volatility surfaces might reproduce each marginal distribution accurately but fail to maintain the economic relationships between them — such as the link between term-structure shifts and changes in implied volatility. Similarly, naive training may violate smoothness or monotonicity constraints that are fundamental to financial realism.

Different architectures offer complementary strengths, as it will be better explained in the following section; for example:

- *Variational Autoencoder*s (VAEs), which can be viewed as the non-linear and non-parametric extension (or competitor) of Principal Component Analysis (PCA), can encode multi-dimensional structures into latent spaces that preserve relationships among term-structure, volatility, and liquidity factors, provided that the latent design reflects the geometry of the underlying data;
- *Generative Adversarial Networks* (GANs) can simulate complex shapes such as yield or volatility surfaces but require careful regularization to avoid mode collapse or unrealistic discontinuities;
- *Physics-informed or constraint-aware networks* can embed structural priors — such as no-arbitrage or curve-smoothness conditions — directly into the training objective, providing a more financially consistent output.

Ultimately, effective multidimensional modeling depends not only on data abundance but also on accurate architectural alignment with financial structure. Models must learn the internal geometry of each risk class (i.e., how its factors evolve jointly and respect financial constraints) before they can reliably support scenario generation or pricing applications.

Inter-Class Dependence

Beyond the internal consistency of each risk class lies the broader challenge of capturing dependencies among them. Financial markets operate as interconnected systems; e.g., movements in interest rates affect equity valuations through discount factors; credit spreads respond to both rates and macroeconomic stress; commodities influence inflation expectations, which in turn shape monetary policy. These linkages evolve dynamically and often strengthen under stress, amplifying systemic risk.

A generative model must therefore go beyond marginal (single-asset) accuracy to reproduce this complex web of co-movements, especially under stress conditions where correlations often spike and behave non-linearly. Standard correlation or copula frameworks provide well-established solutions, while deep generative models must infer such dependencies implicitly from data. This creates recurring potential problems as:

- *Hidden vs. explicit correlation modeling*: most generative models, particularly deep learning architectures, do not explicitly model correlation. Instead, they rely on the joint training of data to implicitly learn dependencies, which may lead to an under/over-estimation of weak or nonlinear links.
- *Overfitting to marginal structure*: generative models often optimize for marginal fidelity (e.g., minimizing reconstruction loss or adversarial divergence), sometimes at the expense of capturing joint behavior.

Asset Class	Statistical Characteristics	Main Driver	Temporal Dynamics
Equities	High volatility, clustering, fat tails.	Investor sentiment, earnings, market cycles.	High-frequency movements.
Rates	Lower volatility, mean-reversion.	Central bank policy, macroeconomic trends.	Slower-moving, medium to long term.
Credit	Sensitive to default risk and liquidity premiums.	Firm fundamentals, systemic risk.	Discontinuous jumps, stress events.
FX	Very high-frequency noise, heavy tails.	Trade flows, interest rate differentials.	Tick-level data, intraday patterns.
Commodities	Seasonality, regime shifts.	Supply/demand shocks, geopolitics.	Mixed: intraday to seasonal cycles.

TABLE 2: Main features of different asset classes.

Asset Class A	Asset Class B	Typical Dependency
Interest Rates	**Government Bonds**	Rising interest rates lead to lower bond prices due to the inverse yield-price relationship.
Interest Rates	**Equities**	Higher rates increase the discount factor for cash flows, generally reducing stock valuations.
Interest Rates	**Credit Spread**	Credit spreads tend to widen when rates rise or in periods of economic uncertainty, reflecting higher default risk.
Commodities (e.g. Oil)	**Inflation**	Energy prices are often positively correlated with inflation expectations, influencing central bank policy.
Equities	**Commodities (Oil)**	Rising oil prices can increase costs for firms, reducing profits and exerting downward pressure on equity markets.
FX (USD)	**Commodities (Oil, Gold)**	Many commodities are priced in USD; a stronger dollar often leads to weaker commodity prices.
Equities	**Credit Markets**	Deterioration in credit conditions or widening spreads often coincides with declining equity valuations.
Volatility index (VIX)	**Equities**	Equities & VIX tends to rise sharply during equity market downturns, reflecting increased risk aversion.

TABLE 3: Examples of cross-asset dependencies observed in financial markets.

- *Insensitivity to systemic events*: rare but crucial market regimes, such as coordinated crashes or liquidity freezes, are underrepresented in training data. Without tailored mechanisms, models may fail to generalize correlation behavior in stress conditions.

No Arbitrage Techniques	
No-Arbitrage Tecniques by Design	Pre- or Post-Fixing Techniques
Embedding economic constraints directly into the model architecture or training objective.	Detecting and correcting arbitrage either during training (as a regularization term) or after training (post-processing).
Examples: - Arbitrage-free neural networks; - Physics-Informed Neural Network; - Generative models trained with economic consistency conditions.	Examples: - Adding no-arbitrage penalties to the loss function during training; - Post-processing scenario outputs to remove arbitrage opportunities.

TABLE 4: Two main paradigms of incorporating no-arbitrage into machine learning models.

Model	Multi Risk Class	Risk Class Dependency	Arbitrage-free
GAN	Needs homogeneous assets.	Capture nonlinear dependencies.	Complex to check.
VAE	Only for multi-latent architectures.	Encode dependencies.	Possible with external constraints.
RNN (e.g. LSTM)	Works with time series of single typology.	Learn time-varying and sequential dependencies	Not easy.
GMM	Need normalization.	Represent dependencies via covariance structures.	Can be implemented.

FIGURE 9: Some AI models and their behavior with respect to the aforementioned challenges in scenario generation.

Arbitrage-free Scenario

One of the core requirements in financial modeling is the absence of arbitrage, the impossibility of obtaining a riskless profit with zero investment [45]. This principle underlies derivative pricing, market equilibrium, and risk management. Any realistic scenario generation method must therefore respect fundamental financial constraints such as no-arbitrage, pricing consistency, and risk-neutral valuation. However, the growing use of machine learning, especially deep neural networks, introduces challenges to these principles. Purely data-driven models optimize statistical loss functions without intrinsic awareness of financial structure, and may thus generate economically inconsistent scenarios, embedding artificial arbitrage opportunities. Typical examples include negative swap spreads, unjustified yield-curve inversions, or violations of put–call parity in option pricing outcomes that, while statistically plausible, are financially invalid. To mitigate such issues, the literature distinguishes **two main approaches**. The first *enforces no-arbitrage by design*, embedding financial structure directly into the model or its loss function. The second *detects and corrects violations ex-post*, either through regularization or post-processing adjustments. Both paradigms aim to align data-driven models with the fundamental economic logic required for credible financial scenario generation and data augmentation.

No-arbitrage by design approaches offer the most principled path toward economically

valid generative models. One strategy is the development of ad hoc architectures, where the model's output space is constrained to only produce arbitrage-free outputs. A notable example of no-arbitrage by design is provided by Ning et al. (2022) [36], who propose a VAE framework specifically tailored to generate arbitrage-free implied volatility surfaces. An even more rigorous variant of the by-design approach leverages **physics-informed neural networks** (PINNs), both methods will be discussed deeply in the next chapter.

In contrast, pre/post fixing methods offer more flexible but less theoretically elengant solutions. Regularization-based techniques add penalty terms to the training loss that softly discourage arbitrage. For example, a term may penalize deviations from put-call parity, or penalize option price surfaces that violate convexity or monotonicity. These constraints act as guardrails during training, encouraging the model to stay close to the arbitrage-free region of function space without imposing hard constraints that might limit generalization. For instance, Hyndman (2021) adopts a pre-fixing approach, embedding arbitrage constraints directly into the training process through projection-based regularization, thus ensuring yield curve forecasts remain economically consistent [25]. Conversely, Vuletić (2023) proposes **VolGAN**, a post-fixing framework where a generative adversarial network produces implied volatility surfaces that are subsequently corrected via projection procedures to remove butterfly and calendar arbitrage, restoring arbitrage-freeness after generation [43]. Both approaches will be discussed in greater detail in the following section.

While these methods cannot eliminate all arbitrage violations, they provide a practical balance when architectural constraints or complex loss formulations are computationally infeasible. Each approach has trade-offs: no-arbitrage architectures may be overly restrictive in high-dimensional settings, PINNs offer theoretical rigor but high computational cost, and regularization or post-processing methods are flexible yet lack formal guarantees. The appropriate choice depends on the application context — strict enforcement for regulatory risk use versus approximate correction for exploratory analyses.

In summary, maintaining arbitrage-freeness remains a central challenge as deep learning advances in finance. To ensure reliability and interpretability, data-driven models must remain anchored in financial theory, integrating no-arbitrage logic — whether through architecture, physics-informed losses, or corrective mechanisms — as a fundamental component of credible AI-based scenario generation consideration.

Approaches in Literature

This chapter provides an overview of the analyzed approaches to market scenario generation as found in the academic and practitioner literature. Starting from the concepts recalled on the classical statistical techniques, here we survey emerging machine learning models, highlighting how each method conceptualizes uncertainty, represents dependence structures, and handles regime dynamics. Particular attention is given to the ways in which different approaches address practical requirements such as stress realism, computational feasibility, and the integration of expert judgment.

The reader can find, for each studied approach, a brief description, the main results drawn from the original papers, and a closing summary of the key takeaways. Likewise, the whole chapter concludes with an overall synthesis of these methods, leaving the reader the choice of where and how to dig deeper.

Overall, through this lens, our aim is to clarify the current state-of-the-art and identify open questions that continue to shape this evolving field.

Scenario Generation Using GAN by Flaig

A notable contribution to the application of deep learning in scenario generation is the work of Flaig [17], who proposes a framework for generating realistic market scenarios using generative neural networks. The study focuses on the use of GANs as tools to model the joint distribution of financial risk factors, emphasizing their *ability to capture non-Gaussian features, tail behavior, and complex dependencies across multiple assets*. In particular, all simulations evaluated by the author are based on multiple benchmark portfolios, with a different ratio of risk classes, as shown in Figure 10.
One of the main strengths of this approach is that it tries to model multiple risk classes at once, learning directly from historical observations without relying on restrictive parametric assumptions. Figure 12 illustrates the generation of new scenarios using a GAN. The generated samples closely replicate the structure of the historical data, while also extending beyond the original observations, thus creating plausible but novel scenarios. By contrast, a simple resampling approach would be confined within the boundaries of the original data cloud, failing to generate such new variations.

The paper also offers a comparative analysis with traditional techniques, demonstrating that neural approaches can outperform both historical simulation and simple parametric models in capturing the marginal and joint dynamics of financial variables. In particular, Flaig's GAN-based framework is benchmarked against the risk scenarios produced by Economic Scenario Generators (ESGs) of European approved Solvency II internal models, drawing on the results of the MCRCS study (available on EIOPA's homepage [46]). The comparison focuses on five key risk factor categories: corporate credit spreads, sovereign credit spreads, equities, and interest rates (up and down shocks), the results are shown in Figures 13, 11 and 14.

This work exemplifies the growing intersection between machine learning and financial risk management, offering a scalable and adaptive alternative to classical scenario generation methods.

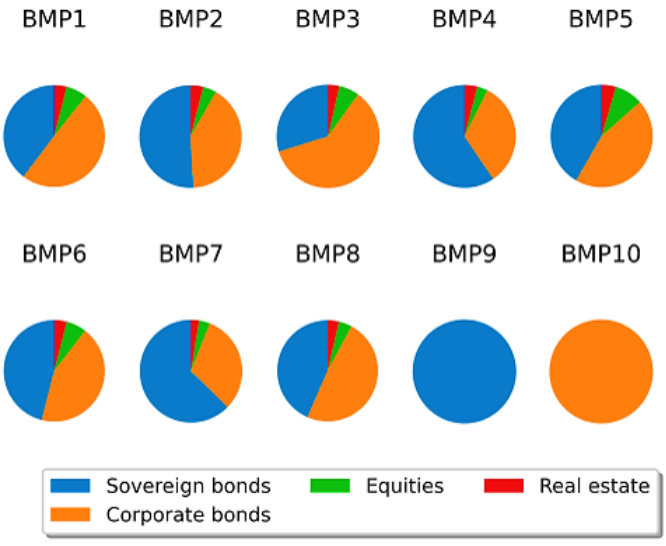

FIGURE 10: *Different composition of the benchmark portfolios used for Flaig analysis [17]. Each of the 10 portfolios are made as a ratio of 4 main risk classes: Sovereign bonds (blue), Equities (green), Real Estate (red) and Corporate Bonds (orange).*

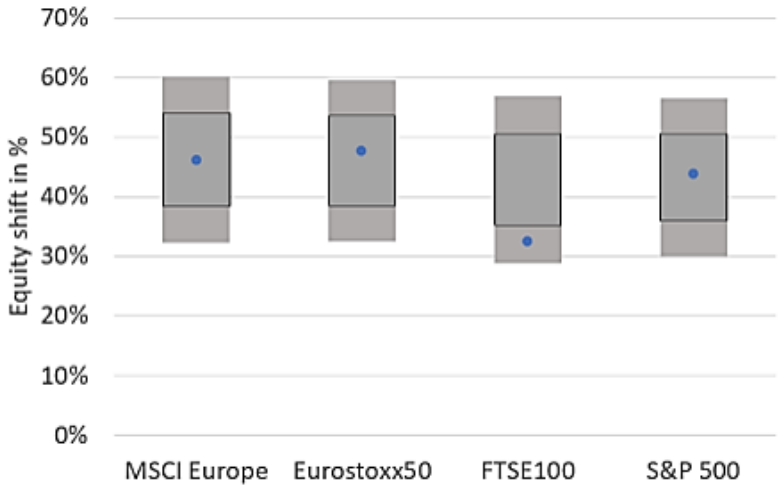

FIGURE 11: Comparison of the simulated shifts for equity risk factors, representation based on own results (blue dots) and EIOPA (gray boxes) [17]. The plot shows that the shifts are broadly consistent across most risk factors. For the FTSE100, the GAN generates less severe shocks compared to the other models. This outcome reflects the characteristics of the training data, as the FTSE100 exhibited lower volatility than the other indices over the period considered. Consequently, the GAN produces results that remain both plausible and data-driven.

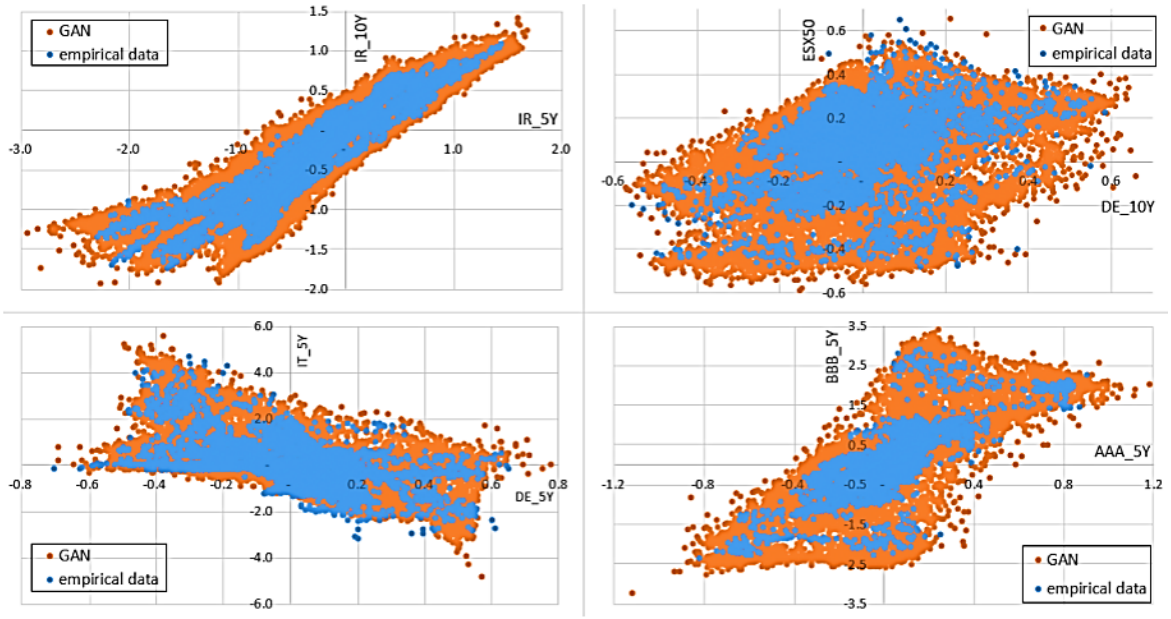

FIGURE 12: Scatterplots of four different risk factor pairs, empirical (blue dots) vs generated data by GAN (orange dots) [17]. The scatterplots display four pairs of risk factors: 5-year vs. 10-year interest rates, Eurostoxx50 vs. German government bond spreads, Italian vs. German government bond spreads, and AAA vs. BBB corporate credit spreads. In each graph, the blue dots correspond to the 4,330 empirical data points used for training, while the orange dots represent 50,000 scenarios generated by the GAN. As shown, the points generated by the GAN exhibit a distribution closely aligned with that of the original empirical data, while also introducing a greater diversity of samples. These newly generated data points lie within the same underlying distribution but were not present in the original training set, demonstrating the model's ability to generalize beyond observed data.

GAN Applied to Synthetic Financial Scenario Generation by Rizzato

Another important contribution to the field of generative modeling for financial applications is provided by Marco Rizzato in his work "Generative Adversarial Networks Applied to Synthetic Financial Scenarios Generation" [38]. The study investigates the use of GANs to simulate synthetic fi-

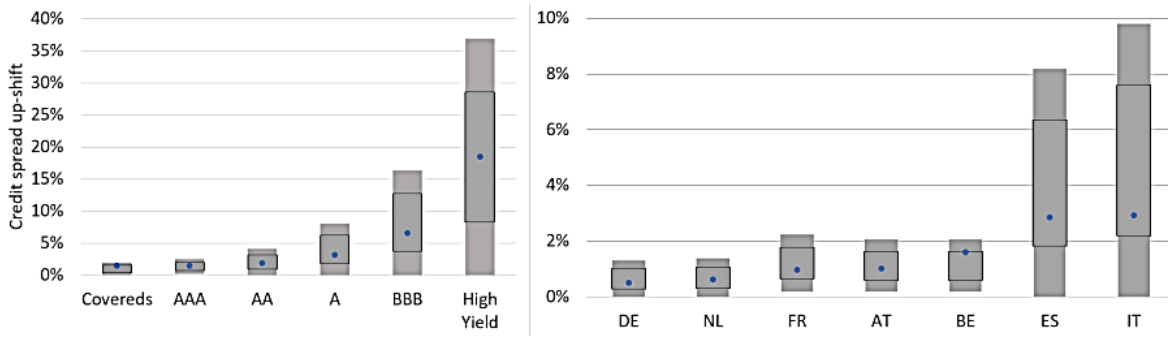

FIGURE 13: Comparison of the simulated shifts for the corporate and sovereign credit spread risk factors, representation based on own results (blue dots) and EIOPA (gray boxes) [17]. As shown, the GAN-based model aligns closely with the approved internal models, especially for corporate and sovereign spreads.

