

Chapter 8 | Financial Risk Management

“It's not whether you're right or wrong that's important, but how much money you make when you're right and how much you lose when you're wrong.” ~ George Soros



8.1 Financial Risk Management

8.1.1 What is Financial Risk Management?

Financial risk management (FRM) prevents very large and unexpected financial losses. It is essential for investors, financial institutions, and any groups or organizations with financial investments and exposures. It can be critical in some market environments, for instance, when a market bubble or financial crisis occurs. Financial crises happen with sometimes devastating consequences for clients, employees, companies, and investors, as during the market crash of 1987, the Asian financial crisis in 1997, the dot-com bubble in 1997-2000, the global financial crisis of 2007-2008, and the European debt crisis in 2009-2014.

Some institutions can suffer significant losses because of poor risk management, like the hedge fund Long Term Capital Management with very large exposure to liquidity risk in 1998 or, more recently, the Silicon Valley Bank in 2023 with significant exposure to interest rate risk. Sometimes, risk can come from individuals who are taking incommensurate risk, such as Nick Leeson, a derivatives trader at Barings in the 1990s that led to a trading loss of 1.4 billion dollars and the collapse of the bank or Jérôme Kerviel at Société Générale in 2008 with a loss of 4.9 billion euros.

Portfolio theory (Chapter 2) provides some tools for FRM but is not always sufficient. One reason is that tail risks are not always modeled accurately because of normal distribution assumptions in which extreme events are very unlikely or the limitation of historical data that do not contain severe losses. A second reason is that some losses are caused by moral hazard due to the actors who are supposed to do the risk calculation. When a trader only captures the upside of his trading profits and losses, the incentive is powerful enough to take too much risk. At worst, he will be dismissed with limited financial liabilities (though he could end up in prison, Nick Leeson served four years and four months in a Singapore prison, Kerviel served three years in prison, and both did not have to recoup the financial losses).

8.1.2 Types of Financial Risk

Financial risk takes many forms. It includes market, credit, liquidity, operational, legal, and regulatory risks.

8.1.2.1 Market risk

Market risk is the risk of loss due to adverse market price or interest rate changes. A trader holding a long (short) position in a financial asset is exposed to a drop (rise) in price. Depending on the asset, the market risk can be an interest rate risk, credit spread risk, equity risk, foreign exchange (FX) risk, or commodities risk. A bond would lose value when interest rates or credit spreads increase. A multinational would suffer losses on foreign revenues if the foreign currency depreciates. Commodities can be affected by supply and demand. Leverage obtained through borrowing or trading derivatives can exacerbate market risk by increasing exposure to significant market changes.

8.1.2.2 Credit risk

Credit or default risk is the risk of a counterparty defaulting on its financial obligations. It is the risk of a financial obligor to jump to default and stop its repayments. A commercial bank, a bond investor, and even a corporation serving customers are exposed to credit risk when a borrower, a bond issuer (e.g., a corporation, a municipality, a public agency, or a country), or a customer defaults.

8.1.2.3 Liquidity risk

Liquidity risk is the risk of not being able to transact an asset at a reasonable cost. Very liquid assets are very easy to sell, even during periods of high market stress. Illiquid assets might not find buyers or need to be sold at a very steep discount (a fire sale). Liquidity stress scenarios could include a loss of short-term retail deposits for a commercial bank or a loss of access to funding facilities and short-term financing.

8.1.2.4 Operational risk

Operational risk is the “risk of loss resulting from inadequate or failed internal processes, people and systems or external events” (BCBS, 2021). There are different types of operational risk (BCBS, 2001):

- **Internal fraud:** *Losses due to acts of a type intended to defraud, misappropriate property or circumvent regulations, the law or company policy, excluding diversity/ discrimination events, which involves at least one internal party*
- **External fraud:** *Losses due to acts of a type intended to defraud, misappropriate property or circumvent the law by a third party*

- **Employment Practices and Workplace Safety:** *Losses arising from acts inconsistent with employment, health or safety laws or agreements, from payment of personal injury claims, or from diversity / discrimination events*
- **Clients, Products, & Business Practices:** *Losses arising from an unintentional or negligent failure to meet a professional obligation to specific clients (including fiduciary and suitability requirements) or from the nature or design of a product.*
- **Damage to Physical Assets:** *Losses arising from loss or damage to physical assets from natural disaster or other events.*
- **Business disruption and system failures:** *Losses arising from disruption of business or system failures*
- **Execution, Delivery & Process Management:** *Losses from failed transaction processing or process management, from relations with trade counterparties and vendors*

The scope of operational risk is, therefore, quite significant. It can cover diverse activities such as accounting or tax fraud, insider trading, market manipulation, natural disasters, health and safety, and cybersecurity.

8.1.3 Identification and Measure of Financial Risk

8.1.3.1 Market Risk

In portfolio theory, market risk can be calculated as the standard deviation of the portfolio. Another helpful measure is the Value-at-Risk (VaR), defined as the maximum loss suffered over a given time horizon, say one day or ten days, for a given confidence level, 95% or 99%, for instance (Figure 8.1). The higher the confidence level or the longer the time horizon, the higher the VaR. In back-testing, losses more elevated than the VaR are counted as VaR exceptions, and their frequencies must be consistent with the confidence level (5% or 1%). A multiplier can be applied to the VaR if the number of exceptions is too high, for instance, if the VaR is calculated under normal distribution assumptions but asset returns have fat-tails.




Figure 8.1. Portfolio Value-at-Risk 95%

A limitation of VaR is that it does not inform about the losses beyond the VaR. Expected Shortfall (ES, also called Conditional VaR or CVaR) is the average loss when the loss is beyond the VaR (Figure 8.2). If a distribution has a fatter tail, the ES will be higher. It is essential for very skewed distributions, like when one sells options or insurance, collecting a premium with a risk of a significant loss.




Figure 8.2. Portfolio Value-at-Risk and Expected Shortfall 95%

Stressed VaR (SVaR) is another measure that tries to capture extreme losses. The VaR is calculated during stressful environments with high market volatility, downturns, adverse events, and low liquidity.




Figure 8.3. Portfolio Value-at-Risk and Stressed VaR 95%

8.1.3.2 Credit Risk

Credit risk is the risk of credit downgrades and the risk of default. Credit risk is partially captured by the ratings of Credit Rating Agencies (CRA), ranging from AAA (lowest risk) to BBB (highest risk) in investment grade and BB (lowest risk) or C (highest risk) in high-yield. Credit migration risk is the risk of moving from one rating to a lower one.

Other valuable measures are incremental and comprehensive risks used to regulate banks' capital. According to the (BCBS, 2009), incremental risk (charge) "represents an estimate of the default and migration risks of unsecuritized credit products over a one-year capital horizon at a 99.9 percent confidence level, taking into account the liquidity horizons of individual positions or sets of positions." Comprehensive risk is the risk of correlation portfolios due to price risk and default risk.

Each credit rating has a historical default rate over a credit cycle. However, credit risk can also be modeled more dynamically by credit models, as described in Chapter 5. For instance, in a structural credit model, credit risk will increase with leverage and asset volatility.

Credit risk is then estimated by combining a probability of default (PD) with a measure of loss-given-default (LGD) in percentage and an exposure-at-default (EAD) (in dollars). The expected credit loss would then be the product $PD \times LGD \times EAD$. For a credit portfolio, default correlation is also important as it affects its diversification and the total expected loss.

8.1.3.3 Liquidity Risk

According to the BIS, “Liquidity is the ability of a bank to fund increases in assets and meet obligations as they come due, without incurring unacceptable losses” (BIS, 2008). Liquidity risk can be measured at the security level with the bid-ask spread or by the effect of the trading volume on the price change. It can also be measured at an institutional level by comparing the amount of liquid assets (cash deposits, government and agency securities) to short-term obligations and cash flows required to finance its operations. If there are some liabilities, managing the mismatch between assets and liabilities is also critical to ensure enough funding liquidity to pay the liabilities when assets are very illiquid.

8.1.3.4 Operation Risk

Operation risk is measured by the impact of various risk factors on an organization’s operations; it can be for a bank, the impact on interest rate income and expenses, operating expenses and revenues, and financial profits and losses. It would combine operational risk exposure with historical loss data. This is a standardized framework adopted to regulate banks (BCBS, 2016). Typically, the impact can be evaluated with statistical models or scenario analysis.

8.1.4 Factor Risk Models

Estimating volatilities and correlations for all assets in a portfolio is not realistic. With N assets, $N(N-1)/2$ correlation coefficients need to be estimated. A large bank can easily have one million assets to follow, meaning half a trillion coefficients to estimate.

A common approach to address this problem in risk modeling is to use factor models in the style of the asset pricing models in Chapter 1. The simplest model would be a

CAPM style model, with one single market factor and some residual independent specific risk ϵ :

$$R_{t+1} = R_f + \beta_1 RF_1 + \epsilon$$

There are only factor exposure coefficients and specific risks to estimate, which grow linearly with the number of assets. The first term might measure systematic risk but is likely a simplification. (Brunnermeier et al., 2010) suggest that “it is unlikely that a single measure of systemic risk will suffice. We anticipate that the variety of inputs ranging from leverage and liquidity to codependence, concentration, and connectedness will all be revealing,” so it is still an ongoing research topic.

In a multifactor model, an asset return is decomposed into factor exposures, factor returns, and some residual specific risk ϵ :

$$R_{t+1} = R_f + \beta_1 RF_1 + \dots + \beta_n RF_n + \epsilon$$

The factor returns can be identified; for instance, they could be industry, country, style (e.g., earnings, value, momentum) index returns, and the betas are then estimated by regression analysis. The betas could also be identified a priori; for instance, they could be some asset characteristics such as duration and convexity for a bond; the factor returns are then estimated by a cross-section regression analysis. Some cross-sectional regression on some asset characteristics could derive the specific risk. In that case, there are fewer numbers to estimate.

Most leading risk models, such as MSCI Barra (Menchero, 2010), have this type of structure. For instance, the Barra Global Equity Model has a World factor, 77 country factors, 34 industry factors, and 11 style factors.

$$R = F_w + \sum_{c=1}^{77} \beta_c F_c + \sum_{ind=1}^{34} \beta_{ind} F_{ind} + \sum_{s=1}^{11} \beta_s F_s + \epsilon$$

The exposure is 1 to the world factor, the respective country and industry factors, and 0 otherwise. The factors are estimated by weighted least-square regression analysis. After normalization, the world factor is the world portfolio return.

In fixed income portfolios, factors are related to currency, yield curve, swap spreads, volatility, investment grade spreads, treasury spreads, credit and agency spreads, MBS and securitized, CMBS, ABS spreads, and emerging market spreads (Dynkin et al., 2006).

Once the model is identified, if a specific return is independent of the factor returns, then the risk of the asset measured by its variance can be calculated as:

$$\sigma^2 = \beta' \Sigma \beta + \sigma_\epsilon^2$$

The risk is decomposed into some factor risk and some idiosyncratic risk. To apply the model, the covariance matrix Σ , the factor exposures and the idiosyncratic variance must be estimated.

These factor models can project future risk based on regression forecasting and historical time series. This approach is backward-looking and relies heavily on the representativeness of historical data. Scenario analysis, stress tests, and Monte Carlo simulations often complement it.

8.1.5 Scenario Analysis

Scenario analysis projects future changes in economic and financial indicators such as unemployment rate, GDP, stock market index, and interest rates and looks at the effect on asset values. For instance, The Federal Reserve publishes stress test scenarios yearly that large banks have to run to show that they have enough capital.

In 2023, the (Board of Governors of the Federal Reserve, 2023) announced their stress test scenario:

- *The severely adverse scenario is characterized by a severe global recession accompanied by a period of heightened stress in both commercial and residential real estate markets, as well as in corporate debt markets.*
- *The U.S. unemployment rate rises nearly 6-1/2 percentage points from the starting point of the scenario in the fourth quarter of 2022 to its peak of 10 percent in the third quarter of 2024.*

- *The sharp decline in economic activity is also accompanied by an increase in market volatility, widening corporate bond spreads, and a collapse in asset prices, including a 38 percent decline in house prices and a 40 percent decline in commercial real estate prices.*
- *The international portion of the scenario features recessions in four countries or country blocs, with heightened stress in advanced economies, followed by declines in inflation and an appreciation in the value of the U.S. dollar against all countries and country bloc's currencies, except for the Japanese yen.*

Each market can have its stress scenarios, such as an increase in interest rates, widening of credit spreads, decline in asset prices, currency depreciation, increase in equity volatility, increase in default rates, and payment delinquencies.

Human domain experts can also propose stress test scenarios. They will often refer to past crises (the crash of 1929, the 1982 Latin American debt crisis, the crash of 1987, the 1997 Asian financial crisis, the LTCM bankruptcy, and the 2007-2008 great financial crisis) and try to project realistic scenarios inspired by them.

8.1.6 Monte-Carlo Simulations

The linear factor risk approach is not appropriate for securities that have non-linear dependencies with the factors. Derivatives, for instance, have convex or concave payoffs and can have significant positive or negative jumps of payoffs when the underlying asset changes value.

