COMILLAS
UNIVERSIDAD PONTIFICIA

ICADE – Facultad de Ciencias Económicas y Empresariales

BAYESIAN AND ROBUST OPTIMISATION FOR EFFICIENT PORTFOLIO SELECTION

Clave: 201803276

MADRID | May 2023

INDEX

1. Introduction

Over seventy years ago, Nobel laureate Markowitz (1952) published his pioneering work on portfolio optimisation, giving rise to Modern Portfolio Theory (MPT). The predicted return and investment risk of a portfolio could then be quantified thanks to this idea, which completely changed the field of portfolio management. The biggest innovation was the change in emphasis from the risk of individual assets to the overall risk of a portfolio. By combining risky assets, Markowitz showed that it was possible to produce a portfolio with an expected return that was similar to that of its constituent parts but with significantly less risk. In other words, he introduced the idea that a portfolio can be built with a risk that is lower than the total of all of its parts.

The ultimate purpose of portfolio management theory is to optimally allocate investments among the range of available assets. By considering the trade-off between a portfolio's risk and return, mean-variance optimisation is a quantitative approach that makes this allocation easier. When deciding how to distribute their wealth, investors can consider their preferences for risk and projected return thanks to the so-called mean-variance portfolio optimization. The method—which gave rise to the fundamental ideas of modern portfolio theory—involves choosing portfolios that, while adhering to a specified level of risk, maximize the portfolio's expected return or, alternatively, that, while adhering to a specified level of expected return, minimise variance.

In the financial industry, Markowitz mean-variance analysis is a crucial tool for both practitioners and scholars. Markowitz mean-variance (MV) optimization has been the de facto norm for efficient portfolio construction for the past few decades. The Markowitz approach is the foundation of almost all commercial portfolio optimizers for asset allocation and equity portfolio management. Because it is simple to execute and directly presents a tool for measuring risk and return, the MV model has gained wide acceptance among academics and practitioners. However, it has also received criticism from sceptics in both academia and financial institutions. Although MPT remains as the main theoretical framework for portfolio construction in the modern era, its adoption by investment professionals is less common than one might expect.

Practitioners' main criticism of MPT is that the *optimal* portfolios generated by the mean-variance approach are usually illogical, counterintuitive, and very sensitive to input parameters. The model does not account for variables like the investor's subjective understanding or particular market behaviour. As a result, the investor obtains portfolios that place too much emphasis on overvalued assets and less on cheap assets (Britten-Jones, 1999). Moreover, asset weights are particularly sensitive to estimation error, with small changes in the values of return and risk estimators leading to significant changes in portfolio values (Black and Litterman, 1992).

Many academics have examined this mean-variance model issue, analysing the portfolios' sensitivity to changes in expected return estimations (Best and Grauer, 1991; Broadie, 1993). The term "error maximisation" has been used to describe portfolio optimisation because of the detrimental consequences of estimation errors on ideal portfolios (Michaud, 1989). According to Michaud, the assets that are overweighted by mean-variance optimisation are precisely those that have a high ratio of expected return to estimated variance and may have significant estimation errors. The majority of the estimation bias in optimum portfolios is believed to be caused by errors in expected return estimations rather than errors in risk estimates. As shown by Chopra and Ziemba (1993), even small changes to estimates of expected returns or risk can lead to optimal portfolios with noticeably different mean variances. As a result, recent developments in the portfolio optimization problem have highlighted the need for upgrading the mean estimation technique.

Alternative risk estimators that perform better when employed within the framework of the Mean Variance model have been explored in an effort to produce better and more stable mean-variance optimal portfolios in order to address the acknowledged shortcomings of the Markowitz model. Since the 1990s, a large number of risk indicators have been put forth, many of which are computationally appealing since they lead to linear programming (LP) issues for discrete random variables. The mean absolute deviation (MAD) model, introduced in 1991 by Konno and Yamazaki, received a great deal of interest about thirty years ago, accelerating the creation of other LP models and giving rise to a great deal of research. The introduction of CVaR models by Rockafellar and Uryasev in 2000 significantly influenced new developments in risk measurement in finance throughout the first decade of the twenty-first century. In order to choose the optimum portfolio, portfolio management also involves optimisation strategies. Over time, portfolio optimization techniques have also advanced from simpler approaches based on heuristics and metaheuristics, such as mean-variance (MV) by Markowitz (1952), mean-absolute deviation (MAD) by Konno and Yamazaki (1991), value-at-risk (VaR) by Jorion (1996), conditional VaR by Rockafellar and Uryasev (2000), and Minimax (MM) by Park et al. (1998). In recent years, techniques based on Machine Learning (ML) by Alpaydin (2014) have gained popularity and have been applied to the field of portfolio optimisation. Bayesian, support vector machine-based, robust, neural network and quantum approaches have been proposed, among others, by Bai et al. (2020), Lee and Yoo (2020), Choi (2018) and Goel et al. (2020) respectively.

This paper provides an in-depth look into the behaviour and limitations of the traditional Markowitz model as well as its classical alternatives. We then present two of the most novel optimisation models that are currently available, Bayesian and Robust optimisation, and apply them to the field of portfolio optimisation. The ultimate goal of this research is to demonstrate how the new techniques could help to overcome the traditional problems of optimal portfolio selection.

2. Methodology

As indicated above, the aim of this research is to analyse in depth the Modern Portfolio Theory, whose cornerstone is the Markowitz Mean-Variance (MV) Model, and to discuss to what extent the application of two state-of-the-art Machine Learning approaches can overcome the limitations that the model has presented so far. To this end, a review of the existing literature on the classical framework, as well as the application of Bayesian Optimisation and Robust Optimisation to these traditional models, is carried out beforehand. Section three provides an overview of the fundamentals of portfolio management and section four analyses the classical theory of portfolio management, explaining the traditional Markowitz model. This section also includes a discussion of the main limitations that the classical theory has encountered over the years, as well as a review of the most widespread alternative models developed to address these shortcomings. The section concludes with a discussion of the relevant market factors that are not covered by the traditional framework. Section five focuses on the study of modern portfolio optimisation, exposing Robust and Bayesian optimisation. Section six shows our experimental implementation of the classical Markowitz model as well as BO and RO to the selection of optimal portfolios for five stocks of the S&P 500 and in section seven we discuss the results obtained and compare the three methods. Finally, in section eight the advantages that our two techniques offer to the classical model are reviewed and suggestions for future areas of research are presented.

3. Portfolio management

From a financial perspective, a portfolio is a pool of individual stocks or investments. By definition, a portfolio is a collection of financial investments in the form of securities, bonds, cash, assets, and other types of commodities that are meant to provide a future return of some kind. At a given time *t*, a portfolio can be depicted as

$$P(t) = \sum_{j=1}^{n} w_j(t) s_j \qquad (1)$$

Where $s_1, \dots, s_j$ describe the securities that compose our portfolio and $w_1, \dots, w_j$, the share of the total amount of money invested in our portfolio, assigned to each of the assets at a time *t*. Markets are dynamic and prices and risks vary constantly, so it is likely that we will need to expand existing investment or withdraw existing stocks. The choice involves the allocation or reallocation of own resources and requires dynamic timing, multiple goals and objective consideration. Successful portfolio management must take into account, among many other things, future events and opportunities, dynamic decision-making capacity and the limitations of allocated resources.

In their daily lives, entities and individuals face choices among two or more alternatives. The idea of utility functions is used in the economic choice theory to show how agents choose among possibilities. All of the options that are offered to the subject are given a numerical value via a utility function. The utility gained from a certain decision increases as its value rises. In terms of the limits the entity must work within, the option chosen offers the most utility. Similar choices exist for institutions when it comes to portfolio management strategies. For each portfolio, the expected return and risk are different.

Generally speaking, risk rises as expected return rises. In cases where the return is favourable but the risk is not, entities must choose which portfolio to invest in. Institutions benefit in variable degrees from different risk-return combinations. The preferences of the institutions' utility function over the combinations of expected return and perceived risk represent the value obtained from any conceivable risk-return combination. Maximizing the value of the portfolio and reducing the risk to which he or she is exposed are typically the two goals of a portfolio manager. According to Markowitz (1952), diversification of investments is a key technique to lower investment risk because of shifting market dynamics.