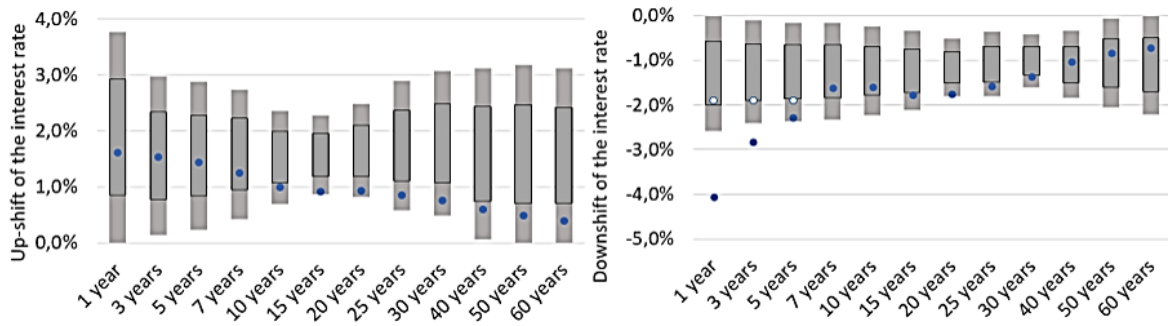

FIGURE 14: Comparison of the simulated shifts (left: up, right: down) for the interest rate risk factors, representation based on own results (blue dots) and EIOPA (gray boxes). The up-shifts generated by the GAN-based model fall within the reference boxes across all maturities, though for longer tenors they tend to cluster toward the lower bound. This pattern reflects the training period, during which interest rates were predominantly declining. For down-shifts, a different behavior emerges: short-term rates lie below the boxes, while medium- and long-term rates remain within them. This outcome can be explained by the sharp decline in short-term rates observed in the training data, in contrast to the relative stability of longer maturities. The GAN replicates this behavior, whereas traditional ESGs, calibrated over longer horizons and often supplemented with expert judgement, typically impose a floor on how negative rates are allowed to become.

Flaig (2023)		
Goal of the Paper	**Results**	**About Challenges**
To extend GAN-based scenario generation to a full internal market-risk model for insurance (one-year horizon under Solvency II) and compare with regulatory approved internal models.	Generated scenarios replicate empirical structures and extend beyond historical data; strong alignment with EIOPA benchmark models across multiple asset classes.	• The paper works with several risk categories, but they still not cover the full multidimensional space of risk classes ; • It does not focus in depth on intra-risk-factor correlation dynamics (the dependencies are learned via GAN) ; • It does not explicitly treat free-arbitrage scenario generation.

TABLE 5: Flaig (2023).

nancial time series aiming at *replicating real market dynamics*. The central idea is to train

a GAN on historical asset returns so that it learns their distributional properties and can generate new synthetic paths that preserve both marginal features and cross-asset dependencies. The neural network architectures evaluated in the paper are illustrated in Figure 15.

Rizzato's work highlights the flexibility of GANs in capturing nonlinear, high-dimensional distributions that are difficult to model with traditional statistical approaches. The GANs successfully reproduce the shape of the empirical distributions, as shown in Figures 16 and 17, and this similarity is quantitatively assessed using performance metrics such as the Kolmogorov–Smirnov test and Principal Component Analysis. From a statistical perspective, these tools provide standard methods for evaluating the closeness of two multivariate distributions. Although GANs have shown excellent performance, Rizzato points out that in modeling high-dimensional data distributions across many domains, they are not universally applicable generative models. In fact, GANs do not provide an exact representation of the underlying data distribution and furthermore GANs with Gaussian priors can only generate sub-Gaussian distributions, highlighting a fundamental theoretical limitation.

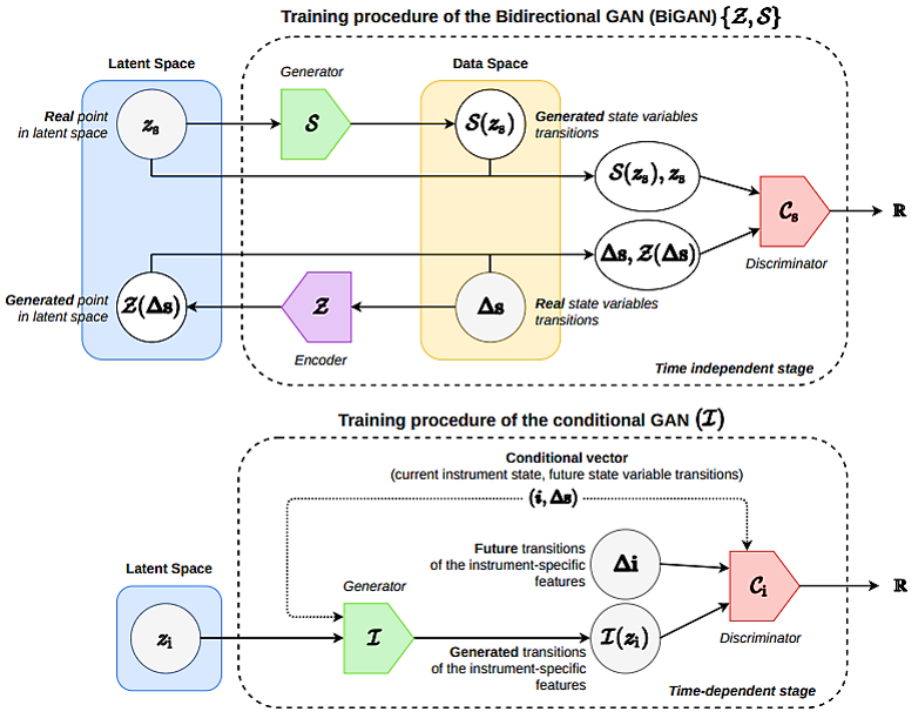

FIGURE 15: Graphical description of the two GAN used in the paper by Rizzato [38]. The BiGAN generator is trained on the historical dataset, and its performance is assessed by comparing the historical test set against an equally sized batch of synthetic data. The Conditional GAN, on the other hand, operates within a conditional probability framework, offering a more nuanced evaluation.

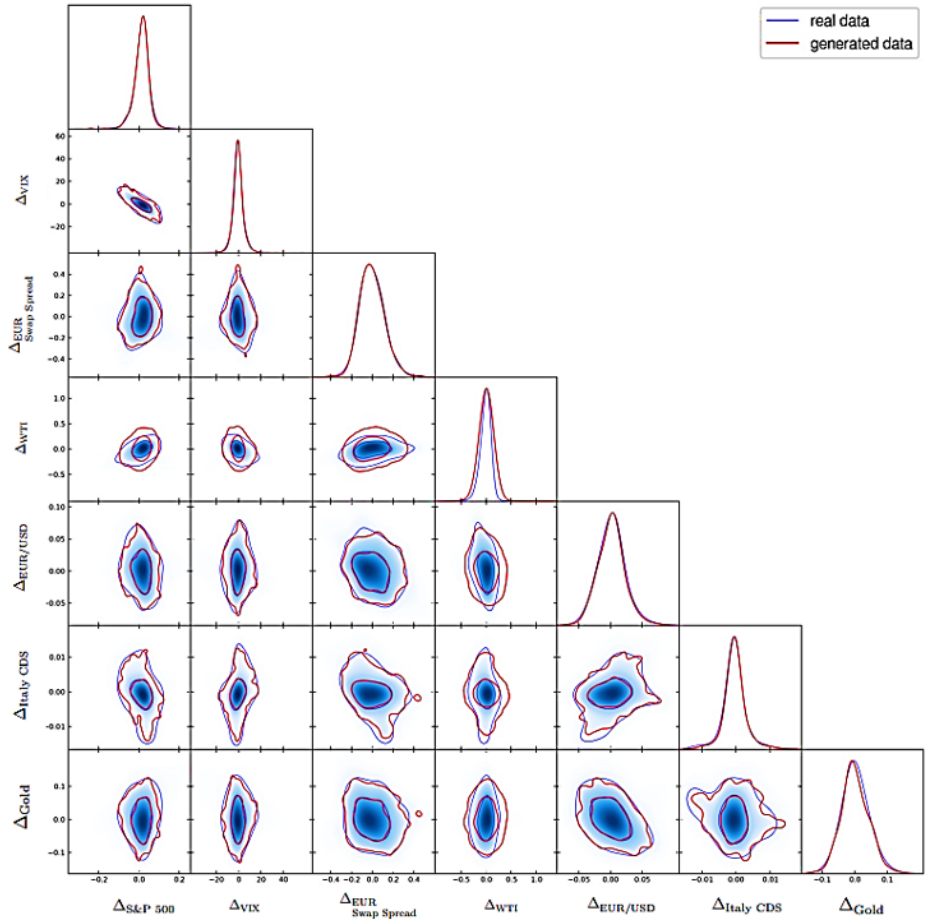

FIGURE 16: *Triangle plot for the evaluation of the BIGAN generator S. Comparison between the real (blue) and the generated (red) state variable transitions. This visualization allows to compare one-dimensional (diagonal panels) and two-dimensional marginal distributions (off-diagonal panels) sorting them by couples of variables. In each subplot, the 68% and the 95% confidence intervals are proposed [38]. As shown in the figure, for each of the asset classes considered, the data generated by the BiGAN exhibit a distribution that closely mirrors that of the real data — both in the univariate case and in the bivariate setting, when examining the joint behavior of two classes.*

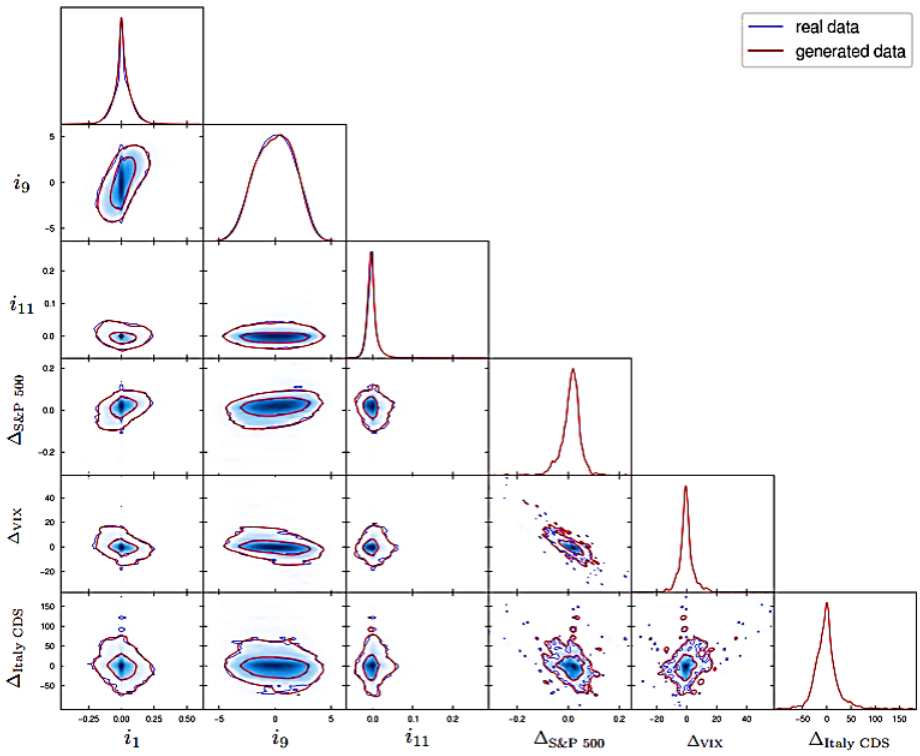

FIGURE 17: *Triangle plot for the evaluation of the cGAN generator I. Comparison between the real (blue) and generated (red) data. This visualization allows to compare one-dimensional (diagonal panels) and two-dimensional marginal distributions (off-diagonal panels) sorting them by couples of features. In each subplot, the 68% and the 95% confidence intervals are proposed [38]. As shown in the figure, for each of the asset classes considered, the data generated by the cGAN exhibit a distribution that closely mirrors that of the real data — both in the univariate case and in the bivariate setting, when examining the joint behavior of two classes.*

Generation of Realistic Synthetic Financial Time-Series by Dogariu

The article "Generation of Realistic Synthetic Financial Time-Series" by Dogariu, Ştefan, Boteanu, Lamba, Kim, and Ionescu [13] addresses the challenge of *producing synthetic financial datasets* that faithfully reproduce key stylized facts such as heavy-tail distributions, volatility clustering, autocorrelation, and cross-asset correlations, while enabling variable-length outputs and better data augmentation for downstream tasks. Rather than relying on sequence modeling networks alone, the authors explore multiple neural architectures, including fully connected GANs, convolutional GANs, VAEs, and Generative Moment Matching Networks (GMMNs), to generate realistic synthetics from U.S. equity market data, showing how such models handle cross-correlation between different assets and fixed-to-variable length time series. A significant contribution of the paper is its emphasis on evaluation methodology: beyond visual similarity, the authors introduce a portfolio trend prediction framework and quantitative metrics to validate the effectiveness of generated series in portfolio forecasting tasks. They demonstrate that their synthetic data preserves both statistical and predictive characteristics when applied to real trading scenarios.

Dogariu et al (2023)		
Goal of the Paper	**Results**	**About Challenges**
The goal of the paper is to develop generative-model techniques (GANs, VAEs) to synthesise realistic financial time-series data for problem of data augmentation.	Synthetic data preserve heavy tails, volatility clustering, and cross-asset correlation.	• The paper focus mainly on multiple stocks, avoiding other type of risk classes; • It addresses cross-correlations between different stocks, but it does not appear to more deeply model intra-risk-factor correlation structures; • The work does not explicitly enforce arbitrage-free constraints.

TABLE 6: Dogariu et al. (2023).

Time-Casual Variational Autoencoder Approach by Acciaio

A noteworthy recent contribution in the field is the **Time-Causal Variational Autoencoder** (TC-VAE) developed by Acciaio, Eckstein, and Hou [1], which addresses the overlooked issue of temporal causality in synthetic financial time-series generation.
The model's loss function provides an upper bound on the causal Wasserstein distance between true market distributions and generated sequences, offering theoretical guarantees on the closeness of statistical performance in decision-making tasks such as pricing, hedging, and portfolio optimization. Numerical experiments on synthetic models, such as Black-Scholes, are shown in Figure 19.
The same paper also provides additional comparisons with the Black–Scholes model, focusing on log-path and volatility distributions for both real and synthetic paths. A key metric employed is the sliced Wasserstein distance, which offers a principled approach to jointly comparing all one-dimensional projections between two measures, together with *Gaussian-kernel MMD* [20] and *Signature MMD* [10]. Figure 21 illustrates that real and generated paths remain relatively close under all three metrics.
All these tests show that TC-VAE effectively reproduces stylized market features (heavy tails, volatility clustering, skewness, kurtosis, and low autocorrelations) while offering strong performance on downstream tasks evaluated against enhanced statistical distance. This approach constitutes an important step forward in financial scenario generation with AI by combining a principled time-causal structure, rigorous mathematical control via adapted Wasserstein metrics, and empirical validation, making TC-VAE a possible robust and reliable tool for generating synthetic financial data.

Synthetic Financial Time Series Generation with LSTM by Schwarz

A recent and highly relevant contribution, if not properly solely generative modeling, is Schwarz contribution (2024) on financial time series with his work titled "Interpretable GenAI: Synthetic Financial Time Series Generation with Probabilistic LSTM" [40]. Despite Generative models which draw samples from the underlying data distribution to produce a range of plausible and coherent financial scenarios, LSTM based approaches are *primarily designed for forecasting, focusing on learning temporal dependencies to predict future values*.
The proposed model (see the neural network design in Figure 22) addresses the dual challenge of producing realistic mar-

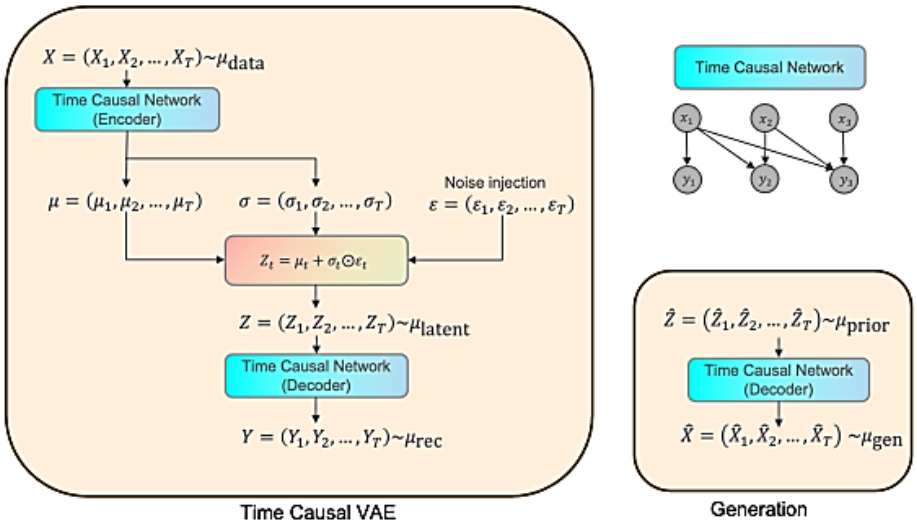

FIGURE 18: Time-casual variational autoencoder and generation. Unlike standard VAEs, TC-VAE enforces causality constraints on both encoder and decoder, ensuring that each modeled time step depends only on past observations, thus faithfully preserving the chronological structure of financial data.

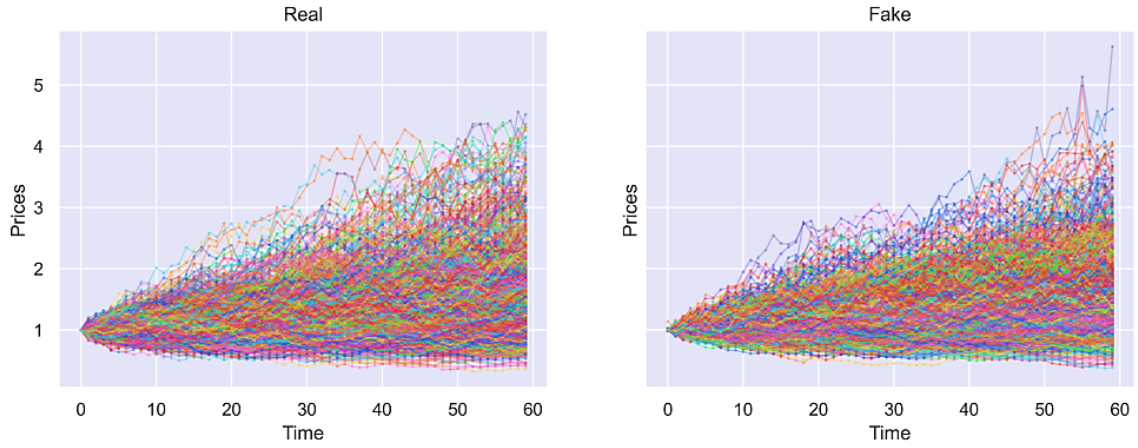

FIGURE 19: Illustration of real paths from a siscretized Black-Scholes model (left) compared to fake paths generated from TC-VAE model (right) [1]. As shown, the TC-VAE is able to reproduce paths that closely resemble those observed in the training data, successfully capturing a comparable level of depth and variability to that found in real-world cases.

ket scenarios while maintaining interpretability, two attributes that are often at odds in deep learning architectures.