Scenario analysis and stress tests can help but can become burdensome when there are too many securities, there is no guarantee that all the scenarios will cover all the future events, and there is consistency between the assets in the same scenario.

Monte-Carlo simulations would create sequences of underlying asset and factor prices that would feed into pricing or simple present value models. The models can be advanced with stochastic volatility or auto-regressive volatility or variance. For instance, a $GARCH(p, q)$ model would assume that the variance of the error term in a time series depends on past variances and past squared values of the error term:

$$R_t = \beta'X_t + \varepsilon_t$$

$$\varepsilon_t \sim N(0, \sigma_t^2)$$

$$\sigma_t^2 = \omega + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j \sigma_{t-j}^2$$

There are some drawbacks to simulations. First, model risk is related to the assumed distribution used for simulations. The distribution could be historical, but it assumes that history has seen all the important events. Second, there are compute costs, especially to simulate rare events, though techniques like important sampling where we can shift the distributions can help.

8.1.7 Risk Mitigation

Once risks are identified, there are many strategies to address them. Risk can be avoided, for instance, by not holding some risky securities. Investment-grade bond investors would avoid high-yield securities and unrated bonds.

Risk can be reduced and controlled by reducing positions in risk assets and increasing diversification. There can be a cap on the share of a security, an obligor, a sector, or a country in a portfolio. Risk can be controlled by risk budgeting by ensuring that total risk or active risk compared to a benchmark is limited and that the risk contribution of each asset is balanced.

Risk can be hedged by shorting the asset, buying insurance, or entering a derivatives contract such as a put option, a futures contract, a swap contract, or a credit default swap that would protect against a drop in the value of the reference asset.

Furthermore, risk can be transferred to a third party with a swap contract or by structuring and securitizing risk. A portfolio of risky assets can be divided into tranches of different risk and seniority and sold to different classes of investors. Some investors would favor the most senior tranches to avoid credit risk but earn better returns than riskless securities. Others will search for yield and higher returns to meet some liabilities.

8.2 The Practice of Financial Risk Management

8.2.1 Governance

Because risk management is not all quantitative and involves an important human dimension, it often starts with a governance framework. For instance, at Goldman

Sachs (Goldman Sachs, 2023), the Board of Directors, with the help of specialized committees such as the Risk Committee and the Audit Committee, supervises risk management and approves the firm's risk appetite. A chief risk officer, head of risk oversight and control, reports to the CEO and the risk committee and informs them of the firm's risk levels and types. The firm has three lines of defense to protect against risk: the business units, the independent risk group that includes Compliance, Conflicts Resolution, Controllers, Legal, Risk and Tax, and the internal audit team.

8.2.2 Processes

At Goldman Sachs, processes involve (i) risk identification and control assessment, (ii) risk appetite, limit and threshold setting, (iii) risk metrics, reporting and monitoring, and (iv) risk decision-making.

8.2.2.1 Risk identification and control assessment

There is a comprehensive process for collecting data and reporting risk events involving all employees. The approach is dynamic, forward-looking, and adaptable to changing risk profiles and business environments. The control assessment, led by risk management, reviews and challenges the control environment to support the strategic business plan. There is a daily practice of marking inventory to current market levels, valuing it at fair value, and reflecting valuation changes in risk management systems and net revenues. Stress testing is crucial to their risk management process, quantifying exposure to tail risks, assessing potential loss concentrations, and analyzing risk positions. Regular firmwide stress tests combine financial and nonfinancial risks into a combined scenario, including credit, market, liquidity, operational, compliance, strategic, systemic, and emerging risks. Ad hoc stress tests are also conducted in anticipation of market events. Stress tests evaluate capital adequacy as part of capital planning and stress testing.

8.2.2.2 Risk appetite, limit, and threshold setting

A framework is applied to control and monitor risk across transactions, products, businesses, and markets. The Board, through its Risk Committee, approves limits and thresholds at firmwide, business, and product levels. The Firmwide Risk Appetite Committee, with delegated authority from the Firmwide Enterprise Risk Committee, is responsible for approving and monitoring the risk limits and thresholds policy. Ongoing dialogue and escalation of risk-related matters are fostered through periodic adjustments of certain limits, encouraging communication among different defense lines, committees, and senior management. Under the delegated authority, Market Risk

and Credit Risk set limits at specific levels such as products, desks, individual counterparties, subsidiaries, industries, and countries. These limits are regularly reviewed and adjusted to align with changing market conditions, business conditions, and risk tolerance.

8.2.2.3 Risk metrics, reporting and monitoring

The firm focuses on the quality and efficiency of its risk systems to ensure accurate and timely information. The risk metrics, reporting, and monitoring processes consider existing and emerging risks, providing insights to risk committees and senior management for informed decision-making. The organization has processes in place to escalate limit and threshold breaches promptly. Additionally, risk profile and business operations changes are evaluated by monitoring risk factors at a firmwide level, including changes in business mix and operating jurisdictions.

8.2.2.4 Risk decision-making

The organization utilizes risk committees that meet regularly to facilitate discussions and mitigate risks. There is a strong emphasis on communication and collaboration among the first and second lines of defense, committees, and senior management. While the first line of defense manages risks, the organization dedicates significant resources to the second line of defense to ensure adequate oversight and segregation of duties. A culture of escalation and accountability across all functions reinforces the firm's commitment to risk management.

8.2.3 people

For Goldman Sachs, the experience and understanding of professionals are crucial in assessing and maintaining risk exposures at prudent levels. The firm reinforces a culture of effective risk management through training and development programs led by senior leaders that emphasize risk management, client relationships, and reputational excellence. Performance evaluations consider reputational excellence, adherence to the code of conduct and compliance policies, and sound risk management judgment. The review and reward processes aim to communicate the connection between behavior and recognition, the focus on clients and reputation, and the importance of acting by high standards.

8.2.4 Market Risk Management

At Goldman Sachs, market risk refers to the potential negative impact on earnings due to changes in market conditions. The firm's market risk primarily arises from positions held for market-making, investing, and financing activities. They employ various risk measures to monitor market risk, including interest rate, equity price, currency rate, and commodity price risks.

The Market Risk department, independent of revenue-producing units, is responsible for assessing and managing market risk across the organization. Managers in revenue-producing units and the Treasury collaborate with Market Risk to discuss market information, positions, and potential loss scenarios. They manage risk within prescribed limits and utilize hedging instruments to mitigate exposures.

The market risk management process involves monitoring compliance with established limits, diversifying exposures, controlling position sizes, and evaluating mitigants such as economic hedges. The firm's market risk management systems enable independent calculations of risk measures such as Value-at-Risk (VaR) and Earnings-at-Risk (EaR). They can generate reports at various levels and conduct ad hoc analyses promptly.