As explained by Kim et al. (2018), portfolios can be categorized as defensive, income, speculative, hybrid, or aggressive based on how much risk they are willing to take on. In order to get higher returns, the aggressive portfolio usually takes on bigger risks. On the other side, the conservative portfolio seeks out minimum risk. An income portfolio prioritizes dividend income or other recurring income, much like a defensive portfolio does. The speculative portfolio, which is sometimes equated to gambling, is designed to take very high risks. With the best amount of risk, the hybrid portfolio aims to provide the best return. It uses a combination of several asset kinds depending on the risk and reward levels. As a result, utility functions differ depending on each investor's profile.

4. Classical approach to portfolio optimization

Since the birth of Markowitz's Mean Variance in the mid twentieth century, numerous authors have proposed alternative models that fall under the umbrella of Modern Portfolio Theory. The conventional model measures risk by variance and optimizes the trade-off between risk and expected return to produce the best possible portfolios. However, the performance of the Markowitz model has been widely challenged due to its sensitivity to input estimation errors. Meanwhile, several authors have proposed alternative approaches to obtain more stable and robust portfolios against the errors generated in the estimation of the risk and return of the assets considered.

4.1. The Mean Variance Model

The original mean-variance model proposed by Markowitz (1952) is a unique objective model whose purpose is to minimize risk for the desired level of return or vice versa. It takes into account a portfolio of m assets with a covariance matrix $C = (\sigma_{ij})$ and uncertain future returns $r' = (r_1, \dots, r_n)$ with expected values $\mu' = (\mu_1, \dots, \mu_n)$. The investor must choose a portfolio depicted by $X' = (X_1, \dots, X_n)$. The anticipated value E and variance V of the portfolio's return, $R = r'X$, can be respectively:

$$E = \mu' X \qquad V = X' C X \qquad (2\ a, b)$$

The portfolio should be chosen in compliance with the following restriction:

$$AX = b, \qquad X \geq 0, \qquad (2\ c, d)$$

Where A, the matrix designating the weight for each of the possible holdings, has a dimension of *m* x *n* and *b*, the resulting pool of assets whose chosen weight is higher than 0, is *m* x 1. Therefore, non-negative X_i should be chosen subject to $m \geq 1$ linear inequalities, avoiding asset short selling. An EV combination is possible if it falls under the *E* and *V* of a feasible portfolio, whereas a portfolio is feasible if it satisfies *(2c)* and *(2d).* To put it another way, a viable EV combination is any one that is created by the combination of all of those feasible portfolios. A feasible portfolio is any portfolio that an investor can establish given the assets that are accessible. Given a set of assets, the feasible set of portfolios is a collection of all feasible portfolios that graphically depicts the risk and expected return combinations that can be obtained by constructing portfolios from all feasible combinations of those assets.

A viable EV combination (E_0, V_0) is inefficient if another feasible EV combination (E_1, V_1) already exists, so that either:

$$(i) \quad E_1 > E_0 \quad and \quad V_1 \leq V_0 \qquad (2\ e, f)$$

$$or$$

$$(ii) \quad V_1 < V_0 \quad and \quad E_1 \geq E_0 \qquad (2\ g, h)$$

Among all possible portfolios, an efficient portfolio is the one that offers the highest anticipated return for a particular amount of risk. For each level of risk, there is thus an efficient portfolio. The collection of all efficient portfolios is known as the efficient set. The set of efficient portfolios is sometimes referred to as the efficient frontier because all efficient portfolios, as depicted graphically in image 1, stand on the edge of the set of possible portfolios that have the best return for a specific level of risk. The efficient frontier's risk-return combinations overwhelm those below it, and there is no way to obtain a risk-return combination above it.

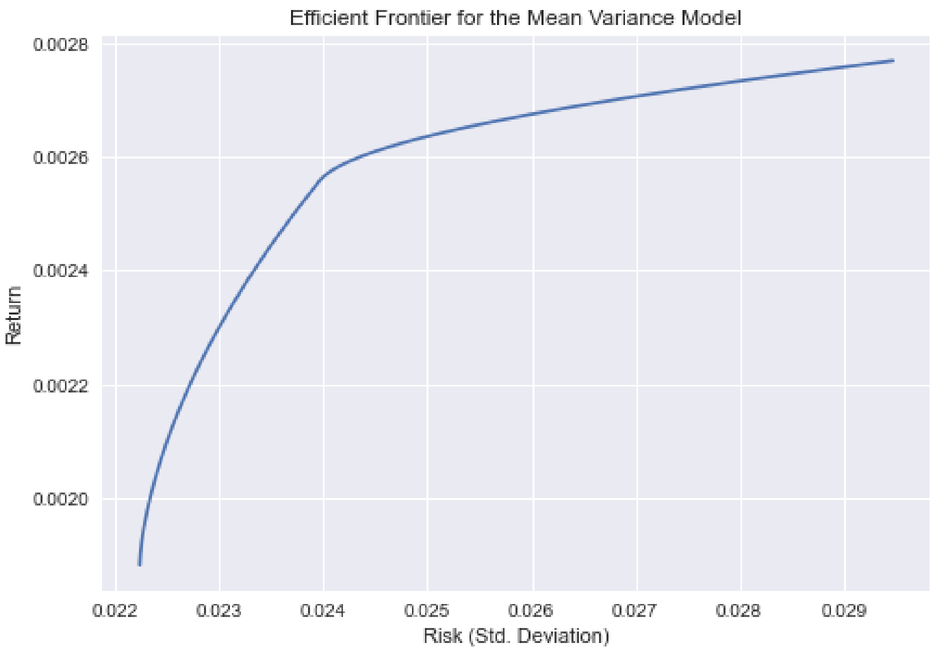

Picture 1: Efficient frontier for the Mean Variance Model applied to a portfolio consisting of Apple, Meta, Tesla, Amazon and Google stocks in the year 2021.

According to the classical MV model proposed by Markowitz (1952), in order to find an efficient portfolio, it is necessary to preselect the level of risk that the investor can bear, or the desired return of the portfolio. This, in fact, may not be entirely possible in real world cases. Thus, to find the efficient portfolio among various combinations of assets in the solution space, instead of considering a single objective, investors must consider all objectives at once. Therefore, the researchers transformed the single-objective model into a multi-objective model. According to Zitzler (1999), the multi-objective mathematical model can be rewritten as follows:

$$\min f(x) = (f_1(x), f_2(x) \dots\ f_p(x))$$

$$\text{subject to:} \quad e(x) = \big(e_1(x), e_2(x) \ldots e_m(x)\big) \leq 0 \qquad (3)$$

$$\text{where:} \quad x = (x_1, x_2, \ldots, x_n) \in X$$

Where the constraints $e(x) \leq 0$ determine the set of feasible solutions, $x = (x_1, x_2, \ldots, x_n)$ is the vector of decision variables (parameters) and X is the decision space, p is equal to 1 in single-objective models and p is greater than one in multi-objective models.

Because tackling MV problem optimization using a single-objective problem-solving structure presents several difficulties due to real-world conditions, multi-objective approaches have been applied. When solving a single-objective maximisation model, it is possible to find a single global optimal solution. As an alternative, the multi-objective approach offers a collection of Pareto optimum solutions, which are the best possible options. The single objective model can be defined in two different ways, as risk minimisation or profit maximisation. Multi-objective models, on the other hand, are based on the idea of simultaneously optimising conflicting objectives, resulting in more authentic depiction of the real market, but computationally more complex.

4.2. Limitations of the Mean Variance Model

The classic MV optimization states that investors favour a portfolio of securities that gives the highest predicted return for a specific level of risk, as measured by the return-variance relationship. The MV optimization approach determines the number of shares of equity portfolio wealth to allocate to each stock based on the estimated means, standard deviations, and return correlations of *N* stocks. The set of specified portfolio weights that is produced describes ideal solutions. The efficient MV frontier is the collection of ideal portfolios for all conceivable degrees of portfolio risk.

The MV model offers significant benefits. Among them, arguably the most substantial one is its application as a means of controlling the portfolio's exposure to various components of risk. It is however well known that Markowitz's optimisation does not work well in practice. This is known as the *Markowitz optimisation conundrum* (Michaud, 1989). According to Zhang et al. (2018), many of the myriad studies on Markowitz optimisation suggest reasons for this conundrum, such as skewness, kurtosis, time series effects, estimation error and estimation risk, and inadequate optimization problems, all of which undeniably have an effect. Had we known accurately the mean and covariance matrix of a set of asset returns, Markowitz optimization would correctly select an optimal portfolio. But we never know them precisely. Hence, we are forced to introduce estimates (Kan and Zhou, 2007). Noise encompasses uncertainty in the real values of the variables on which we want to optimize, such as the mean and variance. And it is complicated by the fact that real assets have time series effects and are not normally distributed.