What sets it apart is its incorporation of interpretability constraints, allowing practitioners to understand how scenarios are formed from a blend of deterministic components (e.g., ARMA/GARCH structures) and learned nonlinear transformations. Overall, Schwarz offers a compelling paradigm: combining generative capability with explainability and statistical structure. This balances performance with understandability, and aims at helping bridge the gap between black-box deep models and classical financial time-series approaches. Shwarz's proposal stands out as a practical tool for producing synthetic scenarios that want to be both realistic and interpretable, suitable for risk analysis, stress testing, algorithmic strategy design, or data augmentation.

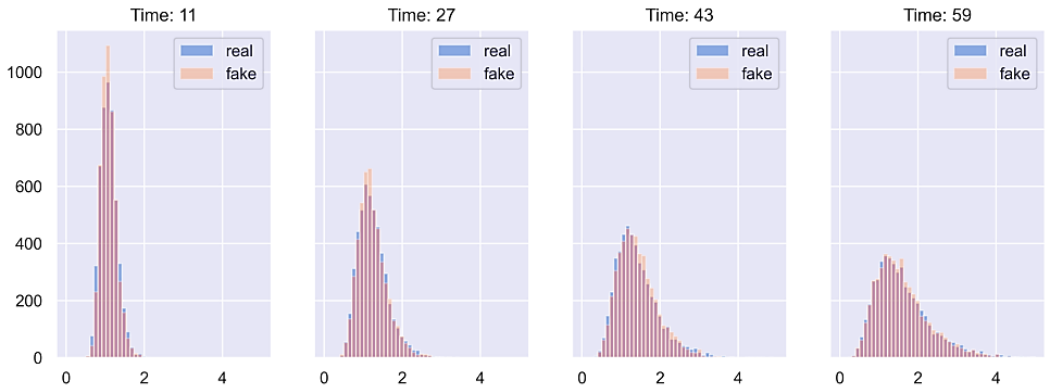

FIGURE 20: Visualization of marginal distributions at different time slices for real paths from a discretized Black-Scholes model (blue) compared to fake paths generated from the TC-VAE model (orange) [1]. As observed, the distributions are nearly overlapping across all the considered time slices, confirming the high reliability and consistency of the TC-VAE model.

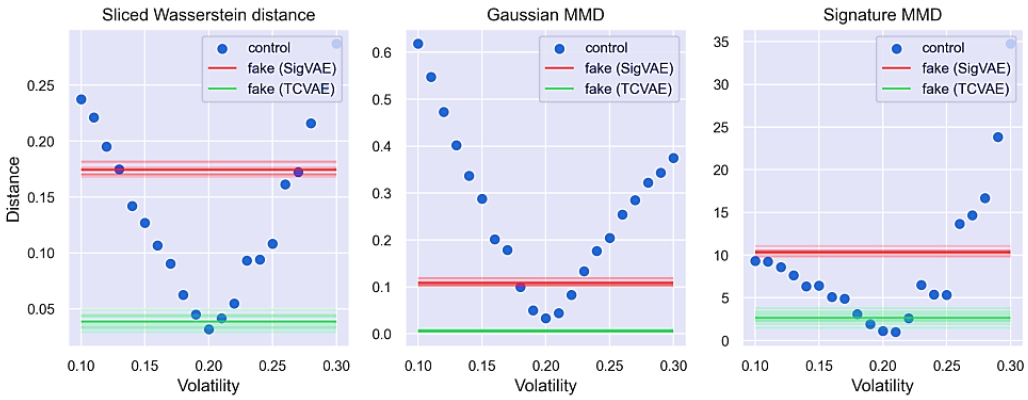

FIGURE 21: From left to right, we visualize the sliced Wasserstein distance, Gaussian MMD and signature MMD. The green (respectively red) lines illustrate distances between real paths of the Black-Scholes model and fake path generated from TC-VAE (respectively Sig-VAE); each line from a different random sees. The blue dots show the distances between real paths and control apths under different volatility levels [1].

Acciaio et al. (2024)		
Goal of the Paper	**Results**	**About Challenges**
To develop a TC-VAE for generating synthetic financial time-series data, enforcing a causality constraint on encoder/decoder networks and bounding the causal Wasserstein distance between real and generated distributions.	Generated paths closely match real ones (evaluated via sliced Wasserstein, Gaussian MMD, Signature MMD); preserves causal and temporal consistency.	• Although it handles time-series generation, it does not focus on a fully specified multidimensional risk-factor framework across distinct risk classes; • The treatment of intra-risk-factor correlation is not emphasized; • The work does not explicitly enforce arbitrage-free constraints.

TABLE 7: Acciaio et al. (2024).

Schwarz (2024)		
Goal of the Paper	**Results**	**About Challenges**
To introduce a probabilistic LSTM-based generative AI model for synthetic financial market time-series data that balances interpretability and performance.	The experiments show that the proposed model: • replicates the probability distribution of real market data; • handles non-linear relationships, market regime changes and external variables; • outperforms traditional ARMA-GARCH models in non-linear settings.	• It does not explicitly tackle a full multidimensional risk-class framework; • It does not emphasize modelling intra-risk-factor correlation structures; • The paper does not claim enforcement of full arbitrage-free constraints

TABLE 8: Schwarz (2024).

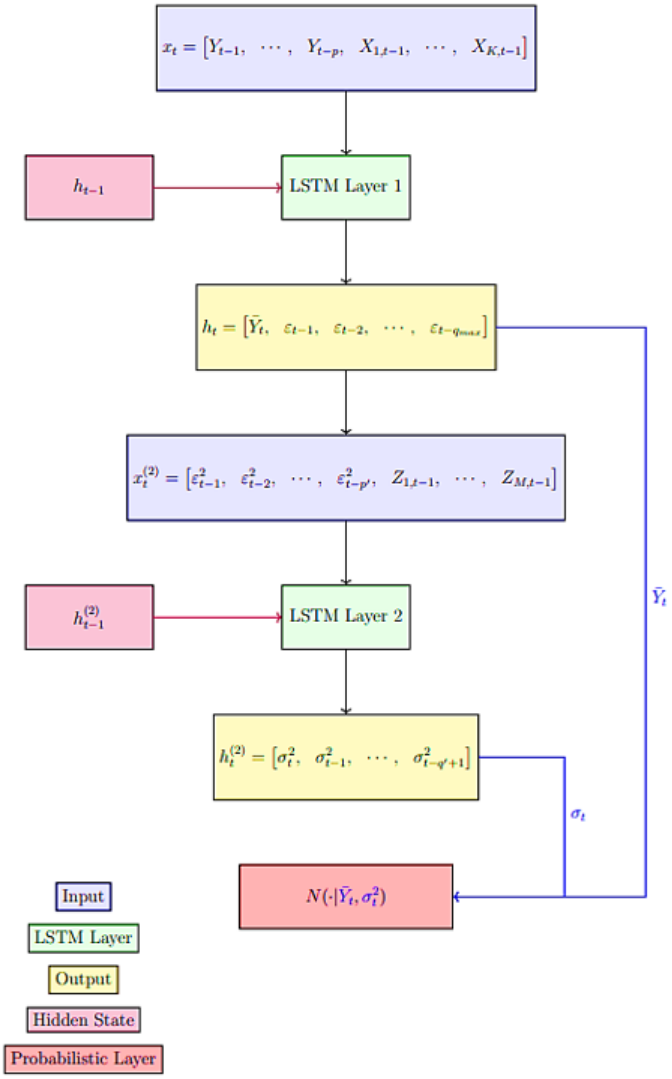

FIGURE 22: Schematic representation of the two LSTM layers feeding into a probabilistic output layer, here illustrated as a conditional Gaussian distribution, but any other distribution could be used [40]. Schwarz's Probabilistic LSTM framework is designed to generate synthetic time-series that capture rich statistical features, such as nonlinear dependencies, regime shifts, and volatility clustering, akin to traditional time series models, but embedded within a generative context.

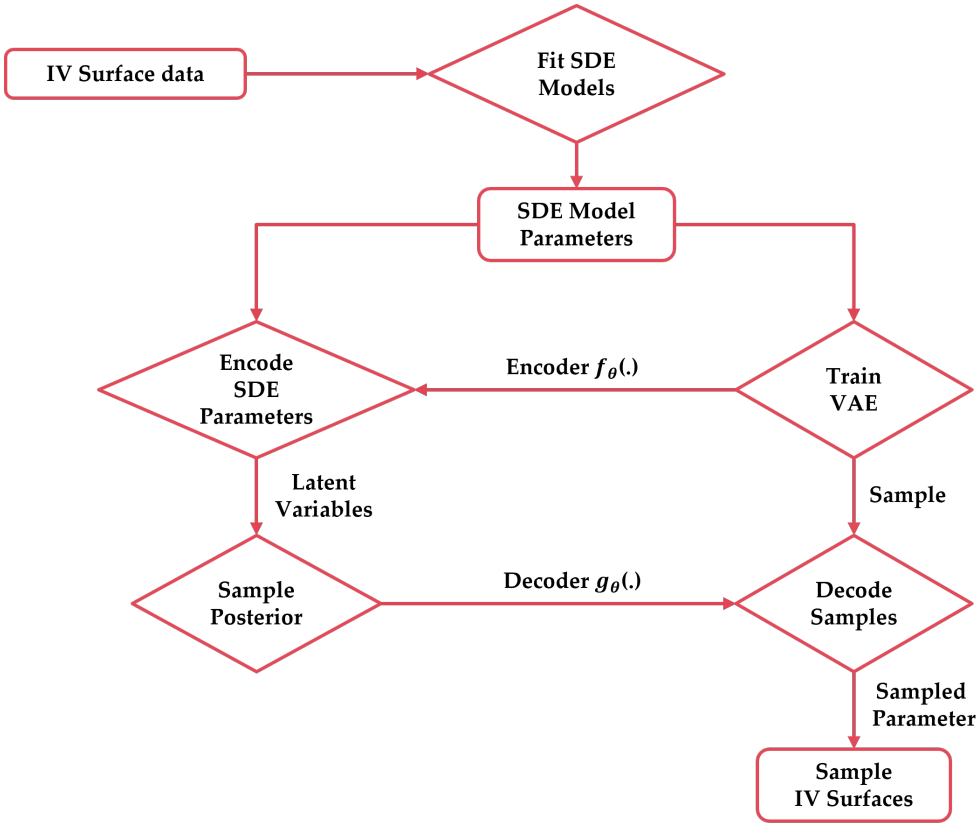

FIGURE 23: Flow chart of algorithm starting from raw surface data to generated surfaces through VAE [36].

Ning et al (2022)		
Goal of the Paper	**Results**	**About Challenges**
To propose a hybrid method combining a variational autoencoder (VAE) with continuous-time SDE-driven models (regime-switching & Lévy additive processes) to generate arbitrage-free implied volatility (IV) surfaces consistent with historical data.	The method demonstrates superior generative out-of-sample performance, producing IV surfaces free of static arbitrage and faithful to historical distributions.	• The paper does not explicitly address multidimensional risk-class scenario generation; • It does not deeply model intra-risk-factor correlation structures; • Free-arbitrage conditions are a core theme.

TABLE 9: Ning et al. (2022).

Arbitrage-Free Implied Volatility with VAE by Ning

No-arbitrage by-design approaches provide the most principled route toward economically consistent generative models. One effective strategy involves developing ad hoc architectures in which the model's output space is intrinsically restricted to arbitrage-free representations. A notable example is offered by Ning et al. (2022) [36], who propose a Variational Autoencoder (VAE) framework specifically *tailored to generate arbitrage-free implied volatility surfaces*. In their design, the decoder network produces the parameters of the **Stochastic Volatility Inspired** (SVI) functional form — an established arbitrage-free parametrization that ensures the absence of both butterfly and calendar — spread arbitrage by construction.
Rather than learning raw option prices or implied volatilities, the VAE captures a low-dimensional latent representation that is mapped to valid volatility surfaces through this parameterization. This architectural constraint guarantees that every generated surface satisfies no-arbitrage conditions without additional regularization or post-processing. The resulting model combines the generative flexibility of VAEs with financial consistency, making it particularly effective for risk management, model calibration, and scenario generation. The workflow of this approach is illustrated in Figure 23.

Calibrating Option Price and Volatility Surface Via Physics-informed Neural Network by Hyeon-ok

An even more rigorous form of the free-arbitrage by-design approach involves the use of Physics-Informed Neural Networks (PINNs), which integrate financial partial differential equations (PDEs) directly into the training objective. Unlike purely data-driven models, PINNs learn simultaneously from data and governing equations, ensuring that outputs adhere to the fundamental dynamics of pricing models. In the context of derivative valuation, the loss function can include the residual of the Black–Scholes PDE, compelling the network to approximate solutions that satisfy the equation across both spatial and temporal domains.
A representative application is presented by Hyeong-ok and Nam Sang-yoon (2023) [24], who employ a *PINN framework to jointly calibrate option prices and implied volatility surfaces*. By embedding the Black–Scholes dynamics directly into the loss function, their model learns arbitrage-free pricing relationships consistent with observed market data. This paradigm elegantly combines the data-driven adaptability of deep learning with the theoretical rigor of financial mathematics, and is increasingly used for high-dimensional pricing and calibration tasks. The network architecture employed for model training is illustrated in Figure 24.

Arbitrage-free Yield Curve and Bond Price Forecasting by Deep Neural Networks by Hyndman

Hyndman (2021) proposes a pre-fixing approach to *enforce arbitrage-free conditions within deep neural network frameworks for yield curve and bond price forecasting*. Rather than correcting arbitrage violations after generation, the model embeds projection-based regularization directly into the training process, ensuring that the outputs remain consistent with the theoretical constraints of fixed-income markets. By incorporating these financial structure conditions into the optimization objective, the network learns to produce yield curves that inherently satisfy no-arbitrage relationships across maturities. This method effectively integrates economic interpretability and machine learning flexibility, offering a robust mechanism for aligning data-driven forecasts with market-consistent dynamics. The resulting model demonstrates improved stability and financial coherence,

making it a promising direction for applications in risk management, pricing, and stress testing [25].

A Generative Model for Arbitrage-Free Implied Volatility Surfaces by Vuletic

Vuletić (2023) introduces *VolGAN*, a generative adversarial network specifically *designed for the synthesis of arbitrage-free implied volatility surfaces*. The framework adopts a post-fixing strategy, in which the model first generates raw volatility surfaces through a data-driven adversarial process, and subsequently applies a set of projection-based correction procedures to enforce financial consistency. These post-processing steps are aimed at minimizing arbitrage violations by projecting the generated outputs back into the admissible, arbitrage-free manifold. The model effectively combines the expressive capacity of GANs with the structural rigor required in quantitative finance, producing volatility surfaces that remain realistic, dynamically coherent, and consistent with market constraints. Beyond arbitrage elimination, VolGAN demonstrates the potential of generative models to capture complex dependencies within the implied volatility space, offering a flexible tool for applications such as option pricing, risk management, and scenario generation [43].
To conclude this section, Table 12 provides a concise summary of all the methodologies discussed, highlighting their main characteristics, achieved results and comparative advantages in the context of financial scenario generation.

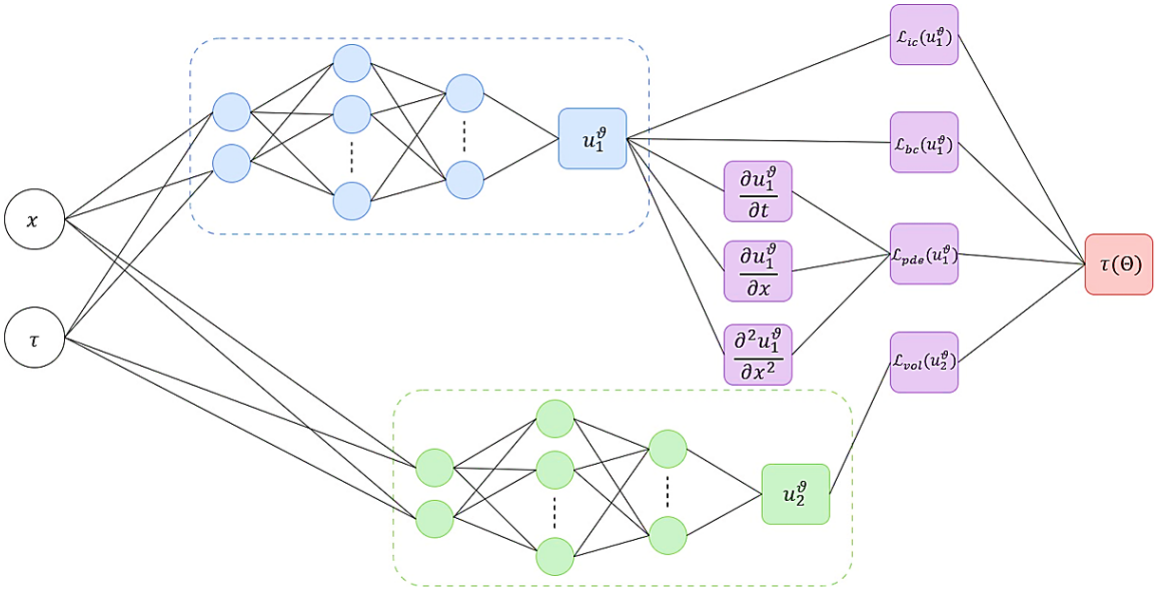

FIGURE 24: *Neural Network schema used by Hyeong-ok and Nam Sang-yoon [24], here x and τ are input data, each are moneyness scaled strike price and time to maturity. u_1^θ and u_2^θ are neural networks parameter. The input used for both the neural network is the same, the second neural network is used as a volatility parameter of the partial differential equation constructed through automatic differentiation for the first one.*

Hyeong-ok B. and Sang-yoon N. (2023)		
Goal of the Paper	**Results**	**About Challenges**
To develop a physics-informed neural network (PINN) framework that calibrates option prices and constructs a local volatility surface by embedding PDE constraints. No generative intent.	The model demonstrates successful estimation of option prices and constructs a local volatility surface consistent with market / model data.	• Does not address a full multidimensional risk-class scenario generation framework; • Does not focus in on intra-risk-factor correlation; • PDE constraints reduce arbitrage-type inconsistencies.

TABLE 10: Hyeong-ok B. and Sang-yoon N. (2023).

Hyndman (2021)		
Goal of the Paper	**Results**	**About Challenges**
To develop a deep-learning framework that forecasts yield curves and coupon-bond prices under the constraint of arbitrage-free regularization	The model shows that introducing an arbitrage-penalty improves forecasting performance.	• The paper does not address multidimensional risk-class scenario generation; • It does not deeply model intra-risk-factor correlation structures; • It introduces arbitrage-free regularization.