8.2.5 Liquidity Risk Management

At Goldman Sachs, liquidity risk refers to being unable to meet funding needs during stressful events. The company has implemented comprehensive and conservative liquidity and funding policies to sustain itself and continue serving clients even in adverse circumstances. The Treasury Department, under the chief financial officer, is responsible for developing and executing the liquidity and funding strategy. An independent unit called Liquidity Risk, reporting to the chief risk officer, oversees and manages liquidity risk through firmwide monitoring, stress testing, and limits frameworks. The firm follows three principles in liquidity risk management: holding sufficient excess liquidity, maintaining appropriate Asset-Liability Management, and having a viable Contingency Funding Plan.

At the macro level, regulators want banks and financial intermediaries to have sufficient access to liquidity, especially during financial crises and periods of high volatility. (Goldberg, 2022) notes, for instance, that access to internal and official liquidity, such as central bank swap lines and repo facilities, can lower the amplitude of global liquidity response to stress events and support the continued provision of credit during systemic shocks. Though banks are subject to strict capital requirements, nonbank financial institutions are much less regulated and can present new liquidity risk vulnerabilities for the financial system.

8.2.6 Credit Risk Management

At Goldman Sachs, credit risk refers to the potential for loss resulting from the default or deterioration in credit quality of a counterparty or issuer of securities. The firm's credit risk primarily arises from client transactions in OTC derivatives, loans, lending commitments, cash placements with banks, securities financing transactions, and customer receivables.

Independent of revenue-producing units, the Credit Risk department is responsible for assessing and managing credit risk across the organization. It oversees credit risk through firmwide oversight and collaborates with the Market Risk department, which monitors and manages credit risks associated with market risk measures and derivatives.

The credit risk management process includes monitoring compliance with established limits, reporting credit exposures and concentrations, establishing or approving underwriting standards, assessing counterparty default likelihood, measuring credit exposure and potential losses, utilizing credit risk mitigants such as collateral and hedging, and maximizing recovery through active workout and restructuring.

Credit analyses are conducted to evaluate the capacity and willingness of counterparties to meet financial obligations. The firm assigns internal credit ratings to counterparties based on evaluations, industry outlook, and economic factors. Senior personnel with industry expertise review and approve credit reviews and internal credit ratings.

Credit risk management systems capture individual and aggregate credit exposures, providing comprehensive information about credit risk by product, internal credit rating, industry, country, and region.

8.2.7 Operational Risk Management

At Goldman Sachs, operational risk is the potential for adverse outcomes from internal processes, people, systems, or external events. It encompasses routine errors and extraordinary incidents such as system failures and legal/regulatory matters. Loss events related to operational risk include process management, system failures, employment practices, fraud, and business disruptions.

The Operational Risk department, independent of revenue-producing units, is responsible for assessing, monitoring, and managing operational risk. Its goal is to maintain operational risk exposure within the organization's risk appetite.

The operational risk management process involves a comprehensive data collection process, firmwide policies and procedures, and a combination of top-down and bottom-up approaches. Senior management assesses operational risk profiles at the firmwide and business levels, while the first and second lines of defense are responsible for day-to-day risk identification and management. Operational risks and events are escalated to senior management.

A control framework is in place to minimize operational risks, overseen by the Firmwide Operational Risk and Resilience Committee. The operational risk management framework complies with regulatory requirements and evolves based on business needs and guidance.

Policies require employees and consultants to report and escalate operational risk events. When events occur, they are documented and analyzed to identify necessary changes in systems and processes for future risk mitigation. Operational risk management applications capture, analyze, aggregate, and report operational risk event data and metrics.

Managers conduct an operational risk and control self-assessment process to identify and rate operational risks and related controls. Results from this process are analyzed to evaluate operational risk exposures and determine areas with heightened levels of risk.

8.2.8 Model Risk Management

According to Goldman Sachs, model risk refers to the potential negative consequences of using incorrect or inappropriate outputs from quantitative models. These models are extensively used in various aspects of their business, such as valuing financial assets, risk management, and regulatory capital measurement.

The Model Risk department, independent of revenue-producing units and other model-related roles, is responsible for assessing, monitoring, and managing model risk. It oversees the global businesses and reports to the chief risk officer, keeping senior management, risk committees, and the Board's Risk Committee updated on model risk.

A governance structure and risk management controls support their model risk management framework. These encompass standards to ensure a comprehensive model inventory, including risk assessment, sound development practices, independent review, and model-specific usage controls. The Firmwide Model Risk Control Committee oversees this framework.

The Model Risk team, composed of quantitative professionals, conducts independent reviews, validations, and approvals of their models. This involves analyzing model documentation, independent testing, assessing methodology appropriateness, and verifying compliance with development and implementation standards.

They continuously refine and enhance their models to account for market and economic changes and align with their business activities. All models undergo an annual review, and new models or significant changes to existing models require approval before implementation.

The model validation process includes reviewing models, trade and risk parameters, and evaluating various scenarios, including extreme conditions. The goal is to critically assess and verify the conceptual soundness, testing strategies, calculation techniques, accuracy in reflecting product characteristics and risks, consistency with similar models, and sensitivity to input parameters and assumptions.

8.2.9 Other Risk Management

Apart from the previously discussed risks, the Goldman Sachs teams also manage other risks, including capital, climate, compliance, and conflicts of interest.

Capital risk management ensures sufficient capital to support their business activities and meet internal and regulatory capital requirements. It has a comprehensive capital management policy and governance structure overseen by the Board and its committees. Risk committees and senior management monitor capital adequacy, evaluate regulatory requirements, and review capital planning and stress test results.

Climate risk management involves addressing physical and transition risks related to climate change. It has integrated climate-related risk oversight into the risk management processes and governance structure, with senior management and the Board's Risk and Public Responsibilities Committees playing key roles. It evaluates climate risk factors in select industries' credit evaluation and underwriting processes.

Compliance risk management focuses on ensuring adherence to applicable laws, regulations, and internal policies. The Compliance Risk Management Program assesses and monitors compliance and regulatory risk, designs controls and procedures, conducts independent testing, and leads responses to regulatory examinations.

Conflict management is crucial for maintaining client relationships and reputation. It has a multilayered approach to identifying and resolving conflicts of interest. Senior management and Conflicts Resolution, Legal, and Compliance teams formulate policies and principles and make judgments regarding conflict resolution. Transaction oversight

committees and internal and external counsel review and address conflicts in various activities.

They regularly assess their policies and procedures to uphold ethical standards and comply with laws and regulations.

8.2.10 Basel III

The Basel Committee developed the Basel III (BCBS, 2017) framework to respond to the global financial crisis. It aims to address the weaknesses of the pre-crisis regulatory framework and establish a resilient banking system. The reforms focus on various aspects of the regulatory framework to ensure a stable and supportive banking system.

The initial phase of Basel III reforms aimed to strengthen the regulatory framework by improving the quality of bank regulatory capital, increasing capital requirements, enhancing risk capture, introducing macroprudential elements, implementing a leverage ratio requirement, and establishing frameworks for liquidity risk and maturity transformation.

The finalized Basel III reforms further enhance the regulatory framework by improving risk-weighted assets (RWAs) calculation and ensuring the comparability of banks' capital ratios. This includes strengthening the standardized approaches for credit risk, credit valuation adjustment (CVA) risk, and operational risk, placing limits on the use of internal model approaches, introducing a leverage ratio buffer for global systemically important banks (G-SIBs), and replacing the existing Basel II output floor with a more robust risk-sensitive floor based on revised standardized approaches.

8.3 AI and Financial Risk Management

AI can be applied to risk measurement, modeling, and forecasting, including tail, systematic, and idiosyncratic risks. AI can identify and accurately measure the relevant risk factors ex-ante that drive risk and model their dynamics and their risk distributions. This could be used for risk budgeting and asset allocation (Chapter 2). Ex-post, AI can be used for return attribution and identifying which factor caused trading or investment losses.

AI can replace traditional credit risk models to forecast rating changes, default risk, and loss-given default. It can leverage traditional financial data such as leverage, interest coverage ratio, cash flow over debt, equity volatility, and interest rates. It can also incorporate new alternative data (Chapter 12), such as job boards or satellite images.

Any indicators correlated with revenue, market demand, and profitability would be helpful to evaluate a firm's credit risk.

AI is beneficial for scenario analysis and simulations. Generative AI, in particular, can simulate realistic risk scenarios that could be used for stress-testing portfolios, trading books, loan books, derivatives books, capital level, margin, and collateral requirements. It is valuable for regulators who can assess the solidity of financial institutions and the financial system in general. Combined with graph neural networks and network analysis, regulators can study the propagation of risk in the economy.

Other risks include operational, liquidity, cybersecurity, and geopolitical risks. Once these risks are appropriately modeled, AI can generate scenarios and risk distributions and manage these risks efficiently by suggesting risk allocation and hedging strategies to improve the resilience of financial portfolios, books, and positions. For instance, AI can model liquidity risk by identifying liquidity factors, characteristics, and the different forms of liquidity risks, such as market liquidity, funding liquidity, and contingent liquidity risks.

8.3.1 Market Risk

Managing market risk with AI is very similar to managing portfolio risk in Chapter 2, the difference being that we are more concerned with tail risk than with the trade-off between risk and return. One important method to manage market risk is to use a factor model that decomposes returns into systematic and idiosyncratic factors, as described in asset pricing (Chapter 1). This allows us to forecast risk and the source of risk.

Usually, it is done in two steps. In the first step, risk factors are identified and estimated. These factors can be grounded in economic theory (stock market portfolio, interest rates, inflation), market long-short portfolios such as the Fama and French portfolios, or implicit factors identified by a Principal Component Analysis (PCA) or autoencoder models. In the second step, the risk of these factors is estimated by a covariance matrix based on historical data or simulation, or a more advanced AI prediction model can be used to forecast the risk.

Another approach is calculating a portfolio's Value-at-Risk (VaR) or Expected-Shortfall (ES). It is an exercise of forecasting the portfolio market loss quantiles that a machine learning or AI model should be able to handle. Typically, a portfolio can be decomposed into its constituents with their risk exposures and weights, and future returns are simulated to generate a portfolio return distribution. The worst $x\%$ of returns will determine the $1-x\%$ VaR level or ES value. More AI-advanced models

would forecast the change in portfolio weights. This would be necessary for hedge ratios to change when market and interest rates vary.

(Fécamp et al., 2020) discuss the challenges of hedging in incomplete financial markets due to factors like transaction costs, illiquidity, and non-tradable risk factors. They explore different criteria for deciding how to share risks between buyers and sellers in such markets, including quantile hedging, utility-based, and moment-based criteria. The authors propose machine learning algorithms to derive optimal hedging strategies, including a global risk minimization approach and local optimization methods. They apply these algorithms to various hedging problems, considering factors like liquidity constraints, transaction costs, and risk criteria. The results indicate that deep neural network algorithms can effectively handle realistic option hedging problems, with the global approach often outperforming local methods in terms of efficiency and accuracy, particularly in higher dimensions. The algorithms also enable the efficient computation of Pareto efficient frontiers for hedging strategies.

8.3.2 Credit Risk

AI and machine learning are natural tools to model credit risk. (Breedon, 2020) suggests three application domains: credit scoring, corporate default, and portfolio forecasting.

8.3.2.1 Credit Scoring

Credit scoring can be enhanced with machine learning models. The FICO score, a widely used credit score, uses payment history (35%), amounts owed (30%), length of credit history (15%), new credit (10%), and credit mix (10%) (“How are FICO Scores Calculated?” 2018). Non-quantitative data could be used (see Chapter 12 on Alternative Data).

(Duan, 2019) uses robust deep neural networks for modeling and predicting peer-to-peer (P2P) lending and highlights the need for effective risk assessment tools in P2P lending, where borrowers provide personal data and interest rates are determined based on this data. The proposed approach involves an innovative Multilayer Perceptron with multiple hidden layers (MLPm) trained using the back-propagation learning algorithm for risk evaluation in P2P lending.

(Sadhvani et al., 2021) use deep learning models to analyze mortgage credit and prepayment risk. Traditional empirical models in mortgage literature often impose linear relationships between risk factors and borrower behavior, which do not capture the nonlinear effects present in the data. They propose a deep learning approach that allows the relationship between risk factors and loan performance to be determined

directly from the data without pre-specified linear assumptions. They use a vast dataset of over 120 million mortgages to study the effects of various risk factors, including borrower-specific variables, macroeconomic factors, and more.

The authors' deep learning model outperforms traditional linear models, providing more accurate forecasts of mortgage risk, improving investment strategies, and enhancing hedge sensitivity estimates. Their approach accommodates complex relationships and variable interactions, making it a valuable tool for understanding mortgage market dynamics.

(Davis et al., 2022) explain the predictions made by ML models, particularly neural networks, in the context of credit risk assessment. Various stakeholders, including loan companies, regulators, loan applicants, and data scientists, have different needs for model explanations.

