# Quantifying Credit Portfolio sensitivity to asset correlations with interpretable generative neural networks

*Sergio Caprioli, Emanuele Cagliero, Riccardo Crupi*

April 2024

INTESA SANPAOLO

# Synthetic data for risk management and model validation



"Excessive reliance **on historical financial data** in model validation can lead to **less robust** risk management in the future." [Marti, 2020].

Generating **realistic synthetic financial data** improves model resilience and accuracy in **back-testing** due to potential biases in historical data.

Examining how assets **correlate in crisis** mode is crucial, so quants use synthetic data to model various market scenarios realistically.

Our analysis provides indications that using **generative neural networks for (correlation) data augmentation** enhances validation and risk management capabilities.

# Generative Neural Networks for correlation analysis



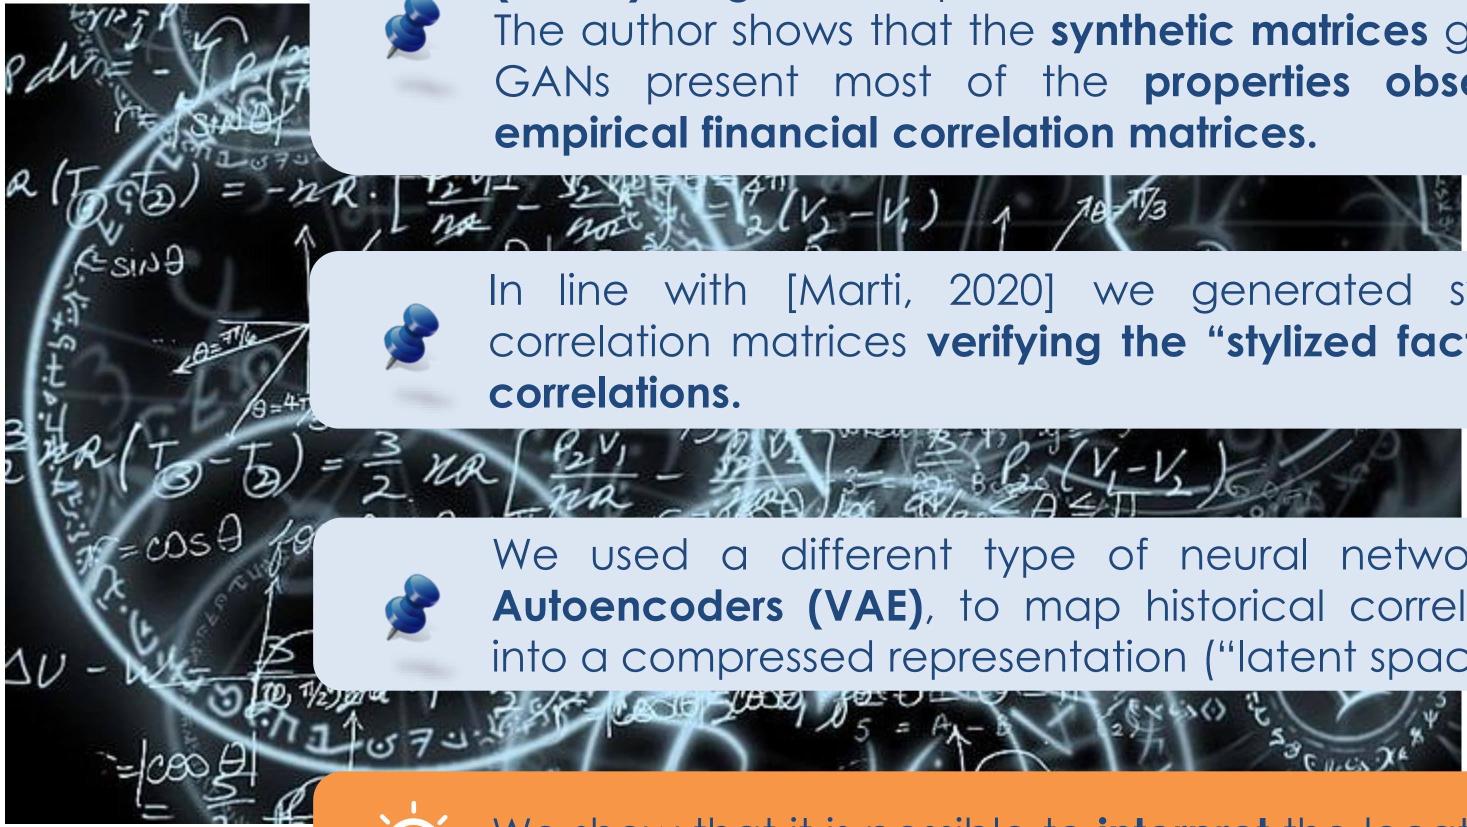

- [Marti, 2020] proposed **Generative Adversarial Networks (GANs)** to generate plausible financial correlation matrices. The author shows that the **synthetic matrices** generated with GANs present most of the **properties observed on the empirical financial correlation matrices.**
- In line with [Marti, 2020] we generated synthetic asset correlation matrices **verifying the "stylized facts" of financial correlations.**
- We used a different type of neural network, **Variational Autoencoders (VAE)**, to map historical correlation matrices into a compressed representation ("latent space").

We show that it is possible to **interpret** the location of points in the latent space, i.e. the **rationale underlying** the mapping.

# Methodological framework



Training **Variational Autoencoder** on historical correlations

Generate synthetic correlations

Analysis of VAE latent space

Check synthetic data “quality”

# Autoencoders

# Auto Encoders



AE is a type of neural network with two parts: an **encoder** and a **decoder**, trained together to reconstruct the original input.

Input: X

Bottleneck:
$C(x) = \sigma(WX + b)$

Reconstruction:
$\hat{X} = f(C(X) = \sigma(W^T c + b)$

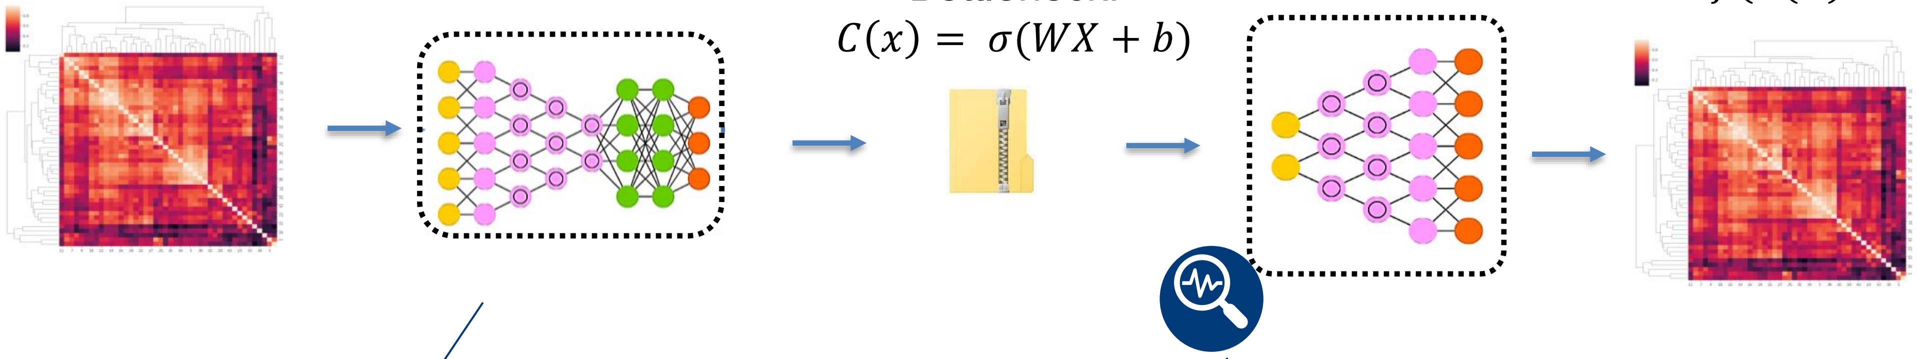

**Encoder**

neural network that maps inputs into a **latent space** (that has few dimensions)

**Decoder**

neural network that maps points in the latent space back to the input space

# Variational Auto Encoders



VAE is a type of neural network with two parts: an **encoder** and a **decoder**, trained together to reconstruct the original image, where **the compressed space is probabilistic** (i.e. it contains a random component).

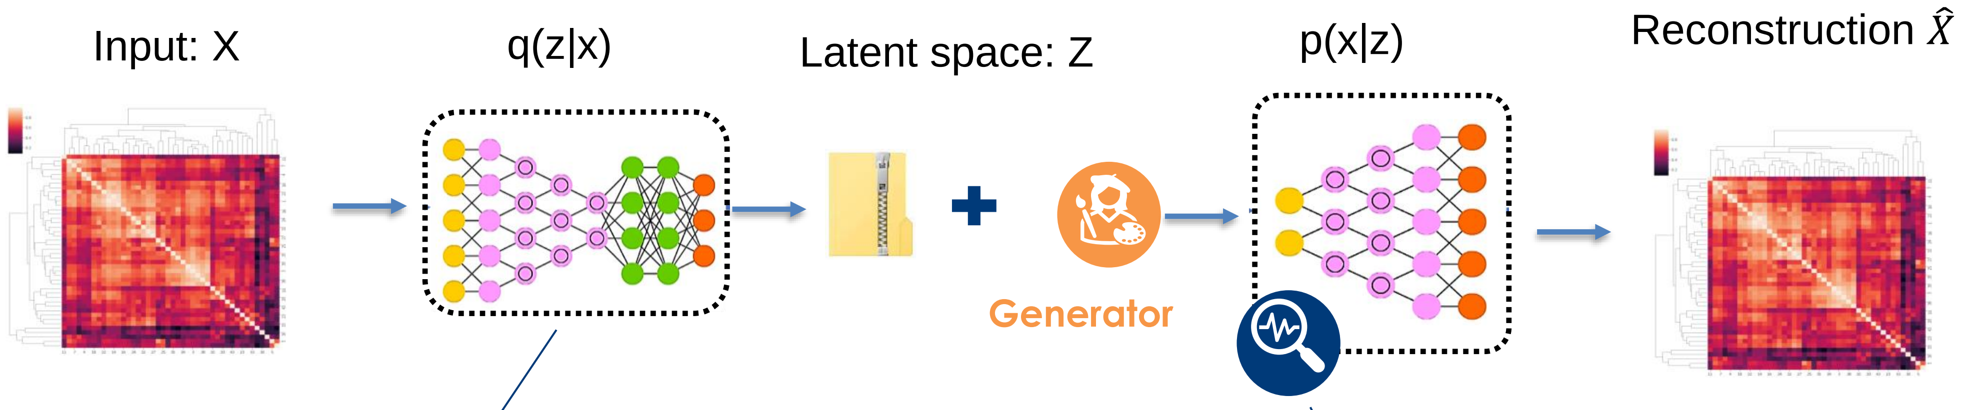


**ENCODER**

neural network that maps inputs into a **latent space** (that has few dimensions)

**DECODER**

neural network that maps points in the latent space back to the input space

# Variational Auto Encoder – Probabilistic rationale



**HIDDEN VARIABLE $Z$**

That generates the observation $X$

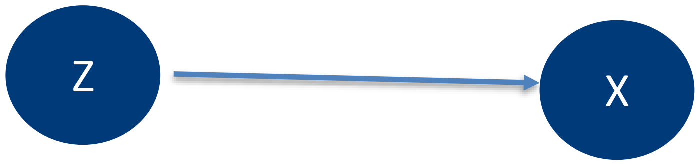


We want to infer the distribution of $Z$ observing $X$

$$p(z|x) = \frac{p(x|z)p(z)}{p(x) = \int p(x|z)p(z)dz}$$

Computing $p(x) = \int p(x|z)p(z)dz$ is difficult.

We approximate $p(z\,|\,x)$ by a **tractable distribution** $q(z\,|\,x)$.

We look for **parameters of** $\boldsymbol{q(z\,|\,x)}$ such that it becomes **very similar to** $\boldsymbol{p(z\,|\,x)}$, so we can use it to perform approximate inference of the intractable distribution.

We can minimize the **Kullback-Leibler** divergence between the two distributions.