The traditional MV approach frequently produces asset allocations and portfolios that are erroneously optimum or financially inefficient. The term *error-maximization* is used by Michaud (1989) to describe the influence of mean error. Due to the fact that MV optimisers are, at their core, estimation error maximisers, many optimized portfolios are often counterintuitive. Estimation error will always be included in risk and return estimates. Stocks with low expected returns, high correlations and small variances are underweighted by Markowitz optimisation. Conversely, those with high expected returns, lower correlations and big variances are largely overweighted. Of course, the likelihood of these overweighted stocks having significant estimation mistakes is highest. Best and Grauer (1991) look at how sensitive optimal portfolios are to changes in expected return estimates. The anticipated efficient frontier overestimates the expected returns of portfolios with varying degrees of estimation adequacy, as demonstrated by Broadie (1993).

Errors in means are almost ten times more significant than errors in variances and covariances combined, according to Kallberg and Ziemba (1984). It has also been demonstrated that equal weighting occasionally performs better than MV optimization. Investors who divide their wealth among individual stocks using the MV framework occasionally set all anticipated returns to zero. Due to the fact that it is often very difficult to generate accurate estimates of expected returns, this can result in a better portfolio allocation. The performance of the VM can be significantly diminished if projections are used that do not accurately reflect the relative expected returns of particular securities. By addressing some of the drawbacks of traditional MV optimizers, novel methods increase the effectiveness of portfolio optimization as an investment.

4.3. Alternative measures of risk

Following the birth of portfolio management theory, a large number of authors have proposed different units of measurement to quantify the risk of a portfolio. As already mentioned, Markowitz used the well-known statistical measures of variance and standard deviation to quantify the idea of risk. The majority of probability density function results—roughly 95% of them—fall within the range of two standard deviations above and below the mean, making the first hypothesis the most logical one. Since the variance is derived using the square of the standard deviation, it is easy to ascertain the variance of a portfolio before calculating its square root to obtain the standard deviation. However, due to the high sensitivity of this model to return, variance and covariance estimation errors, many authors have proposed alternative models for selecting optimal portfolios.

The mean-absolute deviation model (MAD) has been proposed by Konno and Yamazaki (1991) and Konno and Koshizuka (2005) for large-scale, highly diversified portfolio selection challenges. Mathematically, this model can be represented as follows:

$$minimize\ w(x) = E \mid \sum_{j=1}^{n} R_j x_j - E \sum_{j=1}^{n} R_j x_j \mid$$

$$subject\ to \sum_{j=1}^{n} E \mid R_j \mid x_j \geq \rho M_0,$$

$$\sum_{j=1}^{n} x_j = M_0, \tag{4}$$

$$0 \leq x_j \leq u_j, \quad j = 1, \dots, n,$$

Where R_j is the variable depicting the rate of return per period of asset S_j and x_j refers to the share of the total fund M_0 to be invested in the same asset. E is a proxy for the expected value of the variable held within the brackets whereas ρ refers to the minimal rate of return set by the investor. Finally, the maximum share of the funds that can be invested on a certain asset is depicted by u_j. Essentially, the goal of the MAD model consists of minimizing the expected spread between the real return of our portfolio and the forecasted value

Portfolio selection with shortfall limitations was created in 1952 as a result of Roy's safety theory research. In order to restrict the likelihood that the portfolio value will go below a given threshold of catastrophe, this author established the shortfall constraint for the first time. The objective of modeling portfolio selection using shortfall risk criteria is to maximize the probability that the portfolio return will surpass a predetermined minimum threshold level. Maximizing the probability that the portfolio return will exceed a predetermined minimum threshold level is the goal of modeling portfolio selection using shortfall risk metrics. Value-at-Risk, also known as VaR, was first presented by Jorion in 1996 and is likely the most well-known downside risk metric. The VaR of a portfolio can be modelled mathematically as

$$VaR(M) = \mathrm{E}(M) - Q_{\alpha}(M) \tag{5a}$$

Where M refers to the return of the portfolio, E(M), to the expected return of the portfolio and $Q_{\alpha}(M)$ is the α- quantile of return:

$$Q_{\alpha}(M) = \inf\{\upsilon : F(\upsilon) > \alpha\} \tag{5b}$$

This formulation depicts the largest return underperformance regarding the expected return of the portfolio with a confidence level of $1 - \alpha$. Despite VaR is a very effective tool for measuring risk in the financial industry, it has several drawbacks. The VaR optimisation problem is non-convex and not sub additive, thus it does not account for the benefits of diversification. Sub-additivity is one of the key desired properties that define a coherent risk measure for downside risk measures, and for dispersion measures (Serraino & Uryasev, 2013). One of the most well-known coherent risk metrics is the conditional value-at-risk (CVaR), which is the expected value of losses exceeding VaR. It was developed by Rockafellar 13 and Uryasev in 2000. The LP issue that corresponds to the CVaR portfolio optimization model is as follows:

$$\min_{\xi, z_j, x_i} \xi + \sum_{j=1}^{J} \upsilon_j z_j$$

$$subject\ to\ z_j + \sum_{i=1}^{n} r_{ij} x_i + \xi \geq 0 \quad for\ j = 1, \dots, J \tag{6}$$

$$\sum_{i=1}^{n} x_i = 1, \qquad x_i \geq 0 \; for \; i = 1, \dots, n, \qquad z_j \geq 0 \; for \; j = 1, \dots, J$$

where coefficients r_{ij} stand for the return reimbursements, under scenario j *(j = 1, 2,...,* J), for security *i (i = 1, 2,..., n).* The unbounded variable in the portfolio is ξ, and the variables x_i stand for shares of various assets. Coefficients υ_j are expressed as amounts $p_j/(1-\alpha)$, where p_j represents the likelihood of scenario *j* and α represents the level of confidence (CVaR parameter).

A few years later another model based on shortfall constraints was developed; the Minimax model (MM), proposed by Cai et al. (2004), which employs the minimal return as a risk indicator. In situations when asset returns are multivariate and regularly distributed, this strategy, which is based on the work of Park et al. (1998), achieves the same outcome as the Markowitz MV, but it offers certain advantages when returns are not normally distributed, which is the most likely scenario under real worl conditions. Due to its linear programming property, Minimax is quick and can handle more complicated models and restrictions. The model cannot be utilized when there is a lack of historical data, which is a pretty typical occurrence in the actual world, because it is highly sensitive to outliers. The model can be represented mathematically as follows:

$$\begin{aligned} &\max M_p \; subject \; to \sum_{j=1}^{N} w_j y_j - M_p \geq 0, t = 1,2, \dots, T, \\ &\textstyle\sum_{j=1}^{N} w_j \overline{y_j} \geq G, \\ &\sum_{j=1}^{N} w_j \leq W, \\ &w_j \geq 0, j = 1,2, \dots, N \end{aligned} \tag{7}$$

Where w_j is the allocation to asset j, y_j refers to the return of security *j* during the period of time *t* and $\overline{y_j}$, the average return of the same security. *W* accounts for the total allocation, G is the minimum level of return of the portfolio and M_p, the minimum return on portfolio.

The Lower Partial Model (LPM) was created more subsequently, in 1992, when Nawrocki advocated using a set of moments to calculate downside risk in a portfolio. Based on a continuous study, Brogan and Stidham (2008) claim that there could be several types of moments. Therefore, as part of Nawrocki's (1992) LPM, multiple N types of lower moments are analysed. The model can be defined as

$$LPM_{\propto}(\tau, Ri) = \int_{-\infty}^{\tau} (\tau - R)^{\alpha} \partial F(R) \quad (8)$$

Where τ is the desired return, R is the actual return, α, the degree of LPM and $\partial F(R)$, the cumulative distribution function of the return of the security.

The Mean Variance model is still the one that scholars and practitioners most frequently employ, despite the enormous amount of study that has been done in the area of Modern Portfolio Theory since the introduction of Markowitz's work. While its limitations are widely known, none of the alternative models proposed overcome such problems. Earlier, we identified estimation errors in the model's inputs as the trigger for optimisation errors. Out of the three inputs that the Markowitz model requires, it is the errors in the expected return on assets that account for most of the error. The models discussed so far have proposed alternative measures of risk quantification, but each and every one of them requires return estimation as the fundamental measure for portfolio optimisation. Consequently, classical approaches to portfolio optimisation contain a common limitation: their high sensitivity to return estimation errors in the portfolio to be optimised.

4.4. Additional factors

Additionally, we would like to point out that it is often difficult to implement the solutions to the basic portfolio selection problems of the models discussed above in the real world. For example, the optimal MV portfolio of the suggested models can be repeatedly inverted with significant transaction costs. Consequently, although the Markowitz MV model offers a sophisticated theoretical solution to the portfolio selection problem, its effective extensions to real-life applications are hampered by disregard for real-world circumstances. Although many of these elements are not accounted for by portfolio management theory, they have a significant impact on portfolio performance.

In addressing transaction costs, reasonable investors will not trade continuously in the market at all times. The transaction costs related to purchasing and selling assets are disregarded in the original MV model for the sake of simplicity. However, the absence of transaction costs will result in exaggerated portfolio performance. To achieve a portfolio with a specific level of diversification, investors may seek to simultaneously impose high or low limits on the value weighting of each asset in a number of different situations. In addition to these limit restrictions, the building of the portfolio may also be subject to so-called cardinality requirements, which specify the minimum or maximum number of securities to be included in the portfolio and may also call for a specific range. However, some market regulations impose certain real restrictions that an investor cannot ignore. The minimum transaction lot suggests that, in order to trade, you must make an investment in excess of this amount.

These are just a few examples of additional restrictions that must be taken into account in practice. Finally, these are quantitative tools whose results, as we have seen above, can be counter-intuitive. Traditional models are unable to incorporate the investor's prior knowledge or macroeconomic variables. Although the models described above provide a general explanation of the concept of risk-return trade-offs and the possibility of creating optimal portfolios for these variables, the market is much more complex, with noise and other variables that make the initial problem a more complicated task.

5. Modern approach to portfolio optimization

As previously mentioned, it is challenging to gather precise estimates of the data required to determine an optimal portfolio, and estimated mistakes in the forecasts have a major impact on the final portfolio weights. As a result, mean-variance analysis-generated optimal portfolios frequently have high or illogical weights for some assets. However, Fabozzi et al. (2012) point out that in practice, the mean- variance analysis proposed by Markowitz must be modified to achieve reliability, stability, and robustness to model and estimation errors. This is because these examples do not necessarily prove that the theory of portfolio selection is incorrect.

Several scholars have long focused on optimisation under parameter uncertainty (Bertsimas et al. 2011). The two most commonly used methods to solve this problem are stochastic optimisation and resilient optimisation. The foundation of the former is the idea that uncertainty has a probabilistic description, while the latter is a more creative approach in which model uncertainty is joint and determin

istic rather than stochastic. In our study, we present both methods as possible solutions to estimation errors and analyse the suitability of each to the portfolio optimisation problem.

5.1. Robust Optimization

When it comes to robust optimization, the approach is based on a min-regret modelling technique that aims to lessen the effects of impending events in situations where the values of model variables are unknown, fluctuating, and their distributions are inherently unpredictable (Gregory et al., 2011). To properly understand the basis of this model, it is useful to clarify the difference between variability and uncertainty. Vose (2008) explains that variability, usually expressed as a statistical indicator such as variance or standard deviation, refers to an inevitable and unanticipated change in a parameter that cannot be corrected for by data collection. In other words, the performance of the parameter will not change no matter how much information is collected about it. On the contrary, uncertainty can be reduced by collecting more data, as it reflects the lack of prior knowledge about a potential outcome. The difference between variability and uncertainty can be better understood by taking a look at a stock price's random walk, which symbolizes its *natural and unpredictable variation*. One cannot change the random walk, i.e. make it less volatile, no matter how much information you learn or how well you understand its previous behaviour. This is referred to as variability. However, if we estimate a random walk's distribution, learning more about the walk's past actions improves the forecast's accuracy, which is referred to as uncertainty.

In light of the above, we can draw the conclusion that robust optimization entails factoring into the model the uncertainty caused by the estimation errors of the optimisation's parameters (Fabozzi et al., 2012). Gregory et al. (2011) note that by doing this, even if the assumed distributions and parameter estimates are incorrect, the model will guarantee, with a high degree of probability, that, for a given set of uncertainties, the solution will hold for all conceivable results of each unknown variable and that the optimal objective will be achieved or exceeded. According to Pflug & Wozabal (2007) and Fabozzi et al. (2012), the benefit of resilient optimization is that changes to the mathematical model do not alter the fact that the issue is still a quadratic programming one.

Historical data can be used to estimate the uncertainty set when parameter values are unknown. This uncertainty set's definition is largely dependent on subjective opinion and need not include all possible parameter realizations but rather only the most likely values (Gregory et al., 2011). Creating an asset allocation strategy whose performance is optimised under worst-case unknown variables such as returns and covariances is the focus of robust portfolio optimisation when it comes to portfolio management. The robust technique deviates from the traditional approach in that it treats the estimated input parameters for a portfolio allocation problem as reliable and uncertain (Xidonas et al., 2020).

As we have already mentioned, the optimiser's maximization impact causes our portfolio's predicted return to be typically larger than its actual return. As a result, the question of how big this difference can be emerges. Let's think about the biggest discrepancy between $\widetilde{w}$ predicted expected return and its real expected return to get the answer. Ceria and Stubbs (2006) depict this difference as

$$\bar{\alpha}^T\widetilde{w} - \alpha^{*T}\widetilde{w} \qquad (9a)$$

Where $\bar{\alpha}^T\widetilde{w}$ stands for the estimate expected return of the portfolio in period T and $\alpha^{*T}\widetilde{w}$ is the true return of such portfolio. The largest discrepancy between the projected returns on the estimated efficient frontier and the real efficient frontier is determined by solving for the highest difference between equation *9b* within a confidence region of α.

$$\max \bar{\alpha}^T\widetilde{w} - \alpha^{*T}\widetilde{w} \qquad (9b)$$

$$subject\ to\ (\alpha - \bar{\alpha})^T \Sigma^{-1} (\alpha - \bar{\alpha}) \leq K^2$$

where $K^2 = \chi_n^2 (1 - \eta)$ and χ_n^2 is the chi-squared distribution's inverse cumulative distribution function with *n* degrees of freedom. Logically, it is desirable to keep this gap as small as possible. A strategy like this would bring the estimated and actual borders closer to the real efficient frontier while reducing the error-maximization effect. However, the optimal portfolio will be pushed towards a portfolio that would minimize the estimate risk by simply reducing the distance between the two frontiers. That is obviously not the direction we want to go in. Without estimates, it serves no purpose to think about estimating inaccuracy. Instead, we want to maximize portfolio anticipated return while simultaneously reducing estimation risk for a specific amount of expected return. In order to achieve this, we solve an optimization problem where we maximize the lowest feasible value of the portfolio's real expected return. By solving the previously mentioned max-min function optimization problem, we will bring the actual and estimated boundaries closer together.

Most traditional portfolio optimization methods have robust alternatives that share a similar level of computational complexity and need about as much time to solve as the standard alternative issue (Goldfarb and Iyengar, 2003). Fabozzi et al. (2012) claim that resilient Markowitz portfolios outperform conventional mean-variance portfolios outside of samples and are more stable than other portfolios as inputs change. The most common application of robust optimization in finance is portfolio allocation since the mean-variance framework's ideal portfolio weights are very dependent on the estimated input values. As a result, robust optimization has gained popularity as a method for including uncertainty in financial models in the field of portfolio management. The creation of equity portfolios and asset allocation were the early applications. Since robust portfolio optimisation gained prominence as a technique of enhancing traditional portfolio optimisation models approximately two decades ago, significant advancements in the field have taken place. Recent studies have been published in the literature that apply the worst-case approach to bond portfolio construction, currency hedging, and option pricing (Kim et al., 2018).

Yam et al. (2016) recently presented a variant of this resilient model that permits short positions. They examine various reliable formulations of mean-variance problems and discover that the influence of expected return uncertainty has a greater impact on sensitivity control than does uncertainty in the covariance matrix. Asl and Etula (2012) illustrate the usefulness of using robust models to allocate strategic assets. They use a multi-factor model that is appropriate for evaluating expected returns and risk across asset classes to undertake robust optimization as part of their robust asset allocation method. Their empirical studies show that robust portfolios comprised of 15 asset classes exhibit less dispersion, making them more suitable for asset allocation than conventional methods.

The notion that resilient portfolio optimisation outperforms traditional mean-variance optimisation on a large proportion of occasions is supported by several research utilizing both real and simulated market data (Kim et al. 2018). The way in which uncertainty sets are modelled may be the primary area where robust formulations of the portfolio optimization issue diverge (Bertsimas et al. 2011). Finding the ideal mix between robustness and a useful description of uncertainty sets will therefore probably have a big impact on portfolio performance. Furthermore, it appears that robust optimization generally leads to more stable portfolio weights (Xidonas et al., 2020) based on evidence from both simulated and real-world data. Fabozzi et al. (2012) noted that robust optimization typically steers clear of corner solutions. This is due to the fact that in a corner solution, new value is either added or subtracted from the portfolio. Consequently, robust mean-variance optimization leads to smoother and more predictable portfolio returns because it tends to enhance worst-case portfolio performance.

Indeed, the ability of robust efficient portfolios to maintain a relatively constant value over extended periods of time is a key characteristic. Furthermore, applying sound asset allocation procedures can dramatically enhance the worst-case performance of portfolios of various assets, frequently with very minor performance losses in more probable circumstances (Tütüncü and Koenig 2004). Particularly, both the advantages of robust portfolios in worst-case situations and their underperformance in more probable circumstances appear to increase as the size of uncertainty sets increases. This trade-off indicates that it is required to conduct a cost-benefit analysis of the size of uncertainty sets.

5.2. Bayesian Optimization

On the other hand, Bayesian optimisation is a class of machine learning-based optimization techniques focusing on problem solving:

$$\max_{x \in A} f(x) \qquad (10\ a)$$

in which the feasible bundle and the objective function usually have, among others, a number of particular properties. The entry x is in $\mathbb{R}^d$for a value of d that is not excessively large being able to work for less than twenty. The target function $f(x)$, to be optimized, is continuous. This is typically required for modelling any function by Gaussian process regression. The specified function is, in a sense, expensive to evaluate, which limits the number of evaluations that may be conducted, frequently to a few hundred. This limitation typically results from the fact that each evaluation requires a sizable amount of time, but it can also occur because each evaluation has a financial or opportunity cost. One of the most important properties of the objective function is the absence of a well-known unique pattern, such as concavity or linearity, which would make it easier to optimize it using techniques that leverage this pattern to boost efficiency. Since we only observe $f(x)$ and neither first nor second-order derivatives when we evaluate the function, it is referred to as a "black box" in mathematics. This makes it impossible to use first- and second-order methods like Newton's method, gradient descent, or quasi-Newton methods.

The aim of this method is to find a global rather than a local optimum. Summarising these features of the problem, we say that BO is designed for global optimisation without black-box derivatives. This method is very adaptable since it can optimize pricey black-box functions without derivatives. Although it has historically been extensively used for engineering system design, as well as for selecting laboratory experiments in materials and drug design or reinforcement learning (Brochu et al., 2010), it has recently gained a lot of popularity for fitting hyperparameters in machine learning algorithms, especially deep neural networks, as proposed by Snoek et al. (2012). BO was created by Kushner in 1964, but it gained much more popularity after Jones (1998) popularized it. In addition to BayesOpt, it is possible to optimize expensive derivative-free black-box functions using alternative techniques. This broader class of procedures is usually referred to as "surrogate methods". Bayesian optimization distinguishes itself from other surrogate approaches by making use of surrogates produced using Bayesian statistics and selecting where to evaluate the objective using a Bayesian interpretation of these surrogates.

A Bayesian statistical model to describe the objective function and an acquisition function to determine where to sample next constitute the two primary parts of Bayesian optimization. Iteratively, they are used to distribute the remaining funds from a budget of N evaluations of the function after the target has been evaluated in accordance with an initial space-filling experimental design, which frequently consists of points chosen uniformly at random. A Bayesian posterior probability distribution is provided by the statistical model, which is always a Gaussian process, and it describes the possible values of $f(x)$ at a candidate point x. This posterior distribution is changed each time we observe the function at a new point, as seen in figure 2.

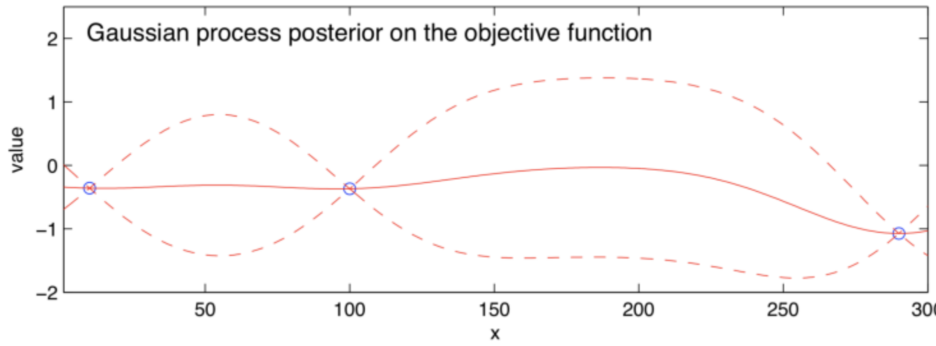

Picture 2: Posterior distribution of an objective function after applying a Gaussian process

According to the current posterior distribution over the initial function, the acquisition function calculates the value that would be produced by the evaluation of the objective function at a new point x (see Figure 3). A posterior probability distribution for each $f(x)$ that is normally distributed and has a mean $\mu_n(x)$ and variance of $\sigma_n^2(x)$ is produced by the Gaussian regression. One interpretation of the mean is as a point estimate of $f(x)$. In frequentist statistics, the credible interval serves as a confidence interval and, according to the posterior distribution, contains $f(x)$ with a 95% percent probability. The previously assessed points are interpolated using the mean. In these locations, the credible interval is at width 0, and as we go away from them, it becomes wider.

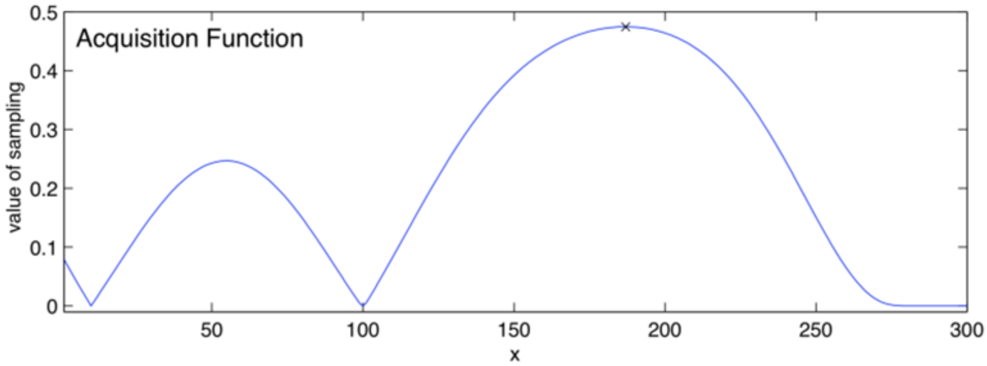

Picture 3: Acquisition function for best-candidate search in Bayesian Optimization

The predicted improvement is the acquisition function that is most frequently employed. Now imagine that at any time we need to carry out a different assessment. We will obtain $f(x)$ if we evaluate x. The value of the best point we have seen after this additional examination will be $f(x)$ (if $f(x) \geq f_n^*$) or f_n^** (if $f(x) \leq f_n^*$). Therefore, if this number is bigger than zero, $f(x) - f_n^*$ is the improvement of the value of the best observed point, and 0 otherwise. This improvement can be expressed more succinctly as $[f(x) - f_n^*]^+ +$, where $a^+ = \max(a, 0)$ denotes the positive part. Despite the fact that we should choose the option with the greatest improvement, $f(x)$ is not known until after evaluation. However, we can choose x to maximize the expected value of this improvement by calculating it. The anticipated development is described as

$$EI_n(x) = E_n[[f(x) - f_n^*]^+] \qquad (10\ b)$$

Where $E_n[\cdot]$ denotes the expectation assigned to the posterior distribution of the function $f(x)$ at $x_1, \ldots, x_n$ and is normally distributed, with mean $\mu_n(x)$ and variance $\sigma_n^2(x)$.

Four benefits of using Bayesian methods over traditional ones can be summed up. First off, long-term investment decisions do take estimation risk into consideration, but the one-period scenario essentially does not. According to Barberis (2000), a long-term investor who disregards it can significantly over-allocate funds to equities. Second, because it may take into account the noise that these parameters experience, the distribution of future returns conditional on the model parameters no longer follows a well-known distribution shape even when the predictors evolve stochastically. The third and fourth benefits are related to the Bayesian investor's capacity to take into account model uncertainty as well as prior perceptions of the level of predictability covered by asset pricing models. We also point out that this method makes it possible to model complex economic magnitudes using quick, simple, and easy-to-use numerical techniques. This strategy's benefits have been proven both conceptually and experimentally. The constructed Bayesian efficient frontier allows us to develop an intelligent technique by performing interval forecasting of the impending updates of the optimal portfolio returns obtained by using the stochastic representation derived from the posterior predictive distribution, according to Bauder et al. (2021), who discuss how to build an optimal portfolio based on the posterior predictive distribution.

The fact that optimal portfolios are quite sensitive to the expected level of returns, which we have explored in prior parts, was noted by Best and Grauer (1991), among other experts. Therefore, enhancing the mean estimation method has recently emerged as a major problem in the portfolio optimisation conundrum. Due to the usage of the acquisition function, Bayesian optimisation techniques are among the most effective methods in terms of the quantity of function evaluations needed. In the 1970s, the earliest portfolio analysis using Bayesian statistics relied solely on non-training or data-driven priors. Comparable to traditional methods of portfolio selection are typically Bayesian approaches based on diffusion a priori. However, Bayesian methods based on the diffusion prior produce different outcomes if certain risky assets have longer histories than others (Stambaugh, 1997). In the spirit of Bayes-Stein contraction priority, Jorion (1996) presented the hyperparametric priority approach,.Contrarily, Black and Litterman (1992) advocated a loosely-defined Bayesian approach supported by equilibrium correlations and economic arguments. Instead of only applying mean-variance optimization, they were able to generate more stable and diversified portfolios by using the Black-Litterman model.

The majority of the cited research used particular a priori model parameter values that were then numerically evaluated in order to assess posterior distributions or asset allocation choices. Posterior beliefs about the values implicit in asset pricing theories were the focus of recent studies Pástor and Stambaugh (2000). These authors specifically looked into investors' portfolio decisions that minimize mean variance and update preconceived notions about risk- or characteristics-based pricing models using sample data. According to Tu and Zhou (2010), who argued that the investment objective provides a useful presumption for portfolio selection, the naive equal-weighted portfolio rule and one of the four sophisticated strategies—the Markowitz (1952) rule, the Jorion (1996) rule, the MacKinlay and Pástor (2000) rule, or the Kan and Zhou (2007) rule—could be combined optimally to improve performance.

The Bayesian environment is similar to how people use information and make decisions in the market. Similar to this, investors base their decisions on their past experiences and recollection, such as historical events or trends. The problem outlined above is the inevitable consequence of the greatest strength of the Bayesian approach: its ability to include the human behaviour of the investor in decision-making. While the Markowitz model is highly theoretical and rather naïve when applied in practice, Bayesian optimisation solves this dilemma and provides us with a radically different view of portfolio optimisation, tailored to the individual investor and able to capture the multitude of factors that affect this process in praxis. According to Avramov and Zhou (2010), the Bayesian framework in portfolio theory may be more appealing from this perspective.

6. Experimental implementation of classical and modern models

Previously, we have explained the benefits that both Robust and Bayesian Optimisation have to offer to the classical Markowitz model. In order to provide a practical example, we have conducted a simple experimental implementation of both methods on a real portfolio optimisation problem, using the traditional Mean-Variance model as a benchmark. Our purpose is to explain how the two proposed methods work in practice as well as the assumed improvement with respect to the classical approach.

6.1. Formulation of the problem

Our objective is to optimise a portfolio composed of stocks listed on the S&P 500. Specifically, our portfolio will include stocks of Apple, Amazon, Google, Tesla and Meta; for the sake of simplicity, we have chosen only five stocks, but our models allow the imputation of as many stocks as desired. Our models will optimise the weights that these five stocks will have in our portfolio based on the evolution of these stocks over the year

2021, having extracted their prices from Yahoo Finance. Once we have obtained the optimal weights for each stock, according to our three models, we create three different portfolios and compare their performance over the year 2022, using the cumulative return as the unit of comparison. The return of our portfolio can be explained by the following mathematical representation:

$$R_p = \sum w_i r_i \tag{11 a}$$

where w_i refers to the weight of stock *i* in our portfolio and r_i , its return. In the same way, the volatility of our portfolio can be explained with the following mathematical expression:

$$\sigma_p = \sqrt{\sum_{i=1}^{n} \sum_{j=1}^{n} w_i w_j \, Cov(r_i r_j)} \tag{11 b}$$

The objective of our models will be to optimise the Sharpe Ratio obtained by our portfolio, measuring the risk-adjusted return of the portfolio:

$$Sharpe\ Ratio = {}^{R_p}/_{\sigma_p} \tag{11 c}$$

In this way, our models will find different optimal weights for our stocks, designing three optimal portfolios based on the return and volatility shown by our five stocks during year 2021 (figure 4).

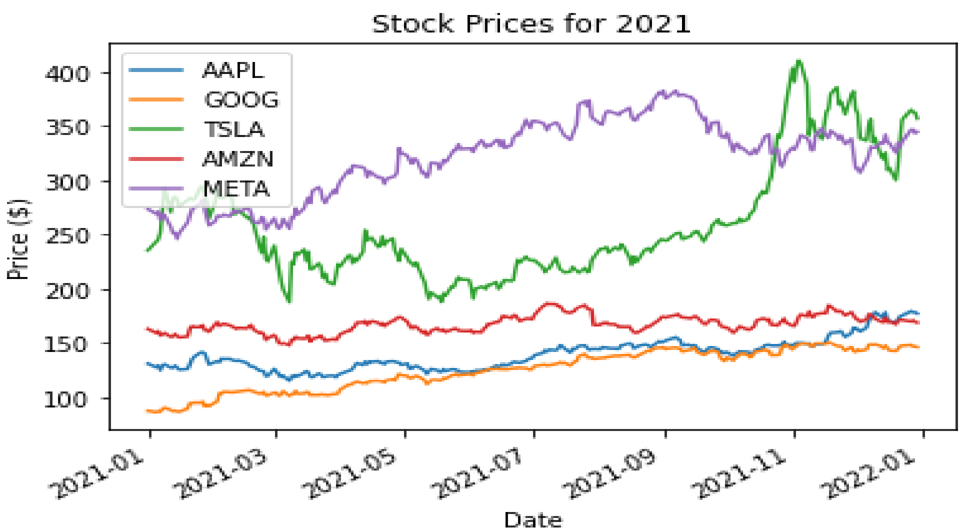

Picture 4: Price evolution for Apple, Amazon, Tesla, Google and Meta stocks in 2021.

6.2. Experimental set up and implementation details

Our Mean-Variance, Robust and Bayesian models have been developed using Python. The Markowitz model is the most straightforward; we define a classical mean-variance optimisation method as the described in section 4. This method is used to optimise a portfolio of stocks by finding the set of weights that minimises the volatility of the portfolio relative to its expected return, in other words, we try to maximise the Sharpe Ratio of our portfolio. Our code initialises the method by setting the number of stocks in the portfolio and the limits of the weights of each stock (between 0 and 1). It then sets the initial estimate of the weights to be the same for all stocks, i.e. we start from a naïve model. The objective function is then defined, which takes the weights as input and calculates the return and volatility of the portfolio based on the input weights. We calculate the deviation between the expected volatility and the actual volatility of the portfolio. Finally, optimisation is performed using the minimise function of the SciPy library, which finds the weights that minimise the objective function while satisfying the constraints given by the bounds. The optimised weights are returned as output.

The Robust Optimisation method, on the other hand, aims to minimise the impact of outliers on asset returns over the final portfolio. It achieves this objective by using a Huber loss function instead of the traditional quadratic loss function. The Huber loss function is a combination of the quadratic loss function for small deviations and the linear loss function for large deviations. This function is less sensitive to extreme values and results in a more robust and stable optimisation process. In our implementation, the function first initialises the variables needed for optimisation, including the number of shares, the bounds for optimisation, which are the same as for the Markowitz model, and the initial assignment of weights. It then defines the objective function, which calculates the return and volatility of the portfolio as a function of the input weights and returns the Sharpe ratio. Finally, the Huber loss function is defined to take the input weights and a Huber constant as a parameter. The Huber loss function returns a loss value based on the difference between the Sharpe ratio and the Huber constant, which is minimised using the L-BFGS-B optimisation method. The resulting weight assignment is returned as a loss value, which we minimise. The resulting weight allocation is returned as the optimal portfolio.

Our Bayesian model will also try to find the optimal weights of a portfolio of stocks. First, we input the number of stocks contained in our portfolio and assign the range of stocks that each weight can take, between zero and one, as in the previous models. Next, we define an objective function that takes the weights as input, calculates the expected return and volatility of the portfolio, and returns the Sharpe ratio (i.e. the return over volatility). We then initialise the Bayesian optimiser by inputting the objective function and the bounds on the weights. The optimiser is set to maximise the objective function and is initialised with a random seed. Finally, we run our optimiser for a certain number of iterations (50 in this case) with a certain number of random initialisations (10 in this case), and the optimal set of weights is returned. Finally, with the weights returned by our three models, we construct our optimal portfolios and contrast their cumulative return over the year 2022.

7. Discussion of the results

The performance of the stocks selected for our portfolio during the test year is shown in Figure 5. Compared to the train values in Figure 4, corresponding to the year 2021, we see that they follow a very different trend. This difference is a fundamental characteristic of the real market and allows us to realistically compare the performance of these three models in a real scenario, with noise, volatility and uncertain patterns.

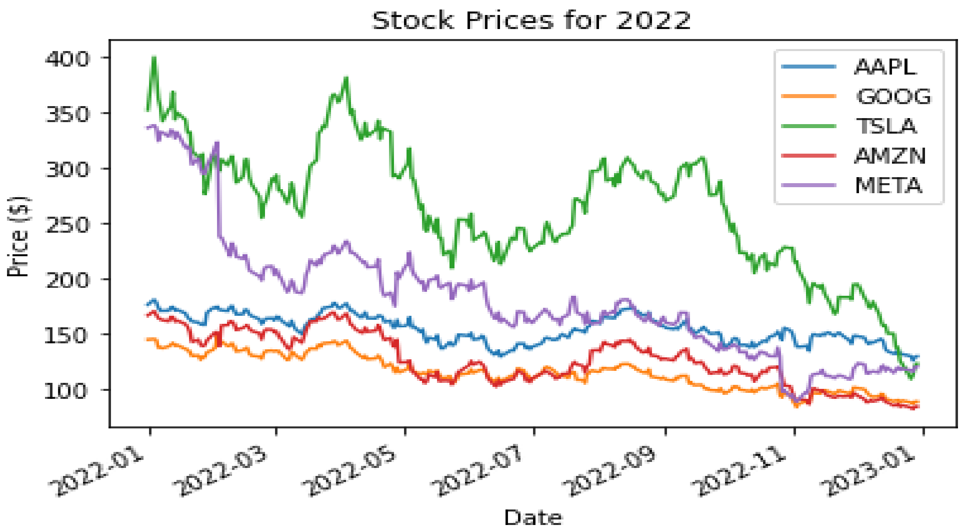

Picture 5: Price evolution for Apple, Amazon, Tesla, Google and Meta stocks in 2022.

Based on the returns for our stocks in the year 2021, our three different methods have designed three different optimal portfolios whose weights for each stock are shown in Figure 6.

	Mean Variance Model	**Robust Model**	**Bayesian Model**
AAPL	0.00	0.00	1.000000e-14
AMZN	8.956761e-01	0.00	9.985640e-01
GOOG	6.955954e-02	0.114224	7.862607e-02
META	2.785404e-17	8.969980e-01	1.000000e-14
TSLA	0.00	0.026479	1.000000e-14

Picture 6: Optimal weights for Mean-Variance, Robust and Bayes based portfolios.

Accordingly, we have constructed our optimal portfolios with the weights provided above and in Figure 7 we show the evolution of the cumulative return of our portfolios over the year 2022.

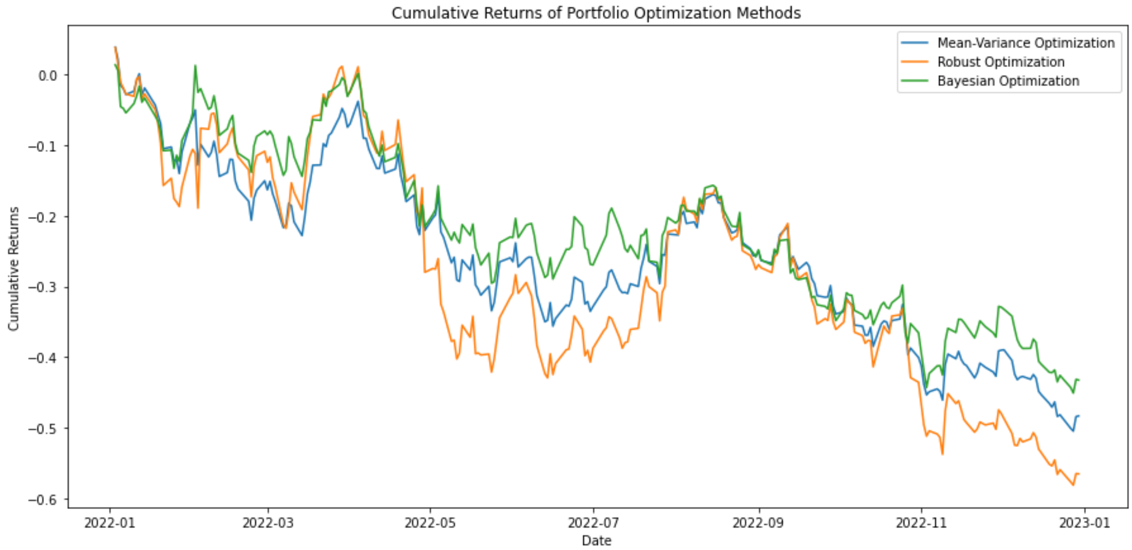

Picture 7: Cumulative return of Mean-Variance, Robust and Bayesian optimal portfolios for 2022.

During 2022, the price of our five stocks followed a negative line. The return of our three portfolios consequently followed this trend, but their returns were different. Overall, the Bayesian portfolio was the best performer and the robust portfolio yielded the largest losses. The traditional Markowitz portfolio remained broadly in the middle ground between the two. In periods when returns are positive, the MV portfolio performs similarly to the Bayesian portfolio, but in declining periods, its returns are much lower and even match the Bayesian portfolio.

Analysis of the return of our portfolios			
	Mean Variance Model	**Robust Model**	**Bayesian Model**
Maximum	0.092928	0.139409	0.083273
Mean	-0.002273	-0.002693	-0.001908
Median	-0.001439	-0.004905	-0.002702
Minimum	-0.082233	-0.142050	-0.095565
Volatility	0.026527	0.035176	0.026321
Final cumulative return	-0.483147	-0.564106	-0.432446

Picture 8: Statistics of the returns of our Mean-Variance, Robust and Bayesian portfolios.

In image number eight, we take a closer look at the statistics of our portfolios. The robust model yields the highest volatility and the most negative returns; both its mean and median are noticeably below the other two. The most stable portfolio turns out to be the Bayesian one, which also provides the highest cumulative return and the highest mean. The median of the MV portfolio is slightly higher than the Bayesian one and it is precisely the Markowitz portfolio that offers the highest return. However, it is a more volatile portfolio whose performance over the annual period analysed is below our Bayesian approach. From the results we can conclude that the performance of the Bayesian portfolio is the most stable and positive as a rule. In the worst cases when stock prices fall, its return remains above its two alternatives and, in the medium and long term, its low volatility makes it the most profitable portfolio on average.

8. Conclusions and further areas of research

Since the birth of Modern Portfolio Theory in 1952, the science of portfolio optimisation has evolved dramatically. In the wake of Markowitz's Mean-Variance model, numerous alternatives have emerged, providing tools for investors to analyse the behaviour of portfolios as a whole. However, the market in which these portfolios move is subject to supply and demand and follows irrational patterns that, to date, cannot be explained by any model. Portfolio optimisation is therefore a complex and continuously evolving science and a field of application for all state-of-the-art tools focused on the optimisation of functions that follow unknown patterns but for which a massive amount of noisy data is available.

It is precisely the irrational behaviour of the components of a portfolio that creates the need to estimate values such as return and risk in order to establish an optimisation strategy. Naturally, any estimation is subject to error. Thus, the models of classical theory, whose basis is the optimisation of the risk-return trade-off from estimates obtained by regression and other traditional methods, show high estimation errors that in many cases can be highly detrimental to the investor. In order to avoid precisely the most damaging cases, the Robust Model tries to impute this uncertainty present in all estimation by means of the inclusion of uncertainty sets. In this way, it tries to avoid those cases in which the portfolio can provide disastrous results for the investor, i.e. it avoids the riskiest scenarios, which can generally also be the most profitable ones, eliminating the outliers and reducing the risk to which the investor is exposed. In our experimental implementation we have found that the Robust Model is the worst performing model. While it is true that our model is simple in order to describe how each method works, the Robust Model errs in excluding precisely those portfolios with a high risk-return ratio. This model creates precisely robust portfolios that avoid both best- and worst-case scenarios, making it more suitable for conservative investors who wish to avoid catastrophic downturns at all costs to the detriment of higher returns when the market is on the upside.

The Bayesian Model, on the other hand, includes the investor's prior knowledge, being able to incorporate that subjectivity present in the market. It is a method that is somehow close to the subjective thinking behind the investment and capable of modelling functions that follow an irrational pattern and hardly explainable by a classical function. In our experimental implementation, this model is the one that shows the best performance, outperforming our benchmark model, the traditional Mean-Variance Model. We have explained above that the major limitation of the Markowitz model lies in the estimation of its inputs. In fact, Markowitz's approach is not flawed, but is limited by the error in the estimates of the securities he analyses, as are his alternative models such as MAD, VaR, CVar, Minimax or Lower Partial Moment, discussed in section 4. These tools changed the investment world, providing for the first-time units of measurement for the risk and return of a portfolio as a whole. The emergence of intelligent tools, such as Bayesian Optimization, does not imply their elimination. On the contrary, we propose this approach as the solution to limit classical estimation errors and provide investors with a tool that allows them to analyse the performance of a portfolio in an accurate way, incorporating the investor's expertise and making use of the massive amount of information they have at their disposal nowadays.

REFERENCES

Alpaydin, E. (2014). *Introduction to Machine Learning*. MIT Press.

Asl, F. M., & Etula, E. (2012). Advancing Strategic Asset Allocation in aMulti-Factor World. *The Journal of Portfolio Management*, *39*(1), 59–66. https://doi.org/10.3905/jpm.2012.39.1.059

Avramov, D., & Zhou, G. (2010). Bayesian Portfolio Analysis. *Annual Review of Financial Economics*, *2*(1), 25–47. https://doi.org/10.1146/annurev-financial-120209-133947

Bai, L., Cui, L., Wang, Y., Jiao, Y., & Hancock, E. R. (2020). A Quantum-inspired Entropic Kernel for Multiple Financial Time Series Analysis. *International Joint Conference on Artificial Intelligence*. https://doi.org/10.24963/ijcai.2020/614

Barberis, N. (2000). Investing for the Long Run when Returns Are Predictable. Journal of Finance, 55(1), 225–264. https://doi.org/10.1111/0022-1082.00205

Bauder, D., Bodnar, T., Parolya, N., & Schmid, W. (2021). Bayesian mean–variance analysis: optimal portfolio selection under parameter uncertainty. *Quantitative Finance*, *21*(2), 221–242. https://doi.org/10.1080/14697688.2020.1748214

Bertsimas, D., Brown, D., & Caramanis, C. (2011). Theory and Applications of Robust Optimization. *Siam Review*, *53*(3), 464–501. https://doi.org/10.1137/080734510

Best, M. D., & Grauer, R. T. (1991). On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results. *Review of Financial Studies*, *4*(2), 315–342. https://doi.org/10.1093/rfs/4.2.315

Black, F., & Litterman, R. (1992). Global Portfolio Optimization. *Financial Analysts Journal*, *48*(5), 28–43. https://doi.org/10.2469/faj.v48.n5.28

Britten-Jones, M. (1999). The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights. *Journal of Finance*, *54*(2), 655–671. https://doi.org/10.1111/0022-1082.00120

Broadie, M. (1993). Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, *45*(1), 21–58. https://doi.org/10.1007/bf02282040

Brochu, E., Hoffman, M. K., & De Freitas, N. (2010). Portfolio Allocation for Bayesian Optimization. *ArXiv (Cornell University)*. https://arxiv.org/pdf/1009.5419.pdf

Brogan, A., & Stidham, S. (2008). Non-separation in the mean–lower-partial-moment portfolio optimization problem. *European Journal of Operational Research*, *184*(2), 701–710. https://doi.org/10.1016/j.ejor.2006.11.028

Cai, X., Teo, K. L., Yang, X., & Zhou, X. K. (2004). Minimax portfolio optimization: empirical numerical study. *Journal of the Operational Research Society*, *55*(1), 65–72. https://doi.org/10.1057/palgrave.jors.2601648

Ceria, S., & Stubbs, R. M. (2006). Incorporating estimation errors into portfolio selection: Robust portfolio construction. *Journal of Asset Management*, *7*(2), 109–127. https://doi.org/10.1057/palgrave.jam.2240207

Choi, H. I. (2018). Stock Price Correlation Coefficient Prediction with ARIMA-LSTM Hybrid Model. *RePEc: Research Papers in Economics*. https://EconPapers.repec.org/RePEc:arx:papers:1808.01560

Chopra, V. K., & Ziemba, W. T. (1993). The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice. *The Journal of Portfolio Management*, *19*(2), 6–11. https://doi.org/10.3905/jpm.1993.409440

Fabozzi, F. J., Markowitz, H. M., Kolm, P. N., & Gupta, F. (2012). Mean-Variance Model for Portfolio Selection. *Encyclopedia of Financial Models*. https://doi.org/10.1002/9781118182635.efm0003

Goel, A., Tripathi, V., & Agarwal, M. (2020). Information asymmetry and stock returns. *Journal of Advances in Management Research*, *18*(1), 85–112. https://doi.org/10.1108/jamr-05-2020-0084

Goldfarb, D., & Iyengar, G. (2003). Robust Portfolio Selection Problems. *Mathematics of Operations Research*, *28*(1), 1–38. https://doi.org/10.1287/moor.28.1.1.14260

Gregory, C., Darby-Dowman, K., & Mitra, G. (2011). Robust optimization and portfolio selection: The cost of robustness. *European Journal of Operational Research*, *212*(2), 417–428. https://doi.org/10.1016/j.ejor.2011.02.015

Jones, C. W. (1998). A Simple Bayesian Method for the Analysis of Diffusion Processes. *Social Science Research Network*. https://doi.org/10.2139/ssrn.111488

Jorion, P. (1996). Measuring the Risk in Value at Risk. *Financial Analysts Journal*, *52*(6), 47–56. https://doi.org/10.2469/faj.v52.n6.2039

Kallberg, J. G., & Ziemba, W. T. (1984). Mis-Specifications in Portfolio Selection Problems. *Springer Berlin Heidelberg EBooks*, 74–87. https://doi.org/10.1007/978-3-642-45569-8_7

Kan, R., & Zhou, G. (2007). Optimal Portfolio Choice with Parameter Uncertainty. *Journal of Financial and Quantitative Analysis*, *42*(3), 621–656. https://doi.org/10.1017/s0022109000004129

Kim, J. K., Kim, W. H., & Fabozzi, F. J. (2018). Recent advancements in robust optimization for investment management. *Annals of Operations Research*, *266*(1–2), 183–198. https://doi.org/10.1007/s10479-017-2573-5

Konno, H., & Koshizuka, T. (2005). Mean-absolute deviation model. *Iie Transactions*, *37*(10), 893–900. https://doi.org/10.1080/07408170591007786

Konno, H., & Yamazaki, H. (1991). Mean-Absolute Deviation Portfolio Optimization Model and Its Applications to Tokyo Stock Market. *Management Science*, *37*(5), 519–531. https://doi.org/10.1287/mnsc.37.5.519

Kushner, H. J. (1964). A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, *86*(1), 97–106. https://doi.org/10.1115/1.3653121

Lee, S. Y., & Yoo, S. J. (2020). Threshold-based portfolio: the role of the threshold and its applications. *The Journal of Supercomputing*, *76*(10), 8040–8057. https://doi.org/10.1007/s11227-018-2577-1

MacKinlay, A., & Pástor, L. (2000). