Chapter 2 | Portfolio Choice and Asset Allocation

> "Since there were two criteria - expected return and risk – the natural approach for an economics student was to imagine the investor selecting a point from the set of Pareto optimal expected return, variance of return combinations, now known as the efficient frontier. These were the basic elements of portfolio theory which appeared one day while reading Williams." ~ Harry Markowitz

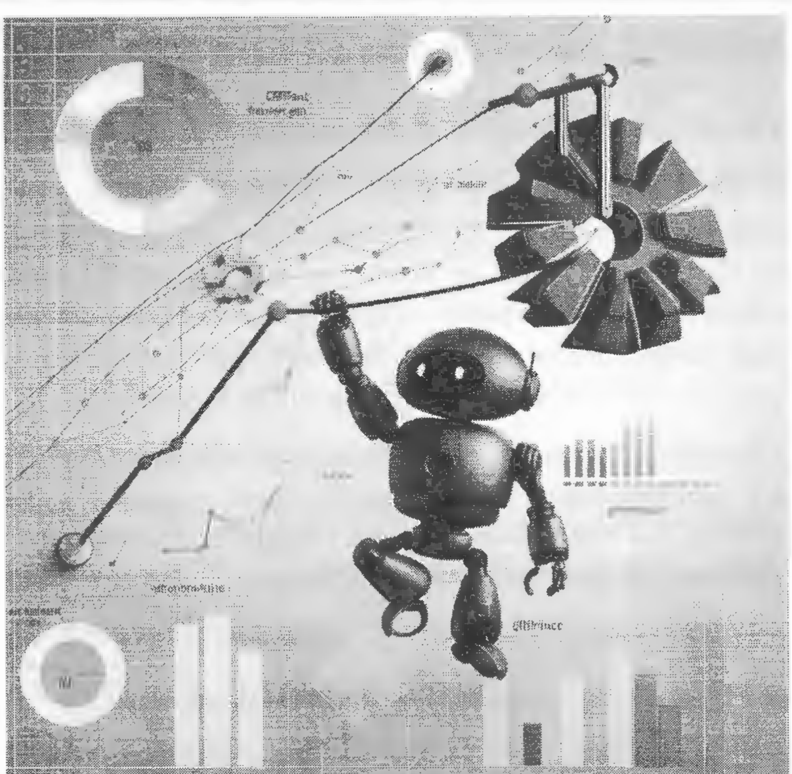

2.1 Theory of Portfolio Choice and Asset Allocation

Portfolio choice and asset allocation consist of deciding how much share of a portfolio to allocate to different individual securities or assets such as stocks, bonds, and commodities. Different heuristic rules exist, such as equal weighting (the same financial amount in each asset) or market weighting (proportional to the outstanding amount in the market). More advanced allocations account for expected returns, asset return volatility, and return correlations.

2.1.1 Preferences, Risk, and Expected Utility

Starting with a couple of principles, the axiom of continuity and the axiom of independence, it is possible to show that the preferences of an economic agent can be modeled by a Von Neumann-Morgenstern expected utility function of wealth (or well-being) w:

$$Eu(w) = \sum_{i=1}^{N} p_i u(w_i)$$

which is a linear combination of economic outcomes $u(w_i)$ and the weights are the probabilities p_i of these outcomes (Gollier, 2004).

The continuity axiom states that if an agent prefers a risky asset A to a risky asset B to a risky asset C, there is a linear combination of assets A and C such that the agent is indifferent between that combination of A and C and the single asset B. The axiom of independence states that if an agent prefers asset A to B, then she will still prefer a combination of A and C to the same combination of B and C.

The utility function u is typically increasing because an agent prefers more wealth than less, and it is concave because the agent prefers more certainty than less for the same average outcome. It is represented by Jensen's inequality that states the expected utility of an uncertain outcome is lower than the utility of the expected outcome:

$$Eu(w) \leq u[E(w)]$$

This also implies that there exists a positive number π, the risk premium, so that:

$$Eu(w) = u[E(w) - \pi]$$

$E(w) - \pi$ is the certainty equivalent.

If the utility function was linear, we could invert the utility function and the expectation operator, and the risk premium would be zero. There would be no risk premium.

If the wealth takes the form $w = w_0 + y$, where y is a random variable with mean zero, The Arrow-Pratt approximation (Gollier, 2004) shows, using a Taylor expansion around 0, that the risk premium can be written as

$$\pi \approx 0.5E(y^2)A(w_0)$$

Where $A(w_0) = -\frac{u''(w_0)}{u'(w_0)}$ is the coefficient of absolute risk aversion.

If the wealth takes instead the form $w = w_0(1 + y)$, where y is a random variable with mean zero, then:

$$\pi \approx 0.5E(y^2)R(w_0)$$

The higher an agent's relative risk aversion R, the higher the risk premium.

If the coefficient $A(w)$ is constant we can solve for the utility function, it can be shown that the function $u(w) = -\frac{exp(-Aw)}{A}$, also called a Constant Absolute Risk Aversion (CARA), has a constant A.

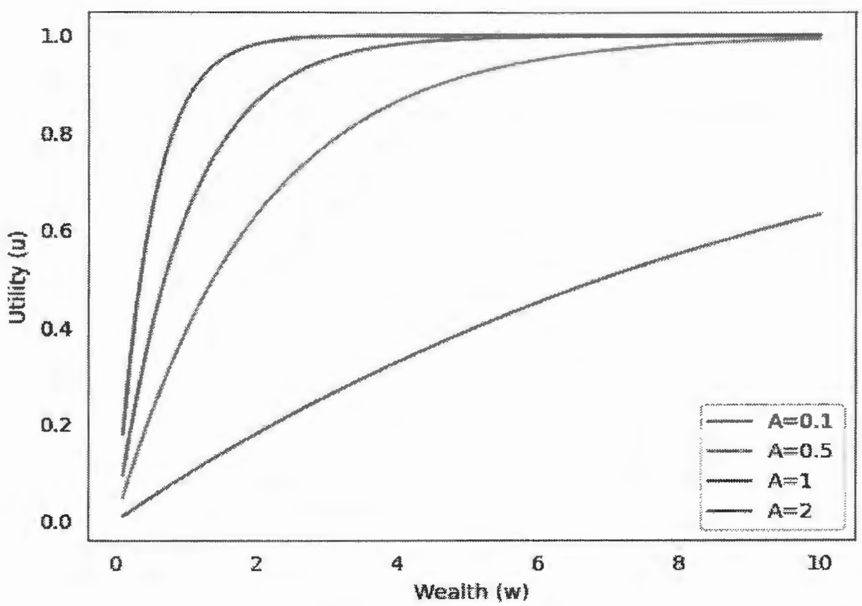

Figure 2.2. CARA utility function

For a constant $R(w)$ then $u(w) = \frac{w^{1-\gamma}}{1-\gamma}$, if $\gamma \neq 1$ or $u(w) = ln(w)$ if $\gamma = 1$. It is the Constant Relative Risk Aversion (CRRA) function.

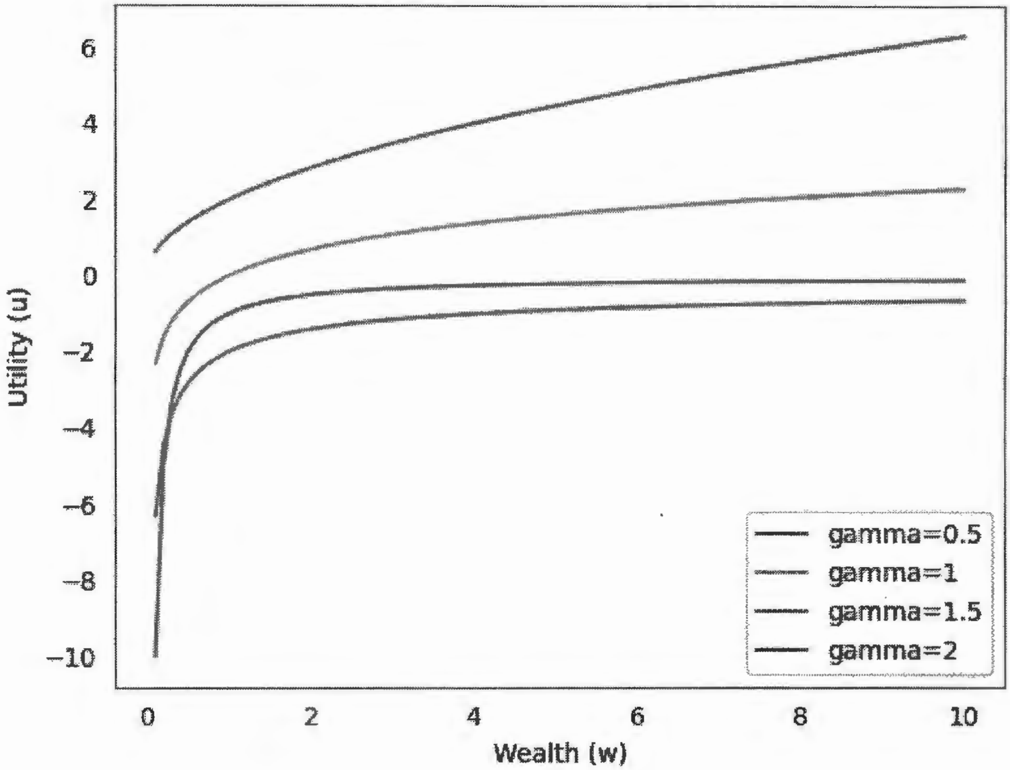

Figure 2.3. CRRA utility function

The utility function can also be quadratic $u(w) = aw - bw^2 + c$, b has to be positive for the function to be concave, and $a - 2bw > 0$ for the function to be increasing with w. This functional form is very common in finance as it simply models the risk-return trade-off (in the return space).

Other specifications exist for the utility function, such as the Epstein-Zin utility function (Epstein and Zin, 1991) that separates the elasticity of intertemporal substitution (willingness to substitute current consumption against future consumption in response to change in interest rates) and the degree of risk aversion (in the case of CRRA, they are inversely related).

2.1.2 Static Portfolio Choice and Asset Allocation

2.1.2.1 Mean-Variance Optimal Portfolio

One way to build a portfolio is to use the mean-variance portfolio analysis of (Markowitz, 1952). We need to estimate the expected returns of each asset, the covariance matrix of the assets, and solve the portfolio weights by constrained quadratic optimization.