# Variational Auto Encoder – Variational lower bound



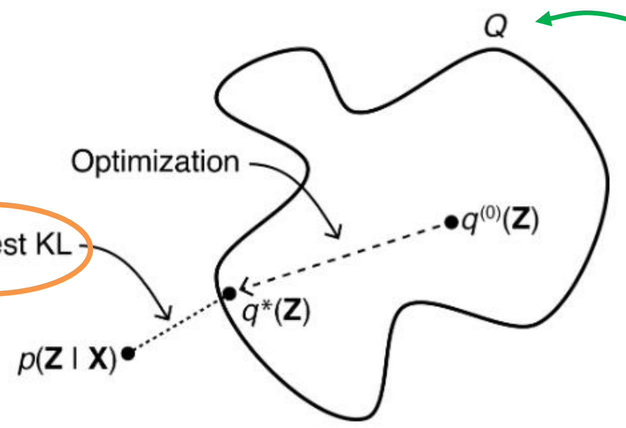


$$
\begin{aligned}
D_{\mathrm{KL}}[q(\mathbf{Z}) \| p(\mathbf{Z}|\mathbf{X})] &= \int_q q(\mathbf{Z}) \log \frac{q(\mathbf{Z})}{p(\mathbf{Z}|\mathbf{X})} \\
&= \mathbb{E}_{q(\mathbf{Z})}\left[\log \frac{q(\mathbf{Z})}{p(\mathbf{Z}|\mathbf{X})}\right] \\
&= \underbrace{\mathbb{E}_{q(\mathbf{Z})}[\log q(\mathbf{Z})] - \mathbb{E}_{q(\mathbf{Z})}[\log p(\mathbf{Z}, \mathbf{X})]}_{-\mathrm{ELBO}(q)} + \log p(\mathbf{X}).
\end{aligned}
$$

Because we cannot compute the desired KL divergence, we **optimize a different objective that is equivalent to this KL divergence up to constant.** This new objective is called the evidence lower bound or ELBO

$$
\mathrm{ELBO}(q) := \mathbb{E}_{q(\mathbf{Z})}[\log p(\mathbf{Z}, \mathbf{X})] - \mathbb{E}_{q(\mathbf{Z})}[\log q(\mathbf{Z})].
$$

$$
\log p(\mathbf{X}) = \mathrm{ELBO}(q) + D_{\mathrm{KL}}[q(\mathbf{Z}) \| p(\mathbf{Z}|\mathbf{X})].
$$

$$
\log p(\mathbf{X}) \geq \mathrm{ELBO}(q).
$$

*Gregory Gundersen, «The ELBO in Variational Inference" https://gregorygundersen.com/blog/2021/04/16/variational-inference/*

# Variational Auto Encoder – Training



The Encoder and Decoder are trained to **minimize a loss** given by two terms:

**reconstruction error term** compares the input $\boldsymbol{x}$ with the output $\hat{\boldsymbol{x}} = D(\boldsymbol{z})$.

**regularization term** is proportional to the Kullback-Leibler divergence of the encoded distribution w.r.t. the target distribution (i.e. to the standard 2D Gaussian).

**Decoder** applied to latent parameters $\boldsymbol{z}$ sampled from the distribution obtained encoding $\boldsymbol{x}$

Gaussian distribution obtained **encoding** $\boldsymbol{x}$

standard Gaussian distribution

correlation matrix in input

regularization strength

$$\text{loss} = \|\boldsymbol{x} - D(\boldsymbol{z})\|^2 + \beta \cdot \text{KL}\left[N\left(\boldsymbol{\mu_x}, \boldsymbol{\sigma_x}\right) \| N(\boldsymbol{0}, \boldsymbol{I})\right]$$

reconstruction error

regularization
(Kullback-Leibler divergence)

# Variational Auto Encoder – Implementation

The loss function **includes the MSE**, that measures the reconstruction error, and the **KL divergence,** that measures the distribution mismatch in the latent space.

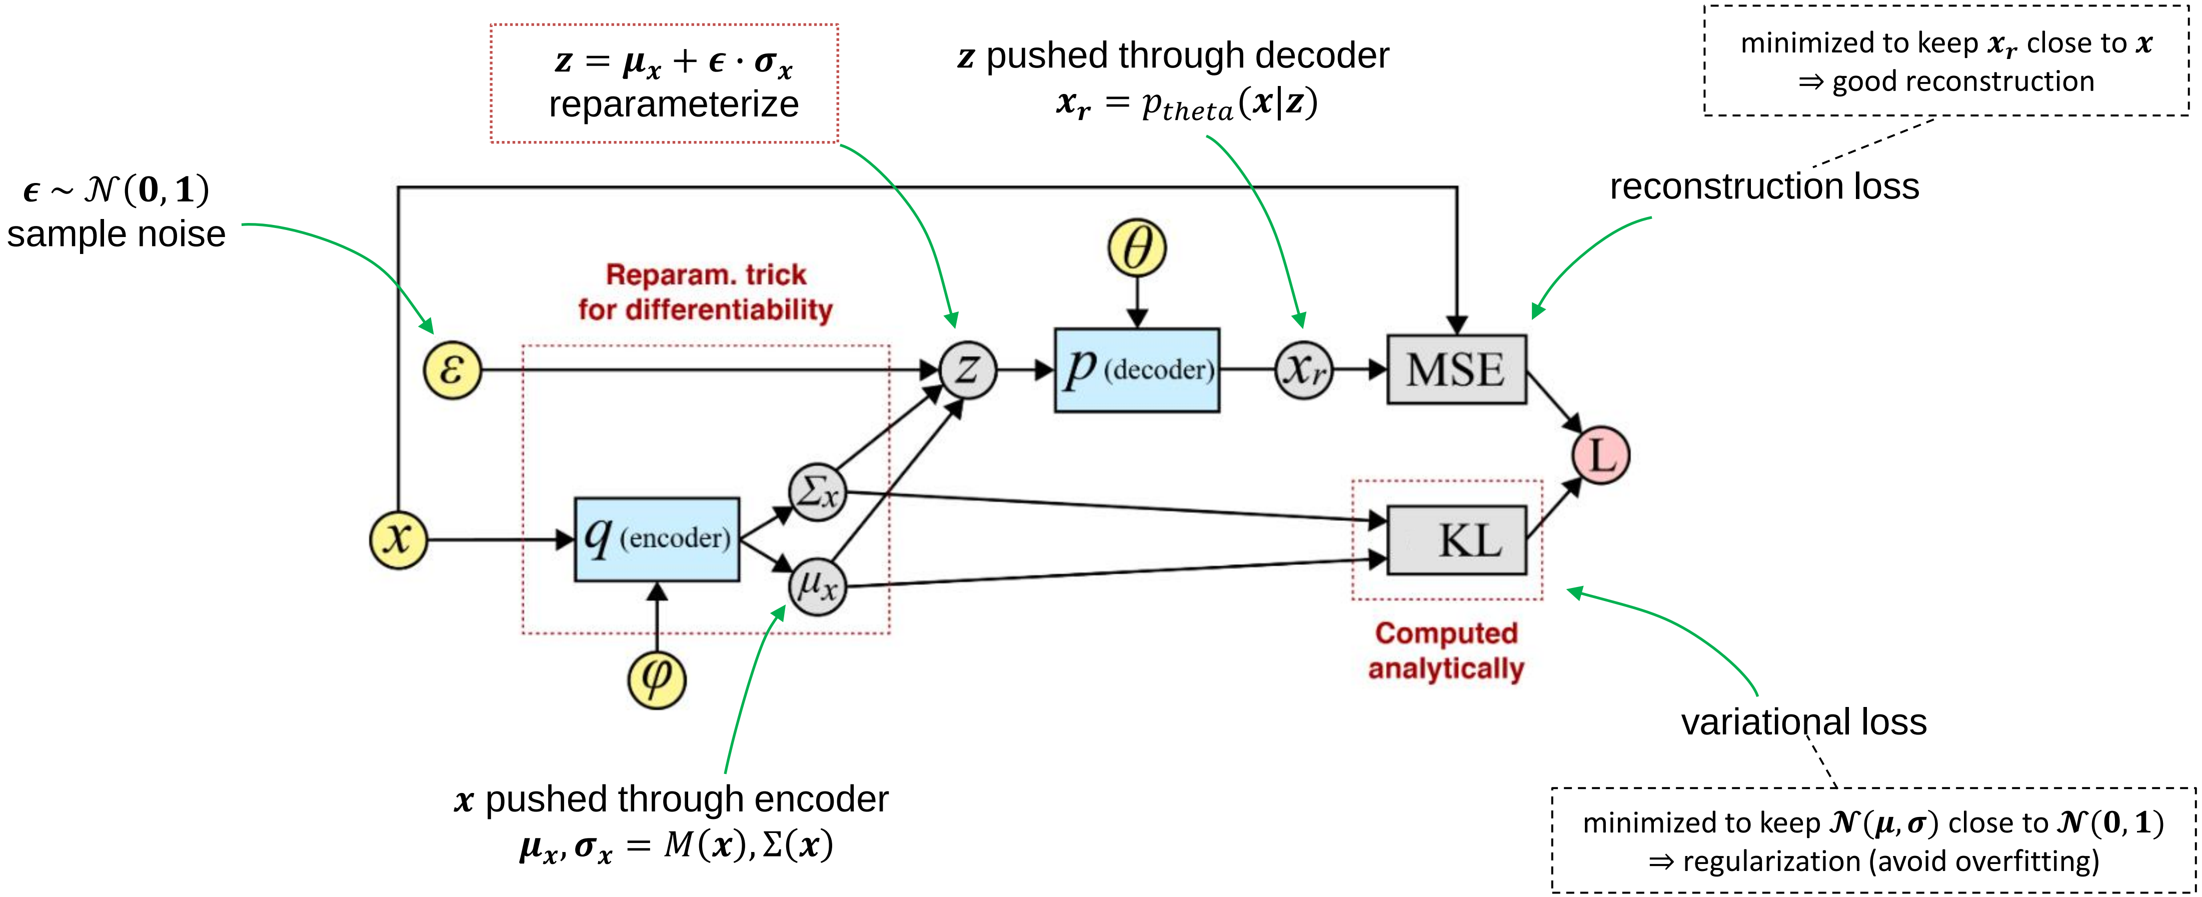


*Gregory Gundersen, The Reparameterization Trick 29 April 2018 https://gregorygundersen.com/blog*

# Variational Auto Encoders – latent attributes



Variational Autoencoders **learn descriptive attributes** (latent attributes) of the input data (such as skin color, smile, glasses, beard, etc.) in an attempt to describe an observation in some **compressed representation**.

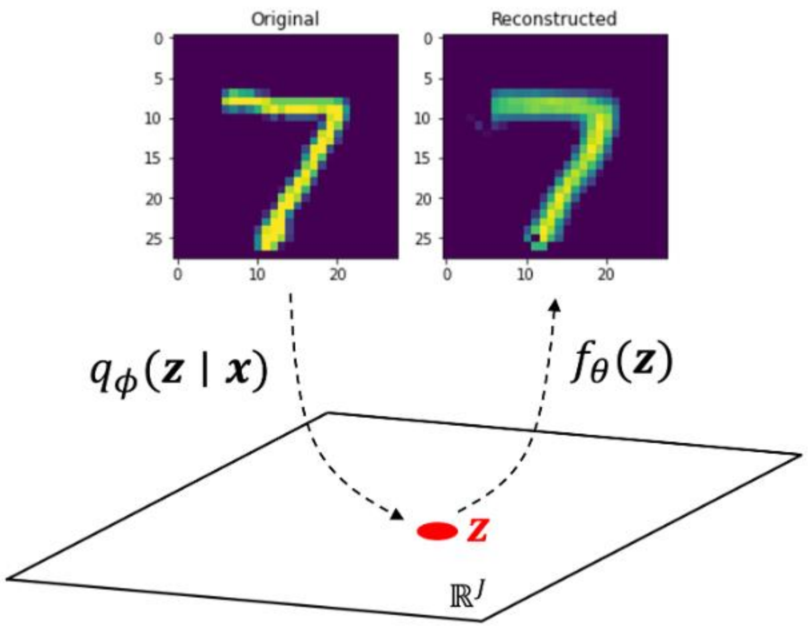


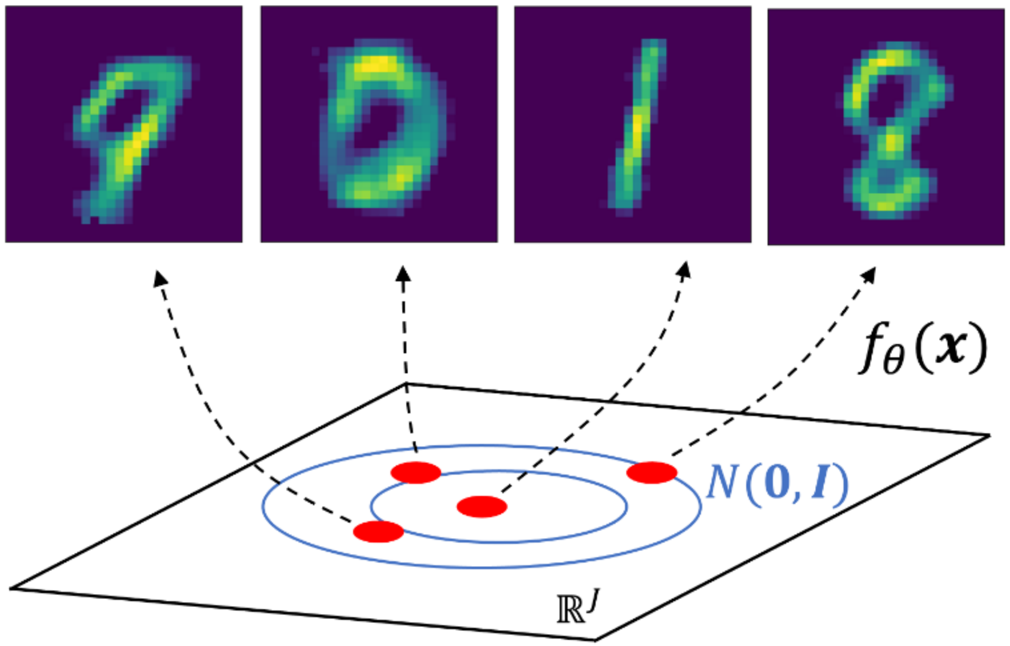


*Matthew N. Bernstein Variational autoencoders March 14, 2023 https://mbernste.github.io/posts/vae/*

# Variational Auto Encoders – probabilistic structure



A Variational Autoencoder represents each latent attribute as a **range of possible values**. It describes **latent attributes in probabilistic terms**.

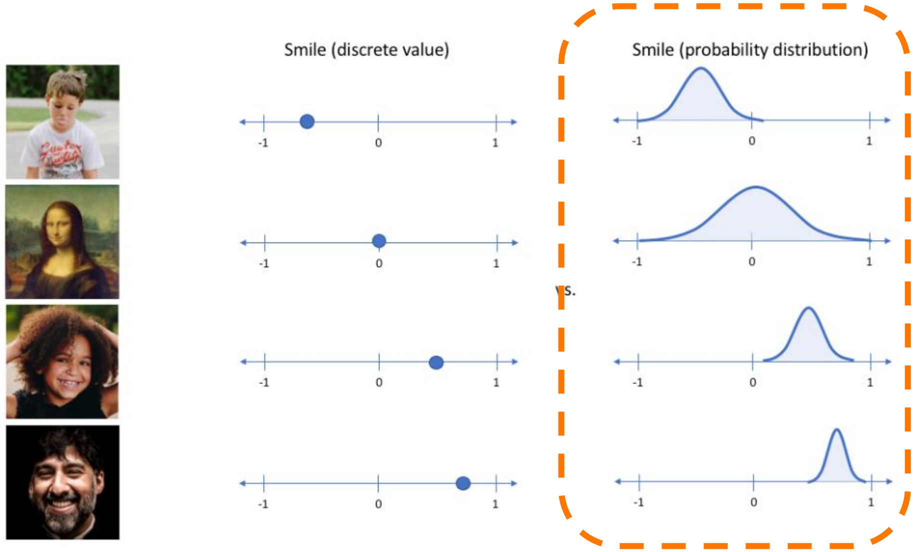


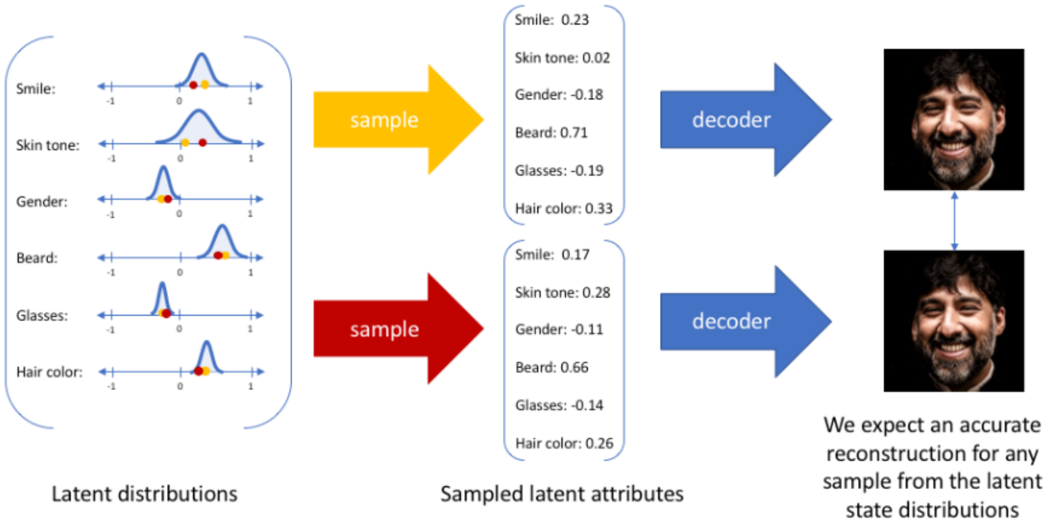


**VAE map images into a set of "latent characteristics".**

**Each feature is expressed as a random variable**

*Jeremy Jordan, Variational Autoencoders https://www.jeremyjordan.me/variational-autoencoders/*

# Variational Auto Encoders – latent space



For any sample from the latent distributions, we expect the decoder to be able to accurately reconstruct the input. **Values that are close to each other in latent space** should correspond to **very similar reconstructions**.

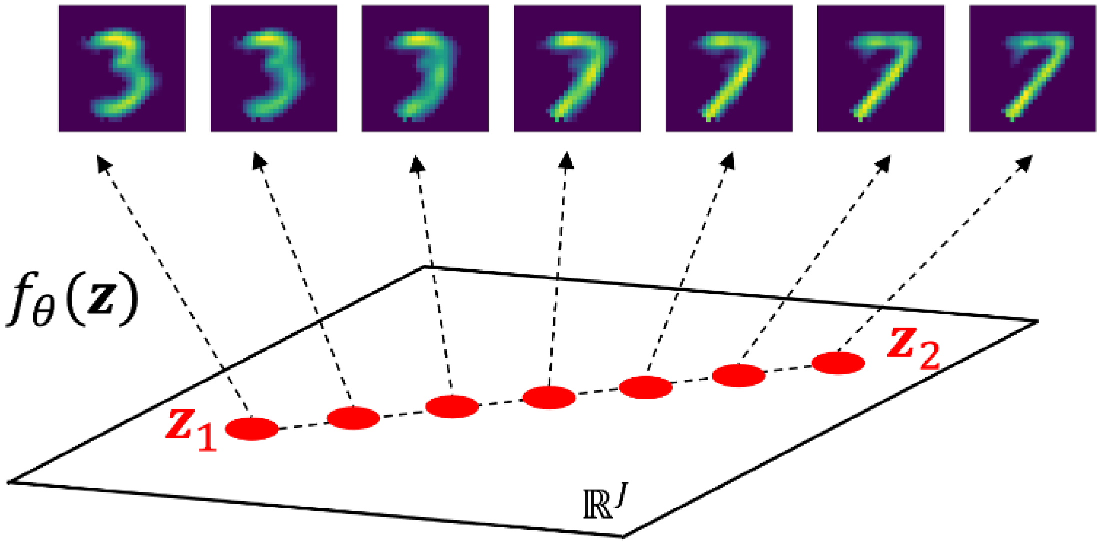


*Matthew N. Bernstein Variational autoencoders March 14, 2023 https://mbernste.github.io/posts/vae/*

# Synthetic data “quality”

# Variational Auto Encoders – Data augmentation



The idea is to use VAE for **data augmentation**, harnessing the **probabilistic nature of the latent space.**

The assumption underlying the latent space is that the data $\boldsymbol{x}$ are **generated from a random process** involving an unobserved continuous random variable $\boldsymbol{z}$ (latent variable).

The **"probabilistic decoder"** is a distribution $\boldsymbol{p(x|z)}$, while **the "probabilistic encoder"** is a distribution $\boldsymbol{q(z|x)}$.

The parameters of these distributions are given by two **neural networks** (encoder and decoder).

# Variational Auto Encoders – data generation



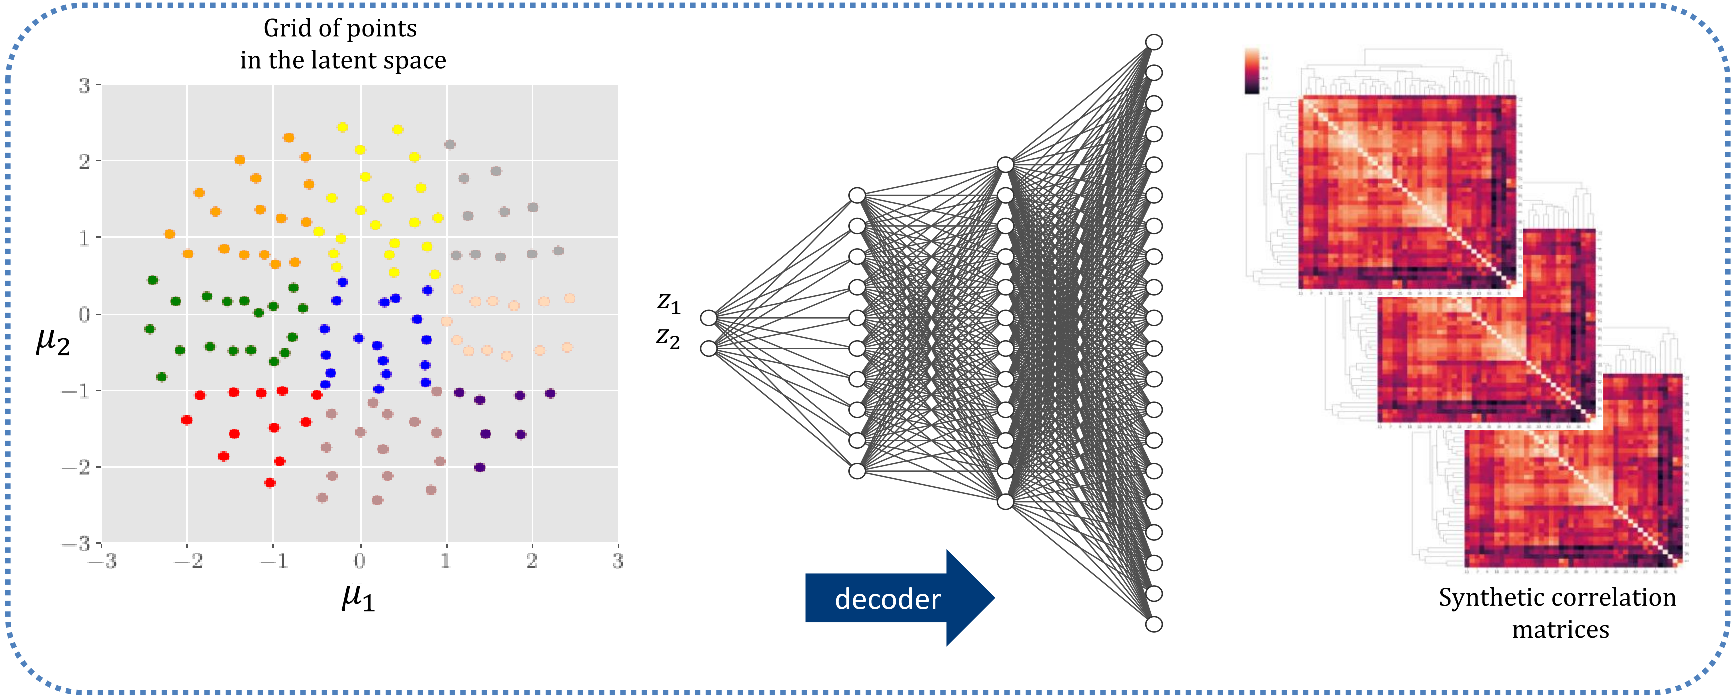


The VAE decoder **generates a "plausible" correlation matrix** from a point in the latent space.

We define a **grid of points of the latent space** and use the decoder to compute **the corresponding correlation matrix**.

We check whether the **stylized facts** of financial correlation matrices hold, for both the historical and the synthetic matrices.

INTESA SANPAOLO

# Stylized facts of financial correlations



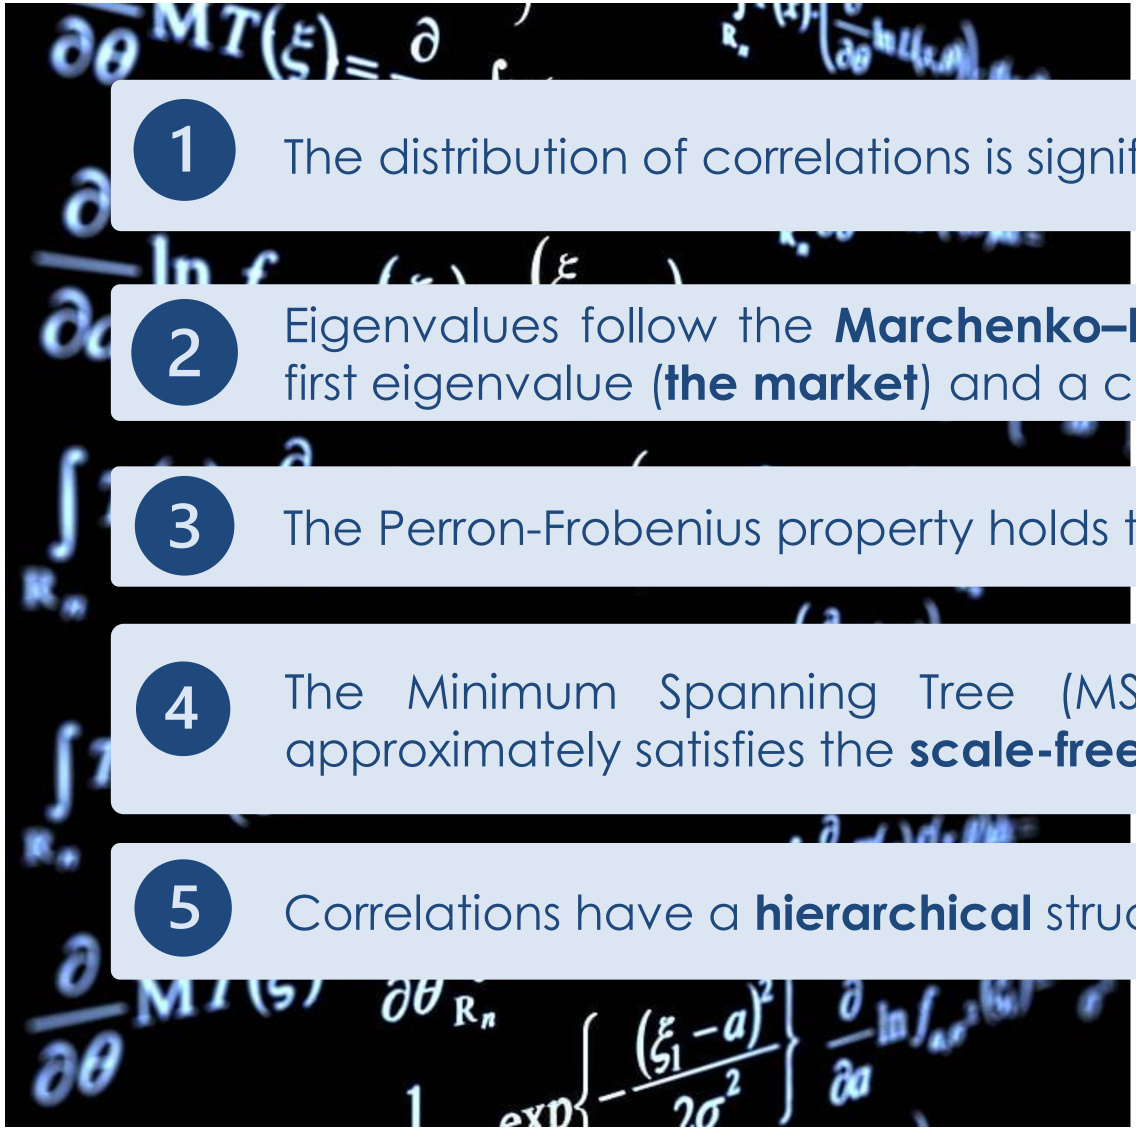

1. The distribution of correlations is significantly shifted towards **positive** values.
2. Eigenvalues follow the **Marchenko–Pastur distribution**, except for a very large first eigenvalue (**the market**) and a couple of other large eigenvalues.
3. The Perron-Frobenius property holds true (**first eigenvector has positive entries**).
4. The Minimum Spanning Tree (MST) obtained from a correlation matrix approximately satisfies the **scale-free property**.
5. Correlations have a **hierarchical** structure.

# Stylized facts : Eigenvalues and eigenvectors



Financial correlation matrices have been extensively studied leveraging on **random matrix theory** (RMT) and **graph analysis**.

Under the lenses of RMT, the **eigenvalues** and **eigenvectors** of the correlation matrix are compared to those of a random matrix, generated from random uncorrelated time-series.

"The **largest eigenvalue** of the correlation matrix is a measure of the **intensity of the correlation** present in the matrix […] Generally this largest eigenvalue is **larger during times of stress and smaller during times of calm**." [Millington and Niranjan, 2021a]

Each **eigenvector** can be viewed as a set of **weights of assets** that defines a new index which is uncorrelated with the other eigenvectors. It follows that a change in eigenvectors can impact **portfolio diversification**. [Nguyen et al., 2018]

# Stylized facts: graph analysis



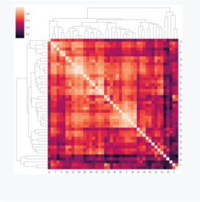

Construct a "**distance matrix**" from the correlation matrix.

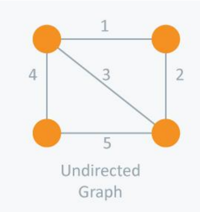


Use this distance matrix as **adjacency matrix** for a "distance graph" **G** whose nodes represent the assets**.**
The distance is used as "weights" of the edges.

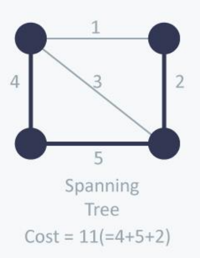


A **spanning tree** of the graph is a **subgraph** that includes every vertex of G and does not contain any loop.

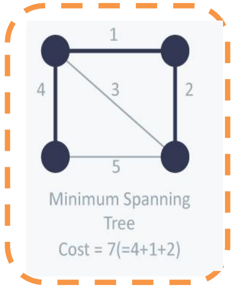


A minimum spanning tree (**MST**) is a spanning tree having the **minimum possible weights** among all possible spanning trees.

**Random network**

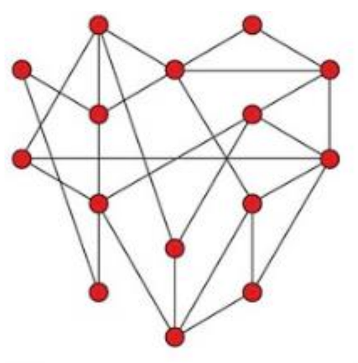

**Scale- free network**

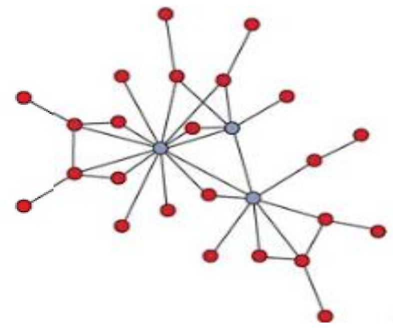

Very few nodes have high degrees while most nodes have degree equal to 1.

# Stylized facts: graph analysis - MST algorithms



1. Determine an arbitrary vertex as the starting vertex of the **Minimum Spanning Tree** (MST).
2. Find edges connecting any tree vertex with the fringe vertices.
3. Find the minimum among these edges.
4. Add the chosen edge to the MST if it does not form any cycle.
5. Repeat steps 2,3,4 till there are vertices that are not included in the MST (known as fringe vertex).

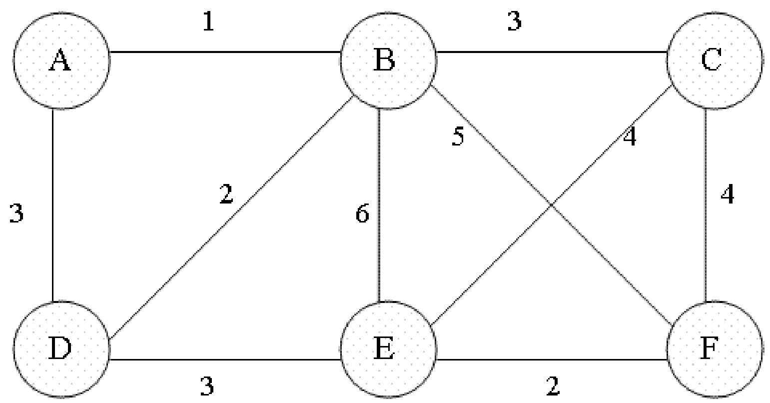


INTESA SANPAOLO

# Hierarchical structure



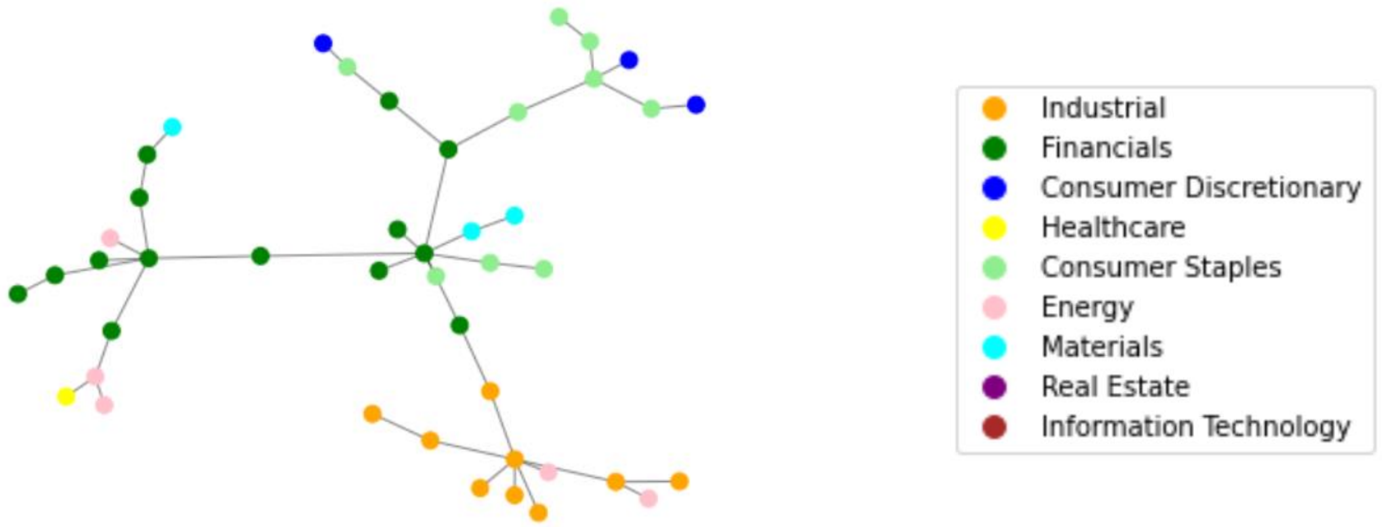


Nodes are colored according to their economic sector membership. Some clustering of nodes is observed: branches of the trees tend to contain **companies in the same sector** and there are several **hub nodes**.

# Application: Credit Concentration Risk

# Literature



## Variational Autoencoders in finance

[Bergeron, Hull et al., 2022] used **Variational Autoencoders** to estimate missing points on partially observed **volatility surfaces**.

[Brugiere and Turinici, 2023] proposed **Variational Autoencoders** to compute an estimator of the **Value at Risk** for a financial asset.

[Sokol 2022] used **Variational Autoencoders** to simulate interest **rates** curves.

# Credit Concentration Risk



*A risk concentration is any* ***single exposure or group of exposures*** *with the potential to produce* ***losses*** *large enough (relative to a bank's capital, total assets, or overall risk level) to threaten a bank's health or ability to maintain its core operations.* ***Risk concentrations*** *are arguably the single most important cause of major problems in banks.*

*Credit risk concentrations, by their nature, are based on* ***common or correlated risk factors****, which, in times of stress, have an adverse effect on the creditworthiness of each of the individual counterparties making up the concentration. Concentration risk arises in both direct exposures to obligors and may also occur through exposures to protection providers. Such concentrations are not addressed in the* ***Pillar 1 capital charge for credit risk****.*

*[..]Banks should explicitly consider the extent of their credit risk concentrations in their assessment of* ***capital adequacy under Pillar 2****.[..]*

*Basel Committee on Banking Supervision Supervisory review process SRP32 Credit Risk 01 January 2023*

# Credit Portfolio Concentration Risk (1/3)



One of the most adopted models to measure the **credit risk of a loan portfolio** was proposed by Vasicek and it is currently a market standard used by regulators for capital requirements. This model provides a **closed-form expression** to measure the risk in the case of **asymptotic single risk factor (ASRF)** portfolios.

A commonly adopted methodology of measuring **concentration risk** is to use a **Monte Carlo simulation of the portfolio loss distribution** according to a more general model of **multiple systematic factors.**

$$V_i = \rho_i Y_j + \sqrt{1-\rho_i^2}\,\varepsilon_i = \rho_i \sum_f \alpha_{j,f} Z_f + \sqrt{1-\rho_i^2}\,\varepsilon_i$$

value of the $i$-th counterparty

systematic factor $Y_j$ belonging to **a set of macroeconomic factors** $\{\boldsymbol{Y_j}\}$, with correlation matrix $\boldsymbol{\Sigma}$

idiosyncratic independent process

**correlation matrix**
$\boldsymbol{\alpha\,\alpha' = \Sigma}$

independent factors

# Credit Portfolio Concentration Risk (2/3)



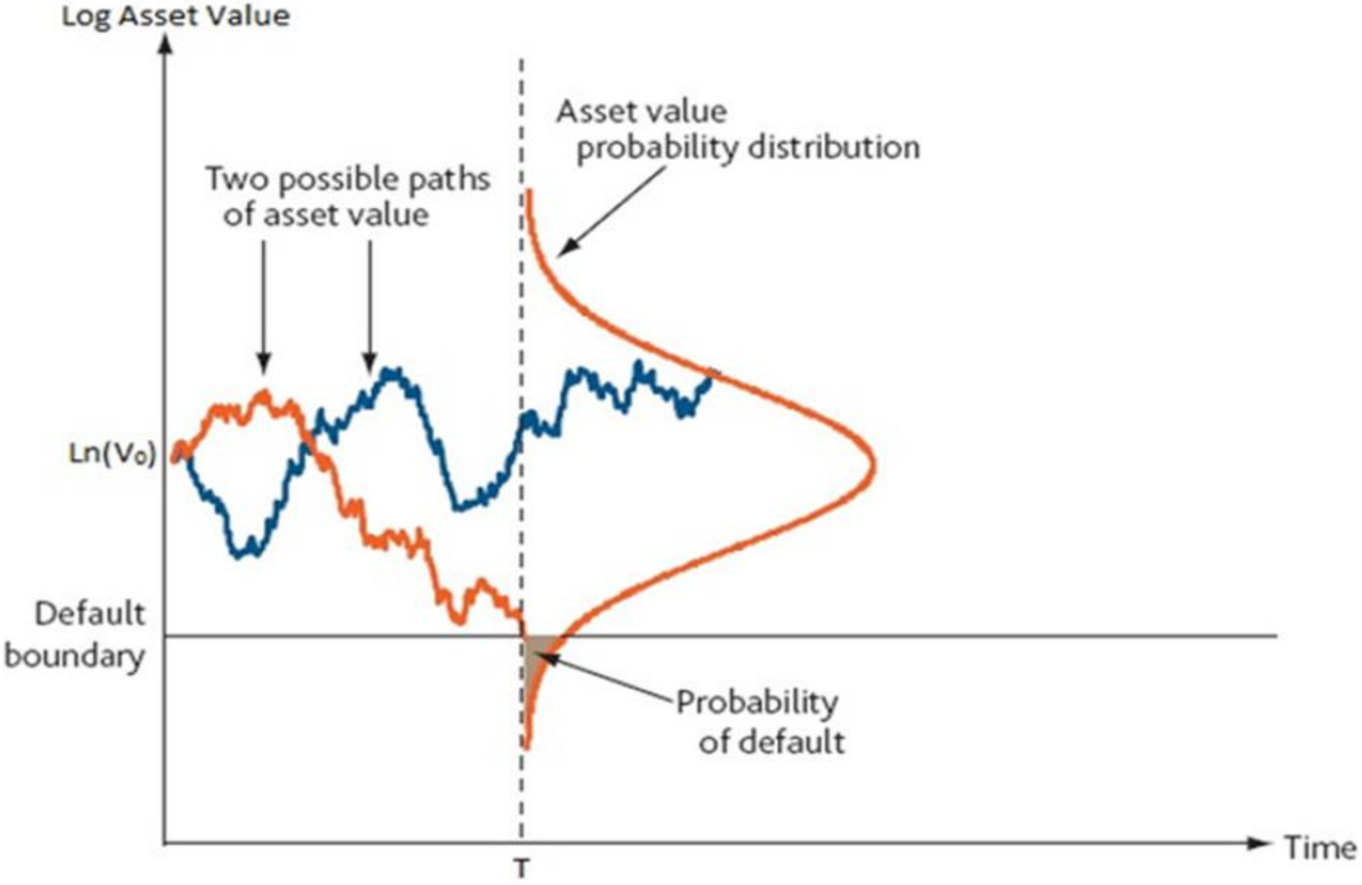

# Credit Portfolio Concentration Risk (3/3)



The bank's portfolio is clustered into **sub-portfolios** with homogeneous **risk characteristics** (i.e. economic sector, geographical area, rating class or counterparty size).

A **distribution of losses** is simulated for each sub-portfolio and the **Value at Risk (VaR)** is calculated on the aggregated loss.

The asset **correlation matrix Σ** is a critical parameter for the estimation of the sub-portfolio loss distribution, that is the **core component** for the estimation of the **concentration risk**.

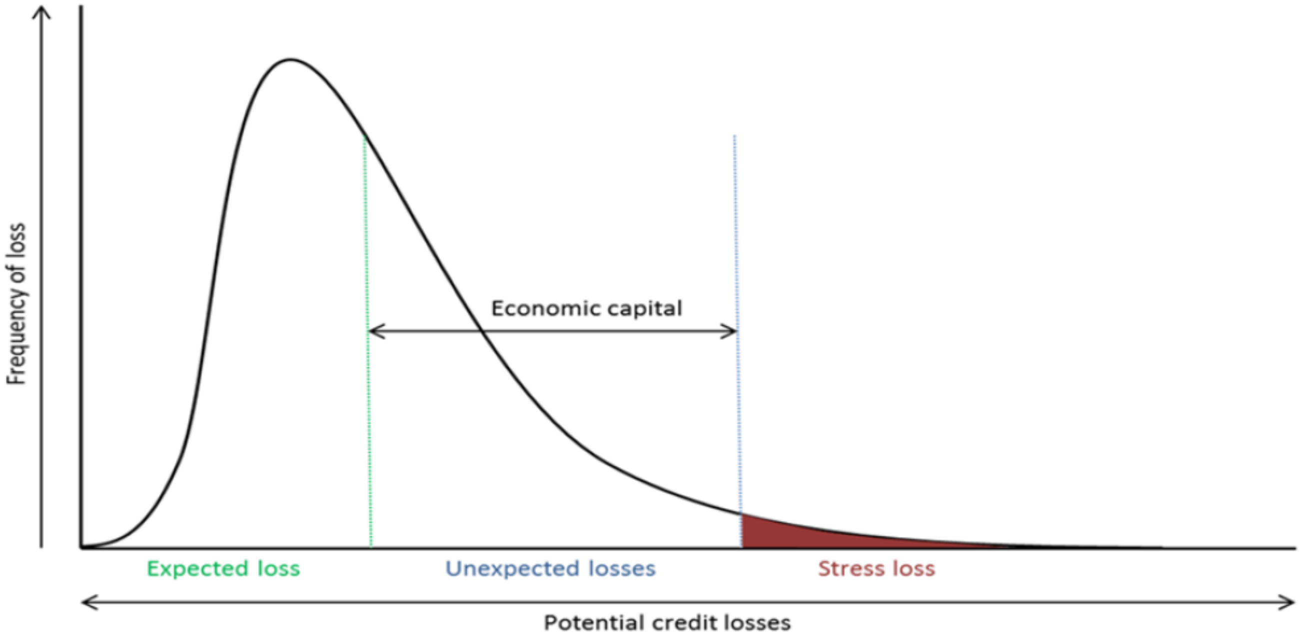

# The dataset - historical correlation matrices



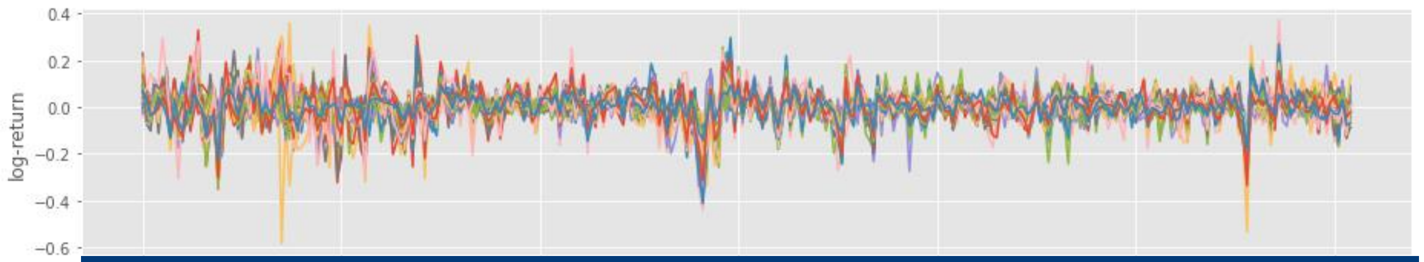


We randomly split (70% - 30%) the dataset in a training sample, used to train the network, and a validation set, used to evaluate the performance.

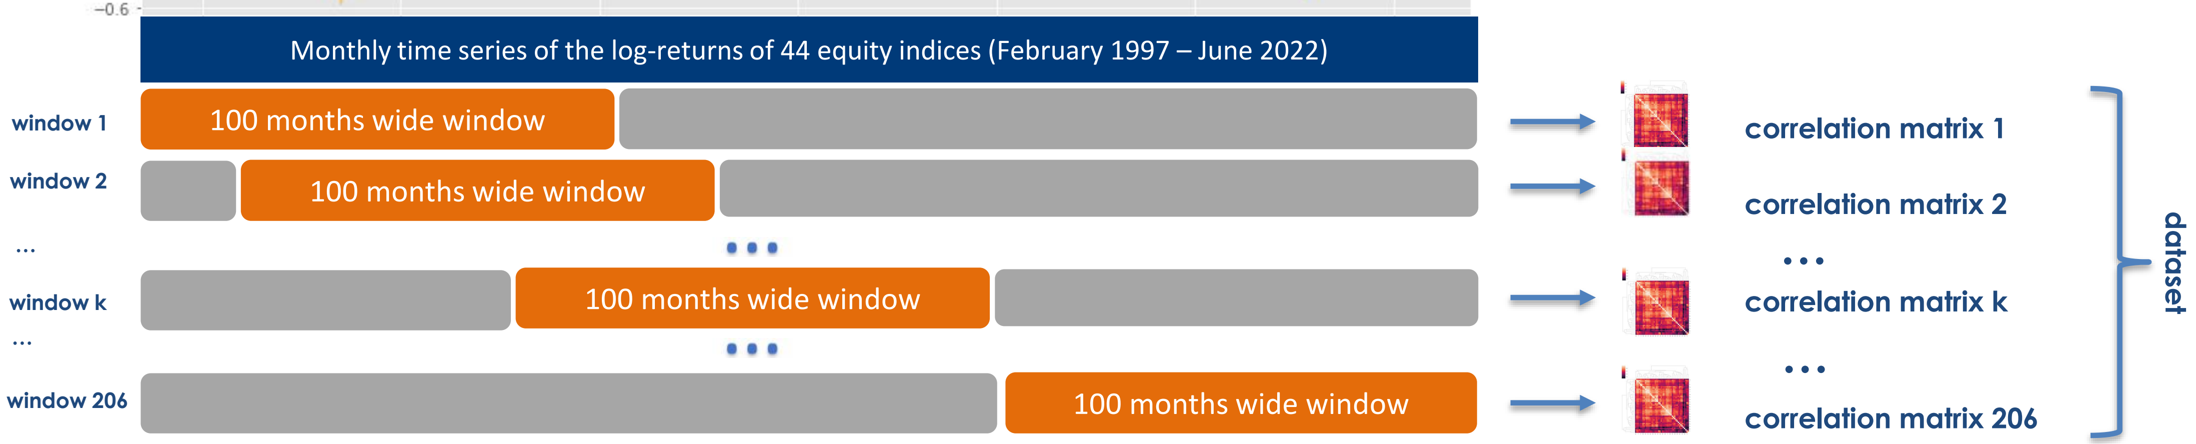


The dataset contains **206 correlation matrices** of the monthly log-returns of **44 equity indices**(*), calculated on their monthly time series from February 1997 to June 2022, using overlapping rolling windows of size 100 months.

*(*) Historical time series considered are Total Market (Italy, Europe, US and Emerging Markets) and their related sector indices (Consumer Discretionary, Energy, Basic Materials, Industrials, Consumer Staples, Telecom, Utilities, Technology, Financials, Health Care), the source is Datastream.*

# Variational Auto Encoders – Correlation Matrix



## Encoder

Encoding layers compress the values of the $44 \times 44$ correlation matrix into the values of just 4 nodes (that represent the mean and variance of a bivariate gaussian distribution).

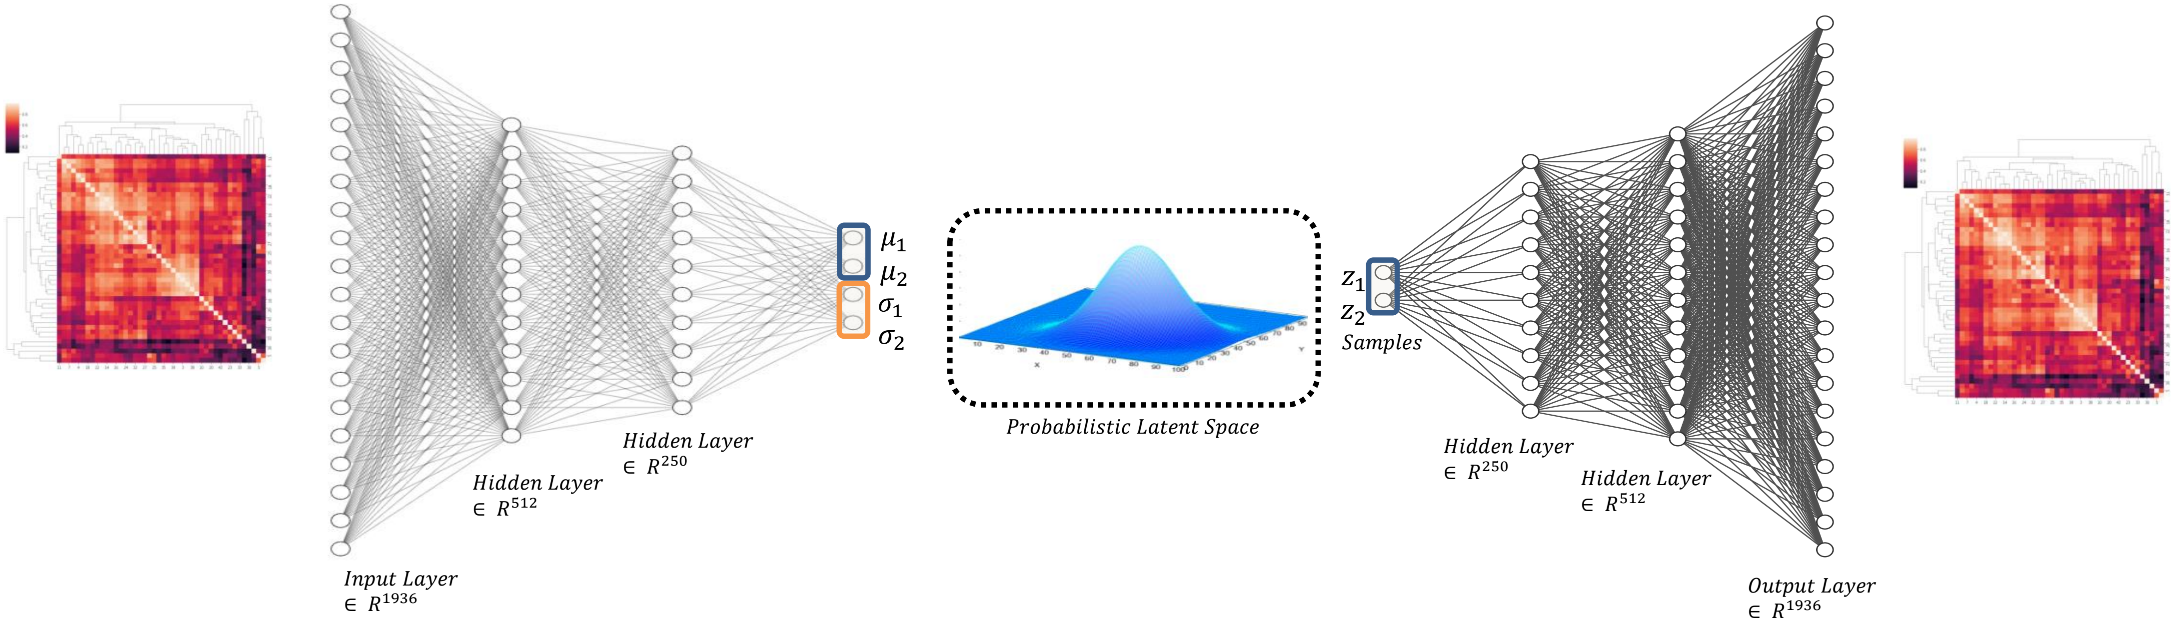


## Latent space

The input matrix is encoded into a 2D Gaussian distribution. A sample from this 2D Gaussian is drawn and it is fed to the decoder.

## Decoder

Decoding layers recover the $44 \times 44$ values representing the output matrix.

# Comparison with linear models

# Comparison with linear models



We compare the performances of the Variational Autoencoder with the **deterministic Autoencoder** (AE) and with **Principal Component Analysis (PCA).**

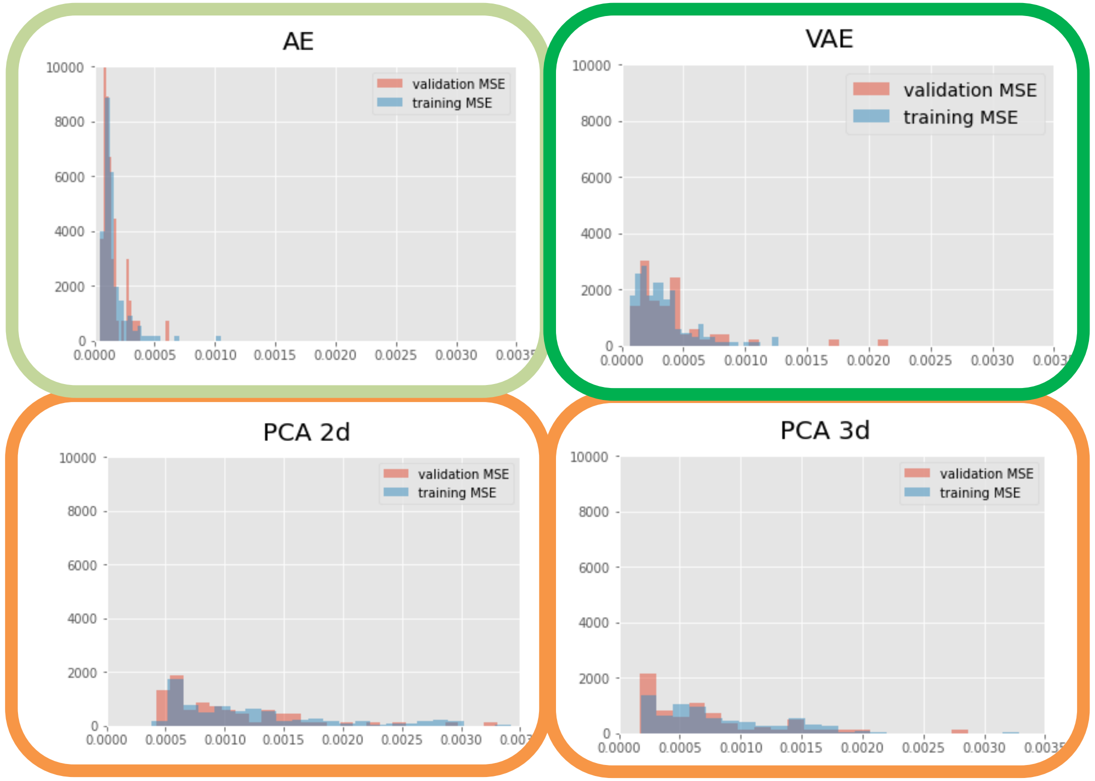


- **Linear models (PCA) have lower performances** in terms of Mean Squared Error both in sample and out of sample, even increasing the dimensions of the latent space.
- Hence, **neural networks actually bring an improvement** in minimizing the reconstruction error.
- The generative probabilistic component of VAE decreases the performance when compared to a **deterministic autoencoder (AE)**. On the other hand, it **allows to generate new correlation matrices** (that are realistic in the sense of stylized facts).

The **linear autoencoder** is equivalent to applying **PCA** to the input data in the sense that its output is a projection of the data onto the low dimensional principal subspace [Plaut, 2018].

# Latent space comparison



The Variational Autoencoder maps each one of the 206 historical correlation matrices into a bivariate normal distribution in a two-dimensional probabilistic latent space. Here we show the location of the encoded distributions.

VAE

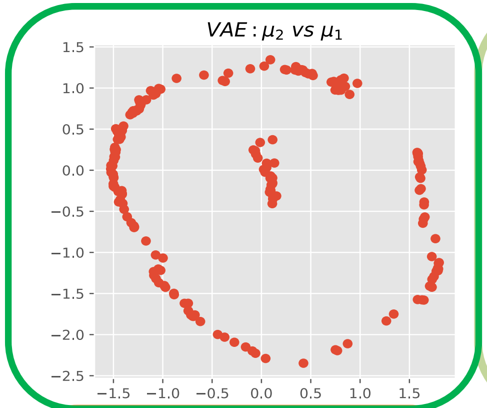


AE

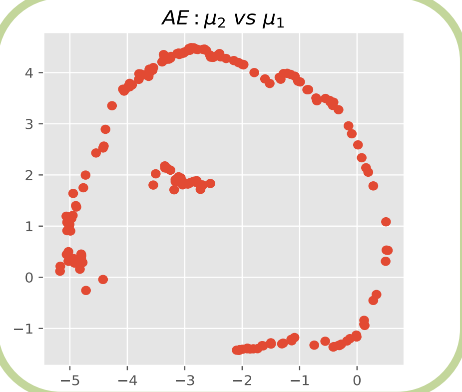


2D PCA

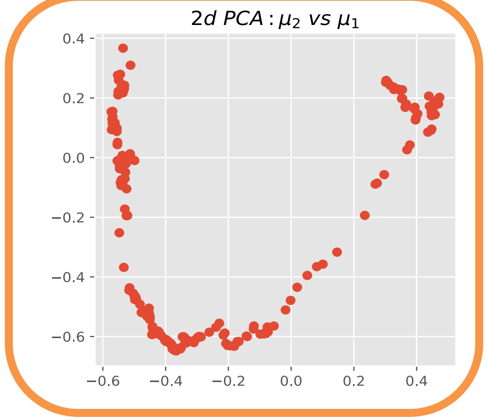


3D PCA

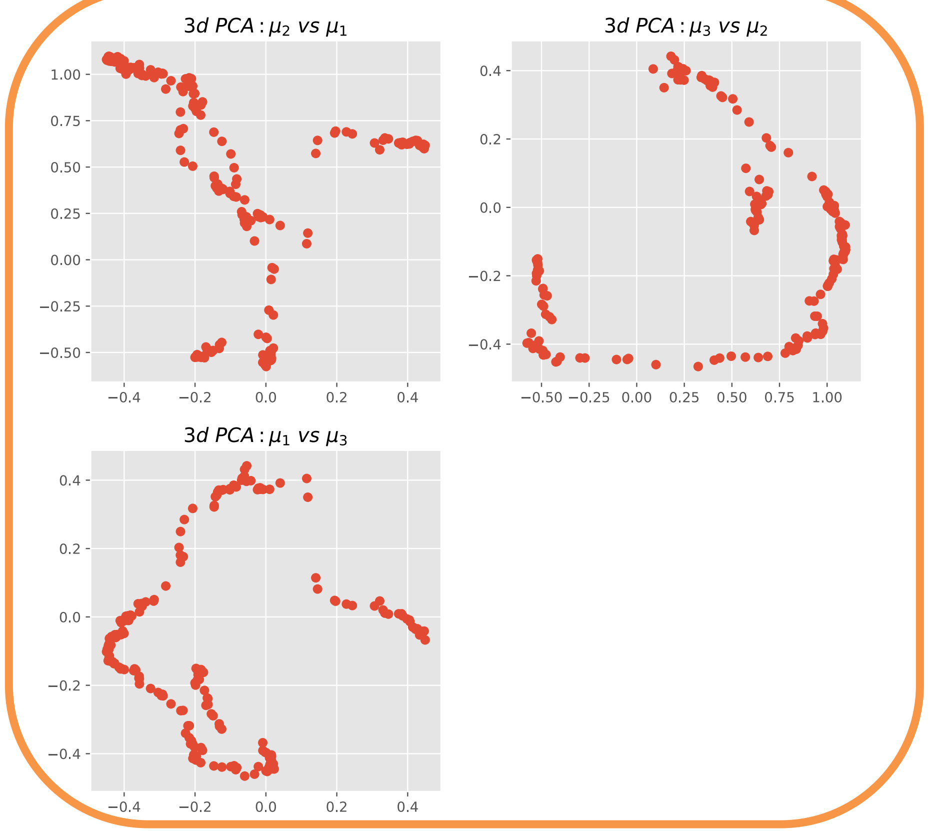


The latent spaces generated by the VAE and AE are similar, while among the linear autoencoders (PCA) **only the 3-dimensional one highlights the cluster of points in the middle**.

# Latent space interpretability

# Artificial Intelligence and Risk Management



NIST NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY U.S. DEPARTMENT OF COMMERCE

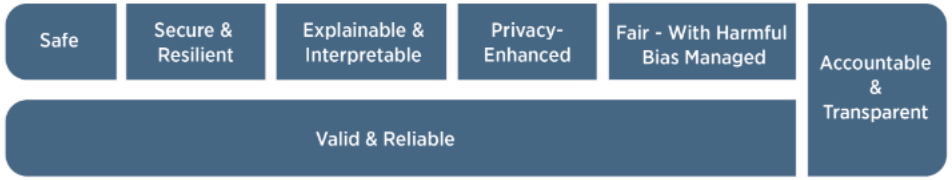


FIGURE 4: CHARACTERISTICS OF TRUSTWORTHY AI SYSTEMS. VALID & RELIABLE IS A NECESSARY CONDITION OF TRUSTWORTHINESS AND IS SHOWN AS THE BASE FOR OTHER TRUSTWORTHINESS CHARACTERISTICS. ACCOUNTABLE & TRANSPARENT IS SHOWN AS A VERTICAL BOX BECAUSE IT RELATES TO ALL OTHER CHARACTERISTICS.

***Transparency*** *can answer the question of* ***"what happened"*** *in the system.*

***Interpretability*** *can answer the question of* ***"why"*** *a decision was made by the system and its meaning or context to the user.*

***Explainability*** *can answer the question of* ***"how"*** *a decision was made in the system.*

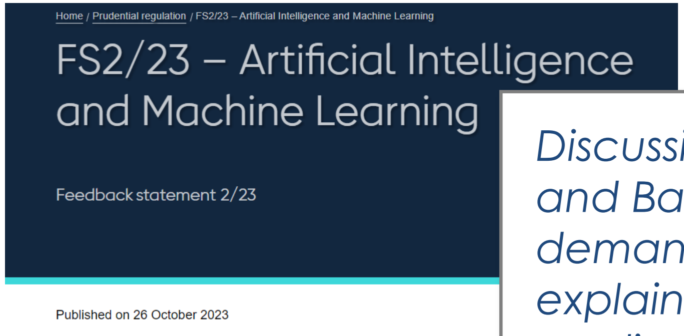


*Discussion paper 5/22 by FCA and Bank of England highlights demand for "more clarity on explainability for AI applications".*


*"ML models can quickly become "**black boxes**", opaque systems for which the internal behaviour cannot be easily understood, and for which therefore it is not easy to understand (and verify) how a model has reached a certain conclusion or prediction."*

# Interpretability vs explainability



## Definitions

According to [Miller, 2019] and [Lipton, 2018]:

**Interpretable** is a model such that an observer can **understand** the cause of a decision.

**Explanation** is one mode in which an observer may obtain understanding in the latent space, for instance, building a simple **surrogate model** that mimics the original model to gain a better understanding of the original model's underlying mechanics.

For the sake of our analysis, we refer to the "**interpretability**" of the VAE as the possibility to understand the reason underlying the responses produced by the algorithm in the **latent space**.

To understand the rationales underlying such representation, we analysed the relationship of the **latent parameters** with the **eigenvectors and eigenvalues** of the original correlation matrices, that are known to be linked to the variance and diversification features of the portfolio.

# Correlations impact on VaR and VAE latent space



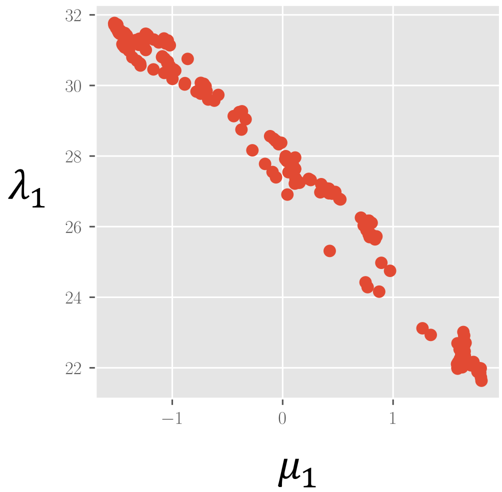


Linear dependence between the first **eigenvalue ($\lambda_1$)** of the correlation matrix and the first latent parameter of the VAE ($\mu_1$).

We expect the VaR to grow with the first eigenvalue and indeed it has a distinct negative linear dependence on $\mu_1$.

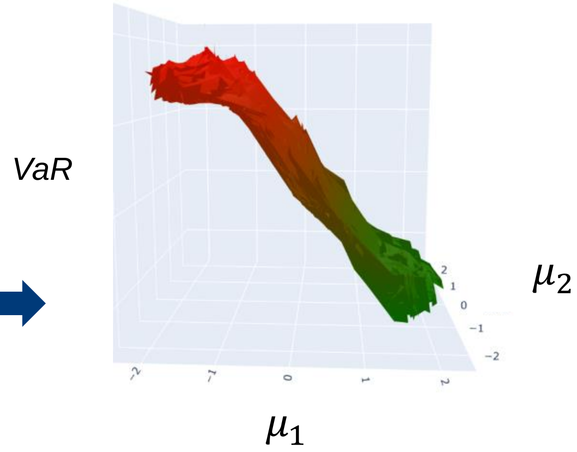


The first latent parameter $\mu_1$ is inversely proportional to the overall correlation, i.e. it expresses the "**diversification opportunities**" on the market.

# Variational Auto Encoders – Interpretability (1/3)



To understand the second latent space dimension, we consider the **cosine similarity $\boldsymbol{\alpha_{i,t}}$** between each eigenvector $\boldsymbol{\nu_{i,t}}$ at a specific time and its average over time $\boldsymbol{\bar{\nu}_i}$.

$$\alpha_{i,t} \equiv \frac{\nu_{i,t} \cdot \bar{\nu}_i}{\|\nu_{i,t}\| \, \|\bar{\nu}_i\|}$$

$$\bar{\nu}_i \equiv \sum_{t=1}^{n} \frac{1}{n} \nu_{i,t}$$

*$i$ is the number of the eigenvector and $t$ is the index of the matrix in the dataset.*

The groups of points in the space $(\alpha_1, \alpha_2, \lambda_1)$ are mapped by the VAE into groups of points in the latent space $(\mu_1, \mu_2)$.
Hence **the "financial proximity" of the correlation matrices is mapped into spatial proximity in the latent space (with a nonlinear mapping)**.

# Variational Auto Encoders – Interpretability (2/3)

Historical

**Nine regions** identified by different colors.

LATENT VARIABLES

portfolio weights

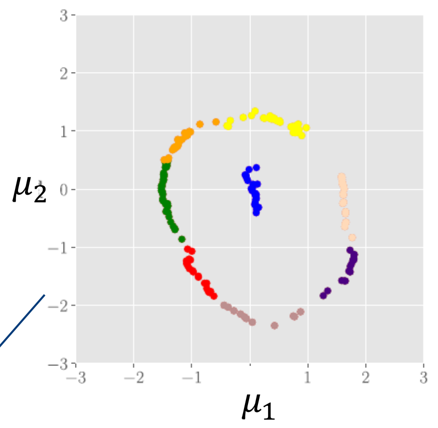


diversification opportunities

FINANCIAL FEATURES

eigenvectors' variation

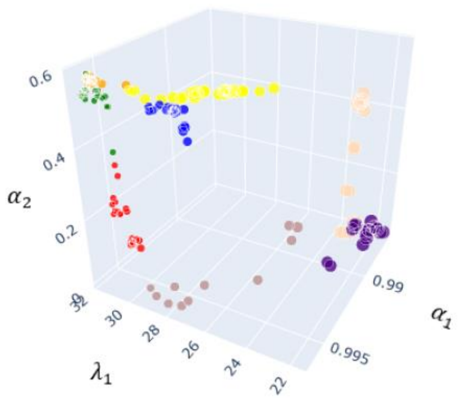


1st eigenvalue

It turns out that groups of points in the space $(\alpha_1, \alpha_2, \lambda_1)$ are mapped by the VAE into **coherent groups** of points in the latent space $(\mu_1, \mu_2)$.

Hence, we find that **the VAE mapping is continuous in meaningful financial features related to diversification.**

# Variational Auto Encoders – Interpretability (3/3)



Sampled

LATENT VARIABLES

portfolio weights

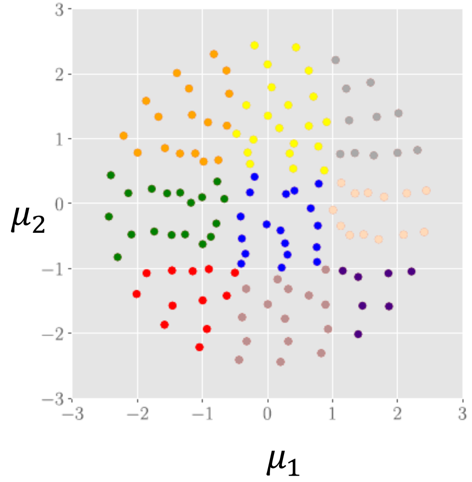


diversification opportunities

FINANCIAL FEATURES

eigenvectors’ variation

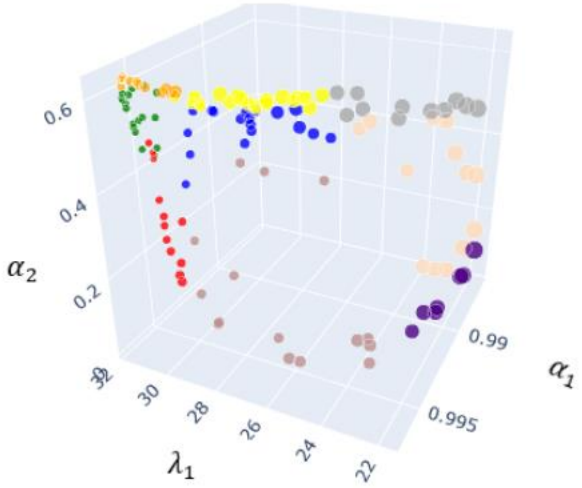


1[st] eigenvalue

Also **for a grid of random samples**, the distances of the first two eigenvectors from their time average $(\alpha_1, \alpha_2)$ and the first eigenvalue ( $\lambda_1$) characterize the regions of the latent space: the **VAE’s mapping preserves the spatial coherence of the groups of points** in the “financial space”.

# Quantifying the sensitivity to asset correlations

# Variational Auto Encoders – Value at Risk estimation



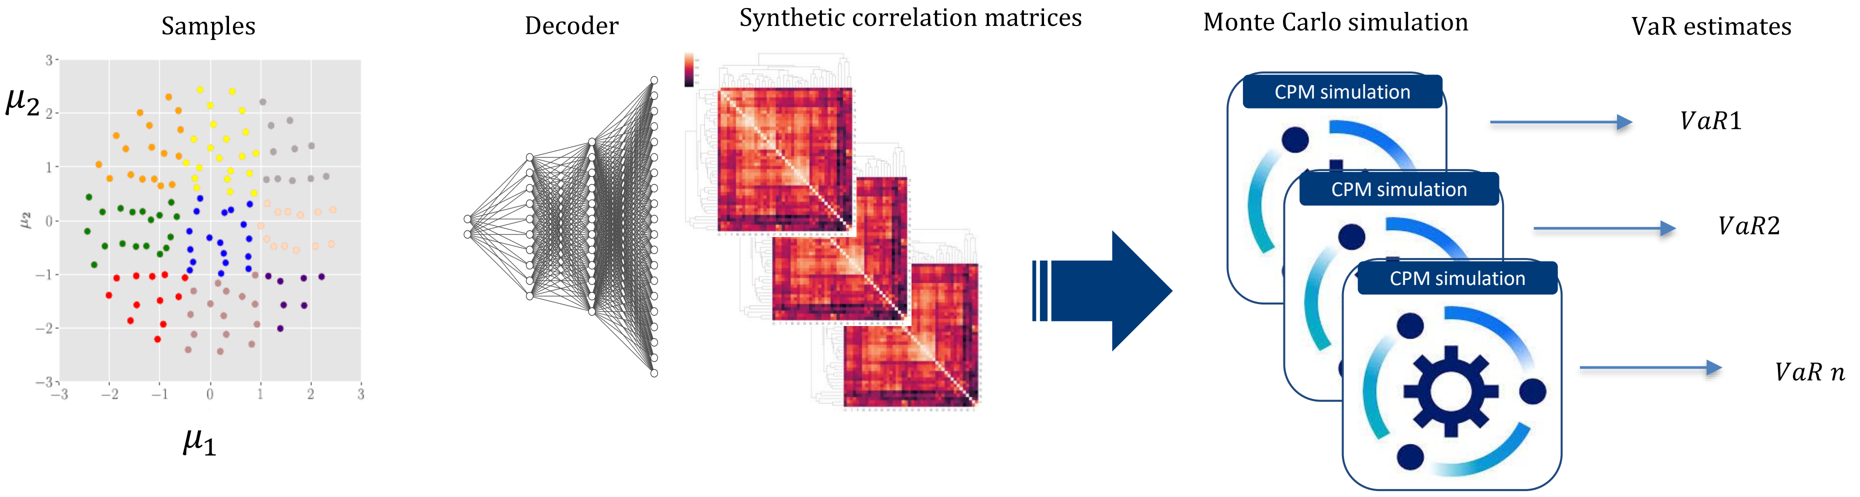


**The CPM (Vasicek multi-factor model)** cannot be solved in closed form, hence running a **Monte Carlo simulation for each generated matrix** is necessary.

In each MC simulation, all the parameters of the model are held **constant except the asset correlation matrix**.

# Variational Auto Encoders – Value at Risk estimation

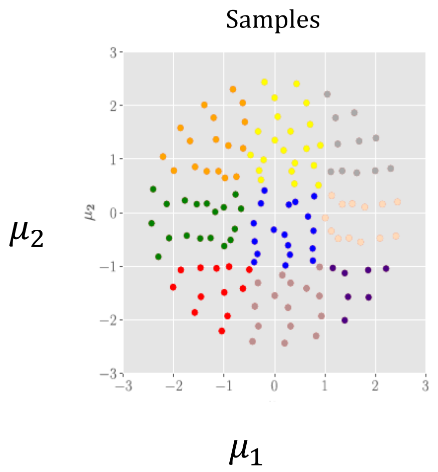


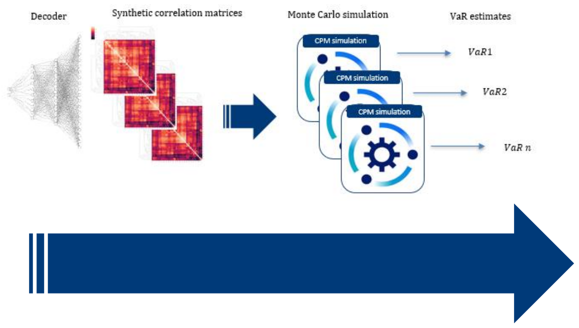


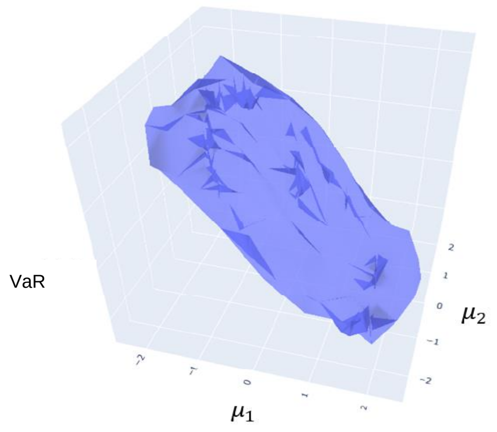


Interpolated VaR surface

Running the MC simulation for every **sampled point** of the latent space and interpolating we obtain the **VaR surface**.

Using the interpolated surface, we can estimate VaR for any given synthetic correlation matrix **avoiding Monte Carlo simulation**.

# Time series - historical correlation matrices

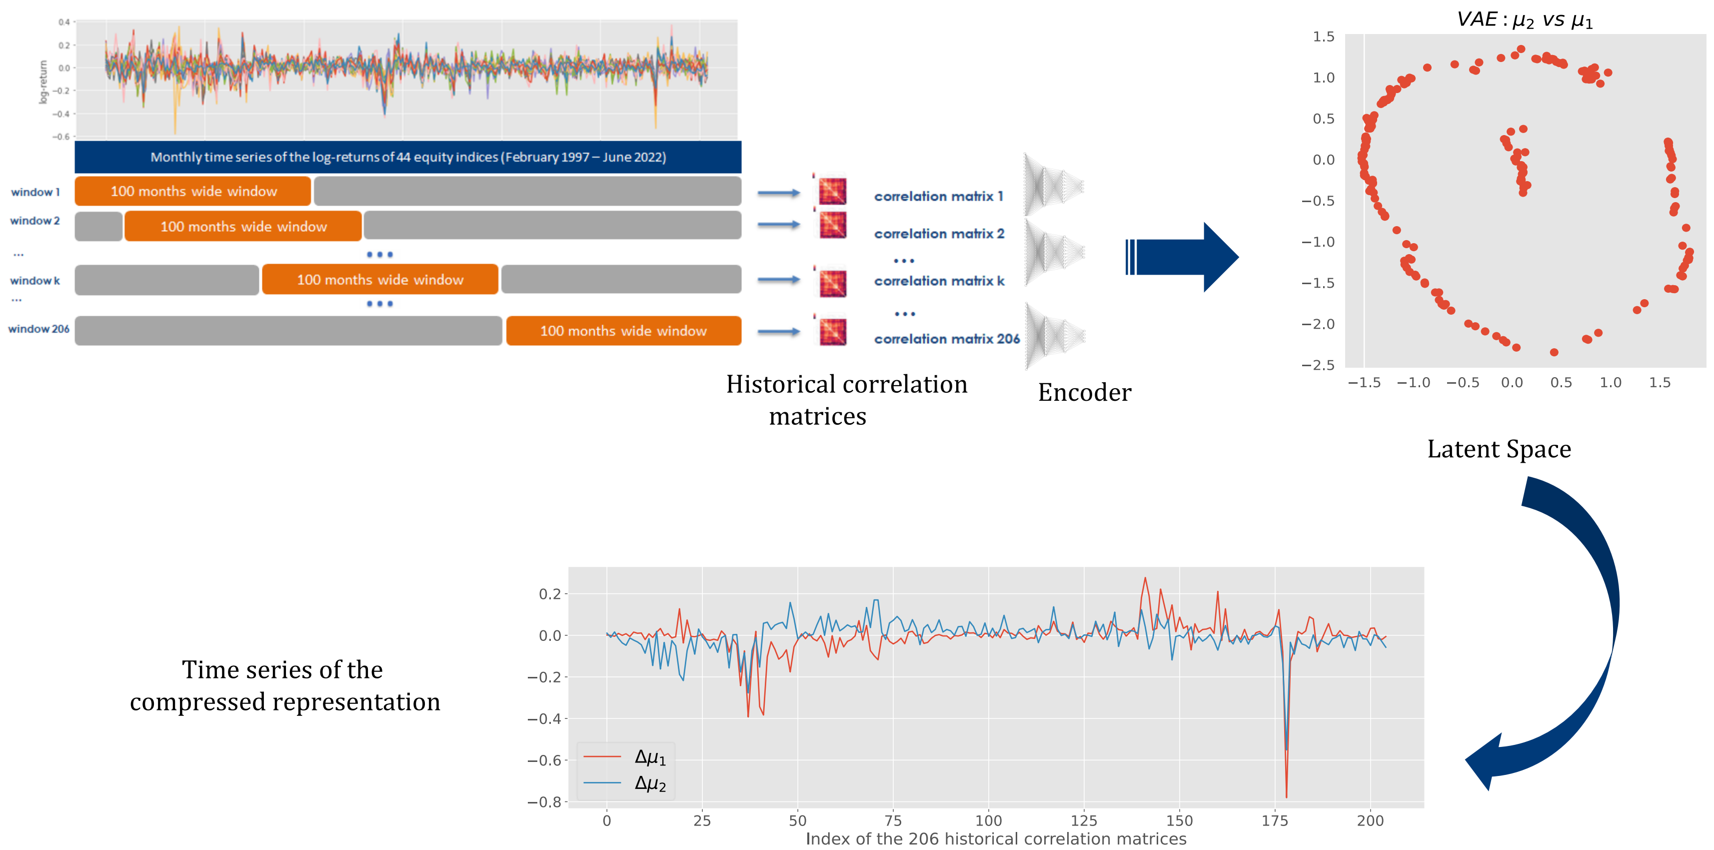


Using the "compressed" representation of the historical correlation matrices, it is possible to analyze a **bivariate time series** that represents correlation changes.

# Variational Auto Encoders – Value at Risk estimation



Bootstrapping the time series of $(\Delta\mu_1, \Delta\mu_2)$ we simulate the **distribution of the variations of the correlation matrix** in a 1-year time horizon.

Time series of variations of latent variables

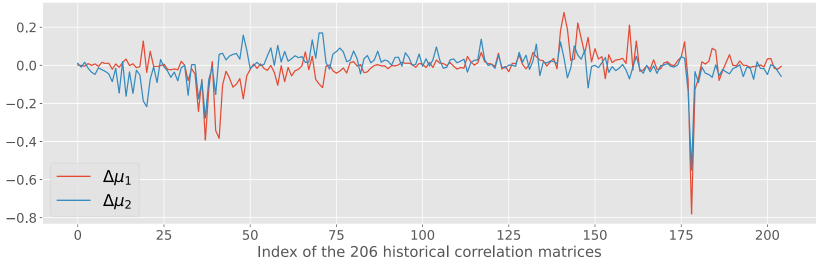


Bootstrap N points in the VAE latent space (1y time horizon)

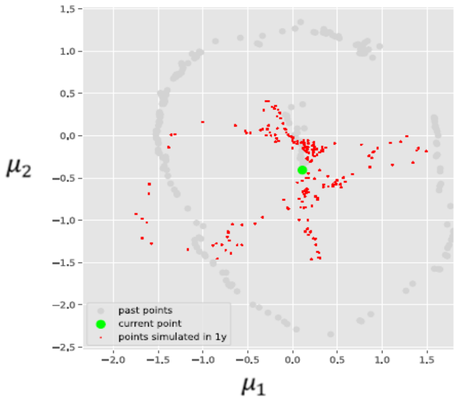


INTESA SANPAOLO

# Variational Auto Encoders – Value at Risk estimation



Harnessing the interpolated VaR surface, we derive the corresponding **VaR distribution**.

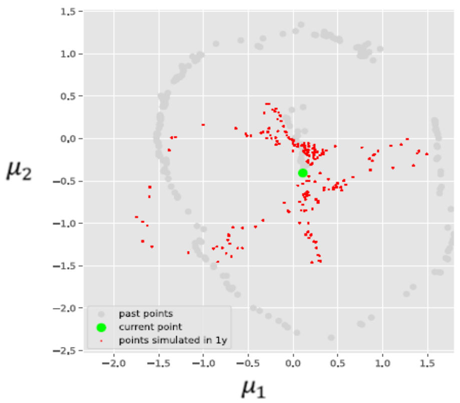


Bootstrap N points in the VAE latent space (1y time horizon)

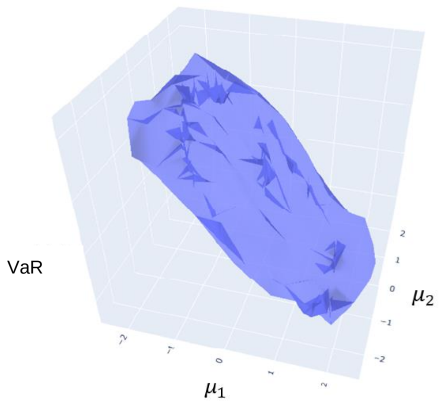


The VaR surface gives N corresponding VaR (without running Monte Carlo simulation N times)

# Variational Auto Encoders – Value at Risk estimation



Bootstrap N points in the VAE latent space (1y time horizon)

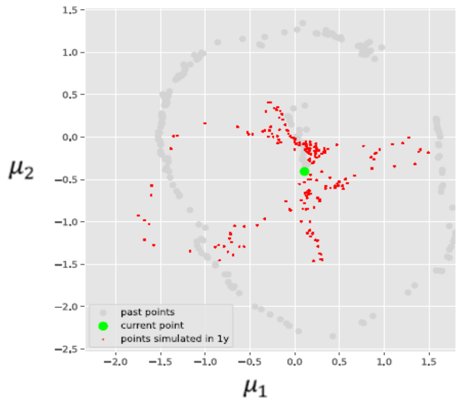


The VaR surface gives N corresponding VaR (without running Monte Carlo simulation N times)

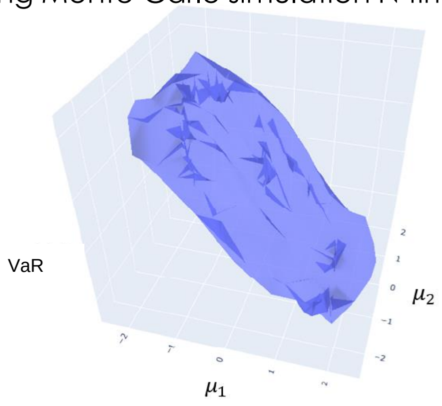


Distribution of VaR over the N synthetic correlation matrices

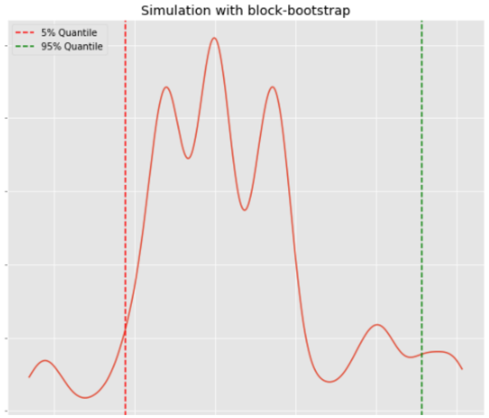


Hence, we obtain an estimation of the **CPM Value at Risk sensitivity to the variations of the correlation matrix.**

# Conclusions



We use a Variational Autoencoder (VAE) to generate **realistic** correlation matrices for assessing **credit portfolio concentration risk** within a multi-factor Vasicek model.

A **VAE** is trained on a dataset of correlation matrices derived from **equity indices time series**.

The VAE's latent space turns out to be **interpretable** in the context of **portfolio diversification**.

The **realistic data augmentation** capability of the VAE is exploited allowing to quickly estimate credit portfolio risk and to assess its **sensitivity to the correlation matrix**.

# References




[EU, 2015] (2015). Regulation (EU) no 575/2013 of the European Parliament and of the Council of 26 June 2013 on prudential requirements for credit institutions and investment firms and amending Regulation (EU) no 648/2012. *Official Journal of the European Union.*

[Abney, 2002] Abney, S. (2002). Bootstrapping. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 360–367.

[Bergeron et al., 2022] Bergeron, M., Fung, N., Hull, J., Poulos, Z., and Veneris, A. (2022). Variational autoencoders: A hands-off approach to volatility. *The Journal of Financial Data Science.*

[Brugière and Turinici, 2023] Brugière, P. and Turinici, G. (2023). Deep learning of value at risk through generative neural network models: The case of the variational auto encoder. *MethodsX*, 10:102192.

[Doersch, 2016] Doersch, C. (2016). Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908.*

[García-Céspedes and Moreno, 2021] García-Céspedes, R. and Moreno, M. (2021). The vasicek credit risk model: A machine learning approach. *Available at SSRN 3836719.*

[Grippa and Gornicka, 2016] Grippa, P. and Gornicka, L. (2016). *Measuring concentration risk-A partial portfolio approach.* International Monetary Fund.

[Kingma and Welling, 2013] Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114.*

[Lipton, 2018] Lipton, Z. C. (2018). The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57.

[Mader et al., 2013] Mader, M., Mader, W., Sommerlade, L., Timmer, J., and Schelter, B. (2013). Block-bootstrapping for noisy data. *Journal of neuroscience methods*, 219(2):285–291.

[Mantegna, 1999] Mantegna, R. N. (1999). Hierarchical structure in financial markets. *The European Physical Journal B-Condensed Matter and Complex Systems*, 11:193–197.

[Marti, 2020] Marti, G. (2020). Corrgan: Sampling realistic financial correlation matrices using generative adversarial networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8459–8463. IEEE.

[Miller, 2019] Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267:1–38.

[Millington and Niranjan, 2021a] Millington, T. and Niranjan, M. (2021a). Construction of minimum spanning trees from financial returns using rank correlation. *Physica A: Statistical Mechanics and its Applications*, 566:125605.

[Millington and Niranjan, 2021b] Millington, T. and Niranjan, M. (2021b). Construction of minimum spanning trees from financial returns using rank correlation. *Physica A: Statistical Mechanics and its Applications*, 566:125605.

[Nguyen et al., 2018] Nguyen, H. T., Tran, P. N., and Nguyen, Q. (2018). An analysis of eigenvectors of a stock market cross-correlation matrix. In *Econometrics for Financial Applications*, pages 504–513. Springer.

[of European Banking Supervisors, 2010] of European Banking Supervisors, C. (2010). Cebs guidelines on the management of concentration risk under the supervisory review process (gl31).

[Papenbrock et al., 2021] Papenbrock, J., Schwendner, P., Jaeger, M., and Krügel, S. (2021). Matrix evolutions: synthetic correlations and explainable machine learning for constructing robust investment portfolios. *The Journal of Financial Data Science*, 3(2):51–69.

[Plaut, 2018] Plaut, E. (2018). From principal subspaces to principal components with linear autoencoders. *arXiv preprint arXiv:1804.10253.*

[Vasicek, 1987] Vasicek, O. A. (1987). *Probability of loss on loan portfolio.* KMV.