TABLE 11: Hyndman (2021).

Criteria / Model	GAN	VAE	GMM	RNN
Intra-class Heterogeneity	Has some difficulties on univariate/multivariate time series	Curve/structure natively supported	It can deal with spot, curves, surfaces and tensors	Limited, requires adaptations
Inter-class Dependence	Complex for homogeneous assets	Only with multi-latent architectures	Possible, but requires normalization	Complex for single-type time series
Arbitrage-Free Scenario	It can be applied externally	It can be applied externally	Arbitrage-free approach supported	It can be applied externally
Computational Cost	High	High	Low	High
Interpretability	Deep network, hard to read	Latent space interpretable, but not directly	GPC interpretable, similar to PCA	Black-box

TABLE 12: Comparison of the main models (GAN, VAE, GMM, and RNN) according to key analytical criteria.

Open Challenges

While *Generative AI* presents a powerful alternative, its application to financial scenario generation is far from straightforward. The transition from academic "proof-of-concept" to a robust, production-grade system introduces several formidable chal-

Vuletic (2023)		
Goal of the Paper	**Results**	**About Challenges**
To introduce a GAN-based generative model (VolGAN) for simulating arbitrage-free implied volatility surfaces (IVS), trained on historical IVS and underlying price data.	Demonstrated that the model learns covariance structure of implied vol co-movements, generates realistic dynamics for the underlying index and VIX.	• The paper does not address multidimensional risk-class scenario generation; • It does not deeply model intra-risk-factor correlation structures; • It targets specifically the theme of no-arbitrage for IV surfaces.

TABLE 13: Vuletic (2023).

TABLE 14: Comparison of main conditional generative models for financial scenario generation, highlighting their integration mechanisms, strengths, and challenges.

lenges that must be addressed.

First, the "black-box" nature of many generative architectures (GANs, VAEs, Transformers) can be a primary obstacle for their adoption in risk management. For regulatory compliance and internal validation, model interpretability is not optional. It is not trivial to validate why a model produced, for example, a specific tail event. Often this low, if not lack of, transparency can push for the choice of more interpretable alternatives like *Gaussian Mixture Models* (GMMs), which allow their factors (*Gaussian Principal Components*) to be analyzed, much like traditional **PCA**.

Second, the calibration and training complexity is not to be underestimated. These models are not "plug-and-play." Training deep generative networks can be notoriously unstable (e.g., **mode collapse** in GANs) and computationally expensive. The process requires significant data engineering and hyperparameter tuning, rivaling (and sometimes exceeding) the complexity of calibrating sophisticated classical models. On the other hand, once a model is calibrated, the scenario generation can be extremely fast, hence exploitable also for (at least near-to) real-time simulations.

Third, the scope of existing research is often too narrow for a real-world financial institution. The literature frequently presents models calibrated on a single asset (e.g., S&P 500) or a single, isolated risk class (e.g., interest rate curves). A bank or insurer, however, must model the entire market — the complex, joint distribution of multiple, heterogeneous risk classes (equities, rates, FX, credit, and commodities) simultaneously. Scaling these models to such high dimensions while preserving delicate inter-class correlations is a massive, and largely unsolved, technical leap. To do so, novel literature is trying to blend classical approaches with AI, as in the "Copula Variational LSTM" propsed by Xu and CAo [44].

Fourth, the enforcement of financial consistency, most notably the absence of arbitrage, is not an inherent property of neural networks. Without specialized architectures (like *Autoencoder Market Models* (AEMM)) or complex, "physics-informed" loss functions, a model may produce statistically plausible but economically invalid scenarios, such as violations of put-call parity or arbitrageable yield curves.

Finally, one important and still largely overlooked dimension in the field of AI-driven

scenario generation for finance is **conditional generation**, the ability to generate future paths or market states that are explicitly dependent on given initial conditions, macroeconomic assumptions, or market configurations. Despite the increasing sophistication of generative models, surprisingly little attention has been paid in the literature to this conditional setting, especially within the context of multi-asset, risk-sensitive financial forecasting. Yet, **conditional generation** has significant practical relevance: it allows risk managers and decision-makers to ask meaningful "what-if" questions, such as how markets might evolve given a credit spread widening, a spike in oil prices, or a central bank intervention. It enables the production of coherent and targeted scenarios that reflect the distributional implications of specific conditions, rather than relying on unconditional or average-case dynamics.

In practice, conditional generation provides a bridge between supervised inference and generative simulation. It gives institutions the ability to simulate plausible market evolutions under regulatory stress scenarios, or under internally defined stress assumptions, in a data-driven and probabilistically consistent manner. This ability is particularly useful in **forward-looking risk assessments** such as *ICAAP*, *CCAR*, or *Solvency II* projections, where the scenarios must not only reflect plausible tail behaviors but also remain consistent with certain macroeconomic narratives or initial market configurations. Conditional generation thus enhances scenario relevance and realism while retaining statistical diversity.

Different neural network architectures can be adapted to tackle this problem, although none of them has been explicitly designed for this task in standard financial literature. One natural starting point is the **Conditional VAE** (CVAE), where the conditioning variable is passed both to the encoder and the decoder, allowing the latent space to be learned in relation to a specified set of market conditions. This setup enables the generation of scenarios that are statistically consistent with observed distributions, yet anchored to specific assumptions or shocks. CVAEs are particularly well-suited when the goal is to map a structured condition — such as a macroeconomic vector or a set of factor values — to a wide distribution of possible outcomes.

In the case of GAN, the extension to conditional generation is achieved through **Conditional GANs** (cGANs), where both the generator and the discriminator are fed with conditioning information. This setup forces the generator to learn how to produce samples that are not just realistic, but also coherent with the provided condition. While powerful in theory, training stability and mode collapse can become even more acute in the conditional case, especially when dealing with high-dimensional financial data with sparse conditional examples. Nonetheless, conditional GANs offer a valuable framework for generating scenario distributions that reflect specific stress drivers or forward-looking views.

Recurrent neural networks (RNNs) and particularly **Long Short-Term Memory networks** (LSTMs) can also be adapted for conditional generation by incorporating the conditioning input into the initial state or by concatenating the condition with each input timestep. For example, an LSTM-based model might take as input a time series of historical prices along with a constant conditioning vector (e.g., a shock to interest rates), and generate a sequence of future values that unfold under that assumption. This approach is particularly appealing in time-series applications where temporal coherence is critical. However, most LSTM-based models in finance are trained unconditionally, missing the opportunity to leverage structured scenario guidance.

Another relevant but underutilized approach is the mixture density network (MDN) or conditional **Gaussian Mixture Model** (GMM), where the model outputs the parameters of a distribution conditional on input features. Such models are natu-

rally capable of representing multi-modal outcomes under specific initial states or shocks, and they allow a high degree of interpretability. However, their use in generative financial modeling, at the best of our knowledge, it remains still extremely studied.
The scarcity of literature on conditional generation in financial contexts is perhaps due to the complexity of obtaining consistent, labeled conditioning data, or the difficulty in defining proper training objectives under partial observability. Nonetheless, as the use of machine learning in risk management matures, the need for scenario generation tools that respond to precise assumptions or stress conditions will only grow. Conditional generation offers a principled and flexible solution to this need, making it a critical area for future development.

Our Contribution

As presented in the previous chapters, there is no single, universally accepted technique for generating synthetic financial data. The landscape is a complex trade-off between the interpretability and proven effectiveness for most cases of classical models and the high-fidelity, data-driven power promised by AI, especially GenAI. Each approach comes with its own assumptions, strengths, and limitations. For this reason, this paper's primary contribution is not to propose yet another novel architecture, but to provide a methodological blueprint for the selection, hybridization, and calibration of these tools, specifically tailored for a production setting in a financial institution.
The successful transition from theory to a production-ready system is a complex task requiring deep domain expertise. Our contribution is to provide the specialized methodological support to design, implement, test, and govern bespoke solutions at each stage, ensuring a blend of AI innovation and financial-grade robustness.

Scenario Classification and Data Framework Design

- **The Core Task**: this initial stage involves transforming vast, heterogeneous market data into a curated, model-ready "Scenario Catalog".
- **Our Value**: this is not simple "tagging". We partner with clients to architect the data framework. Our expertise helps:
  - *Define the Problem Space*: we identify the relevant risk factors, historical regimes (e.g., "inflation-shock", "flight-to-quality"), and "semantic contexts" required for a specific institutional need (e.g., Market Risk VaR, CCR, etc.).
  - *Data Augmentation Strategy*: we design and implement the preprocessing pipeline, addressing missing data, different time scales, and heterogeneous data structures (e.g., curves, surfaces).

Model Selection, Hybrid Design, and Calibration

- **The Core Task**: selecting, building, and training a generative model that is both statistically powerful and financially consistent.
- **Our Value**: this is the primary technical hurdle where generic approaches fail. We design and fine-tune bespoke hybrid models by navigating the critical trade-offs:
  - *Model Selection*: we can guide the choice between different models and architectures, navigating the complexities and trade-offs among high-fidelity, interpretable latent-space, and data-efficient, transparent factor models.
  - *Hybridization*: we can architect solutions that blend classical approaches with AI, such as using

a VAE to learn a non-linear representation of a yield curve, and then applying a classical (and auditable) stochastic process to its latent factors.

– *Constraint Enforcement*: our core expertise in financial modeling can help in injecting financial domain knowledge directly into the model — e.g., arbitrage-free conditions and causality constraints to ensure the model's outputs are not just plausible, but economically valid and robust for dynamic problems).

Validation, Governance, and Integration

- **The Core Task**: ensuring the calibrated model is trustworthy, maintainable, and usable within the bank's existing infrastructure.
- **Our Value**: a model is only production-ready once it is governed. We provide a complete validation and governance framework:
  - *Rigorous Testing*: we can move beyond standard loss metrics to perform comprehensive statistical validation of the synthetic data, ensuring it correctly reproduces all critical stylized facts (heavy tails, volatility clustering) and, most importantly, preserves complex cross-asset tail dependencies.
  - *Governance & MLOps*: we provide the clear documentation required for internal model validation teams and regulators. We design the MLOps pipeline for efficient retraining, monitoring, and versioning of the generative engine.
  - *Integration*: we support the final integration of the generative model's output into the institution's existing risk engines, ensuring the solution is scalable, auditable, and truly operational.

Conclusion and Future Works

This work has reviewed the evolution of financial scenario generation, from the foundational classical techniques to the Generative AI models entering the scene. Traditional approaches, such as historical simulation and Monte Carlo methods, remain widely used and have been consistently improved in the decades to overcome constraints on possibly reductive assumptions. To tackle differently critical market stylized facts like non-linear dependencies, heavy tails, and volatility clustering, and blend-in unstructured data, the academia and the industry started to look for more flexible, data-driven frameworks.

Our analysis detailed leading GenAI alternatives — GANs, VAEs, and GMMs with their variants — along with LSTMs, and surveyed key literature presenting their potential. These models promise to offer a powerful new toolkit for data augmentation and for generating high-fidelity scenarios by learning complex market structures directly from data. The primary advantage of these AI-driven approaches lies in their *flexibility* and *potential for customization*. For instance, an Autoencoder Market Model (AEMM) can provide a parsimonious, data-driven representation of yield curves, while a GMM offers an interpretable factor model via Gaussian Principal Components (GPCs).

However, our review concludes that these advanced models, even if worthy to be studied deeper and tested in use cases with financial institutions, are not yet "production-ready" off-the-shelf. Critical and practical challenges remain. The "black-box" nature of many deep learning models poses a significant hurdle for internal validation and regulatory approval; we acknowledge indeed that these approaches can be used mainly, if not only, for managerial purposes.

Furthermore, the calibration and training complexity of these models is substantial, often rivaling classical Monte Carlo in its demand for expert tuning and computational power.
Most importantly, two key gaps seems to persist:

1. **Scalability**: the majority of existing research remains narrowly focused on single assets or risk classes. A production-ready system must be capable of modeling the joint distribution of all heterogeneous, multi-asset risk classes in a bank's portfolio.

2. **Financial Consistency**: ensuring that generated scenarios are, for example, arbitrage-free is a non-trivial requirement that demands specialized architectures or complex loss functions.

Looking ahead, we see the future not in a pure AI-only solution, but in *hybrid approaches* that combine the statistical rigor of classical models with the adaptive power of GenAI. In this context, Large Language Models (LLMs) are poised to play a crucial, dual role. First, we see their primary strength not in generating the quantitative scenarios themselves, but in managing unstructured data, processing news, sentiment, and geopolitical reports to create the rich, qualitative inputs for conditional scenario generation. Second, LLMs can serve as a powerful natural-language interface, providing a "chat-like" entry point that allows risk managers and business people to translate financial intuition into robust, calibrated simulations. Ultimately, the path to adoption relies on a practical framework for selection, hybridization, and governance.
By carefully integrating these new technologies, financial institutions can leverage the power of GenAI to build scenario generation tools that are not only data-driven and realistic but also transparent, operationally feasible, and trusted for real-world risk management and business applications.

References

[1] **Acciaio B., Eckstein S. and Hou S.** *Time-Causal VAE: Robust Financial Time Series Generator.* November 2024.

[2] **Atkins A., Niranjan M. and Gerding E.** *Financial news predicts stock market volatility better than close price*, The Journal of Finance and Data Science, Vol. 4, Issue 2, pp. 120-137, June 2018.

[3] **Bagattini G. et al.** *Leverage Large Language Models in Finance: Pathways to Responsible Adoption*. ESMA ILB, The Alan Turing Institute, 2025.

[4] **Baillie R. T., Bollorslev T. and Mikkelsen H. O.** *Fractionally integrated generalized autoregressive conditional heteroskedasticity*. Journal of Econometrics, Vol. 74, pp. 3-30, September 1996.

[5] **Barone-Adesi G. and Giannopoulos K.** *Non parametric VaR Techniques. Myths and Realities*. July 2001.

[6] **Bollerslev T.** *Generalized autoregressive conditional heteroskedasticity*. Journal of Econometrics, Vol. 31, pp. 307-327, April 1986.

[7] **Bonollo M., Damato V. and Luce F.** *FHS-VaR and Financial Risks*, MAF, 2025.

[8] **Brigo D. and Mercurio F.** *Interest Rate Models — Theory and Practice*. 2006.

[9] **Cheng Z., Sjarif N.N.A. and Roslina I.** *Deep learning models for price forecasting of financial time series: A review of recent advancements: 2020–2022*. September 2023.

[10] **Chevyrev I. and Oberhauser H.** *Signature moments to characterize laws of stochastic processes*. Journal of Machine Learning Research, Vol. 23, pp. 1-42, October 2018.

[11] **Demarta S. and McNeil A.** *The t copula and related copulas*. May 2004.

[12] **Divate M. et al.** *Harnessing LLMs for Financial Forecasting: A Systematic Review of Advances in Stock Market*. International Journal for Research in Applied Science and Engineering Technology, Vol. 12, pp. 1101-1105, November 2024.

[13] **Dogariu M., Stefan L. and Boteanu B. A.** *Generation of Realistic Synthetic Financial Time-series*. Association for Computing Machinery, March 2022.

[14] **Embrechts P., Kluppelberg C. and Mikosch T.** *Modelling Extremal Events for Insurance and Finance*. February 1997.

[15] **Embrechts P., McNeil A. and Straumann D.** *Correlation and Dependence in Risk Management: Properties and Pitfalls. In Risk Management: Value at Risk and Beyond.* 2002.

[16] **Engle R. F.** *Dynamic Conditional Correlation: A Simple Class of Multivariate GARCH Models*. Journal of Business and Economic Statistics, Vol. 20, pp. 339-350, January 2012.

[17] **Flaig S. and Junike G.** *Scenario generation for market risk models using generative neural networks*. Risks, August 2023.

[18] **Fritzsh S. and Timphus M.** *Marginals versus copulas: Which account for more model risk in multivariate risk forecasting?*. Journal of Banking & Finance, Vol. 158, January 2024.

[19] **Goodfellow I. J. et al.** *Generative Adversarial Networks*. June 2014.

[20] **Gretton A. et al.** *A kernel two-sample test*. The Journal of Machine Learning

Research, Vol. 13, pp. 723-773, December 2012.

[21] **Harrison J. and Pliska S.** *Martingales and stochastic integrals in the theory of continuous trading*. Stochastic Processes and their Applications, Vol. 11, pp. 215-260, August 1981.

[22] **Haugh M.** *An introduction to copulas*. 2016.

[23] **Hochreiter S.** *Long Short-Term Memory*. Neural Computation, Vol. 9, pp. 1735-1780, 1997.

[24] **Hyeong-ok B. and Sang-yoon N.** *Calibrating option price and volatility surface via physics-informed neural network*. The Graduate School, Ajou University, February 2023.

[25] **Hyndman C.** *Arbitrage-free yield curve and bond price forecasting by deep neural network*. September 2021.

[26] **Jolliffe I. T.** *Principal Component Analysis*. 2002.

[27] **Kingma D. P. and Welling M.** *An Introduction to Variational Autoencoders*. Foundations and Trends in Machine Learning, Vol. 12, pp. 307-392, December 2019.

[28] **Liu Y. and Yun L.** *The Development of Large Language Models in the Financial Field*. Proceedings of Business and Economic Studies Vol. 8, pp. 49-54, April 2025.

[29] **Longin F.** *From Value at Risk to Stress Testing: The Extreme Value Approach*. Journal of Banking & Finance, Vol. 24, pp. 1097-1130, July 2000.

[30] **Lu C. and Sester J.** *Generative model for financial time series trained with MMD using a signature kernel*. July 2024.

[31] **Mandelbrot B.** *The variation of certain speculative prices*. The Journal of Business, Vol. 36, pp. 394-419, December 2009.

[32] **McLachlan G. J.** *Finite Mixture Models*. June 2019.

[33] **McNeil A. and Frey R.** *Estimation of tail-related risk measures for heteroscedastic financial time series: an EVT approach.* June 2000.

[34] **Mirza M. and Osindero S.** *Conditional Generative Adversarial Nets*. November 2014.

[35] **Mohammadabadi S. et al.** *A Survey of Large Language Models: Evolution, Architectures, Adaptation, Benchmarking, Applications, Challenges, and Societal Implications*. Electronics, Vol. 14, September 2025.

[36] **Niang B., Jaimunga S. and Bergeron M.** *Arbitrage-Free Implied Volatility Surface Generation with Variational Autoencoders*. January 2022.

[37] **Reynolds D.** *Gaussian Mixture Models*. Encyclopedia of Biometrics, pp. 659–663, 2009.

[38] **Rizzato M. et al.** *Generative Adversarial Networks applied to synthetic financial scenarios generation*. Physica A: Statistical Mechanics and its Applications, May 2024.

[39] **Schmitt T. A. et al.** *Non-stationarity in financial time series: Generic features and tail behavior*. May 2013.

[40] **Schwarz C.** *Interpretable GenAI: Synthetic Financial Time Series Generation with Probabilistic LSTM*. June 2024.

[41] **Tunnicliffe G.** *Time Series Analysis: Forecasting and Control*. Journal of Time Series Analysis, Vol. 37, March 2016.

[42] **Vaswani, A. et al.** *Attention Is All You Need*. August 2023.

[43] **Vuletić M. and Cont R.** *VolGAN: A Generative Model for Arbitrage-Free Implied Volatility Surfaces*. Applied Mathematical Finance, Vol. 31, pp. 203–238, December 2023.

[44] **Xu J. and Cao L.** *Copula Variational LSTM for High-dimensional Cross-market Multivariate Dependence Modeling*. May 2023.

Sitography

[45] **Arbitrage.** *Website*.

[46] **EIOPA. YE2019 Comparative Study on Market & Credit Risk Modelling, 2021a.** *Website*.

Annex

This chapter presents a structured overview of these methods, highlighting both classical techniques and more recent data-driven frameworks. From traditional parametric models to deep generative networks, the landscape of scenario generation has grown increasingly rich and diverse — reflecting the complex, nonlinear, and high-dimensional nature of financial markets. Understanding the underlying principles, assumptions, and trade-offs of each approach is essential for selecting the most appropriate tool for a given risk management or forecasting application.

Generative Adversarial Network

Generative Adversarial Networks (GANs), introduced by Ian Goodfellow et al. in 2014 [19], represent a groundbreaking framework for training generative models. GANs leverage game theory to model a system in which two neural networks — a generator and a discriminator — are trained simultaneously in a minimax game. This approach enables the generator to learn the underlying data distribution and produce highly realistic synthetic data, without requiring explicit probability models or likelihood functions.
At the core of GANs is the interaction between two models:

- Generator (G): Takes as input a random noise vector $z \sim p_z(z)$ (typically sampled from a uniform or Gaussian distribution) and transforms it into a sample $G(z)$ that mimics the data distribution;
- Discriminator (D): Receives either a real data sample $x \sim p_{data}(x)$ or a generetad sample $G(z)$ and outputs the probability that the input is real (i.e., from the true data distribution).

The generator and discriminator are trained in opposition: the generator aims to fool the discriminator, while the discriminator aims to correctly classify real versus generated data, in Figure 25 can be seen an example structure of a GAN. This dynamic is formalized as a two-player minimax game:

$$\min_G \max_D V(D,G) = \mathbf{E}_{x\sim p_{data}(x)}[\log D(x)] + \mathbf{E}_{z\sim p_z(z)}[\log(1 - D(G(z)))].$$

Here, $D(x)$ is the probability assigned by the discriminator that x is a real sample, and $G(z)$ is the synthetic sample produced by the generator from noise z. Goodfellow et al. demonstrated that under optimal conditions, the generator recovers the true data distribution p_{data}. Specifically, for a fixed generator G, the optimal discriminato r $D*(x)$ is:

$$D*(x) = \frac{p_{data}(x)}{p_{data}}, \tag{59}$$

where p_g is the model distribution induced by the generator G. Plugging $D*$ back into the value function yields:

$$\begin{aligned} C(G) &= \max_D V(D,G) = \mathbb{E}_{x\sim p_{\text{data}}}\big[\log D^*(x)\big] + \mathbb{E}_{z\sim p_z(z)}\big[\log(1 - D^*(G(z)))\big] \\ &= -\log 4 + 2\cdot \text{JSD}(p_{\text{data}} \parallel p_g). \end{aligned} \tag{60}$$

This shows that minimizing the GAN objective is equivalent to minimizing the Jensen–Shannon divergence (JSD) between the real and generated data distributions — a symmetric, bounded

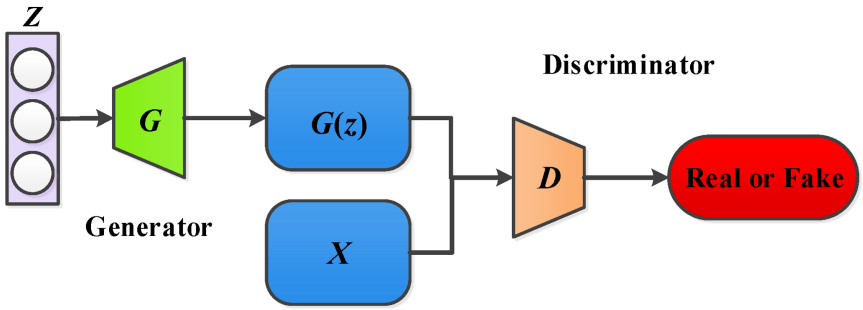

FIGURE 25: In the figure above there is an example of the architecture of a GAN. At the beginning a noised input is fed to the generator to produce fake sample $G(z)$. Both the real data X and the generated sample are then given to the discriminator D, which learns to distinguish real from fake.

divergence measure.

GAN Extension: Conditional GAN

CGANs extend traditional GANs by introducing a conditioning vector y that influences both the generator and the discriminator. This allows data x to be generated in a conditional way, leading to more targeted and controllable outputs. The objective function becomes:

$$\min_G \max_D V(D,G) = \mathbb{E}_{x\sim p_{\text{data}}(x)}\big[\log D(x,y)\big] + \mathbb{E}_{z\sim p_z(z),\, y\sim p_{\text{data}}}\big[\log(1 - D(G(z,y),y))\big]. \tag{61}$$

- The generator G receives both noise z and condition y as input, typically concatenated before the upsampling stages;
- The discriminator D evaluates pairs (x,y) versus $(G(z,y),y)$.

Gan Extension: Time GAN

TimeGAN (Time-series Generative Adversarial Network) is a powerful model designed to generate realistic synthetic time series data while preserving both temporal dynamics and realistic feature distributions. Traditional GANs struggle with time series because they do not model temporal dependencies explicitly. TimeGAN addresses this by combining supervised and unsupervised learning in a unified framework that captures both: step-wise temporal dynamics and latent representations of multivariate data.
TimeGAN blends components of a GAN, an autoencoder and a recurrent network (usually LSTM). Its architecture consists of:

- Embedding Network: learns a latent representation h of the original data x, typically via a RNN encoder;
- Recovery network: reconstruncts the input from embedding $\hat{x} = R(h)$;
- Generator G: generates latent sequences from random noise and conditions;
- Discriminator D: tries to distinguish real vs synthetic latent sequances;
- Supervisor S: predict the next latent state from the current one, helping model temporal transitions.

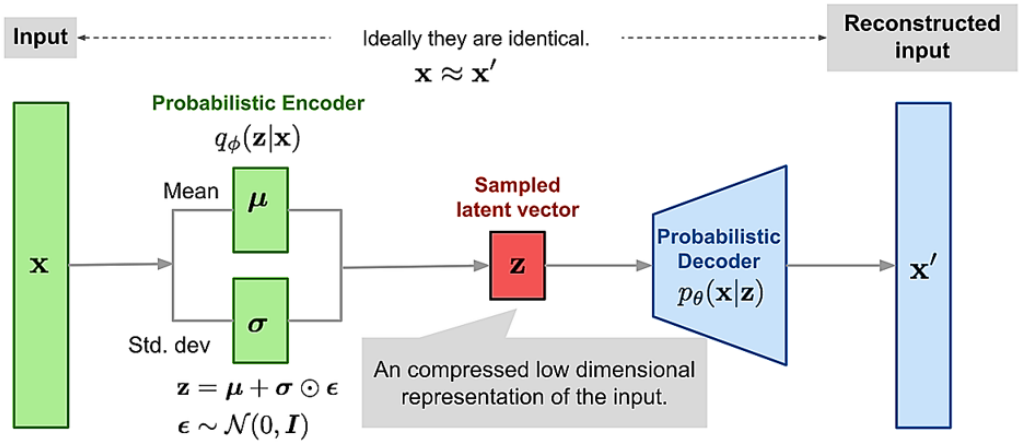

FIGURE 26: Example of a Variational Autoencoder (VAE). The input x is encoded into a mean μ and a standard deviation σ, from which a latent vector z is sampled as a compressed low-dimensional representation of x. A probabilistic decoder then reconstructs the data, generating x'.

And the training includes three losses, Reconstruction, Supervised and Adversaral Loss as follow:

$$\mathcal{L}_{tot} = \mathcal{L}_{rec} + \mathcal{L}_{sup} + \mathcal{L}_{adv} = \mathbf{E}[||x - \hat{x}||^2] + \mathbf{E}[||h_{t+1} - \hat{h}_{t+1}||^2] + \min_G \max_D \mathbf{E}[\log D(h)] + \mathbf{E}[\log(1 - D(\hat{h}))]. \tag{62}$$

Variational Autoencoder

A Variational Autoencoder (VAE) is a type of generative model in machine learning that learns a compressed representation of data while also being able to generate new, synthetic data samples. A VAE consists of two neural networks:

- Encoder/Inference model $q_\phi(z|x)$: encodes input x to a latent distribution;
- Decoder/Generative model $p_\theta(x|z)$: reconstructs x given latent z.

This setup implements amortized inference by sharing encoder parameters across data points — far more efficient than traditional per-sample optimization. In Figure 26 is possible to see an example of the architecture of a VAE.
Training maximizes the ELBO, which serves as a tractable lower bound on the data log-likelihood $\ln p_\theta(x)$:

$$\mathcal{L}(\theta, \phi; x) = \mathbf{E}_{z \sim q_{phi}(z|x)}[\ln p_\theta(x|z)] - D_{KL}(q_\phi(z|x)||p(z)). \tag{63}$$

- Reconstruction term encourages $p_\theta(x|z)$ to match true x;
- KL divergence regularizes latent z to stay close to the prior $p(z)$, typically $\mathcal{N}(0, I)$.

Maximizing ELBO achieves a trade-off between accurate reconstructions and latent space regularization.
To backpropagate through stochastic sampling, the reparameterization trick is used:

$$z = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon, \epsilon \sim \mathcal{N}(0, I). \tag{64}$$

This transforms z-sampling into a differentiable operation, allowing efficient gradient estimation with encoder parameters.

VAE Extension: Conditionale VAE

Conditional Variational Autoencoders (CVAEs) are powerful models for learning conditional generative processes. They extend standard VAEs by introducing a conditioning variable y, allowing generation of data x conditioned on labels, attributes, or context. While the idea is conceptually simple, practical challenges arise in training CVAEs effectively. In this chapter, we revisit CVAEs through the lens of Christopher Beckham's analysis, offering a theoretically sound and empirically motivated perspective on their structure, training, and trade-offs.
At the core of CVAEs lies the Evidence Lower Bound (ELBO), which becomes:

$$\log p_\theta(x|y) \geq \mathbf{E}_{q_\phi(z|x,y)}[\log p_\theta(x|z,y)] - D_{KL}(q_\theta(z|x,y)||p(z|y)). \tag{65}$$

A key modeling decision in CVAEs is whether to use a conditional prior $p(z|y)$ or a fixed, indipendent prior $p(z)$. The dependent prior allows the latent space to shift in a way that reflects the semantics of y, potentially enriching expressiveness.

VAE Extension: Time VAE

The Time-Variational Autoencoder (TimeVAE) framework combines probabilistic latent variable modeling with temporal encoding mechanisms, enabling expressive and interpretable time-series generation. TimeVAE can be viewed as an extension of the Variational Autoencoder (VAE) that integrates temporal models like RNNs, GRUs, or transformers into the encoder–decoder architecture. This allows the model to learn not only latent structure but also temporal dynamics of sequences.
The model is trained to maximize the ELBO:

$$\mathcal{L}_{ELBO} = \sum_{t=1}^{T} \mathbf{E}_{q_\phi(z_t|x_{\leq t},z_{<t})}[\log p_\theta(x_t|z_t)] - D_{KL}(q_\phi(z_t|x_{\leq t},z_{\leq t})||p_\theta(z_t|z_{<t}). \tag{66}$$

Recurrent Neural Network

In many real-world applications such as natural language processing, time-series forecasting, and speech recognition, data is inherently sequential. Traditional feedforward neural networks are inadequate for modeling such sequences, as they assume all inputs are independent of each other. Recurrent Neural Networks (RNNs) address this limitation by introducing cycles within the network architecture, enabling information to persist across time steps.
An RNN processes input sequences one element at a time, maintaining a hidden state that captures information from previous inputs. Given a sequence $x = (x_1, x_2, ..., x_T)$, the RNN updates its hidden state h_t and outputs y_t at each time step using the following equations:

$$h_t = \phi(W_{hh}h_{t-1} + W_{xh}x_t + b_h); \tag{67}$$

$$y_t = W_{hy}h_t + b_y. \tag{68}$$

Here, ϕ is typically a non-linear activation function such as tanh or ReLU. The key idea is that the hidden state h_t carries temporal information across the sequence.
To address the shortcomings of vanilla RNNs, Hochreiter and Schmidhuber (1997) proposed the Long Short-Term Memory network. LSTMs introduce a memory cell and gating

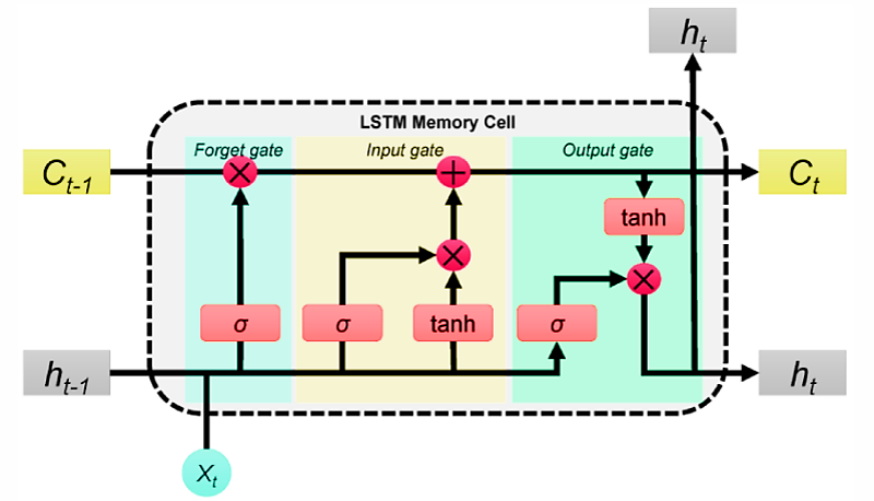

FIGURE 27: Illustration of a Long Short-Term Memory (LSTM) unit. It contains three main gates: the forget gate, the input gate, and the output gate, which regulate the flow of information and help capture long-term dependencies in sequences.

mechanisms that regulate the flow of information, enabling the network to selectively retain or discard information over long periods.
An LSTM unit consists of the following gates:

- Forget gate f_t: Decides what information to discard from the cell state;
- Input gate i_t: Decides what new information to store;
- Cell candidate $\tilde{C}_t$: Potential values added to the cell state;
- Output gate o_t: Determines the output and hidden state.

In Figure 27 is showed an example architecture of a LSTM.
The computations are as follows:

$$f_t = \sigma(W_f[h_{t+1}, x_t] + b_f; \tag{69}$$

$$i_t = \sigma(W_i[h_{t-1}, x_t] + b_i); \tag{70}$$

$$\tilde{C}_t = \tanh(W_C[h_{t-1}, x_t] + b_C); \tag{71}$$

$$C_t = f_t \odot C_{t-1} + i_t \odot \tilde{C}_t; \tag{72}$$

$$o_t = \sigma(W_o[h_{t-1}, x_t] + b_o; \tag{73}$$

$$h_t = o_t \odot \tanh C_t, \tag{74}$$

where σ is the sigmoid activation function, $\odot$ denotes element-wise multiplication, C_t is the cell state at time t and h_t is the hidden state output at time t.
LSTM networks are capable of learning dependencies over long time horizons by preserving gradients and maintaining stable learning dynamics. They have become the de facto standard in sequence modeling tasks prior to the rise of Transformer-based architectures.

Gaussian Mixture Models

Gaussian Mixture Models (GMMs) provide a flexible, unsupervised method to model such heterogeneity by assuming that data are generated from a mixture of multiple Gaussian

distributions. Unlike k-means clustering, GMMs offer a soft clustering approach, assigning probabilities for a data point to belong to each cluster.
A Gaussian Mixture Model defines the probability density function (pdf) of a data point $\mathbf{x} \in \mathbf{R}^D$ as a convex combination of K Gaussian components:

$$p(\mathbf{x}) = \sum_{k=1}^{K} \pi_k \mathcal{N}(\mathbf{x}|\mu_k, \mathbf{\Sigma_k}), \tag{75}$$

where:

- $\pi_k \in [0, 1]$ are the mixing coefficients, such that $\sum_{k=1}^{K} \pi_k = 1$;
- $\mu_k \in \mathbf{R}^D$: mean of the $\mathbf{k}^{th}$ Gaussian;
- $\Sigma_k \in \mathbf{R}^{DxD}$: covariance matrix of the $\mathbf{k}^{th}$ Gaussian;
- $\mathcal{N}$;$(\mathbf{x} \mid \mu_k, \mathbf{\Sigma_k})$: multivariate Gaussian PDF.

GMM can be viewed ad a latent variable model where each data point $\mathbf{x}_i$ is associated with a latent cluster label $z_i \in \{1, ..., K\}$. The generative process firstly sample a cluster index $z_i \sim Categorical(\pi_1, ..., \pi_K)$ and the sample the data point $\mathbf{x}_i \sim \mathcal{N}(\bar{}_{z_i}, \mathbf{\Sigma}_{z_i})$.

A Common Collateral Pooling Arrangement for Corporate Lending

About the Author

Antonio Castagna:
Founder and Managing Partner
Antonio Castagna is currently managing partner and founder of the consulting company iason. He previously was in Banca IMI, Milan, from the 1999 to 2006: there, he first worked as a market maker of cap/floor's and swaptions; then he set up the FX options desk and ran the book of plain vanilla and exotic options on the major currencies, being also responsible for the entire FX volatility trading. He started his carrier in the investment banking in the 1997 in IMI Bank, in Luxemborug, as a financial analyst in the Risk Control Department. He graduated in Finance at LUISS University in Rome in 1995, with a thesis on American options and the numerical procedures for their valuation. He wrote papers on different topics, including credit risk, derivative pricing, collateral management, managing of exotic options risks and volatility smiles. He is also author of the books "FX options and smile risk" and "Measuring and Managing Liquidity Risk", both published by Wiley.

We present a framework for establishing a mutual collateral scheme that a group of corporate debtors can use to enhance the overall credit quality of the loan portfolio extended to them by an investor. We analyze the extent of credit protection provided by the pooled collateral and how the debtors can be compensated for it, based on actuarial and financial fairness principles.

We analysed in a previous paper (see Castagna [3]) a risk-sharing mechanism for a group of companies selling their accounts receivable to an investor. In this work, we extend the idea by analysing a collateral pooling agreement that allows part of the investor's credit risk to be covered. The agreement should facilitate loan approval decisions and potentially lead to an enhancement of the interest rate paid by the borrowers.

The concept hinges on the old idea of mutual aid, *i.e.*: an arrangement where resources are voluntarily shared for the mutual benefit of a group of peers. Contemporary forms of mutual aid can be more easily implemented through modern technologies that enable effective organisation and operation of such schemes, often referred to as P2P risk sharing.