Suppose w is a vector of portfolio weights, μ the vector of expected returns, and Σ the variance-covariance matrix. In that case, the investor wants to maximize the portfolio return under some risk constraint or minimize risk under some return constraint. She wants to solve the following:

$$max_w w'\mu$$
$$w'\Sigma w \leq \sigma^2$$

Or equivalently

$$max_w w'\mu - \lambda w'\Sigma w$$

The solution is $w = \frac{1}{2\lambda}\Sigma^{-1}\mu$ and the portfolio weights have to sum to a value W then $\frac{1}{2\lambda}1'\Sigma^{-1}\mu = W$ and $1'$ is a vector on 1s. Then

$$w = \frac{\Sigma^{-1}\mu}{1'\Sigma^{-1}\mu}W$$

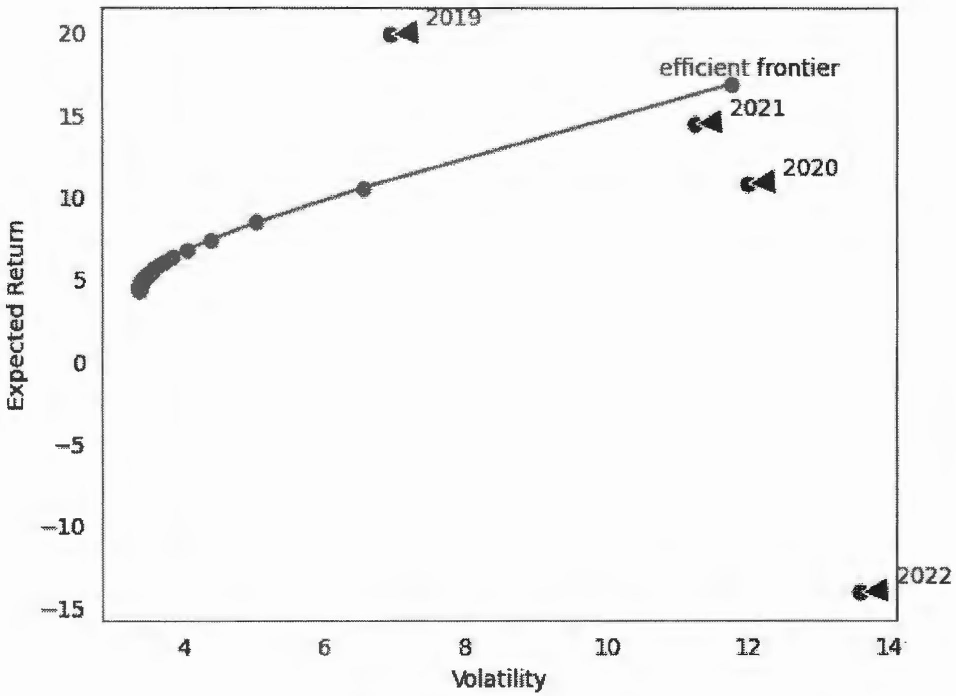

Figure 2.4. Mean-variance efficient frontier for Norges Bank fund and performance in 2019-2022

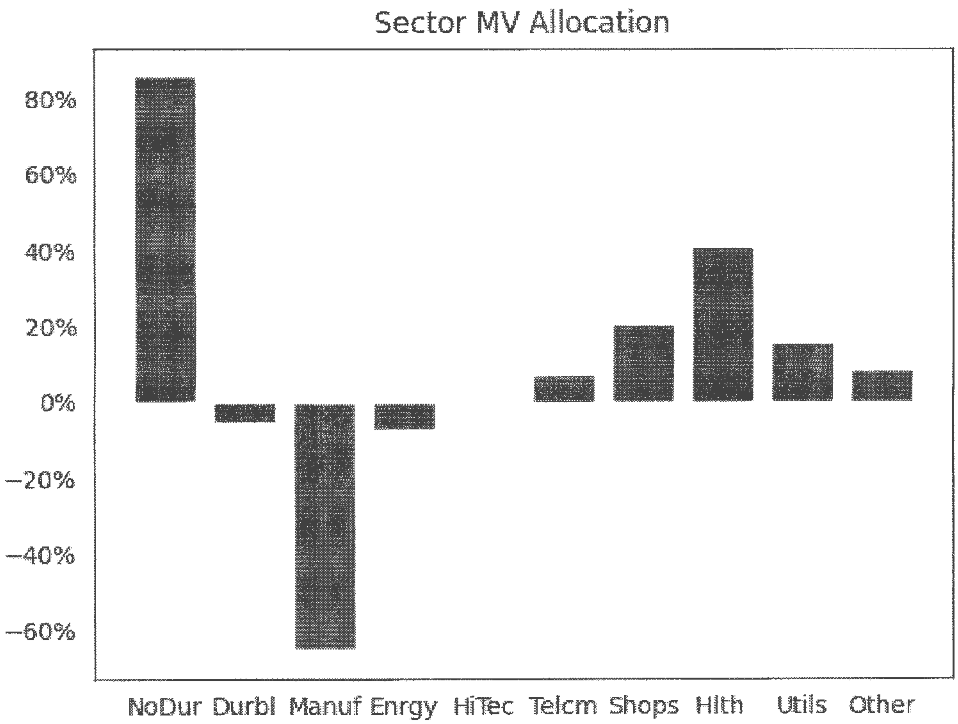

Figure 2.5. Mean-variance sector portfolio weights

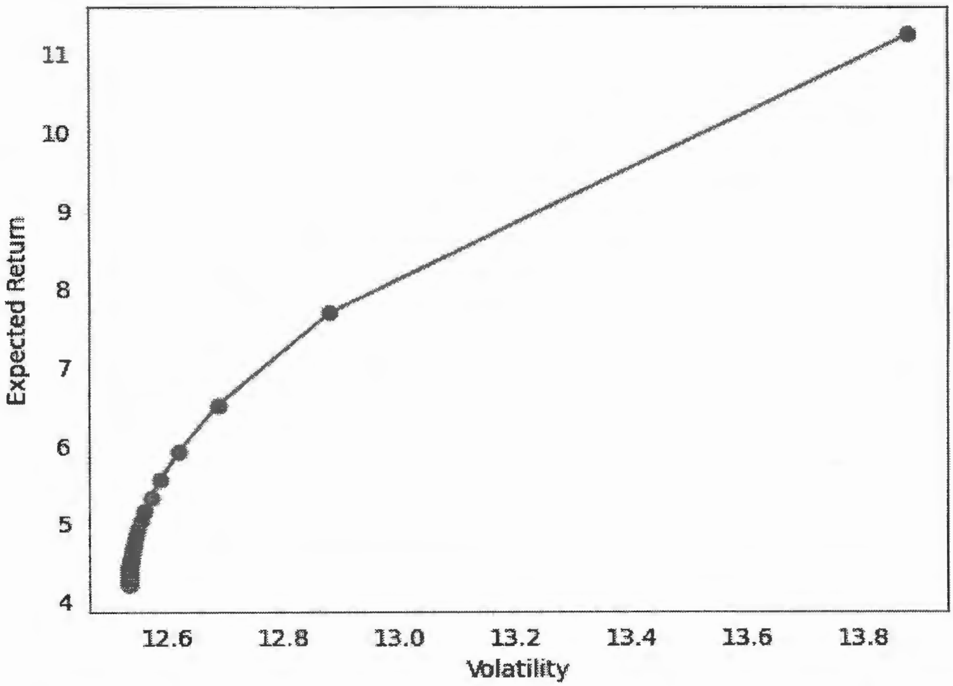

Figure 2.6. Mean-variance efficient frontier for Fama French sectors

Therefore, we need to calculate the expected returns μ and the covariance matrix Σ. A long time series of return observations is required to estimate returns accurately when they are stationary. If expected returns are time-varying because risk premia are time-varying (high in bearish markets and low in bullish markets), for instance, then precise estimation might be extremely difficult. (Merton, 1980) points out that estimating expected returns should go beyond the simple arithmetic average of historical returns but include changes in risk and adjusted for heteroscedasticity.

Similar problems appear with the covariance matrix. Risk can also be time-varying (high in economic recession and low in expansion). Estimation errors in the covariance matrix can lead to very imprecise estimates of portfolio weights. For instance, if the assets were independent, an asset with minimal volatility would have a large weight depending on the expected return. One possible regularization technique is to shrink the covariance matrix towards the identity matrix as in (Ledoit and Wolf, 2004).

Another method is to impose no short-sale constraints (the portfolio weights cannot be negative). It is a natural condition for long-only portfolios. Without such constraints, a mean-variance portfolio can have extremely positive and negative positions, especially for highly correlated assets. (Green and Hollifield, 1992) show that this situation is common because of the presence of (at least) one single factor in the covariance matrix.

(Jagannathan and Ma, 2003) suggest that no short-sale constraints improve the efficiency of optimal portfolios because it reduces the covariance of each asset with the others and upward estimation errors. (Best and Grauer, 1991) show that the weights of a positively weighted MV-efficient portfolio can be highly sensitive to changes in the means of individual assets. Even a slight increase in the mean of one asset can drive half the securities from the portfolio. In contrast, the expected return and standard deviation of the portfolio remain nearly unchanged.

(DeMiguel et al., 2009a) present another framework for constructing portfolios that perform well out-of-sample in the presence of estimation error. The framework solves the minimum-variance problem subject to the constraint that the norm of the portfolio-weight vector is smaller than a given threshold (a form of regularization). The new portfolios are compared to 10 strategies in the literature across five data sets, and the norm-constrained portfolios have a higher Sharpe ratio than existing strategies.

2.1.2.2 Minimum-Variance Portfolio

A particular case of the mean-variance optimal portfolio is the minimum-variance portfolio. We solve for the portfolio weights by minimizing the volatility (or variance) of the portfolio (choose σ^2 to be the smallest possible):

$$min_w w'\Sigma w$$

$$1'w = W$$

The result is $w = \frac{\Sigma^{-1}1}{1'\Sigma^{-1}1} W$ since implicitly it assumes the returns are the same across assets and focuses only on the covariance.

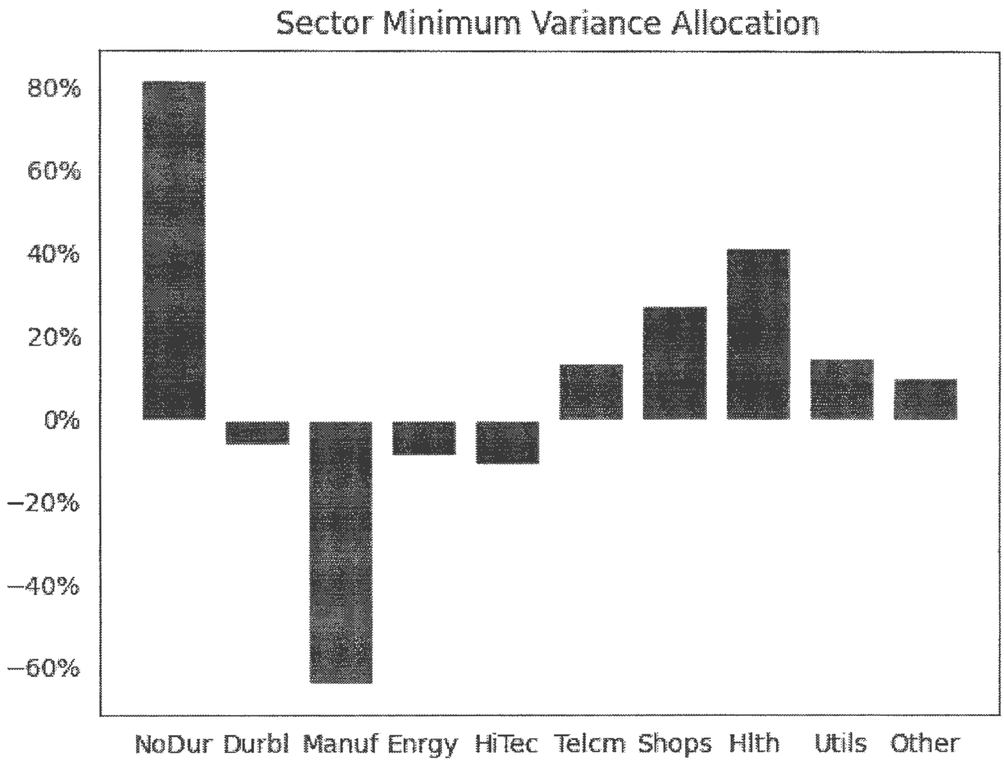

Figure 2.7. Portfolio weights the minimum variance sector portfolio

The advantage of this portfolio is that we do not need to estimate the expected return of each asset. The vector of expected returns μ does not appear anymore.