For loan companies, they describe using post hoc explainability tools like LIME and SHAP to provide reasons for loan denials. They also highlight the importance of generating explanations that reveal the relationships learned by the model, such as the impact of an applicant's credit history on their creditworthiness. For regulators, they suggest methods to investigate the fairness of ML models and perform stress tests to analyze their behavior in extreme scenarios. They give suggestions for loan applicants to help individuals improve their creditworthiness, even when black-box ML models are used for loan approval decisions.

They find that explainable AI tools can make the functionality of black-box ML models accessible to various stakeholders, enabling them to make informed decisions, conduct stress tests, and answer "what-if" questions. They highlight the potential of explainable AI in traditional fields like credit modeling.

8.3.2.2 Corporate Default

Ed Altman, the creator of the Altman Z-Score, has examined the use of machine learning models for predicting corporate bankruptcy one year in advance (Barboza et al., 2017), comparing their performance with traditional statistical methods. The research utilizes data from 1985 to 2013 on North American firms, encompassing over 10,000 firm-year observations. The authors compare traditional statistical methods like discriminant analysis and logistic regression with machine learning models, including support vector machines, bagging, boosting, and random forest.

They find that the latter show a substantial improvement in prediction accuracy compared to traditional models. On average, they find that machine learning models achieve around 10% higher accuracy and that random forest performs the best with

the highest accuracy of 87%, while logistic regression and linear discriminant analysis achieve 69% and 50% accuracy, respectively, in the testing sample. Financial indicators, such as operating margin, change in return-on-equity, and growth measures, significantly enhance all models' accuracy.

8.3.2.3 Portfolio Forecasting and Optimization

With default probabilities, loss given default, and a model of joint default behavior, it is possible to calculate portfolio loss distribution using analytical or Monte Carlo methods. Some methods can involve copulas, such as the Gaussian copula.

(Sirignano and Giesecke, 2019) develop efficient numerical methods for analyzing these large, heterogeneous pools of loans, particularly in the context of dynamic, discrete-time models of loan-level risk. These models consider various factors, including loan-level features, common risk factors, and historical loan performance data. The study also proves certain mathematical limit theorems and constructs Monte Carlo approximations for assessing risk, pricing, and hedging asset-backed securities backed by these loan pools. The proposed methods are tested using a dataset of over 25 million prime and subprime mortgages. They provide accurate results at significantly lower computational costs compared to brute-force simulations.

8.3.3 Liquidity Risk

Machine learning methods can quantify liquidity risk. It can be explored in the stock market, corporate bond market, or derivative market.

(Crego et al., 2023) focus on the limitations of using realized returns as a proxy for expected returns and introduces the concept of "yield" as an alternative estimator for expected returns. They combine the option-based yield estimator with machine learning methods to improve the accuracy of predicting expected returns. Their goals include extending the applicability of yield beyond S&P 500 stocks and identifying consistent characteristics that predict expected returns more effectively.

They employ various statistical techniques, including the Kalman filter and neural networks, to predict yields and assess their ability to explain differences in average stock returns. The authors find that predicted yields spread average returns, with the difference in returns between the highest and lowest portfolios being statistically significant. Furthermore, they highlight the importance of certain stock characteristics, including liquidity, momentum, volatility, market frictions, firm quality, and measures of intangible assets, in predicting expected returns. These characteristics consistently appear as influential factors across different samples.

(Müller et al., 2023) predict liquidity in the corporate bond market and present a machine learning-based model for forecasting bond liquidity. They introduce a prediction model that uses machine learning techniques to forecast a bond's liquidity for the next month. They measure expected illiquidity and downside liquidity risk for individual bonds. The model utilizes a wide range of liquidity predictors based on previous research in the corporate bond market. These predictors include current bond illiquidity, illiquidity of similar bonds, interest rate risk, and credit risk. The predictive model's performance is compared to a naive approach that assumes a bond's current liquidity is the same in the future. The predictive model outperforms the naive approach, reducing forecasting errors by approximately 21%.

(Ghysels et al., 2020) apply machine learning algorithms in forecasting uncertainty shocks in the financial domain, explicitly focusing on variance risk premiums and synthetic variance swap payoffs. They employ machine learning algorithms such as LASSO, Elastic Net, Ridge, Random Forest, Gradient Boosting Regression Tree, Feedforward Neural Network, and Sparse Group LASSO. They use daily options prices, implied volatilities, or implied variances from approximately 1,500 contracts each month to predict next month's synthetic variance swap payoff.

The research demonstrates that machine learning methods outperform a simple AR(1) model in predicting synthetic variance swap payoffs. The highest out-of-sample R-squared (OOS R²) is achieved with Neural Networks (NN), followed by other machine learning methods. Furthermore, the study examines the impact of liquidity on prediction accuracy. The study finds that penalized linear models like LASSO and Elastic Net outperform other methods by pre-filtering input data based on liquidity measures. Liquidity-guided machine learning approaches show that these models perform well even when selecting a subset of contracts with higher liquidity.

8.3.4 Operational Risk, Cybersecurity Risk, and Geopolitical Risk

Operational risk includes cyber risk, a form of external fraud risk. (H. Jiang et al., 2020) study the influence of cyber risk on stock returns; they employ machine learning algorithms to develop a real-time estimate of the likelihood of individual firms experiencing a cyberattack shortly. This estimate is derived from various firm characteristics, such as industry competition and linguistic analysis of risk disclosures in 10-K filings. The cyber risk measure developed using machine learning techniques significantly enhances the prediction of future cyberattacks compared to traditional models. A meaningful and statistically significant relationship exists between cyber risk and stock returns. Stocks with higher cyber risk tend to have higher average returns. A

portfolio analysis reveals that stocks with higher cyber risk consistently outperform those with lower risk, even after adjusting for market exposure and other factors.

Operational risk includes natural disasters which can damage physical assets. (Zeng and Bertsimas, 2023) introduce a novel approach to global, multi-year flood prediction using machine learning, precisely a multimodal approach combining text-based geographical information with statistical features. They use a unique multimodal framework integrating text-based geographical information with time-series statistical features for global flood prediction, leveraging advanced natural language processing techniques. They obtain strong results in multi-year flood risk forecasting, with the best model achieving a 75%-77% ROC-AUC score for predicting floods in the next 1-5 years. Multimodal models combining text and statistics outperform single-modal models.

Operational risk includes geopolitical risk, which can lead to severe business disruptions. (Bondarenko et al., 2023) address the tracking and analysis of geopolitical risk (GPR), which encompasses adverse events related to wars, terrorism, and international tensions. The Russo-Ukrainian War exemplifies how GPR can have economic, social, political, and human costs. The authors create a monthly news-based GPR measure based on Russian local news sources. This local GPR measure is compared with those based on English-language, German, and Ukrainian news sources.

8.4 Case Studies

8.4.1 Explainable Machine Learning Models of Consumer Credit Risk

(Davis et al., 2022) use data from FICO on the Home Equity Line of Credit (HELOC) to illustrate how to explain machine learning models in finance. A Home Equity Line of Credit is a pre-approved loan that a homeowner can draw and is backed by the equity value of the home property.




Figure 8.4. US HELOC balance and limit, Credit Card balance and limits.

Source: <https://www.newyorkfed.org/microeconomics/hhdc>

The authors explore four categories of models: Inductive Logic Programming (ILP), Random Forests, Optimal Classification Trees, and Neural Networks.

Inductive Logic Programming is using logical rules to classify the data. If there are n features, there is a probability distribution over N rules and range (rule size) R for each feature X_i (that needs to be normalized to be between 0 and 1) and its negation $1 - X_i$. The rules use clauses with $\hat{\wedge}$ (AND) and $\hat{\vee}$ (OR) and their continuous extensions to build a loss function that can be minimized by gradient descent.

Optimal Classification Trees are variants of Decision Trees but use global mixed integer optimization to build an optimal tree. They are more efficient but less transparent, with more complex rules at each tree node.

Random Forests and Neural networks are also efficient learners but are less explainable because of the number of layers, nodes, and parameters.

FICO provided the data as part of a data science competition and includes 10,459 borrowers from March 2000 to March 2002 with a two-risk classification, creditworthy or not observed in March 2003, and a set of explanatory variables that cover a FICO measure of credit risk, dates regarding lines of credit, delinquency status and information, number of lines of credit, number of requests of new lines of credit,

fraction of lines of credit used, utilization rate, data on revolving and installment credit lines.

They compare the models using traditional measures for classifiers such as accuracy, the area under the curve (AUC), false positive rate (Type I error), and false negative rate (Type II error). They also compare their results with the model that won the competition. The Neural Network model delivers the best overall accuracy at 74.75% and the best AUC at 81.40%.

They use a variation of LIME (Ribeiro et al., 2016) but also consider constraints and correlation on inputs and a variation of SHAP (Lundberg and Lee, 2017) to explain the results. They can order the features of contribution importance to present the results. They can also stress test the results with LIME by replacing the inputs with more extreme values.

To create counterfactual results of the model (how the inputs should change to modify the model result) that would help users and other stakeholders understand the model better, the authors rely on the DiCE (Diverse Counterfactual Explanations) method proposed by (Mohtilal et al., 2020). Several k counterfactuals are found by minimizing a loss function that depends on the deviation from the counterfactual result y , a measure of diversity between the counterfactual inputs $c_1, \dots, c_k$, and a proximity measure so that the counterfactual inputs do not deviate too much from the original inputs x .

They also run a simpler optimal decision tree model with a small depth and few features per node to describe the dataset. They also use ILP with few decision rules to develop a simplified and more transparent model. These approaches that can still explain 70-72% of the outputs have the benefit of being more straightforward to interpret but are not necessarily good approximations of the real, more complex machine learning models.

8.4.2 Deep Learning for Mortgage Risk

(Sadhvani et al., 2021) apply a deep learning model to study the behavior of mortgage borrowers and the performance of their mortgage loans. This information is essential for banks and for investors in mortgage-backed securities who need to assess their risk and expected returns. Losses can occur because some borrowers default on their mortgage or prepay it and save on future interest payments.

The mortgage data comes from CoreLogic and covers around 120 million mortgage loans during nineteen years from January 1995 to June 2014. The data provides information on the features of the loans and their monthly performance.

The features at origination include the FICO score, Original debt-to-income ratio, Original loan-to-value ratio, Original interest rate, Original balance, Original term of loan, Original sale price, Buydown flag, Negative amortization flag, Occupancy Type (Owner occupied, second home, non-owner occupied or investment property, other), Prepayment penalty flag, Product type, Loan purpose (Purchase, Refinance Cash-out, Refinance No Cash Out, Second mortgage, Refinance Cash Out, Unknown, Construction Loan, Debt Consolidation Loan, Home Improvement Loan, Education Loan, Medical Loan, Vehicle Purchase, Reverse Mortgage, Other), Documentation (Full documentation, Low or minimal documentation, No asset or income verification, Other), Lien type (1st Position, 2nd Position, 3rd Position, 4th Position, Other), Channel (Retail Branch, Wholesale Bulk, Mortgage Broker, Realtor Originated, Relocation Corporate, Relocation Mortgage Broker, Builder, Direct Mail, Other Direct, Internet, Other Retail, Mortgage Banker, Correspondent, Other), Loan type (Conventional Loan, VA Loan, FHA Loan, Other Government Loan, Affordable Housing Loan, Pledged Asset Loan, Other), Number of units (1,2,3,4,5), Appraised value (is it < sale price?), Initial Investor Code (Portfolio Held, Securitized Other, GNMA/Ginnie Mae, GSE), Interest Only Flag, Margin for ARM mortgages, Periodic rate cap, Periodic rate floor, Periodic pay cap, Periodic pay floor, Lifetime rate cap, Lifetime rate floor, Rate reset frequency (1,2,3, . . . (months)), Pay reset frequency (1,2,3, . . . (months)), First-rate reset period (1,2,3, . . . (months since origination)), Convertible flag, Pool insurance flag, Alt-A flag, Prime flag, Subprime flag, Geographic state (CA, FL, NY, MA, etc.), Vintage (origination year) (less than 1995, 1995, 1996, .., 2014).

The performance data are Current status (seven states: Current, 30 days delinquent, 60 days delinquent, 90+ days delinquent, Foreclosed, Real Estate Owned (by the bank during foreclosure), paid off), Number of days late, Current interest rate, Current interest rate - national mortgage rate, Time since origination, Current balance, Scheduled principal payment, Scheduled principal + interest payment, Number of months the mortgage's interest has been less than the national mortgage rate and the mortgage did not prepay, Number of occurrences of current in past 12 months, Number of occurrences of 30 days delinquent in past 12 months, Number of occurrences of 60 days late in past 12 months, Number of occurrences of 90+ days outstanding in past 12 months, Number of occurrences of Foreclosed in past 12 months.

The model's objective is to predict the mortgage monthly performance based on the features and past behavior. Mortgages can transition to seven states (Current to paid off). Many features and performance data should inform the mortgage's future performance. We would expect, for instance, that a mortgage that has been delinquent for more than 90 days is more likely to be foreclosed or that the mortgage on a non-owner-occupied home is riskier than a rental property.

Current/Next	Current	30 days	60 days	90+ days	Foreclosure	REO	Paid Off
Current	96.7	1.6	0	0	.002	.0004	1.61
30 days	34.2	44.5	19.3	0	0.02	.003	1.84
60 days	12	16.7	34.5	33.8	1.9	.009	1.1
90+ days	4.1	1.4	2.5	80.4	9.9	.3	1.3
Foreclosure	1.9	.3	.1	6.8	86.8	2.6	1.3
REO	0	0	0	0	0	100	0
Paid Off	0	0	0	0	0	0	100

Figure 8.5. Unconditional transition matrix (%)

A mortgage n at time t has a state U_t^n that can be one of the seven states. The states are components of a vector $U_t = (U_t^1, \dots, U_t^N)$. The explanatory variables are described by a vector X_t^n that includes U_t^n , loan-level features, past behaviors of all the mortgages, as well as economic factors V_t^n (Monthly zip code median house price per square feet, Monthly zip code average house price, zip code house price change since origination, Monthly state unemployment rate, Yearly county unemployment rate, Original interest rate - National mortgage rate, Original interest rate - National mortgage rate at origination, Number of months where interest rate is < national mortgage rate, Median income in same zip code, Total number of prime mortgages in same zip code, Lagged subprime default rate in same zip code, Lagged prime default rate in same zip code, Lagged prime paid off rate in same zip code). They note $X_t = (X_t^1, \dots, X_t^N)$.

The model estimates the conditional transition probabilities (F_{t-1} is the appropriate information filtration):

$$P[U_t^n = u | F_{t-1}] = h_\theta(u, X_{t-1}^n)$$

h_{θ} is represented by a neural network and estimated by Maximum Likelihood. Regularization, dropout, and ensemble methods are used to limit the overfitting of the model.




Figure 8.6. The neural network to estimate transition probability to the seven states

The model is estimated on training data from 1995 to May 2012. The validation data are from May 2012 to October 2012. The test data are from November 2012 to May 2014. The goodness of fit is measured with the average log-likelihood.

Variable sensitivity is calculated using finite difference to approximate the derivatives of h_{θ} with respect to the components of X . For different transitions, the authors identify the highest sensitivities. For instance, for the transition from current to 30 days delinquent, the most critical factors are Number of Times 30 Days Delinquent in Last 12 Months, FICO Score, Number of Times 60 Days Delinquent in Last 12 Months, Current Outstanding Balance, Original Loan Balance, Original Interest Rate, Zillow Zip Code Housing Price Change Since Origination, Original Interest Rate - National Mortgage Rate, Number of Times 90+ Days Delinquent in Last 12 Months, Lagged Prime Default Rate in Same Zip Code, Number of Times Foreclosed in Last 12 Months, and Zillow zip code median house price change since origination.

The authors also study the nonlinear effects between mortgage performance and the most critical variables. The prepayment rate is not monotonic with the FICO score; it increases a lot then flattens after ten months with the Time since origination, increases then decreases to 20 LTV (Loan-To-Value ratio), flattens then drops again at 90 LTV.

They examine their out-of-sample pool-level prepayment predictions using their 5-layer neural networks and a logit regression model. The mortgages are divided into portfolios of 1,000 loans of similar probability of being current after 12 months. For each portfolio, they compare the realized prepayment rate in the 12 months and the model predictions. They find that the neural network results are more aligned on the 45-degree lines (predicted number of prepayments vs observed number of prepayments). The logit model tends to overpredict the number of prepayments.

8.4.3 Quant GANs: Deep Generation of Financial Time Series

With Quant GANs, (Wiese et al., 2020) can generate financial time series with accurate distributional and dependence properties, including volatility clusters, leverage effects, and serial autocorrelations. The authors use Generative Adversarial Networks (GANs) and a time series model called Temporal Convolutional Network (TCN).

A TCN is a neural network that works for time series. The time series observations form an input layer are grouped into K (the kernel size) observations and then feed into nodes of a hidden layer, where the nodes are separated by a dilation factor D . A layer is a causal convolutional layer that uses dilated causal convolutional operator noted $*_D$ to transform the input.

If the input X is of dimension $N_I \times T$, and the output of dimension $N_O \times T$, with $m \in \{1, \dots, N_O\}$

$$(W *_D X)_{m,t} = \sum_{i=1}^K \sum_{j=1}^{N_I} W_{i,j,m} \cdot X_{j,t-D(K-i)}$$

The causal convolutional layer applies the function.

$$w(X)_{m,t} = (W *_D X)_{m,t} + b_m$$

(N_I, N_H, K, D) are arguments of the causal convolutional layer.




Figure 8.7. TCN with $K=2$ and $D=2$

These TCNs are then used in a GAN with a generator (or actor) and a discriminator (or critic). The generator is a network that creates a time series, and the discriminator is a network that wants to differentiate between a generated time series and an original time series. These networks are trained by gradient descent and are trained alternatively.

Quant GAN is used to model the daily log returns of the S&P 500 index (from May 2009 to December 2018) and is compared to the results with a GARCH(1,1) model and a Constrained Log Return Neural Process (C-SVNN), and a GARCH(1,1). It is implemented with Python, PyTorch, and GPU (RTX 2070 from NVIDIA), which were used for model implementation and training. Preprocessing and evaluation were done using NumPy and SciPy.

The models are compared using various metrics, including Earth Mover Distance (EMD) (how much probability mass has to be moved to transform one return distribution into another, return over t periods), DY metric (difference of the empirical probability density functions of the historical and generated t -differenced log path, $t=1, 5, 20, 100$), Autocorrelation Function (ACF) (dependence properties of the historical and the generated time series of log returns, their absolute values and their squared), and Leverage Effect score (using the correlation of the lagged, squared log returns and the log returns) (Figure 8.8).

	TCN	C-SVNN with drift	GARCH(1,1)
EMD(1)	0.0039	0.004	0.0199
EMD(5)	0.0039	0.004	0.0145
EMD(20)	0.004	0.0069	0.0276
EMD(100)	0.0154	0.0464	0.0935
DY(1)	19.1199	19.8523	32.709
DY(5)	21.1167	21.2445	27.476
DY(20)	26.3294	25.0464	39.3796
DY(100)	28.1315	25.8081	46.4779
ACF(id)	0.0212	0.022	0.0223
ACF(.)	0.0248	0.0287	0.0291
ACF(.,.)2	0.0214	0.0245	0.0253
Leverage Effect	0.3291	0.3351	0.4636

Figure 8.8. Metrics (Green indicates the best model)

Quant GANs, particularly the TCN model, outperform the GARCH(1,1) model in modeling the distribution and dependence properties of the S&P 500 log returns. The C-SVNN performs comparably to the TCN but slightly worse in most metrics due to the model's constraints. The GARCH(1,1) model performs the worst in distributional and dependence properties, particularly in capturing the tails of the log return distribution.