Current research on P2P risk sharing includes Abdikerimova and Feng [1], who analyse the concept of P2P risk transfer and allocation networks; Denuit and Dhaene [4], who introduce a conditional mean risk-sharing mechanism to achieve Pareto-optimality by leveraging the risk-reducing properties of conditional expectations with respect to convex ordering; building on this approach, Denuit and Robert [5] formalise the three business models dominating peer-to-peer (P2P) property and casualty insurance: the self-governing model, the broker model, and the carrier model; Feng et al. [9] propose a framework for P2P risk-sharing pools based on Pareto optimality and actuarial fairness principles, providing an exact solution for allocation ratios in the unconstrained optimal P2P risk-sharing setting, with an application to catastrophe risk pooling and several P2P alternatives for flood risk management.

The work is organised as follows: we present the general setup of a portfolio of loans and the contract terms they share; then, we analyse how collateral can be posted by the borrowers and its dynamics, also considering the credit losses suffered by the lender. We then study how the premium for the protection offered by the pooled collateral is determined, and how this premium should be allocated to individual borrowers based on actuarial and financial fairness conditions. Finally, we investigate how to calculate the premium through a numerical procedure and present a practical implementation of the scheme.

Credit Risk and Collateral Agreement

Let us assume that at time $t = 0$, a group of companies seeks financing from one or more lending entities. Each company i ($i \in \{1, \dots, I\}$) requires funding for a notional amount $K_i(0)$, with maturity at $t = T$, and agrees to pay an annual interest rate r_i. We assume that the maturity T is common to all companies; a more general case will be addressed later in this work.

The interval $[0, T]$ is divided into discrete time points at which interest and principal payments occur, *i.e.*, $t \in \{0, 1, \dots, T\}$.

Each loan is repaid according to an amortisation schedule $A_i(a)$, with $a \in \{1, \dots, T\}$, such that:

$$\sum_{a=1}^{T} A_i(a) = K_i(0).$$

The outstanding amount at any time t is then defined as:

$$K_i(t) = K_i(0) - \sum_{a \le t} A_i(a).$$

The debtor companies are subject to possible default: let τ_i be the time of default of debtor i. If $t \leq \tau_i \leq T$, the lender suffers a loss $L_i(\tau_i)$, which is a percentage $\mathbf{Lgd}_i$ (assumed fixed) of the exposure at default $\mathbf{EAD}_i(\tau_i)$, *i.e.*, the outstanding amount of the loan at time τ_i. Thus:[4]

$$\begin{aligned} L_i(\tau_i) &= \mathbf{EAD}_i(\tau_i) \times \mathbf{Lgd}_i \\ &= \left[K_i(0) - \sum_{a=1}^{\tau_i - 1} A_i(a) \right] \times \mathbf{Lgd}_i. \end{aligned} \tag{76}$$

Let $D_i(t,s) \in \{0,1\}$ be the indicator function of default ($D_i(t,s) = 1$) for debtor i between times t and s ($t < s$). The probability that debtor i defaults between t and s is denoted by $\mathbf{PD}_i(t,s)$, which may be abbreviated as $\mathbf{PD}_i$ when the context clearly refers to the interval $[0,T]$. The corresponding survival probability is $\mathbf{SP}_i(t,s) = 1 - \mathbf{PD}_i(t,s)$. Although default can occur at any instant τ_i, it is only observed at discrete times within the interval $[0,T]$. We assume that defaults among debtors are independently distributed events.

The credit risk borne by the lender can complicate the borrowing process, both in terms of the amount actually lent - which may be less than needed - and in terms of the interest rate required to compensate for the risk. To mitigate this issue, the debtor companies agree to post collateral to reduce the lender's credit risk. To minimise the total collateral posted and maximise its effectiveness in reducing credit risk, the companies agree to pool the collateral contributed by each of them and use it to cover any credit losses suffered by the lender, regardless of which companies actually default.

Collateral Dynamics and Net Credit Loss

Depending on the arrangement between the debtor companies and the lender(s), we can determine the dynamics of the collateral pool, assuming we are at time $t = 0$. In the most general agreement, each company i agrees to deposit as collateral an amount of cash C_i equal to a percentage x_i of the initial borrowed amount $K_i(0)$. The deposit is made in $M \leq T$ instalments $C_i^q(t)$ at one or more times t, so that the total collateral contributed by company i at any time $t \geq 0$ is:

$$C_i(t) = \sum_{m=0}^{t \vee M} C_i^q(m) \cdot (1 - D_i(0,m)). \tag{77}$$

The collateral contributed by company i depends on the accumulation plan and the survival of the company up to each instalment time m. If company i survives up to the last instalment time M, then for any time $t > M$ we have:

$$C_i(t) = \sum_{m=0}^{M} C_i^q(m) = x_i K_i(0).$$

In a typical arrangement, the accumulation schedule is the same for all debtors, and the total collateral contributed by each of them is a common percentage of the borrowed amount.

The expected collateral contributed up to time t, calculated at time 0, is:

$$EC_i(t) = \mathbb{E}[C_i(t)] = \sum_{m=0}^{M} C_i^q(m) \cdot \mathbf{SP}_i(0,m). \tag{78}$$

The pooled collateral $C_P(t)$ contributed by all debtor companies is kept in an escrow account until maturity T. At any time t, it is simply the sum of the collateral posted by each company:

$$C_P(t) = \sum_{i=1}^{I} C_i(t). \tag{79}$$

Similarly, the expected pooled collateral at time t, calculated at time 0, is:

$$EC_P(t) = \mathbb{E}\left[\sum_{i=1}^{I} C_i(t) \right].$$

[4] We do not include in the credit loss the possible missed interest payments due from the debtor. These can be easily incorporated into the analysis without substantially altering the structure of the arrangement, aside from the definition of credit loss.

Over the interval $[0, T]$, the cumulative total credit loss suffered by the lender is:

$$L_P(T) = \sum_{i=1}^{I} \sum_{t=1}^{T} L_i(t) \cdot D_i(t-1, t). \quad (80)$$

The expected credit loss $EL_P(T)$ is then:

$$\begin{aligned} EL_P(T) &= \mathbb{E}[L_P(T)] \\ &= \sum_{i=1}^{I} \sum_{t=0}^{T} L_i(t) \cdot \mathbf{PD}_i(t-1, t). \end{aligned} \quad (81)$$

This loss can be mitigated by the available collateral.
It is also useful to compute the present value of $EL_P(T)$. In general terms, let $\mathcal{D}(t, T)$ be the discount factor used to compute the present value at time t of a cash flow occurring at time T. Denoting by r_t the instantaneous risk-free interest rate at time t, we have:

$$\mathcal{D}(t, T) = \exp\left(-\int_t^T r_s \, ds\right).$$

In a stochastic interest rate framework, the discount factor is defined as:

$$\begin{aligned} P(t, T) &= \mathbb{E}^Q\left[\mathcal{D}(t, T)\right] \\ &= \mathbb{E}^Q\left[\exp\left(-\int_t^T r_s \, ds\right)\right], \end{aligned}$$

where the expectation is taken under the risk-neutral measure Q. We also assume that r_s is independent of any other stochastic variable in the model we are presenting.
The present value of the expected total credit loss is:

$$\begin{aligned} EL_P^D(T) &= \mathbb{E}\left[\sum_{i=1}^{I} \sum_{t=0}^{T} \mathcal{D}(0, t) L_i(t) D_i(t-1, t)\right] \\ &= \sum_{i=1}^{I} \sum_{t=0}^{T} P(0, t) L_i(t) \mathbf{PD}_i(t-1, t). \end{aligned} \quad (82)$$

Whenever a debtor defaults, the collateral is used to compensate the lender for the outstanding amount of the loan, *i.e.*, $\mathbf{EAD}_i(t) = K_i(0) - \sum_{a<t} A_i(a)$. The credit loss suffered by the lender upon the default of debtor i is zero only if a sufficient residual amount of pooled collateral is available, considering previous defaults. Since the collateral is accumulated according to a predefined schedule, it can be used to cover both future and past credit losses.
At time T, the present value of the cumulative total credit loss suffered by the lender, net of the protection offered by the pooled collateral $C_P(T)$, denoted by $NL_P^D(T)$, is:

$$\begin{aligned} NL_P^D(T) = &\sum_{i=1}^{I} \sum_{t=1}^{T} \mathcal{D}(0, t) L_i(t) D_i(t-1, t) \\ &- \sum_{t=1}^{T} \mathcal{D}(0, t) \Delta R(t), \end{aligned} \quad (83)$$

where $\Delta R(t) = \min[L_P(t), C_P(t)] - \min[L_P(t-1), C_P(t-1)]$ represents the compensation for credit losses at each time t, including possible uncovered past losses. The protection from credit losses is capped at the available pooled collateral, which may be fully exhausted by time T, or partially used, leaving residual collateral to be redistributed to the debtor companies according to a predefined rule.
The collateral can be interpreted as a form of insurance against credit losses, with a cap at the total pooled amount funded according to the instalment plan set at time 0. The expected value of the present value of the net loss is:

$$\begin{aligned} ENL_P(T) &= \mathbb{E}[NL_P^D(T)] \\ &= EL_P^D(T) - \mathbb{E}\left[\sum_{t=1}^{T} \mathcal{D}(0, t) \Delta R(t)\right]. \end{aligned} \quad (84)$$

The first term on the right-hand side of Equation (84) is the expected loss of the loan portfolio; the second term is the expected coverage provided by the available collateral, capped at its total pooled value.
Under fair actuarial principles, the premium Π that the lender would pay to purchase the protection offered by the pooled collateral arrangement equals the expected

present value of the credit loss coverage:

$$\Pi = \sum_{t=1}^{T} P(0,t) \cdot \mathbb{E}\left[\min[L_P(t), C_P(t)] - \min[L_P(t-1), C_P(t-1)]\right]. \quad (85)$$

This premium should be paid to the individual debtors and depends, in addition to the planned total amount to be posted, also on the probability that the pooled collateral is not fully contributed by each company, in cases where the collateral is posted according to a schedule rather than upfront. There is a risk that a debtor may default before contributing the full amount. Possible approaches to the valuation of Π are presented in the following sections.

Fair Allocation of the Premium and of the Remaining Collateral

Since the pooled collateral can be interpreted as the full funding of a capped insurance against credit losses suffered by the lender, debtor companies should be fairly remunerated for their expected contribution to covering actual losses (within the limit of the posted collateral). In this section, we outline the rules for making the entire pooled collateral agreement fair for all debtors.

We consider a setup in which every debtor company participates in the reimbursement of the remaining collateral pool, regardless of whether it defaulted before time T. This setup allows for a tractable - though not closed-form - solution to the problem.

Let us consider company i: it contributes an expected amount $\mathbb{E}[C_i(T)]$ to the collateral pool according to the contribution schedule. For this contribution, it should receive a fraction $0 \leq \alpha_i \leq 1$ (with $\sum_i \alpha_i = 1$) of the total insurance premium Π paid by the lender for the credit protection provided by the collateral pool. Additionally, if some collateral remains after covering the credit losses at time T, company i receives a fraction $0 \leq \beta_i \leq 1$ of this residual amount.

At maturity T, the total outcome for company i resulting from the collateral agreement is:

$$R_i^A(T) = \alpha_i \Pi + \beta_i \max\left[\sum_{i=1}^{I}\sum_{t=1}^{T} C_i(t) - \sum_{i=1}^{I}\sum_{t=1}^{T} L_i(t) D_i(t-1,t), 0\right] - C_i(T). \quad (86)$$

The quantities α_i and β_i represent what company i receives, respectively, as a fraction of the premium and of the residual collateral after the credit losses are deducted (this quantity is clearly floored at zero), while $C_i(T)$ is its actual contribution to the collateral pool. It should be noted that Equation (86) does not include any discounting of cash flows and should be interpreted purely from an actuarial fairness perspective.

We can rewrite the argument of the max operator in Equation (86) as:

$$\max\left[\sum_{i=1}^{I}\sum_{t=1}^{T} C_i(t) - \sum_{i=1}^{I}\sum_{t=1}^{T} L_i(t) D_i(t-1,t), 0\right] = \left[\sum_{i=1}^{I}\sum_{t=1}^{T} C_i(t) - \min\left[\sum_{i=1}^{I}\sum_{t=1}^{T} L_i(t) D_i(t-1,t), \sum_{i=1}^{I}\sum_{t=1}^{T} C_i(t)\right]\right].$$

Taking the expected value of $R_i^A(T)$, and using Equation (85), we obtain:

$$\mathbb{E}[R_i^A(T)] = \alpha_i \Pi + \beta_i \left[\mathbb{E}[C_P(T)] - \Pi\right] - \mathbb{E}[C_i(T)]. \quad (87)$$

Equation (87) explicitly shows the expected net collateral return (the term in square brackets) as the difference between the expected collateral pool at time T and the expected credit losses, which equal the fair premium Π. Actuarial fairness requires that $\mathbb{E}[R_i^A(T)] = 0$, and thus we must define rules for setting α_i and β_i to satisfy this condition.

Let us start with the share of residual pooled collateral β_i. The rule we aim to de-

fine must satisfy the following constraints:

$$0 \leq \beta_i \leq 1,$$
$$\sum_i \beta_i = 1,$$

so that the residual collateral is fully returned to the debtor companies, and none of them is required to make any additional payment. A simple and sensible rule is to return to each company a share proportional to its expected contribution to the collateral pool up to time T, *i.e.*,

$$\beta_i = \beta_i^* = \frac{EC_i(T)}{EC_P(T)}. \tag{88}$$

It is easy to verify that this rule satisfies both constraints.
Plugging Equation (88) into Equation (87), and computing the expected value, we obtain:

$$\mathbb{E}[R_i^A(T)] = \alpha_i \Pi - \frac{EC_i(T)}{EC_P(T)}\Pi. \tag{89}$$

From Equation (89), the fair value of α_i that ensures $\mathbb{E}[R_i^A(T)] = 0$ is:

$$\alpha_i = \alpha_i^* = \frac{EC_i(T)}{EC_P(T)}. \tag{90}$$

This choice also satisfies the constraints on α_i.
While these conditions ensure actuarial fairness, they do not necessarily guarantee financial fairness. To address this, we must consider the timing of the various cash flows and apply appropriate discounting.
Discounting the cash flows in Equation (86) can be done under different assumptions regarding when the premium Π is paid to each borrower. Without loss of generality, we assume the premium is paid in instalments over the dates $t \in \{0, \ldots, T\}$, with each instalment equal to a fraction ξ of Π. The discounted return from the collateral agreement is then:

$$\begin{aligned} &R_i^F(T) \\ &= \alpha_i^* \xi \Pi \sum_{t=0}^{T} \mathcal{D}(0,t) + \beta_i^* \mathcal{D}(0,T)\left[C_P(T) - \Pi\right] \\ &- \sum_{m=0}^{M} \mathcal{D}(0,m) C_i^q(m)(1 - D_i(0,m)). \end{aligned} \tag{91}$$

Taking expectations, we obtain:

$$\begin{aligned} &\mathbb{E}[R_i^F(T)] = \\ &= \alpha_i^* \xi \Pi \sum_{t=0}^{T} P(0,t) + \beta_i^* P(0,T)\left[EC_P(T) - \Pi\right] \\ &- \sum_{m=0}^{M} P(0,m) C_i^q(m) \cdot SP_i(0,m), \end{aligned} \tag{92}$$

where ξ is such that $\alpha_i^* \xi \Pi \sum_{t=0}^{T} P(0,t) = \alpha_i^* \Pi$. The introduction of discounting generally implies $\mathbb{E}[R_i^F(T)] \neq 0$, so we introduce an additional payment Φ_i from the lender to each borrower, paid with the same schedule and fraction ξ as the premium Π, to restore financial fairness. That is:

$$\begin{aligned} &\mathbb{E}[R_i^F(T)] \\ &= \alpha_i^* \xi (\Pi + \Phi_i) \sum_{t=0}^{T} P(0,t) \\ &+ \beta_i^* P(0,T)\left[EC_P(T) - \Pi\right] \\ &- \sum_{m=0}^{M} P(0,m) C_i^q(m) \cdot SP_i(0,m) = 0. \end{aligned} \tag{93}$$

Let $EC_i^D = \sum_{m=0}^{M} P(0,m) C_i^q(m) \cdot SP_i(0,m)$. Then the additional payment is making nil the financial return $R_i^F(T)$ is:

$$\begin{aligned} &\Phi_i = \Phi_i^* = \\ &\frac{EC_i^D - \beta_i^* P(0,T)\left[EC_P(T) - \Pi\right]}{\alpha_i^* \xi \sum_{t=0}^{T} P(0,t)} \\ &- \frac{\alpha_i^* \xi \Pi \sum_{t=0}^{T} P(0,t)}{\alpha_i^* \xi \sum_{t=0}^{T} P(0,t)}. \end{aligned}$$

This value will typically be negative. These payments by the lender may be netted, on each date, with the payments the debtor must make for interest and amortisation of the loan.

Application of the Framework to a Homogeneous Portfolio

We implement the model under the assumption of a homogeneous portfolio of debtor companies and loans. More realistic assumptions will be addressed later by approximating the formulae derived for the homogeneous case.

This portfolio consists of I debtor companies, each with an identical probability of default, *i.e.*, $\mathbf{PD}_i(t,s) = \mathbf{PD}(t,s)$ for $s \in [0,T]$, $t \leq s$. It is important to stress that all defaults are independent events.

Additionally, each company borrows the same initial amount $K_i(0) = K(0)$, follows the same amortisation schedule $\sum_{a=0}^{T} A_i(a) = \sum_{a=0}^{T} A(a)$, and has the same loss given default, so that $L_i(\tau_i) = L(\tau_i) = \mathbf{EAD}(\tau_i) \times \mathbf{Lgd}$. All companies also post the same amount of collateral according to the same payment schedule, *i.e.*, $C_i(m) = x_i K_i(0) = C(m) = xK(0)$, with $\sum_{m=0}^{M} C_i^q(m) = \sum_{m=0}^{M} C^q(m) = xK(0)$.

Under these assumptions, the number of defaults between times t and s follows a binomial distribution $\mathcal{B}_{t,s}(I, \mathbf{PD}(t,s))$, with I trials. The cumulative total credit loss distribution is directly linked to the default distribution. At any time t, the default of a company in the interval $[0,t]$ produces a loss equal to the fraction of the outstanding portfolio $L(\tau_i)$, due to the homogeneity of the loans and amortisation schedules.

The total credit loss density function between $t-1$ and t is:

$$L_P(t-1,t) = L(t) \times \mathcal{B}'(i, I, \mathbf{PD}(t-1,t)), \tag{94}$$

where $\mathcal{B}'(i, I, \mathbf{PD})$ is the binomial density function evaluated at i, *i.e.*, the probability that exactly i defaults occur.

The expected credit loss is:

$$\begin{aligned} EL_P(t-1,t) &= \mathbb{E}[L_P(t-1,t)] \\ &= \sum_{i=1}^{I} i \times L(t) \times \mathcal{B}'(i, I, \mathbf{PD}(t-1,t)) \\ &= I \times L(t) \times \mathbf{PD}(t-1,t). \end{aligned} \tag{95}$$

The standard deviation (volatility) of the total credit loss is:

$$\begin{aligned} &\mathbf{Vol}(L_P(t-1,t)) = \\ &I \times L^2(t) \times \mathbf{PD}(t-1,t)(1 - \mathbf{PD}(t-1,t)). \end{aligned} \tag{96}$$

Unfortunately, even under the homogeneous portfolio assumption, the premium Π (*i.e.*, the expected protection offered by the collateral pool) cannot be computed in closed form. We propose a numerical scheme inspired by methods from Duffie and Garleanu [7], Duffie and Singleton [8], Duffie [6], and Brigo, Pallavicini, and Torresetti [2]. The procedure exploits the homogeneous portfolio assumption:

Procedure 1 (Premium Valuation under Homogeneous Portfolio)**.**

1. *Let S be the total number of simulations and $sd(t)$ the number of surviving debtors at time t.*
2. *For each of the S simulations, set $sd(0) = I$, $C_P(0) = sd(0) \times C^q(0)$, $L_P(0) = 0$, $\Pi_s(t) = 0$, and $\Pi_p = 0$.*
3. *For each time $t \in \{1, \ldots, T\}$, draw the number of defaults $nd(t)$ from a binomial distribution $\mathcal{B}(sd(t-1), \mathbf{PD}(t-1,t))$.*
4. *Update the total credit loss: $L_P(t) = L_P(t-1) + nd(t) \times \mathbf{Lgd} \times \mathbf{EAD}(t-1)$.*
5. *Update the number of surviving debtors: $sd(t) = sd(t-1) - nd(t)$.*
6. *Update the posted collateral: $C_P(t) = C_P(t) + sd(t) \times C^q(t)$.*
7. *Update the premium:*

$$\begin{aligned} \Pi_s(t) = \Pi_s(t-1) &+ P(0,t)\Big[\min\left(L_P(t), C_P(t)\right) \\ &- \min\left(L_P(t-1), C_P(t-1)\right)\Big]. \end{aligned}$$

8. *Store the premium: $\Pi_p = \Pi_p + \Pi_s(T)$.*
9. *After completing all S simulations, compute the premium: $\Pi = \Pi_p / S$.*

EC_P		
1 Inst.	**5 Inst. Scen. 1**	**5 Inst. Scen. 2**
€ 28,120	€ 26,188	€ 25,269

TABLE 15: Total collateral contributed under different probability of default scenarios and collateral contribution schedules.

	Year					
	0	**1**	**2**	**3**	**4**	**5**
Average Loss		€ 6,014	€ 4,881	€ 3,715	€ 2,513	€ 1,275
Loss M²		€ 12,068,171	€ 7,950,128	€ 4,603,759	€ 2,106,732	€ 542,365
Loss Variance		€ 3,474	€ 2,820	€ 2,146	€ 1,451	€ 736
PD*		3.56%	3.56%	3.56%	3.56%	3.56%
EAD*	€ 3,468	€ 2,815	€ 2,142	€ 1,449	€ 735	€ -
I*	100	81	81	81	81	81

TABLE 16: Approximation by MMT to a homogeneous portfolio in the first scenario for probability of default.

It is important to note that this approximation is used only to compute the premium in Equation (85), while all other quantities must be computed using the original input data.
Procedure 1 must be adapted to carefully handle the number of debtors and the posting of collateral. We propose the following procedure, which extends the original one designed for the homogeneous portfolio:

Procedure 2 (Premium Valuation under Approximated Homogeneous Portfolio)**.**

1. *For each time $t \in \{1, \dots, T\}$, compute* $\mathbf{PD}^*(t-1,t)$, $\mathbf{EAD}^*(t)$, *and* $I^*(t)$ *from the actual data.*

2. *From the actual collateral instalment schedule, compute the equivalent instalment:*

$$C^{q*}(m) = \frac{1}{I^*(m)} \sum_{i=1}^{I} C_i^q(m).$$

3. *Let S be the total number of simulations and $sd(t)$ the number of surviving debtors at time t.*

4. *For each of the S simulations, set:* $sd(0) = I^*(0)$, $C_P(0) = sd(0) \cdot C^{q*}(0)$, $L_P(0) = 0$, $\Pi_s(t) = 0$, $\Pi_p = 0$, $cnd\%(0) = 0$.

5. *For each time $t \in \{1, \dots, T\}$, draw the number of defaults $nd(t)$ from a binomial distribution $\mathcal{B}(sd(t-1), \mathbf{PD}^*(t-1,t))$.*

6. *Update the cumulative default percentage:*

$$cnd\%(t) = cnd\%(t-1) + \frac{nd(t)}{sd(t-1)},$$

if $sd(t-1) = 0$ then $cnd\%(t) = 1$.

7. *Update the total credit loss:*

$$L_P(t) = L_P(t-1) + nd(t) \cdot \mathbf{Lgd} \cdot \mathbf{EAD}^*(t-1).$$

8. *Update the posted collateral:*

$$C_P(t) = C_P(t) + [sd(t-1) - nd(t)] \cdot C^{q*}(t).$$

9. *Update the number of surviving debtors:*

$$sd(t) = \lfloor I^*(t) \cdot [1 - cnd\%(t)] \rfloor.$$

10. *Update the premium:*

$$\begin{aligned}\Pi_s(t) &= \Pi_s(t-1)\\ &+ P(0,t)\Big[\min\big(L_P(t), C_P(t)\big)\\ &- \min\big(L_P(t-1), C_P(t-1)\big)\Big].\end{aligned}$$

11. *Store the premium:*

$$\Pi_p = \Pi_p + \Pi_s(T).$$

12. *After completing all S simulations, compute the final premium:*

$$\Pi = \frac{\Pi_p}{S}.$$

Procedure 3 is quite efficient and easily implemented.

Approximation to a Homogeneous Portfolio

We resort to an approximation that makes the actual portfolio as similar as possible to an equivalent homogeneous portfolio.[5] This approach allows us to derive, for each time $t \in \{1,\dots,T\}$, the equivalent default probability $\mathbf{PD}^*(t-1,t)$, the outstanding loan amount $\mathbf{EAD}^*(t)$, and the equivalent number of debtors $I^*(t)$ of a homogeneous portfolio that approximates the actual one, assuming that the loss given default is the same for all debtors, *i.e.*, $\mathbf{Lgd}_i = \mathbf{Lgd}$. We apply a Moment Matching Technique (MMT), which equates the first and second moments of the credit loss, *i.e.*: the expected value and variance of the actual portfolio over the period $[t-1,t]$, to those of the equivalent homogeneous portfolio. This is done under the constraint that the total value of the outstanding loans in both portfolios is the same. Formally:

$$\begin{cases}\mathbb{E}[L(t-1,t)] = \\ \sum_{i=1}^{I} L_i(t)\cdot \mathbf{PD}_i(t-1,t) = I^*(t)\cdot L\cdot \mathbf{PD}^*(t-1,t)\\ \\ \mathrm{Var}[L(t-1,t)] = \\ \sum_{i=1}^{I} (L_i^*(t))^2\cdot \mathbf{PD}_i(t-1,t)\cdot(1-\mathbf{PD}_i(t-1,t)) = \\ I^*(t)\cdot (L_i^*(t))^2\cdot \mathbf{PD}^*(t-1,t)\cdot(1-\mathbf{PD}^*(t-1,t))\\ \\ I^*(t)\cdot \mathbf{EAD}^*(t) = K\end{cases},$$

where $K = \sum_i \mathbf{EAD}_i(t)$ is the total value of the outstanding loans at time t. The solution to the system () is:

$$\begin{cases}\mathbf{PD}^*(t-1,t) = \frac{\sum_{i=1}^{I}\mathbf{EAD}_i(t)\cdot\mathbf{PD}_i(t-1,t)}{K}\\ \\ \mathbf{EAD}^*(t) = \\ \frac{\sum_{i=1}^{I}\mathbf{EAD}_i^2(t)\cdot\mathbf{PD}_i(t-1,t)\cdot(1-\mathbf{PD}_i(t-1,t))}{K\cdot\mathbf{PD}^*(t-1,t)\cdot(1-\mathbf{PD}^*(t-1,t))}\\ \\ I^*(t) = \frac{K}{\mathbf{EAD}^*(t)}\end{cases} \quad (97)$$

It is important to note that this approximation is used only to compute the premium in Equation (85), while all other quantities must be computed using the original input data.
Procedure 1 must be adapted to carefully handle the number of debtors and the posting of collateral. We propose the following procedure, which extends the original one designed for the homogeneous portfolio:

Procedure 3 (Premium Valuation under Approximated Homogeneous Portfolio)**.**

1. *For each time $t \in \{1,\dots,T\}$, compute $\mathbf{PD}^*(t-1,t)$, $\mathbf{EAD}^*(t)$, and $I^*(t)$ from the actual data.*

2. *From the actual collateral instalment schedule, compute the equivalent instalment:*

$$C^{q*}(m) = \frac{1}{I^*(m)}\sum_{i=1}^{I} C_i^q(m).$$

3. *Let S be the total number of simulations and $sd(t)$ the number of surviving*

[5] See Castagna [3] and the references therein for the same approach applied to a portfolio of receivables.

debtors at time t.

4. *For each of the S simulations, set:* $sd(0) = I^*(0)$, $C_P(0) = sd(0) \cdot C^{q*}(0)$, $L_P(0) = 0$, $\Pi_s(t) = 0$, $\Pi_p = 0$, $cnd\%(0) = 0$.

5. *For each time $t \in \{1, \ldots, T\}$, draw the number of defaults $nd(t)$ from a binomial distribution $\mathcal{B}(sd(t-1), \mathbf{PD}^*(t-1,t))$.*

6. *Update the cumulative default percentage:*

$$cnd\%(t) = cnd\%(t-1) + \frac{nd(t)}{sd(t-1)},$$

if $sd(t-1) = 0$ then $cnd\%(t) = 1$.

7. *Update the total credit loss:*

$$L_P(t) = L_P(t-1) + nd(t) \cdot \mathbf{Lgd} \cdot \mathbf{EAD}^*(t-1).$$

8. *Update the posted collateral:*

$$C_P(t) = C_P(t) + [sd(t-1) - nd(t)] \cdot C^{q*}(t).$$

9. *Update the number of surviving debtors:*

$$sd(t) = \lfloor I^*(t) \cdot [1 - cnd\%(t)] \rfloor.$$

10. *Update the premium:*

$$\Pi_s(t) = \Pi_s(t-1) + P(0,t) [\min(L_P(t), C_P(t)) - \min(L_P(t-1), C_P(t-1))].$$

11. *Store the premium:*

$$\Pi_p = \Pi_p + \Pi_s(T).$$

12. *After completing all S simulations, compute the final premium:*

$$\Pi = \frac{\Pi_p}{S}.$$

Procedure 3 is quite efficient and easily implemented.

	Year					
	0	**1**	**2**	**3**	**4**	**5**
Average Loss		€ 9,103	€ 7,388	€ 5,622	€ 3,803	€ 1,930
Loss M²		€ 17,895,414	€ 11,788,931	€ 6,826,732	€ 3,123,989	€ 804,251
Loss Variance		€ 4,230	€ 3,434	€ 2,613	€ 1,767	€ 897
PD*		5.40%	5.40%	5.40%	5.40%	5.40%
EAD*	€ 3,463	€ 2,811	€ 2,139	€ 1,447	€ 734	€ -
I*	100	81	81	81	81	81

TABLE 17: *Approximation by MMT to a homogeneous portfolio in the second scenario for probability of default.*

Application to a Realistic Portfolio

We apply the framework sketched above to a realistic portfolio consisting of 100 loans. The details of each loan are shown in Appendix "Loan Portfolio Details". Table 21 provides, for each debtor, the initial amount and the amortisation schedule. All loans start at time $t = 0$ and mature at time $t = 5$. In Table 22, we report the probability of default **PD** for each debtor under two different scenarios: the first with moderately low probabilities, and the second with higher ones. We assume that **PD** is constant over each of the five years of the loan duration. The (weighted) average **PD** for the entire portfolio is 3.56% in the first scenario and 5.40% in the second.

For each probability scenario, we also consider the expected collateral contributed by each debtor under two assumptions:

- A single contribution made fully at the inception of the loan;
- A contribution made in five instalments, one at the beginning of each year.

In the first case, the possibility of default has no impact on the contributed amount. In both cases, the total contributed collateral (if the debtor does not default before completing the instalment schedule) is 10% of the initial borrowed amount.

Given the expected contributed collateral $EC_i(5)$, we compute the parameters $\beta_i = \alpha_i$ that make the expected return in Equation (87) equal to zero for each debtor. The total collateral for each probability of default scenario and contribution schedule is reported in Table 19. It is clear that the smaller expected amount of collateral (25,269) refers to an accumulation plan in 5 instalments in the higher probability of deafult scenario.

We now compute the total premium Π, which should be paid by the lender to the debtors, under two different probability of default scenarios and five assumptions regarding the collateral contribution schedule, ranging from 1 to 5 instalments.

For the computation, we first apply the Moment Matching Technique (MMT) described in Section "Approximation to a Homogeneous Portfolio" to approximate the actual portfolio with an equivalent homogeneous portfolio in each of the five years. For the first **PD** scenario, the expected loss, second moment, and variance-along with the three quantities solving the system in Equation (97) - are reported in Table 16. The same results for the second **PD** scenario are shown in Table 17.

We now implement Procedure 3 to compute the total premium Π, which should be paid by the lender to the debtor companies. The computation is performed under the two different **PD** scenarios and across five assumptions regarding the collateral contribution schedule, ranging from 1 to 5 instalments.

The results for the total premium Π in absolute terms are reported in Table 18, while the same results expressed as a percentage of the initial amount of the loan portfolio are shown in Table 20.

It can be noted that the level of the premium is, not surprisingly, linked to the average level of the default probabilities of the debtors. In the first case (average of 3.56% p.a.), the premium is around 5.70% of the initial notional of the loans, for a total theoretical collateral pool of 10% of the same value. In the second case (5.40% p.a.), the premium varies from around 8.00% to 7.50%.

More unexpectedly, the dependence of the premium on the collateral contribution schedule is very mild: it is practically negligible in the first **PD** scenario, and around 0.5% in the second scenario.

As a final application, we show how to calculate the total premium including the extra payment Φ that ensures financial fairness. Consider debtor 1, with all details taken from the tables in Appendix "Loan Portfolio Details, for the first **PD** scenario. We assume that both the fraction of the premium $\alpha_1\Pi$ and the extra payment Φ_1 are

	Installments				
PD Scen.	1	2	3	4	5
3.56%	€ 16,129	€ 16,067	€ 16,031	€ 16,011	€ 15,970
5.40%	€ 22,691	€ 22,354	€ 21,981	€ 21,624	€ 21,128

TABLE 18: Premium Π for the credit protection, in each of the two scenarios of default probabilities and collateral contribution schedule with different number of instalments. Absolute values.

EC_P		
1 Inst.	**5 Inst. Scen. 1**	**5 Inst. Scen. 2**
€ 28,120	€ 26,188	€ 25,269

TABLE 19: Total collateral contributed under different probability of default scenarios and collateral contribution schedules.

paid at the start of each year of the loan's duration, regardless of whether the debtor survives until expiry or defaults earlier.
Let us assume that the annual risk-free interest rate is constant at 3.5%. The discount factors are given by:

$$P(0,n) = \exp(-0.035 \times n).$$

Additionally, we assume that the collateral is contributed by debtor 1 in 5 instalments at the beginning of each year of the loan's duration. The probability of default is taken from the first scenario.
First, we compute the quantity ξ in Equation (93), and we have:

$$\xi = \frac{\alpha_i^* \Pi}{\alpha_i^* \Pi \sum_{t=0}^{T} P(0,t)}.$$

Since $\alpha_1 \Pi = 242.18$, we get:

$$\xi = 0.21.$$

Calculation of the extra payment Φ_i is now straightforward from Equation ():

$$\Phi_i = -51.60.$$

which is negative, as expected.

	Installments				
PD Scen.	1	2	3	4	5
3.56%	5.74%	5.71%	5.70%	5.69%	5.68%
5.40%	8.07%	7.95%	7.82%	7.69%	7.51%

TABLE 20: *Premium Π for the credit protection, in each of the two scenarios of default probabilities and collateral contribution schedule with different number of instalments. In percentage of the initial value of the portfolio of loans.*

Conclusion

We proposed a collateral pooling scheme that offers credit protection to the lender of a portfolio of loans to corporate debtors. The framework is flexible enough to allow for the contribution of collateral in instalments, even during the life of the loans.
The value of the credit protection offered by the collateral, and the corresponding premium to be paid to the debtors, has been derived. The allocation of this premium and of the remaining collateral at the expiry of the loan portfolio is determined based on actuarial and financial fairness conditions.
More sophisticated schemes and alternative options for allocating the residual collateral - and consequently the premium - can be designed starting from the present framework. Also a credit model that allows for a correlation of default events would be make the framework more effective.

References

[1] **Abdikerimova, S. and Feng, R.** *Peer-to-peer multi-risk insurance and mutual aid*. European Journal of Operational Research, Vol. 299, Issue 2, pp. 735-749, June 2022.

[2] **Brigo, D., Pallavicini, A. and Torresetti, R.** *Calibration of CDO Tranches with the Dynamical Generalized-Poisson Loss Model*. SSRN Electronic Journal, May 2007.

[3] **Castagna, A.** *Risk Sharing Insurance Schemes for Invoice Discounting Platforms*. Research Paper Series N. 42. February 2022.

[4] **Denuit, M. and Dhaeneb, J.** *Convex order and comonotonic conditional mean risk sharing*. Insurance: Mathematics and Economics, Vol. 51, Issue 2, pp. 265-270, September 2012.

[5] **Denuit, M. and Robert, C.** *Risk sharing under the dominant peer-to-peer property and casualty insurance business models*. Risk Management and Insurance Review, Vol. 24, Issue 2, pp. 181-205, June 2021.

[6] **Duffie, D.** *First-to-Default Valuation*. Stanford University, May 1998.

[7] **Duffie, D. and Garleanu, N.** *Risk and Valuation of Collateralized Debt Obligations*. Stanford University, September 2001.

[8] **Duffie, D. and Singleton, K.** *Simulating Correlated Defaults*. Stanford University, May 1999.

[9] **Feng, R., Liu, C. and Taylor, S.** *Peer-to-Peer Risk Sharing with an Application to Flood Risk Pooling*. Annals of Operations Research, Vol. 321, pp. 813-842, July 2022.

Annex

Loan Portfolio Details

Table shows the details of the portfolio of loans. Each loan starts on the same date $t = 0$ and expires on $t = 5$.

	Notional (K(0)) at Year	Amortisation (A(t)) at Year				
ID Debtor	**0**	**1**	**2**	**3**	**4**	**5**
1	€ 4,200	€ 791	€ 815	€ 839	€ 864	€ 890
2	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
3	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
4	€ 4,700	€ 885	€ 912	€ 939	€ 967	€ 996
5	€ 3,700	€ 697	€ 718	€ 739	€ 762	€ 784
6	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
7	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
8	€ 2,900	€ 546	€ 563	€ 579	€ 597	€ 615
9	€ 2,700	€ 509	€ 524	€ 540	€ 556	€ 572
10	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
11	€ 2,000	€ 377	€ 388	€ 400	€ 412	€ 424
12	€ 4,900	€ 923	€ 951	€ 979	€ 1,009	€ 1,039
13	€ 3,000	€ 565	€ 582	€ 599	€ 617	€ 636
14	€ 3,900	€ 735	€ 757	€ 779	€ 803	€ 827
15	€ 3,300	€ 622	€ 640	€ 659	€ 679	€ 700
16	€ 4,100	€ 772	€ 795	€ 819	€ 844	€ 869
17	€ 2,500	€ 471	€ 485	€ 500	€ 515	€ 530
18	€ 3,400	€ 640	€ 660	€ 679	€ 700	€ 721
19	€ 5,000	€ 942	€ 970	€ 999	€ 1,029	€ 1,060
20	€ 4,800	€ 904	€ 931	€ 959	€ 988	€ 1,018
21	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
22	€ 4,300	€ 810	€ 834	€ 859	€ 885	€ 912
23	€ 4,600	€ 866	€ 892	€ 919	€ 947	€ 975
24	€ 2,900	€ 546	€ 563	€ 579	€ 597	€ 615
25	€ 1,200	€ 226	€ 233	€ 240	€ 247	€ 254
26	€ 1,600	€ 301	€ 310	€ 320	€ 329	€ 339
27	€ 2,700	€ 509	€ 524	€ 540	€ 556	€ 572
28	€ 4,400	€ 829	€ 854	€ 879	€ 906	€ 933
29	€ 800	€ 151	€ 155	€ 160	€ 165	€ 170
30	€ 1,200	€ 226	€ 233	€ 240	€ 247	€ 254

TABLE 21: *Premium Π for the credit protection, in each of the two scenarios of default probabilities and collateral contribution schedule with different number of instalments. In percentage of the initial value of the portfolio of loans.*

	Notional (K(0)) at Year	Amortisation (A(t)) at Year				
ID Debtor	0	1	2	3	4	5
31	€ 2,100	€ 396	€ 407	€ 420	€ 432	€ 445
32	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
33	€ 2,300	€ 433	€ 446	€ 460	€ 473	€ 488
34	€ 2,400	€ 452	€ 466	€ 480	€ 494	€ 509
35	€ 4,700	€ 885	€ 912	€ 939	€ 967	€ 996
36	€ 3,900	€ 735	€ 757	€ 779	€ 803	€ 827
37	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
38	€ 1,700	€ 320	€ 330	€ 340	€ 350	€ 360
39	€ 2,600	€ 490	€ 504	€ 520	€ 535	€ 551
40	€ 2,500	€ 471	€ 485	€ 500	€ 515	€ 530
41	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
42	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
43	€ 4,300	€ 810	€ 834	€ 859	€ 885	€ 912
44	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
45	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
46	€ 2,500	€ 471	€ 485	€ 500	€ 515	€ 530
47	€ 3,000	€ 565	€ 582	€ 599	€ 617	€ 636
48	€ 3,700	€ 697	€ 718	€ 739	€ 762	€ 784
49	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
50	€ 4,700	€ 885	€ 912	€ 939	€ 967	€ 996
51	€ 2,200	€ 414	€ 427	€ 440	€ 453	€ 466
52	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
53	€ 3,400	€ 640	€ 660	€ 679	€ 700	€ 721
54	€ 1,900	€ 358	€ 369	€ 380	€ 391	€ 403
55	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
56	€ 2,600	€ 490	€ 504	€ 520	€ 535	€ 551
57	€ 600	€ 113	€ 116	€ 120	€ 123	€ 127
58	€ 3,600	€ 678	€ 698	€ 719	€ 741	€ 763
59	€ 3,200	€ 603	€ 621	€ 639	€ 659	€ 678
60	€ 4,600	€ 866	€ 892	€ 919	€ 947	€ 975
61	€ 3,600	€ 678	€ 698	€ 719	€ 741	€ 763
62	€ 4,000	€ 753	€ 776	€ 799	€ 823	€ 848
63	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148
64	€ 4,700	€ 885	€ 912	€ 939	€ 967	€ 996
65	€ 700	€ 132	€ 136	€ 140	€ 144	€ 148

TABLE 21: Premium Π for the credit protection, in each of the two scenarios of default probabilities and collateral contribution schedule with different number of instalments. In percentage of the initial value of the portfolio of loans.

	Notional (K(0)) at Year	Amortisation (A(t)) at Year				
ID Debtor	0	1	2	3	4	5
66	€ 4,000	€ 753	€ 776	€ 799	€ 823	€ 848
67	€ 4,900	€ 923	€ 951	€ 979	€ 1,009	€ 1,039
68	€ 600	€ 113	€ 116	€ 120	€ 123	€ 127
69	€ 2,000	€ 377	€ 388	€ 400	€ 412	€ 424
70	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
72	€ 2,300	€ 433	€ 446	€ 460	€ 473	€ 488
73	€ 2,600	€ 490	€ 504	€ 520	€ 535	€ 551
74	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
75	€ 1,500	€ 283	€ 291	€ 300	€ 309	€ 318
76	€ 900	€ 170	€ 175	€ 180	€ 185	€ 191
77	€ 3,700	€ 697	€ 718	€ 739	€ 762	€ 784
78	€ 2,800	€ 527	€ 543	€ 560	€ 576	€ 594
79	€ 3,500	€ 659	€ 679	€ 699	€ 720	€ 742
80	€ 800	€ 151	€ 155	€ 160	€ 165	€ 170
81	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
82	€ 500	€ 94	€ 97	€ 100	€ 103	€ 106
83	€ 3,900	€ 735	€ 757	€ 779	€ 803	€ 827
84	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
85	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
86	€ 1,700	€ 320	€ 330	€ 340	€ 350	€ 360
87	€ 2,700	€ 509	€ 524	€ 540	€ 556	€ 572
88	€ 3,900	€ 735	€ 757	€ 779	€ 803	€ 827
89	€ 4,500	€ 848	€ 873	€ 899	€ 926	€ 954
90	€ 3,200	€ 603	€ 621	€ 639	€ 659	€ 678
91	€ 3,700	€ 697	€ 718	€ 739	€ 762	€ 784
92	€ 2,500	€ 471	€ 485	€ 500	€ 515	€ 530
93	€ 3,700	€ 697	€ 718	€ 739	€ 762	€ 784
94	€ 1,300	€ 245	€ 252	€ 260	€ 268	€ 276
95	€ 1,300	€ 245	€ 252	€ 260	€ 268	€ 276
96	€ 2,100	€ 396	€ 407	€ 420	€ 432	€ 445
97	€ 4,700	€ 885	€ 912	€ 939	€ 967	€ 996
98	€ 3,600	€ 678	€ 698	€ 719	€ 741	€ 763
99	€ 1,000	€ 188	€ 194	€ 200	€ 206	€ 212
100	€ 2,500	€ 471	€ 485	€ 500	€ 515	€ 530

TABLE 21: *Premium Π for the credit protection, in each of the two scenarios of default probabilities and collateral contribution schedule with different number of instalments. In percentage of the initial value of the portfolio of loans.*

	PD		EC_i			$\beta_i = \alpha_i$		
	Sc.1	Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2
1	2.80%	7.40%	€ 420	€ 397	€ 362	1.49%	1.52%	1.43%
2	3.30%	7.40%	€ 450	€ 421	€ 388	1.60%	1.61%	1.54%
3	2.50%	1.90%	€ 350	€ 333	€ 337	1.24%	1.27%	1.33%
4	3.30%	8.80%	€ 470	€ 440	€ 394	1.67%	1.68%	1.56%
5	4.40%	2.90%	€ 370	€ 339	€ 349	1.32%	1.29%	1.38%
6	4.70%	8.90%	€ 70	€ 64	€ 59	0.25%	0.24%	0.23%
7	3.20%	4.90%	€ 70	€ 66	€ 63	0.25%	0.25%	0.25%
8	5.00%	3.60%	€ 290	€ 262	€ 270	1.03%	1.00%	1.07%
9	2.60%	7.20%	€ 270	€ 256	€ 234	0.96%	0.98%	0.93%
10	3.50%	4.50%	€ 100	€ 93	€ 91	0.36%	0.36%	0.36%
11	2.70%	7.30%	€ 200	€ 189	€ 173	0.71%	0.72%	0.68%
12	5.00%	3.20%	€ 490	€ 443	€ 460	1.74%	1.69%	1.82%
13	2.50%	7.60%	€ 300	€ 285	€ 258	1.07%	1.09%	1.02%
14	3.40%	8.20%	€ 390	€ 364	€ 331	1.39%	1.39%	1.31%
15	3.00%	1.10%	€ 330	€ 311	€ 323	1.17%	1.19%	1.28%
16	2.90%	4.70%	€ 410	€ 387	€ 373	1.46%	1.48%	1.48%
17	5.00%	6.10%	€ 250	€ 226	€ 221	0.89%	0.86%	0.88%
18	2.90%	6.60%	€ 340	€ 321	€ 298	1.21%	1.23%	1.18%
19	4.50%	3.40%	€ 500	€ 457	€ 467	1.78%	1.74%	1.85%
20	3.20%	4.40%	€ 480	€ 450	€ 440	1.71%	1.72%	1.74%
21	4.70%	8.60%	€ 70	€ 64	€ 59	0.25%	0.24%	0.23%
22	3.30%	5.00%	€ 430	€ 403	€ 389	1.53%	1.54%	1.54%
23	2.60%	4.40%	€ 460	€ 437	€ 421	1.64%	1.67%	1.67%
24	2.60%	2.00%	€ 290	€ 275	€ 279	1.03%	1.05%	1.10%
25	3.40%	1.00%	€ 120	€ 112	€ 118	0.43%	0.43%	0.47%
26	4.20%	5.20%	€ 160	€ 147	€ 144	0.57%	0.56%	0.57%
27	2.90%	7.50%	€ 270	€ 255	€ 232	0.96%	0.97%	0.92%
28	3.60%	5.90%	€ 440	€ 409	€ 391	1.56%	1.56%	1.55%
29	3.60%	7.90%	€ 80	€ 74	€ 68	0.28%	0.28%	0.27%
30	2.90%	1.30%	€ 120	€ 113	€ 117	0.43%	0.43%	0.46%
31	2.60%	6.20%	€ 210	€ 199	€ 186	0.75%	0.76%	0.73%
32	3.10%	3.50%	€ 350	€ 329	€ 326	1.24%	1.26%	1.29%

TABLE 22: Probability of default of each debtor, expected collateral contributed in 1 instalment and 5 instalments and parameters $\beta = \alpha$ for each debtor (in the two different probability of default scearios).

	PD		EC_i			$\beta_i = \alpha_i$		
	Sc.1	Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2
33	3.30%	6.00%	€ 230	€ 215	€ 204	0.82%	0.82%	0.81%
34	3.20%	2.90%	€ 240	€ 225	€ 226	0.85%	0.86%	0.90%
35	3.10%	4.50%	€ 470	€ 442	€ 430	1.67%	1.69%	1.70%
36	4.80%	5.50%	€ 390	€ 354	€ 349	1.39%	1.35%	1.38%
37	3.30%	1.80%	€ 70	€ 66	€ 68	0.25%	0.25%	0.27%
38	3.10%	4.40%	€ 170	€ 160	€ 156	0.60%	0.61%	0.62%
39	3.80%	8.40%	€ 260	€ 241	€ 220	0.92%	0.92%	0.87%
40	4.30%	8.70%	€ 250	€ 229	€ 210	0.89%	0.88%	0.83%
41	2.80%	6.30%	€ 450	€ 425	€ 397	1.60%	1.62%	1.57%
42	4.40%	4.30%	€ 100	€ 92	€ 92	0.36%	0.35%	0.36%
43	3.60%	5.80%	€ 430	€ 400	€ 383	1.53%	1.53%	1.52%
44	4.60%	2.40%	€ 70	€ 64	€ 67	0.25%	0.24%	0.26%
45	3.90%	8.10%	€ 450	€ 416	€ 383	1.60%	1.59%	1.51%
46	4.20%	8.10%	€ 250	€ 230	€ 213	0.89%	0.88%	0.84%
47	3.50%	2.40%	€ 300	€ 280	€ 286	1.07%	1.07%	1.13%
48	4.20%	2.70%	€ 370	€ 340	€ 351	1.32%	1.30%	1.39%
49	2.80%	3.20%	€ 350	€ 331	€ 328	1.24%	1.26%	1.30%
50	4.50%	4.10%	€ 470	€ 430	€ 433	1.67%	1.64%	1.71%
51	2.70%	5.30%	€ 220	€ 208	€ 198	0.78%	0.80%	0.78%
52	3.20%	6.00%	€ 350	€ 328	€ 310	1.24%	1.25%	1.23%
53	2.70%	2.60%	€ 340	€ 322	€ 323	1.21%	1.23%	1.28%
54	4.60%	2.70%	€ 190	€ 173	€ 180	0.68%	0.66%	0.71%
55	3.20%	5.10%	€ 100	€ 94	€ 90	0.36%	0.36%	0.36%
56	3.30%	8.80%	€ 260	€ 243	€ 218	0.92%	0.93%	0.86%
57	3.50%	4.90%	€ 60	€ 56	€ 54	0.21%	0.21%	0.22%
58	3.50%	9.00%	€ 360	€ 336	€ 301	1.28%	1.28%	1.19%
59	2.60%	3.10%	€ 320	€ 304	€ 301	1.14%	1.16%	1.19%
60	3.60%	7.90%	€ 460	€ 428	€ 393	1.64%	1.63%	1.55%
61	3.50%	5.30%	€ 360	€ 336	€ 324	1.28%	1.28%	1.28%
62	3.50%	8.70%	€ 400	€ 373	€ 336	1.42%	1.42%	1.33%
63	4.90%	5.70%	€ 70	€ 63	€ 62	0.25%	0.24%	0.25%
64	3.60%	8.20%	€ 470	€ 437	€ 399	1.67%	1.67%	1.58%

TABLE 22: Probability of default of each debtor, expected collateral contributed in 1 instalment and 5 instalments and parameters $\beta = \alpha$ for each debtor (in the two different probability of default scearios).

	PD		EC_i			$\beta_i = \alpha_i$		
	Sc.1	Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2	1 Inst.	5 Inst. Sc.1	5 Inst. Sc.2
65	3.50%	8.50%	€ 70	€ 65	€ 59	0.25%	0.25%	0.23%
66	4.40%	6.80%	€ 400	€ 366	€ 349	1.42%	1.40%	1.38%
67	3.70%	3.10%	€ 490	€ 455	€ 461	1.74%	1.74%	1.82%
68	2.80%	6.30%	€ 60	€ 57	€ 53	0.21%	0.22%	0.21%
69	3.30%	1.90%	€ 200	€ 187	€ 193	0.71%	0.71%	0.76%
70	3.90%	8.80%	€ 350	€ 324	€ 294	1.24%	1.24%	1.16%
71	3.90%	7.30%	€ 130	€ 120	€ 112	0.46%	0.46%	0.44%
72	3.80%	4.30%	€ 230	€ 213	€ 211	0.82%	0.81%	0.84%
73	2.80%	5.50%	€ 260	€ 246	€ 233	0.92%	0.94%	0.92%
74	4.40%	7.10%	€ 450	€ 412	€ 390	1.60%	1.57%	1.55%
75	4.30%	5.50%	€ 150	€ 138	€ 134	0.53%	0.53%	0.53%
76	4.10%	3.60%	€ 90	€ 83	€ 84	0.32%	0.32%	0.33%
77	3.20%	5.70%	€ 370	€ 347	€ 330	1.32%	1.33%	1.31%
78	4.40%	4.20%	€ 280	€ 256	€ 257	1.00%	0.98%	1.02%
79	3.00%	6.70%	€ 350	€ 330	€ 306	1.24%	1.26%	1.21%
80	4.10%	6.80%	€ 80	€ 74	€ 70	0.28%	0.28%	0.28%
81	4.50%	2.00%	€ 450	€ 411	€ 432	1.60%	1.57%	1.71%
82	4.90%	1.30%	€ 50	€ 45	€ 49	0.18%	0.17%	0.19%
83	4.10%	6.80%	€ 390	€ 359	€ 340	1.39%	1.37%	1.35%
84	2.70%	7.30%	€ 100	€ 95	€ 86	0.36%	0.36%	0.34%
85	3.40%	2.40%	€ 100	€ 93	€ 95	0.36%	0.36%	0.38%
86	4.70%	6.70%	€ 170	€ 155	€ 149	0.60%	0.59%	0.59%
87	4.80%	7.20%	€ 270	€ 245	€ 234	0.96%	0.94%	0.93%
88	4.10%	6.00%	€ 390	€ 359	€ 346	1.39%	1.37%	1.37%
89	2.90%	4.10%	€ 450	€ 425	€ 415	1.60%	1.62%	1.64%
90	3.10%	1.60%	€ 320	€ 301	€ 310	1.14%	1.15%	1.23%
91	4.50%	7.40%	€ 370	€ 338	€ 319	1.32%	1.29%	1.26%
92	4.00%	3.50%	€ 250	€ 231	€ 233	0.89%	0.88%	0.92%
93	2.60%	5.70%	€ 370	€ 351	€ 330	1.32%	1.34%	1.31%
94	3.00%	7.10%	€ 130	€ 122	€ 113	0.46%	0.47%	0.45%
95	4.40%	7.20%	€ 130	€ 119	€ 113	0.46%	0.45%	0.45%
96	4.60%	6.80%	€ 210	€ 192	€ 183	0.75%	0.73%	0.73%
97	2.60%	5.40%	€ 470	€ 446	€ 422	1.67%	1.70%	1.67%
98	4.00%	2.80%	€ 360	€ 332	€ 340	1.28%	1.27%	1.35%
99	4.30%	2.80%	€ 100	€ 92	€ 95	0.36%	0.35%	0.37%
100	2.70%	5.70%	€ 250	€ 237	€ 223	0.89%	0.90%	0.88%

TABLE 22: Probability of default of each debtor, expected collateral contributed in 1 instalment and 5 instalments and parameters $\beta = \alpha$ for each debtor (in the two different probability of default scearios).

In the previous issue

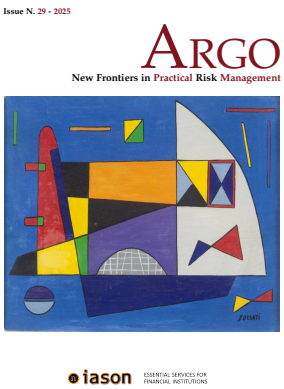

Issue N. 29 - 2025

Insurance Risk

Risks Aggregation, Tail Dependence and Beyond

Market Risk

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

Innovation

AI Risk Management Frameworks

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