There is still the remaining challenge of estimating the variance-covariance matrix, especially when the number of assets is large and when the asset volatilities and correlations are time-varying.

2.1.2.3 Buy-and-hold Portfolio

Figure 2.8. Market capitalization by sector

The buy-and-hold portfolio is the simplest portfolio. It is the market portfolio. It avoids transaction costs but can be volatile. According to the CAPM, everybody should hold a portfolio with the market portfolio and a risk-free asset. The amount allocated between the market portfolio and the risk-free asset depends on the risk aversion of the investor. This has been shown by (Tobin, 1958) for mean-variance investors.

(Kan and Zhou, 2007) demonstrate, however, that when the returns and the variance-covariance matrix are estimated with errors, the previous two-fund separation theorem does not hold anymore. Still, a third fund, the minimum variance portfolio, improves the out-of-sample performance of three fund portfolios (the risk-free asset, the tangency or market portfolio, and the minimum variance portfolio).

In practice, buy-and-hold portfolios can be acquired cheaply through index funds and ETFs. According to the Investment Company Institute (Investment Company Institute, 2022), US index mutual funds and ETFs grew from 21% of assets (2 trillion dollars) in long-term funds in 2011 to 43% (12.6 trillion dollars) in 2021. The average expense ratios of index equity mutual funds were 6 bps in 2021 vs. 68 bps for active equity funds.

Interestingly, there is mixed evidence that the market portfolio itself is mean-variance efficient. Some tests that compare returns of the market portfolio and portfolios on the mean-variance frontier seem to reject it (Brière et al., 2013; Roll, 1977). In theory, there is no reason why the market portfolio should only be restricted to a portfolio of stocks of publicly listed companies; it probably should also include other assets (bonds, real estate. Commodities, private equity) and even foreign assets.

2.1.2.4 Equally-weighted Portfolio

Figure 2.9. Equally sector share

In an equally-weighted portfolio, we hold the different assets with the same nominal weight at the date of formation of the portfolio. If there are N assets, the weights are simply $w_i = 1/N, i = 1..N$.

We can rebalance the portfolio at different frequencies to maintain equal weighting. This rebalancing can be very costly if there are too many assets, say all the stocks in the investment universe, in the portfolio, including very small and illiquid stocks.

Though equal-weighting is a relatively unsophisticated way to build a portfolio, it is often surprisingly effective when the assets are not individual assets but rather portfolios. (DeMiguel et al., 2009b) compare the 1/N portfolio strategy with 14 optimal portfolio strategies using different sets of assets:

- Ten sector portfolios of the S&P 500 index and the US equity market portfolio
- Ten industry portfolios and the US equity market portfolio
- Eight country indexes and the World equity index
- SMB, HML, and US equity market portfolio MKT
- Twenty size and book-to-market portfolios and US equity MKT
- Twenty size and book-to-market portfolios and US equity MKT, SMB, and HML portfolios
- Twenty size and book-to-market portfolios and US equity MKT, SMB, HML, and UMD portfolios
- Simulated data with 10, 25, or 50 portfolios

They find that none of the 14 portfolio strategies consistently outperform the equally weighted strategy regarding Sharpe ratio, certainty-equivalent return, or turnover because of estimation errors.

(Yuan and Zhou, 2022) discuss how estimation risk arises when implementing mean-variance portfolio optimization due to the need for accurate estimates of asset returns and covariance matrices. Several strategies have been proposed to mitigate this risk. The 1/N investment strategy, which allocates equal weights to all assets, is compared with more advanced strategies in estimating risk. Surprisingly, it is found that for portfolios with a large number of assets, the estimation window needed for advanced strategies to outperform the 1/N strategy is impractically long.

They find that the 1/N rule can be optimal when there is one factor that prices all assets with diversifiable risks. In such cases, the 1/N portfolio is equivalent to the factor plus idiosyncratic risks, making it optimal for large portfolios. The theoretical findings are illustrated using empirical examples, including a calibrated economy and applying the 1/N rule to selecting stocks from the S&P 500. They propose tractable portfolio rules that can outperform the 1/N rule under certain conditions. They examine cases where N is smaller relative to the sample size T and cases where N is greater than T.

2.1.2.5 Risk-Parity Portfolio

The risk-parity portfolio allocates the risk to each asset so that the asset contribution to the total portfolio risk is identical. Because portfolio risk is a homogeneous function (a doubling of the portfolio size doubles the risk), we can apply Euler's equation that states that the volatility is equal to the sum of weighted derivatives of risk with respect to the portfolio weights:

$$\sigma = \sum_{i=1}^{N} w_i \frac{\partial \sigma}{\partial w_i}$$

The portfolio risk contribution (PRC) of each asset i is $PRC_i = w_i \frac{\partial \sigma}{\partial w_i}$

$$\sigma = \sum_{i=1}^{N} PRC_i$$

The risk parity portfolio will make all the PRC_i equal to σ/N. It ignores the expected return of each asset. It will leverage up the low-risk assets (with a low marginal contribution to risk) or underweight the riskiest assets (with a high marginal contribution to risk).

In the particular case where all the asset correlations are equal, it is possible to show (Maillard et al., 2010) that the weights are equal to the inverse of the volatility of each asset, normalized by the harmonic mean of the volatilities. Thus:

$$w_i = \frac{1/\sigma_i}{\sum_{j=1}^{N} 1/\sigma_j}$$

This is quite intuitive, if an asset is very volatile, its contribution to risk is high, so it needs to be scaled down. On the contrary, if the asset is not very volatile, it must be scaled down. So typically, the stock market exposure has to be lowered, and government bond exposure has to increase, potentially using leverage. In the general case, no closed-form formula and simple optimization algorithms can be used to find the portfolio weights.

In the case of our sector portfolio, the weights would be as follows:

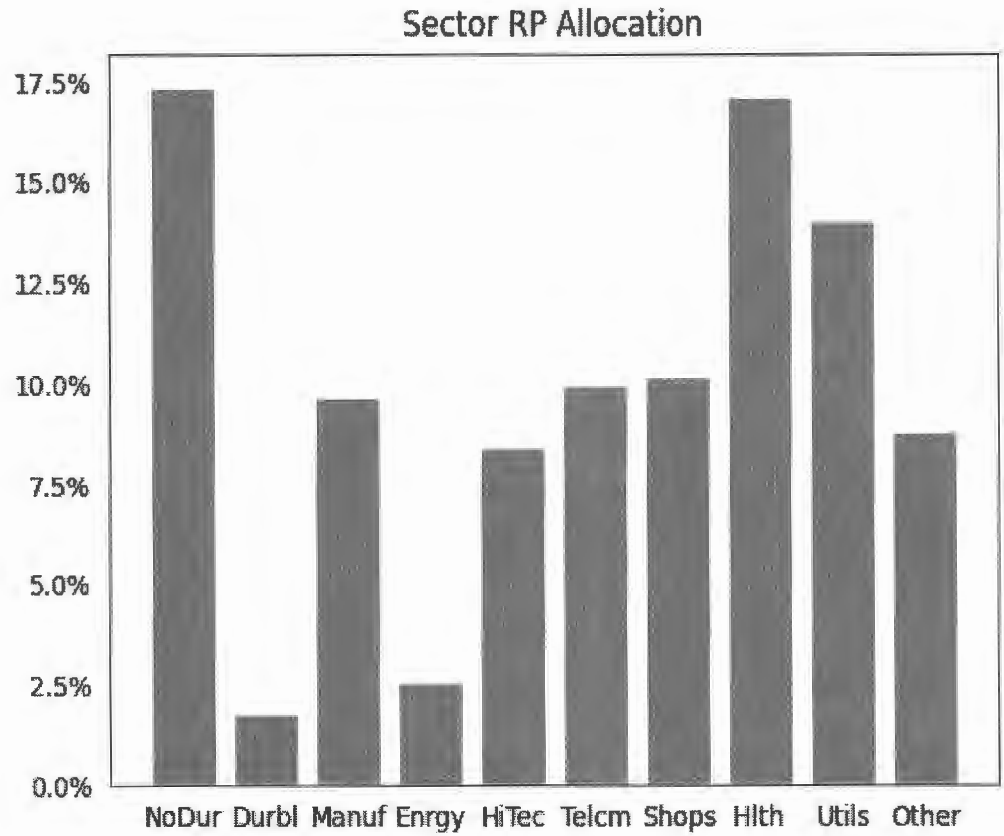

Figure 2.10. Portfolio weights the risk parity sector portfolio (assuming equal correlation)

2.1.2.6 Risk Budgeting

Risk budgeting is traditionally used in the context of a portfolio evaluated against a benchmark index (e.g., the S&P500 for a large-cap equity fund or the Bloomberg U.S. Aggregate Bond Index for a fixed-income fund). The standard deviation of the portfolio excess return vs. the benchmark is called the active risk. A portfolio manager will optimize the active risk by allocating risk across the different securities in his portfolio. This is the risk budgeting allocation procedure.

(Cetingoz et al., 2022) discuss risk budgeting portfolios in the context of modern portfolio theory. They explore mathematical results related to the existence and uniqueness of risk budgeting portfolios for a wide range of risk measures. They demonstrate that, for many of these measures, computing the weights of such portfolios can be achieved using standard stochastic algorithms.

They highlight Equal Risk Contribution (ERC) as a widely adopted Risk Budgeting method, where risk budgets are distributed equally among assets. ERC portfolios are appealing because they are less sensitive to covariance matrix estimation errors.

They present mathematical results demonstrating the existence and uniqueness of Risk Budgeting portfolios for various risk measures. They use a convex minimization problem and show that, up to rescaling, the condition defining Risk Budgeting portfolios can be met. They use standard stochastic gradient descent (SGD) algorithms for computing Risk Budgeting portfolios. Their approach is applied to a wide range of risk measures, including deviation measures like volatility and spectral risk measures like Expected Shortfall.

In a dynamic investment context, (Jaimungal et al., 2023) propose a concept called "dynamic risk contributions," which extends traditional risk contributions and allows for a recursive calculation of risk contributions.

They define dynamic risk contributions using Gâteaux derivatives (a generalization of the directional derivative concept in multivariable calculus), primarily focusing on the class of dynamic coherent distortion risk measures. This approach allows for precisely measuring how each asset contributes to the overall risk over time. It introduces dynamic risk budgeting portfolio strategies, where the risk contributions of assets are predetermined percentages of the future risk of the strategy. The authors prove that any self-financing dynamic risk budgeting strategy with an initial wealth of 1 can be expressed as a scaled version of a unique solution to a sequence of strictly convex optimization problems. To solve the sequence of optimization problems efficiently, they develop an actor-critic approach that leverages deep learning techniques.

2.1.2.7 Bayesian Portfolio

The mean-variance paradigm introduced by Markowitz is a well-accepted framework for portfolio selection. However, there are limitations of classical approaches, especially in handling estimation risk and model uncertainty. (Avramov and Zhou, 2010) review Bayesian approaches to portfolio analysis, focusing on their advantages and performance measures compared to classical methods.

Bayesian methods for portfolio analysis offer advantages such as incorporating prior information, accounting for estimation risk and model uncertainty, and using intuitive numerical algorithms. Bayesian portfolio analysis involves three key components: forming prior beliefs, specifying the law of motion for asset returns, and recovering the predictive distribution of future asset returns. The Bayesian optimal portfolio is obtained by maximizing expected utility with respect to the predictive distribution.

The authors introduce the concept of performance measures in Bayesian portfolio analysis, particularly the Certainty Equivalent Return (CER) and objective utility gain (loss). These measures help assess the benefits of using informative priors versus diffuse priors. Informative priors can significantly impact portfolio decisions, especially when there is limited historical data or certain assets have longer histories than others. Utility gain (loss) is a performance metric to compare portfolio strategies. It measures the expected utility differential between alternative strategies and can help investors make informed choices.

The authors discuss how alternative portfolio estimators can be compared using simulations or analytical methods, even without knowing the true parameters. Bayesian solutions are shown to dominate classical rules in some cases, but informative priors are often necessary to leverage Bayesian methods fully.

Priors can be of different kinds. They can be diffuse or uninformative priors (Barry, 1974; Bawa, Brown, and Klein, 1979). They can be identical priors (Frost and Savarino, 1986). They can use shrinkage estimators (e.g., Jobson, 1979; Jorion, 1986).

Despite its theoretical appeal, the Bayesian portfolio allocation is uncommon in the industry. An exception is the Black and Litterman (1990, 1992) model, which addresses two issues that have limited the adoption of quantitative portfolio modeling by practitioners. First, investors have views on expected returns only for a few assets in a portfolio, and second, standard mean-variance optimization is very sensitive to expected return assumptions. Slightly different expected returns can lead to different weights, especially when short positions are allowed and some assets are highly correlated.

Asset pricing models can be used to improve the performance of portfolio selection algorithms. (Pastor, 2000; Pastor and Stambaugh, 2000) study the problem of portfolio selection in the presence of asset pricing models. They propose a new approach to portfolio selection that is based on the idea of using asset pricing models to estimate the expected returns of assets. They show that the new approach can be used to construct portfolios that have higher expected returns and lower risk than portfolios built using traditional methods. They apply the new approach to a sample of US stocks and find that it produces portfolios with higher expected returns and lower risk than portfolios constructed using traditional methods. The authors also find that the new approach can be used to improve the performance of investment strategies.

2.1.2.8 Robust Portfolios

Financial models rely on many assumptions and estimations of financial variables, such as expected returns and covariance matrices. (Raponi et al., 2021) present a normative theory for constructing mean-variance portfolios that are robust to model misspecification. They identify two types of portfolios: an "alpha" portfolio, which depends on pricing errors, and a "beta" portfolio, which relies on observed factor risk premia. These portfolios, when combined, form mean-variance efficient portfolios. The study argues that misspecification in these portfolios should be addressed differently due to their distinct economic properties.

The key insights from this research led to a significant improvement in out-of-sample portfolio performance, with a particular focus on the role of latent asset demand. The authors emphasize distinguishing between alpha and beta components and addressing misspecification separately. They also demonstrate that as the number of assets in the portfolio increases, the alpha portfolio's elements dominate, making it crucial to handle misspecification in this component effectively.

Their theoretical foundation relies on the Arbitrage Pricing Theory (APT) and emphasizes the significance of asset-specific pricing errors, challenging traditional models focusing solely on pervasive risk factors. Empirical results indicate a substantial enhancement in portfolio performance, especially when accounting for latent asset-specific risk. The study highlights the APT's ability to handle model misspecification and suggests different approaches for alpha and beta portfolios.

2.1.3 Dynamic Portfolio Choice

(Merton, 1971, 1969) solves the problem of optimal investment and consumption in continuous time. The portfolio weights can be adjusted through time, which is probably more realistic than the case of static portfolios. (Campbell and Viceira, 2002) provide an example of the Merton problem with one risky asset with price *P*, one risk-free asset with price *B*, and a state variable *S*:

$$\frac{dP_t}{P_t} = \mu_P(S,t)dt + \sigma_P(S,t)dZ_{P,t}$$

$$\frac{dB_t}{B_t} = r(S,t)dt$$

$$dS_t = \mu_S(S,t)dt + \sigma_S(S,t)dZ_{S,t}$$

The risk asset and the state variables follow geometric Brownian motions like in the Black-Scholes formula (Chapter 4). They each have drift and, except for the risk-free asset, a stochastic component. The investor will aim to maximize the lifetime utility function by controlling her fraction *w* of wealth invested in the risk asset and her consumption *C*:

$$max_{C,w} E_0\left[\int_0^{\infty} U(C,t)dt\right]$$

Total wealth *W* has to follow the dynamic:

$$dW = [(w(\mu_P - r) + r)W - C]dt + wW\sigma_P dZ_P$$

If *J(W,S,t)* is the value function result of the utility maximization problem, it can be shown to verify a Partial Differential Equation that can be solved either in closed form under some conditions or numerically. The optimal allocation to the risky asset will depend on the asset risk premium divided by its volatility and a hedging demand component that depends on the state variable, the correlation with the asset, and the marginal utility of wealth.

(Campbell and Viceira, 2002) show that if there is volatility risk with volatility negatively correlated with stock returns, investors should reduce their allocation to the risky asset (equities), even more so if they are long-term investors. This effect is, however, relatively small compared to the negative hedging demand for stocks due to changes in interest rates or risk premiums.

(Moreira and Muir, 2016) introduce a trading strategy called "volatility managed portfolios" that adjusts investment allocations based on the previous month's realized variance of assets. Their strategy aims to decrease risk exposure when the variance is high and increase it when the variance is low. The volatility-managed portfolios seek to exploit variations in asset volatility over time. When asset volatility is high, the strategy reduces risk exposure, and when volatility is low, it increases exposure. The authors find that this simple strategy generates significant positive risk-adjusted returns (alphas) across various asset pricing factors. These alphas indicate that investors can benefit from timing volatility.

The strategy's effectiveness is demonstrated across asset classes, including equities, currencies, and factors like value, momentum, profitability, etc. The utility gains from using the volatility-managed strategy are substantial, potentially increasing a mean-variance investor's lifetime utility by a significant percentage. The strategy is also applied to portfolios constructed from multiple factors, with positive alphas observed consistently. These portfolios expand the mean-variance frontier. The results are robust across various factors, time periods, and multifactor combinations. The strategy's efficacy is not limited to specific asset classes or timeframes. The strategy is shown to have performed well during various crisis episodes, including the Great Depression, the Great Recession, and the 1987 stock market crash.

In the same line of work, (DeMiguel et al., 2021b) challenge the traditional finance notion of a strong risk-return tradeoff by showing that investors can improve their Sharpe ratios by reducing exposure to risk factors during periods of high volatility. Previous studies have discussed the benefits of managing volatility in individual-factor portfolios, but this study introduces a novel approach by proposing a multifactor portfolio that adjusts its weights on each factor based on market volatility. The authors find this strategy outperforms unconditional multifactor portfolios even when accounting for costs.

Their conditional multifactor portfolio significantly outperforms its unconditional counterpart, indicating that estimation errors and transaction costs do not explain the gains from volatility management. The performance improvement is attributed to trading diversification, optimizing for transaction costs, and varying the relative weights of factors based on market volatility. The risk-return tradeoff for individual factors weakens as market volatility increases, providing a rationale for the effectiveness of the conditional multifactor portfolio. The study explores the asset-pricing implications of these findings, demonstrating that the price of risk for individual factors generally decreases with realized market volatility.

We have seen that CAPM does not hold, and some factors, such as value or momentum, provide some risk premia. (Polk et al., 2022) explore dynamic portfolio choices for long-horizon investors when traditional CAPM-based models do not perform well in predicting asset prices and portfolio allocations. They focus on the value and momentum anomalies and aim to guide the optimal portfolio for long-term investors. They use an equilibrium approach, assuming asset prices are determined as capital moves between underperforming and overperforming market segments.

They find that over short horizons (up to two years), the Sharpe ratios of value and momentum strategies decrease as the horizon increases, reflecting short-horizon positive autocorrelation in returns. Over longer horizons, the Sharpe ratio of momentum remains relatively constant, while the Sharpe ratio of value significantly increases, overtaking momentum for horizons longer than thirteen years. Combining value and momentum strategies provides diversification gains. The correlation between these strategies changes from negative over short horizons to positive over longer horizons.

A typical long-horizon investor is a foundation or an endowment. The Yale endowment model is a benchmark for university endowments (Swensen, 2009). "The endowment model" emphasizes investing in less liquid and efficiently priced assets that could earn a risk premium for a long-term investor. In 2020, domestic US marketable securities make up less than 10% of the portfolio, while foreign equity, private equity, absolute return strategies, and real assets constitute over 90% of the Endowment's investments ("Yale's Strategy," n.d.).

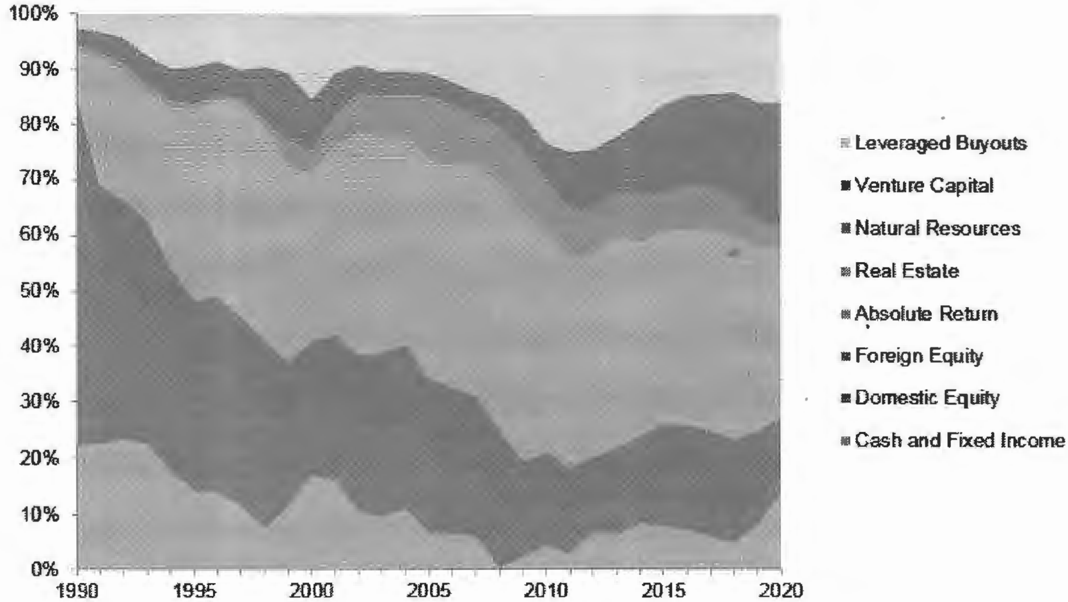

Figure 2.11. Yale University endowment asset allocation. Source: https://investments.yale.edu/about-the-yio

How many illiquid assets should an endowment have in its portfolio? (Dimmock et al., 2021) address this investment strategy and develop a dynamic asset allocation model based on modern portfolio theory (MPT) to assess the endowment model. They incorporate illiquid investment opportunities and key features of alternative assets, such as private equity and hedge funds, into the MPT framework. The model considers illiquidity, transaction costs, and natural liquidity events.

The model reveals that allocations to illiquid alternative assets follow a double-barrier policy, allowing them to rise or fall until they reach predefined rebalancing boundaries. This is in contrast to classic MPT, where allocation ratios are constant over time. The study shows that investors who stagger their investments in alternative assets over time (liquidity diversification) can maintain more stable portfolio allocations and achieve higher allocations to alternatives. The model considers investor characteristics, such as spending flexibility, and demonstrates how these factors affect asset allocations. Investors with higher spending flexibility tend to accept higher allocations to illiquid assets.

2.2 Empirical Portfolio Choice and Asset Allocation

2.2.1 Portfolio Construction and Optimization

The Norwegian Government Pension Fund Global (GPFG), which has gained recognition as a model for managing financial assets, is managed by Norges Bank Investment Management ("Norges Bank") (Chambers et al., 2011). The GPFG, ranked as the largest fund globally, follows a professional, low-cost, transparent, and socially responsible approach to asset management.

We outline the investment universe and benchmark index for the investment portfolio of Norges Bank, specifically focusing on the equity and fixed-income portfolios:

Section 2-1: Investment Universe

The Bank can invest the investment portfolio in financial instruments, real estate, renewable energy infrastructure, and cash deposits, subject to certain limitations.

Permissible investments for each portfolio:

a. Equity portfolio: Equities listed on a regulated and recognized marketplace, listed securities equivalent to equities, depositary receipts for equities, and unlisted companies intending to seek listing.

b. Fixed-income portfolio: Tradable debt instruments and depositary receipts for such debt instruments.

c. Unlisted real estate portfolio: Real estate, equity, and interest-bearing instruments issued by unlisted companies and other legal entities focused on real estate.

d. Unlisted renewable energy infrastructure portfolio: Physical renewable energy infrastructure facilities, equity, and interest-bearing instruments issued by unlisted companies and other legal entities focused on renewable energy infrastructure.

e. Financial derivatives and fund units are inherently linked to the equity, fixed-income, or unlisted portfolios.

f. Other financial instruments acquired through corporate events.

Prohibited investments include securities, real estate, and infrastructure in Norway, certain government bonds under large-scale UN sanctions, certain countries without tax treaties or information sharing, and specific oil and offshore drilling-related equities.

Section 2-2: Benchmark Index for the Fixed-Income Portfolio

The fixed-income portfolio's benchmark index has fixed weights and includes government bonds (70%) and corporate bonds (30%).

The government bond portion consists of bonds issued by developed market countries included in specified Bloomberg indices.

The corporate bond portion comprises securities issued by companies in developed market countries and specified currencies.

Specific formulas determine the weights of countries in the government bond portion based on GDP and market capitalization weights.

Bonds denominated in Norwegian kroner, issued in Norway, or excluded under guidelines cannot be part of the benchmark index.

Section 2-3: Benchmark Index for the Equity Portfolio

The equity portfolio's benchmark index is based on the FTSE Global All Cap Index.

Adjustment factors are assigned based on country affiliation: European developed markets (excl. Norway) get 2.5, USA and Canada get 1, other developed markets and emerging markets get 1.5.

Each country's weight in the benchmark index is calculated based on market capitalization and adjusted for free float.

The benchmark index is adjusted for the Bank's tax position.

Securities excluded under guidelines are not included in the benchmark index, and those reinstated later will be included again.

Equities classified under specific oil and offshore drilling categories are not included in the benchmark index.

2.2.2 Factors Influencing Asset Allocation Decisions

Norges Bank's investment strategy aims for long-term growth and wealth preservation for future generations. The fund diversifies its investments across various asset classes, including international equities, bonds, real estate, and renewable energy infrastructure.

It has a global perspective, investing in markets, countries, and currencies worldwide to capture opportunities for growth and value creation.

Key points regarding this investment strategy and mandate include:

- ***Evolution and Oversight:*** *The strategy evolves over time through expert reviews and analysis, with significant changes requiring parliamentary approval. The goal is to leverage the fund's long-term horizon and substantial size to generate strong returns.*
- ***Mandate:*** *The Ministry of Finance issues a management mandate that specifies the markets, asset classes, and risk parameters within which the fund can invest. The mandate also defines the benchmark index against which the fund's performance is measured.*
- ***Benchmark Index:*** *The fund uses a benchmark index based on indices from FTSE Russell Group and Bloomberg Barclays, with fixed weightings of 70% in equities and 30% in bonds. However, the fund constructs a portfolio that deviates slightly from the benchmark to optimize returns, meet environmental and fiscal criteria, and ensure cost-effective management.*
- ***Equity and Bond Allocations:*** *The fixed-weight allocation between equities and bonds may drift due to differences in price movements. There is a limit for this drift, and rebalancing is required if it exceeds a set threshold.*
- ***Investment Strategies:*** *The fund employs a mix of strategies, including market exposure, security selection, and fund allocation, across equity, fixed income, and real asset management. These strategies aim to manage risk and optimize returns while considering the fund's unique characteristics.*
- ***Equity Investments:*** *The fund, as one of the world's largest shareholders, uses strategies for efficient market exposure and fundamental research. It seeks to avoid mechanical benchmark replication to reduce trading costs and actively manages its equity portfolio.*
- ***Fixed-Income Investments:*** *The fund invests in various government and corporate bonds and tailors its strategies to different bond segments. The goal is to build cost-effective portfolios while capitalizing on opportunities at different levels.*
- ***Real Assets Investments:*** *Real estate and renewable energy infrastructure investments aim to diversify the portfolio. The mandate sets guidelines for these investments, allowing a specific allocation to unlisted real estate and unlisted renewable energy infrastructure.*

- ***Portfolio Goals:*** *The fund aims to gradually build up its real estate and renewable energy portfolios while focusing on projects with stable cash flows and reduced risk.*

Overall, this investment strategy emphasizes diversification, active management, and responsible investing to achieve strong returns within the framework set by the Ministry of Finance.

Economic policy uncertainty (EPU) also plays a vital role in asset allocation. (Alok et al., 2022) examine the response of global institutional funds to EPU in both their home countries and investment destinations. On average, there is a negative relationship between EPU and capital flows for global funds. An increase in EPU in a target investment country is associated with a decrease in fund capital flows into that country. Funds tend not to withdraw capital from their home countries when EPU spikes at home. However, this is primarily observed in G7 economies. Funds domiciled in ex-G7 and developing countries reduce capital allocation in response to increased home EPU. The strongest capital withdrawal response to EPU spikes is observed in emerging markets, followed by ex-G7 advanced countries and then G7 countries.

The study also finds that funds' exposure to EPU shocks in other countries affects their investment decisions in a target country. This effect amplifies the response to the target country's EPU shock. Additionally, it explores the role of cultural, legal, and political factors in mitigating the impact of EPU on fund flows. It suggests that cultural, legal, and political proximity between the home country of a fund and the investee country can influence the fund's response to EPU.

2.2.3 Portfolio Risk Management

Norges Bank has to follow some risk limits defined in its investment management mandate:

Section 2-4: Risk limits for the investment portfolio

- *Equity portfolio should constitute 60-80% of the portfolio.*
- *Fixed-income portfolio should constitute 20-40% of the portfolio.*
- *Unlisted real estate portfolio can be up to 7% of the portfolio.*
- *Unlisted renewable energy infrastructure portfolio can be up to 2% of the portfolio.*

- *Derivatives should be included in the calculations based on net market value and underlying financial exposure.*
- *The standard deviation of excess return from the benchmark index should not exceed 1.25 percentage points.*
- *Diversification of systematic risk factors is sought for equity and fixed-income portfolios.*
- *Fiscal strength differences between countries should be considered for government bond investments.*
- *High-yield debt instruments (below investment grade) should not exceed 5% of the fixed-income portfolio.*
- *Credit ratings are required for debt instrument investments, and all internal credit ratings must be documented.*
- *Emerging market debt instruments can constitute up to 5% of the fixed-income portfolio.*
- *The Bank's holdings in a single company should not exceed 10% of voting shares, except for real estate and renewable energy companies.*
- *Leverage may be used for efficient management execution but not to increase exposure to risky assets.*
- *Cash collateral reinvestment should not increase exposure to risky assets.*
- *Short selling is allowed if the Bank has access to securities through borrowing arrangements.*

Section 2-5: Supplementary risk limits established by the Executive Board

- *Additional risk limits cover factors not adequately captured by expected tracking error, including overlap between equity and fixed-income portfolios, credit risk, liquidity risk, counterparty exposure, leveraging, cash collateral reinvestment, and securities borrowing.*
- *An operational risk limit is established.*
- *Unlisted real estate risk is limited based on investments in countries, sectors, emerging markets, real estate under development, and debt ratios for individual investments.*
- *Unlisted renewable energy infrastructure risk is limited based on investments in countries, emerging markets, projects under development, and debt ratios for individual investments.*

- *Limits are set for the scope of renewable energy infrastructure investments based on representation in unlisted companies and proportion in fund structures.*
- *Exceeding maximum amounts for unlisted real estate and infrastructure investments and investments in companies seeking listing requires approval from the Executive Board.*
- *A limit is set on the Fund's ownership of voting shares in a single listed real estate company.*
- *A limit is established for the portion of the portfolio managed by external managers.*
- *Limits are set for co-investing with any one investment partner in the real estate and infrastructure portfolios.*
- *All limits and changes are to be presented to the Ministry with a minimum three-week notice unless exceptional circumstances allow for a shorter notice period.*

2.2.4 Evaluating Portfolio Performance and Reporting

Norge Bank has to follow various principles and guidelines for measuring and managing risks in the investment portfolio. Here's a summary:

Section 3-2: Valuation and Return Measurement

- *The return on the investment portfolio is measured in the currency composition of the benchmark index and unlisted portfolios.*
- *Net return on unlisted portfolios is calculated after deducting all costs.*
- *The return calculations follow the Global Investment Performance Standards (GIPS) methodology.*
- *Principles for valuation and return measurement of financial instruments are established.*
- *The method used for determining the value of financial instruments should be verifiable and express the fair value as of the measurement date.*
- *External, independent valuation of unlisted real estate and infrastructure investments is commissioned annually.*

Section 3-3: Measurement and Management of Market Risk

- *Principles are established for measuring and managing market risk, including systematic risk factors.*
- *Multiple methods are used to estimate risk, and stress tests are conducted based on historical events and future development scenarios.*
- *Expected tracking error is calculated and approved by the Ministry.*
- *Extreme event risk analyses are integral to risk management.*

Section 3-4: Measurement and Management of Climate Risk

- *Principles are established for measuring and managing climate risk.*
- *Multiple methods are used to estimate risk, and stress tests are conducted based on future development scenarios, including a 1.5°C global warming scenario.*

Section 3-5: Measurement and Management of Credit Risk

- *Principles are established for measuring and managing credit risk.*

Section 3-6: Measurement and Management of Counterparty Exposure

- *Principles are established for measuring and managing counterparty exposure.*
- *The Bank has procedures for selecting and evaluating counterparties, setting exposure limits, credit rating requirements, collateral management, and netting arrangements.*

Section 3-7: Measurement and Management of Leveraging

- *Principles are established for measuring and managing investment portfolio leveraging, including implicit leveraging through derivatives and reinvestment of cash collateral.*

Section 3-8: Securities Borrowing, Securities Lending, and Short Selling

- *Guidelines are established for securities borrowing, securities lending, and short selling.*

Section 3-9: Reinvestment of Cash Collateral Received

- *Guidelines are established for reinvesting cash collateral received.*

Section 3-10: Measurement and Management of Operational Risk

- *Operational risk is defined and assessed in terms of probability and impact.*
- *Operational risk factors are monitored and managed.*

Section 3-11: Instrument Approval and Investment Due Diligence Review

- *The Executive Board must approve all financial instruments, markets, and issuing countries for government bond investments.*
- *Approval is based on the instrument's contribution to the management assignment's efficiency and effectiveness and comprehensive risk management.*
- *Due diligence reviews are conducted before investing in unlisted real estate, unlisted renewable energy infrastructure, and companies intending to seek listing.*
- *The reviews include assessments of various risks, such as market risk, liquidity risk, credit risk, etc.*

Section 3-12: Disposal of Unlisted Shares in Companies Where a Planned Listing Is Not Carried Out

- *Guidelines are established for disposing of shares in unlisted companies if a planned listing does not occur.*
- *The plan includes assessments of appropriate and cost-efficient disposal and ongoing evaluation of investment risks.*
- *These sections provide a comprehensive framework for managing various risks associated with the investment portfolio.*

Chapter 6 outlines the requirements for public reporting by the Bank. The reporting should provide a true, fair, and comprehensive overview of the Bank's execution of the management assignment. It must be published in Norwegian and address various aspects of the investment portfolio.

Section 6-1: Reporting Requirements

- *The Bank must regularly report on its investment strategies for managing the equity portfolio, fixed-income portfolio, unlisted real estate portfolio, and unlisted renewable energy infrastructure portfolio.*
- *Semi-annually, the Bank should report on developments in the investment portfolio's value, performance, and risk, as well as the composition of excess return and utilization of limits.*
- *Annually, additional reporting is required, covering performance in managing unlisted real estate, renewable energy infrastructure, and investment strategies. The Bank must also address risk-return relationships, management costs, climate risk efforts, stress test results, and other factors.*
- *Decisions made based on guidelines for observation and exclusion from the GPFG should also be disclosed.*

Section 6-2: Publication of Calculation Methods and Data

- *The methods and data used in public reporting, including risk measurement and reporting, should be described and made publicly available in a computer-readable format.*
- *If calculation methods change, the reasons for the changes and the effects must be disclosed, along with pro-forma figures based on the previous calculation methods for the next two years.*

Section 6-3: Publication of the Executive Board's Guidelines

- *The principles, guidelines, and limits adopted by the Executive Board as a result of this mandate should be made public.*
- *Overall, the public reporting aims to enhance transparency and provide stakeholders with relevant information about the Bank's investment strategies, performance, risk management, and responsible management efforts.*

2.2.5 The Difference between Theory and Practice

(Cochrane, 2021) highlights the gap between academic portfolio theory and actual investment practices. While academic theory, based on Merton's work, suggests that investors should build portfolios using factors and state variables, real-world investors, including institutions and endowments, tend to rely on more straightforward methods. They often categorize investments into buckets based on asset classes and industry sectors, allocate funds to different managers, and evaluate performance based on short-term returns relative to benchmarks.

Despite the theoretical foundation, Merton's approach has had little impact on portfolio practice. Institutional investors often fail to implement state-variable hedging, and their strategies are not aligned with academic theory. The study also mentions the lack of success in active management despite high fees, which contradicts academic findings.

Cochrane raises questions about the disconnect between theory and practice and suggests that both sides may have shortcomings. He proposes focusing on the stream of dividends and taking a general equilibrium perspective to align theory with practice better and address the challenges of implementing Merton's portfolio theory in the real world. Ultimately, he calls for reevaluating portfolio theories and approaches to bridge the gap between academia and investment practice.

2.3 AI and Portfolio Choice and Asset Allocation

2.3.1 Risk Assessment

AI can be used to select an optimal portfolio and asset allocation. In the simple implementation of portfolio theory, investors need estimates of returns and risk, as measured, for instance, by volatilities and correlation. Forecasting these numbers and adjusting them to changing market and economic conditions are special tasks that an AI is well-designed to do.

Forecasts can be made for long time horizons and be used for strategic allocation. They can be made for weekly, monthly, or yearly horizons for more tactical asset allocation. Some sectors, countries, and assets can be overweight or underweight, depending on the predictions. AI can help select which asset should belong to the portfolio and the costs and benefits of holding these assets.

AI is particularly useful when assets do not follow the simple distributions assumed by financial models, such as a log-normal or normal distribution, or if their distributions are changing. Collectively, the assets can also follow different patterns. Some assets will be more correlated, and the correlation could change with market conditions. For instance, it is well-known that the correlation between risky assets tends to increase in market downturns, probably because investors want to reduce their risks and sell these risky assets simultaneously.

2.3.2 Forecasting Returns

Forecasting returns is covered by the Asset Pricing chapter (Chapter 1). It can be based on factor models, on risk (riskier assets leading to higher expected returns), or valuation-based, for instance, on the projection of future earnings and cash flows.

There can be benefits to predicting factor returns. (Haddad et al., 2020) examine the predictability of asset returns in both the time series and the cross-section, focusing on factor-based investing and factor timing. They determine the value of constructing an optimal factor-timing portfolio that combines market timing and factor investing.

They propose an approach that focuses on the predictability of the largest principal components (PCs) of factors to overcome the challenge of measuring predictability across many returns accurately. They impose the condition that the implied stochastic discount factor (SDF), which prices assets, is not too volatile. The estimated SDF, when considering factor timing, exhibits significantly different properties compared to estimates assuming constant factor premia. It is more volatile, and the benefits of factor timing are time-varying. The predictability of factor returns is substantial, particularly for the dominant principal components. This supports the concept of factor timing and is more significant than predicting the aggregate market return.

The study uses dimension reduction techniques, specifically principal component analysis (PCA), to handle high-dimensional factor spaces, leading to robust predictability estimates. An optimal factor timing portfolio is constructed, demonstrating the investment benefits of timing factors, and the conditional variance of the SDF is found to be substantially larger than implied by static strategies alone. The SDF's variance exhibits fluctuations at shorter business-cycle frequencies and is correlated with various macroeconomic variables, providing insights into the dynamic properties of risk premia.

2.3.3 Portfolio Optimization

An investor does not necessarily have mean-variance preferences, and an AI can help her optimize difference preference objectives. An investor might care about the social impact of her portfolio in addition to the pure financial return. ESG considerations such as climate change (Chapter 9) might be critical to the investor. Using the same logic as recommender systems and leveraging characteristics of the investor (age, wealth, income, retirement objectives, risk tolerance, etc..), an AI could suggest assets to include in a portfolio.

An investor might also be very uncertain about the traditional portfolio models and wish for a more robust optimal portfolio. Model parameters might be highly uncertain, and an AI could guide her toward a better portfolio. Portfolio models that account for uncertainty aversion tend to be more complex, and AI could help in their practical implementation.

In the same spirit, Bayesian portfolio analysis and other more advanced portfolio optimization techniques might be less straightforward. An AI could elicit priors from investors to perform Bayesian updating and calculate portfolio weights consistent with their views. More advanced models than the Black-Litterman model could be implemented without the expertise of Bayesian statistical analysis.

Transaction costs, liquidity, and sometimes taxes are important practical considerations easily incorporated into traditional portfolio and asset allocation models without very simplified assumptions such as constant proportional transaction costs, especially over a long-time horizon. AI can help solve these problems with simulated data and techniques such as reinforcement learning.

Another area of application is the simulation of other portfolio allocations. Demand-based asset pricing has documented the importance of flows buying and selling pressure on assets. One's portfolio is affected by the portfolios of other investors. To the extent that the portfolios can be observed or inferred (by 13F filings, for instance), an AI can simulate or forecast the rebalancing of these portfolios under different scenarios or market conditions.

For large portfolios, (Ao et al., 2019) propose a new methodology for estimating the mean-variance efficient portfolio. They call their method the maximum-Sharpe-ratio estimated and sparse regression (MAXSER) method. The MAXSER method is a general approach that can be applied to various situations in which the number of assets in portfolios is not negligible compared with the sample size. Under mild assumptions, the MAXSER portfolio can asymptotically (1) achieve mean-variance efficiency and (2) satisfy the risk constraint. To the authors' knowledge, this is the first time that both objectives can be simultaneously achieved for large portfolio optimization.

The MAXSER method is based on two key insights. First, the mean-variance optimization problem can be equivalent to a regression problem. Second, the regression problem can be solved using sparse regression methods.

The MAXSER method has several advantages over traditional methods for large portfolio optimization. First, the MAXSER method is more efficient, as it does not require the estimation of a covariance matrix. Second, the MAXSER method is more robust to estimation errors, as it uses a sparse regression method. Third, the MAXSER method can be used to estimate the mean-variance efficient portfolio for many assets. The MAXSER method is effective in empirical studies. The authors show that the MAXSER portfolio outperforms traditional methods for large portfolio optimization in terms of both risk-adjusted returns and tracking errors.

(Zhang et al., 2020) propose a novel approach that utilizes deep learning models to optimize portfolios. Unlike traditional methods that predict expected returns before optimization, this approach directly optimizes the Sharpe ratio, which measures the return per unit of risk.

The authors use Exchange-Traded Funds (ETFs) based on market indices to construct portfolios, reducing the complexity of choosing individual assets. They select four market indices: the US total stock index, the US aggregate bond index, the US commodity index, and the Volatility Index. These indices offer liquidity and low expense ratios and are generally uncorrelated, enhancing diversification.

The methodology of the proposed framework involves using a neural network to model asset allocations in a long-only portfolio. The objective is to maximize the Sharpe ratio, and gradient ascent is used to optimize the model's parameters. The model architecture includes input layers, neural layers (utilizing LSTM), and an output layer with softmax activation to ensure portfolio weights are positive and sum to one.

(Ma et al., 2020) use deep neural networks (DNNs) for short-term financial market prediction and suggest combining these models with portfolio optimization techniques. They specifically focus on DNNs like deep multilayer perceptron (DMLP), long short-term memory (LSTM) neural networks, and convolutional neural networks (CNN) for predicting future stock returns and propose prediction-based portfolio optimization models.

These prediction-based models use DNNs to predict future stock returns and employ metrics like semi-absolute deviation (SAD) to measure the downside risk of portfolios. The goal is to create more efficient portfolio optimization models for short-term investments. The study also mentions experiments comparing these models with equal-weighted portfolios and models based on support vector regression (SVR).

The authors collect a dataset of historical stock prices and other financial data. They train a deep neural network to predict future stock prices and use the trained neural network to identify the most important factors for portfolio optimization and select optimal portfolios in real-time. They evaluate their approach on a real-world dataset of stock prices. They show that their approach outperforms traditional portfolio optimization methods in terms of both return and risk.

(Jensen et al., 2022) integrate machine learning with traditional portfolio optimization techniques. They argue that traditional portfolio optimization techniques are not well-suited for the modern financial markets, which are characterized by high levels of uncertainty and volatility. They say that machine learning can be used to improve the performance of traditional portfolio optimization techniques by providing more accurate forecasts of future returns and by allowing for more flexibility in the portfolio construction process.

They propose a new framework for portfolio optimization called the "implementable efficient frontier." The implementable efficient frontier is a set of portfolios that are efficient (i.e., they maximize expected return for a given level of risk) and implementable (i.e., they can be constructed without incurring excessive trading costs). They show that the implementable efficient frontier can be constructed by integrating machine learning with traditional portfolio optimization techniques. The implementable efficient frontier outperforms the traditional efficient frontier. The authors also show that the implementable efficient frontier is more robust to changes in the market environment.

(Chatigny et al., 2021) also use deep learning but with a focus on attention-guided deep learning. They identify the most relevant stock characteristics that can define a stochastic discount factor (SDF) or pricing kernel, which, in turn, helps construct an optimal mean-variance efficient portfolio. The attention networks in this deep learning architecture are designed to mimic human-like thinking, identifying fundamental and non-redundant characteristics that contribute to long-term portfolio performance.

They use a no-arbitrage assumption to train their neural networks and show that their approach performs well out-of-sample, achieving a high Sharpe ratio and alpha. They also assess the economic gains of the portfolio after accounting for trading costs, demonstrating that significant gains remain, especially when excluding certain stocks based on market capitalization.

2.3.4 Deep Reinforcement Learning

We review reinforcement learning in Appendix A. Deep reinforcement learning (DRL) has wide applications in portfolio selection because it can deal with an agent who wants to maximize an objective such as profit a portfolio by taking some actions such as asset allocation or security selection.

(Zhang et al., 2020) propose a DRL-based approach that can learn to select portfolios that maximize the expected return while minimizing the transaction costs and risk.

The environment is the financial market, modeled as a Markov decision process (MDP). The MDP has a set of states which represent the possible states of the financial market. The MDP also has a set of actions that represent the possible portfolios that can be selected. The MDP has a reward function, which represents the expected return of a portfolio. The agent is a deep neural network that learns to choose portfolios that maximize the expected return while minimizing the transaction costs and risk. The agent is trained using a reinforcement learning algorithm, such as Q-learning.

The agent is trained by interacting with the environment. The agent starts with a random policy and learns to select portfolios that lead to higher returns and lower transaction costs. The agent is trained until it converges to a policy that achieves a satisfactory level of performance. Once the agent is trained, it can be deployed to select portfolios in real-time. The agent can be used to choose portfolios for individual investors or institutional investors.

The authors evaluate their approach on a real-world dataset of stock prices. They show that their approach outperforms traditional portfolio selection methods in terms of both return and risk.

(Almahdi and Yang, 2017) use reinforcement learning to adjust the portfolio allocation dynamically. The system is designed to maximize the expected return of the portfolio while minimizing the risk. The system works by first learning a value function that maps from the current state of the market to the expected return of the portfolio. The value function is learned using a reinforcement learning algorithm called Q-learning (see Appendix A). Once the value function is learned, the system can use it to make trading decisions.

The system makes trading decisions by first identifying the best action to take in the current state of the market. The best action is the one that will lead to the highest expected return. Once the best action is identified, the system executes the trade.

The system is adaptive because it can learn to adjust its trading behavior in response to changes in the market. This is done by updating the value function as new data is observed. The updated value function will then be used to make better trading decisions in the future.

The system was tested on a historical dataset of stock prices. The results show that the system was able to outperform a buy-and-hold strategy. The system is also able to beat several other portfolio trading systems that were tested. The main limitation of the system is that it requires a large amount of data to train the value function. The authors are working on developing a system that can learn to trade with less data.

2.4 Case Studies

2.4.1 AlphaPortfolio: Direct Construction through Reinforcement Learning and Interpretable AI

(Cong et al., 2021) propose an investment strategy based on Reinforcement Learning. They trade US common stocks from Amex, Nasdaq, and NYSE, using data from July 1965 to June 2016 from CRSP (returns data) and Compustat (financial data). They use 51 variables as input features (state variables). The variables are fundamental financial data, changes to these data, and financial ratios such as cash flow to book value of equity or book value of equity to market value of equity.

They use sequence representation extraction models (SREM) to extract information from the state variables and then some cross-asset attention networks (CAANs) to capture interactions between the variables of different stocks. The SREM can take two forms, either a Transformer Encoder (TE) or a Long Short-Term Memory (LSTM) model.

TE works as follows:

An asset is represented by a series of 12 (monthly) vectors of 51 features. These vectors are converted to new representations through an SREM using a transformer or an LSTM model. They model the time series behaviors of these features within an asset.

CAANs work as follows:

The new representations are processed into a CANN that replaces a covariance matrix, and the representation vectors of the assets that come out of the SREM are concatenated into one single vector R. This vector is multiplied with matrices W_V, W_Q, and W_K to form the vectors V, Q, and K.

Figure 2.12. APN Network

The output of this neural network model (the AlphaPortfolio or AP model) is a set of winner scores probabilities of selection in the final portfolio. These scores are used to construct long-short portfolios, long (short) the stocks with the highest (lowest) scores.

The parameters of the AP model are then estimated thanks to Reinforcement Learning. The state variables are the input features, and the actions are the weights of the long-short portfolio. The objective of the portfolio manager is the sum of rewards earned in each time step (trading period). The model parameters (for each trading period) are found by doing a gradient descent on the average reward (across all trading periods) or Sharpe ratio (Return divided by Standard Deviation). Each episode is a sequence of trading periods (e.g., one year for twelve monthly periods).

The study compares AP's performance with alternative models, including a nonparametric (NP) model and a traditional two-step approach. The results show that AP achieves a high OOS Sharpe ratio, indicating good performance. AP outperforms most machine-learning-based strategies, demonstrating the effectiveness of reinforcement learning (RL) and artificial intelligence (AI) in investment.

2.4.2 Long Horizon Multi-Factor Investing with Reinforcement Learning

(Goyenko and Zhang, 2022) study long-term multi-factor investing using firm characteristics and macro factors. They account for volatility and liquidity risks as well as transaction costs. They use a reinforcement learning framework to make portfolio rebalancing decisions.

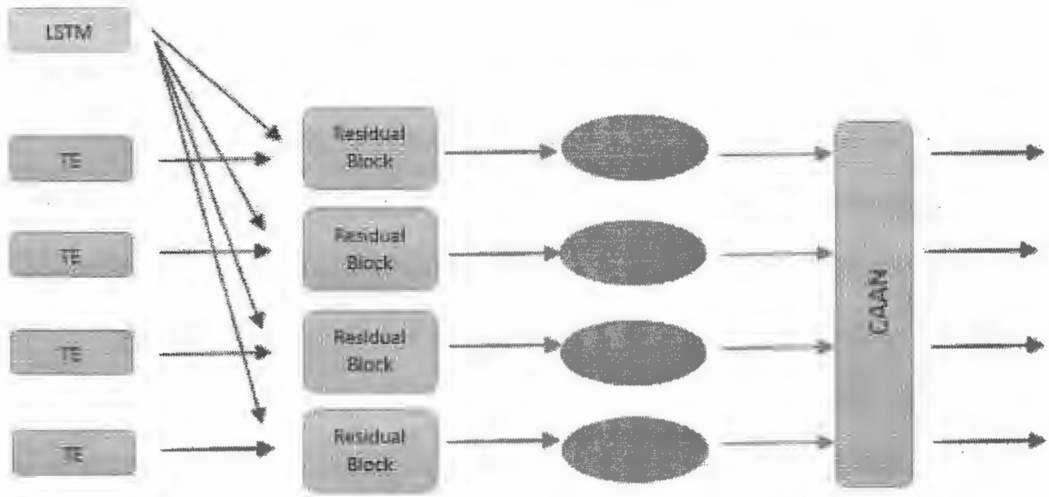

Figure 2.13. Model architecture

Their model uses stock characteristics and macro variables as inputs (Figure 2.13). The vector of 12 months of stock characteristics goes through Transformer Encoders (TE) (Figure 2.13), and the vector of 12 months of macro variables feeds into an LSTM model. For each stock, its characteristics are combined with the macro variables in a residual block and then concatenated to form a stock-level representation. The stock representations then go through a cross-asset attention network architecture (CAAN) (Figure 2.13) and generate portfolio weights as outputs instead of long-short positions as in (Cong et al., 2021).

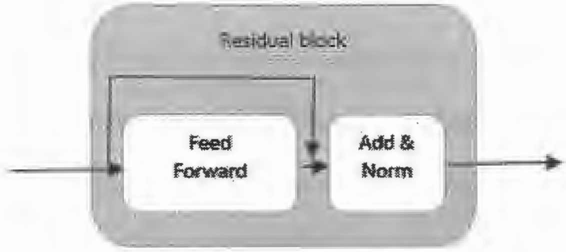

Figure 2.14. Residual block

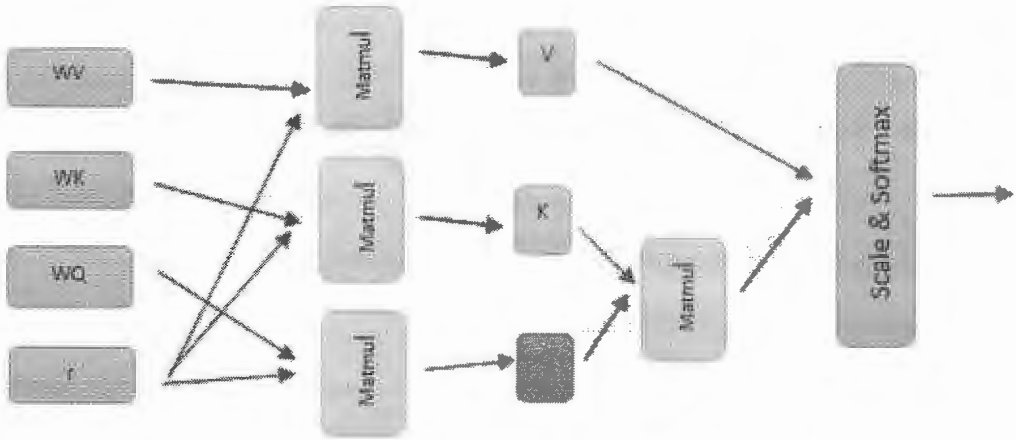

Figure 2.15. CAAN architecture

They consider nine market and factor portfolios: market (MKT), small-minus-big (SMB), high-minus-low (HML), robust-minus-weak (RMW), and conservative- minus-aggressive (CMA), the momentum (MOM) factor, the profitability (ROE) and investment (IA) factors, and the betting-against-beta (BAB). They estimate the transaction costs.

For each factor portfolio, they constructed 153 firm characteristics (Jensen et al., 2021). The variables range from firm age, liquidity of book assets, liquidity of market assets, Amihud measure, book leverage, asset growth, assets-to-market, capital turnover, change in common equity, book-to-market equity, market beta to price momentum, long-term reversal, maximum daily return, total skewness, return volatility, assets turnover, labor force efficiency, sales growth, sales-to-market, revenue surprise, lagged returns, change in short-term investments, and total accruals.

They also use 12 macro variables described in (Welch and Goyal, 2008):

- dividend-price ratio (dp)
- dividend yield (dy)
- Earning price ratio (ep)
- stock variance (svar)
- book-to-market ratio (bm)
- net equity expansion (ntis)

- Treasury-bill rate (tbl)
- long term rate of returns (ltr)
- term spread (tms)
- default spread (dfy)
- default return spread (dfr)
- Consumer Price Index (infl)

They train their model by reinforcement learning. The actions are the portfolio weights. The state variables are the stock and macro variables. The reward is a mean-variance (quadratic) utility net of transaction costs of a risk-averse investor over a holding period T.

They train their model with data from January 1980 to December 2004. They test their model from January 2005 to December 2020. After randomly picking a month, they build the macro variables and stock characteristics and use the past 12-month data to forecast the 13^{th}-month portfolio weights. Then, roll their model forward by one month and repeat the process eleven times. They also look at the quarter, semi-annual, and annual rebalancing. The model weights are re-estimated only annually.

They find that, out-of-sample, the RL strategy delivers a Sharpe ratio of 2.344 for monthly rebalancing better than the conditional mean-variance multifactor portfolio with rebalancing that depends on volatility. The return is 72bps per month, and the alphas are 53bps and 51bps over 6 and 9p-factor benchmark portfolios. They obtain higher Shape ratios for semi-annual (3.051) and annual (2.702) rebalancing. The strategies do not produce extreme weights, leverage, or portfolio turnovers.

(Zhang and Aaraba, 2022) employ the same methodology for bottom-up portfolio construction and find that the same strategies outperform the market out-of-sample with a higher Sharpe ratio of 2.872 from January 2005 to December 2020. The turnover is also low at 14% per month.

2.4.3 A Universal End-to-End Approach to Portfolio Optimization via Deep Learning

(Zhang et al., 2021) develop an end-to-end (E2E) framework to solve the portfolio optimization problem using deep learning. They consider a mean-variance portfolio optimization (MVP), the global minimum variance portfolio (GMVP), and the maximum Sharpe ratio portfolio (MSRP).

The objective of the MVP problem is:

$$max_{w_t} E(r_{p,t+1}) - \frac{\lambda}{2} Var(r_{p,t+1})$$

subject to $\sum_{i=1}^{N} |w_{i,t}| = 1$ if short-selling is allowed

And with $r_{p,t+1} = \sum_{i=1}^{N} w_{i,t} . r_{i,t+1}$

w_t is the vector of portfolio weights, r_t is the vector of returns.

The objective of the GMVP is the minimization of the portfolio volatility or variance:

$$min_{w_t} Var(r_{p,t+1})$$

The objective of MSRP is the maximization of risk-adjusted return or the return divided by the variance (though it is more common to divide by the volatility):

$$max_{w_t} E(r_{p,t+1})/Var(r_{p,t+1})$$

Their model has two connected neural networks. The first one uses as inputs the lagged returns $R_t = (r_{t-L}, .., r_t)$ and produces fitness (investment) signals $s_t = (s_{1,t}, ..., s_{N,t})$. The second one transforms these signals into investment portfolio weights $w_t = (w_{1,t}, ..., w_{N,t})$. The first network h_1 is the Score block. The second network h_2 is the Portfolio block.

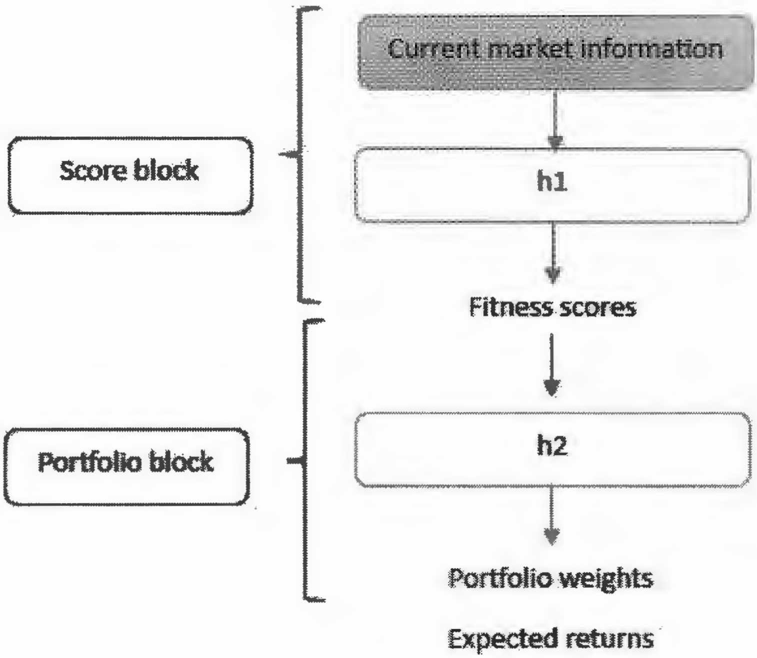

Figure 2.16. Neural network model

The objective is the minimization of a loss function; in this case, it will be the negative of the mean-variance objective function in the case of MVP, the negative of the Sharpe ratio in the case of MSRP, or the portfolio variance in the case of the GMVP.

Portfolio constraints are commonly imposed because of investment mandates or risk management. Positions can be restricted to be long-only (no short-selling is allowed), to have a maximum limit (for diversification or maximum concentration limit), to have a maximum leverage, or to be limited in number (cardinality constraint).

Deep learning techniques can be used if the loss function is differentiable. Fortunately, many of these constraints can be expressed as differentiable transformations of the signals. For instance, for positive weights that add to one, a softmax function can be used:

$$w_{i,t} = \frac{exp(s_{i,t})}{\sum_{j=1}^{N} exp(s_{j,t})}$$

To make the weights smaller than a constant u, the authors suggest a generalized sigmoid function $\phi_a(x) = a + 1/(1 + exp(-x))$ that is increasing in x and bounded by a and $a + 1$ function. If $a = (1 - u)/(Nu - 1)$ the weights $w_{i,t} = \frac{\phi_a(s_{i,t})}{\sum_{j=1}^{N} \phi_a(s_{j,t})}$ are always lower than u.

For the cardinality constraint, they propose to rank the weights and take a fixed number of weights on the long and short side so their number adds up to the cardinality limit. Using a softmax operation, it is possible to make this step differentiable.

To test their model, the authors use simulated and real daily data of 735 US stocks data from the Russell 3000 index from March 1984 to July 2021. The training set is between 1984 and 1999, the validation set is in 2000, and testing is done on 2001 data. The training set is expanded annually, and the testing is done from 2001 to 2021.

To benchmark their end-to-end model, they compare with the classical setup (CS) of estimating the expected returns, the covariance matrix, and then solving for the portfolio weights. Forecasting returns are done with sample data (SAMPLE), a linear model (LM), a multilayer perceptron (MLP, one neural network with 64 units), a long short-term memory model (LSTM with 64 units), and a convolutional neural network (CNN, 3 layers of one-dimensional convolutional layer with filter size 32, 64, and 128, and kernel size (3,1), and a final LSTM layer with 64 units).

As performance metrics, they report $E(R)$, the annualized expected return, $Std(R)$, the annualized standard deviation return, $Sharpe$, the annualized Sharpe ratio, $DD(R)$, the annualized downside deviation (like the standard deviation but using only negative returns), $Sortino$, the annualized Sortino ratio (expected annualized return divided by the downside deviation), MDD, the maximum drawdown (loss from the maximum gain), $\%\ of\ +Ret$, the percentage of positive returns, $Frobenius\ norm$, the distance between estimated and optimal portfolio weights (for the simulation case only), $Turnover$, the portfolio turnover rate, and $Beta$, the beta against the S&P 500 index.

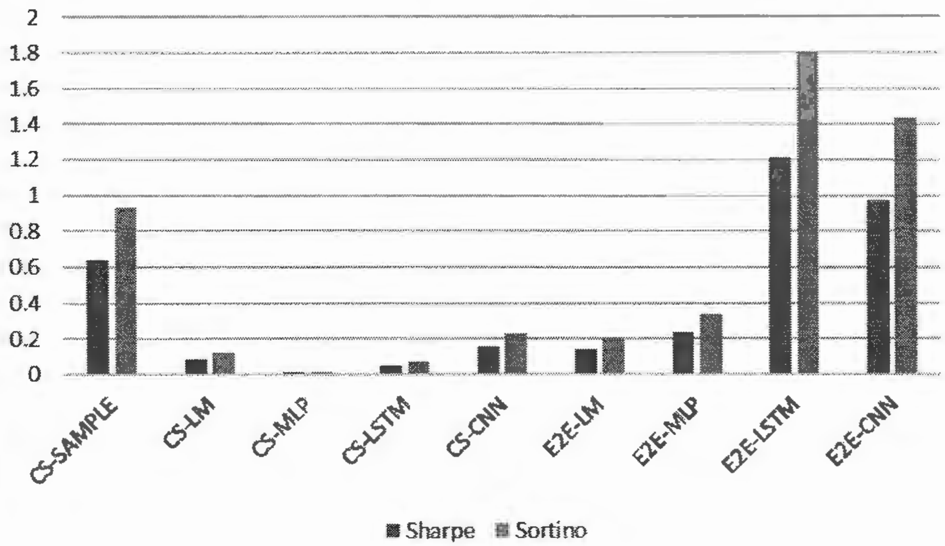

Figure 2.17. Sharpe and Sortino ratios by models

They find that the model E2E-LSTM (end-to-end with LSTM to predict returns) delivers the highest Sharpe ratio and Sortino index and outperforms all the baseline models when trying to solve the MSRP problem (maximum Sharpe ratio). Results are similar to the simulated data (distribution calibrated on real data).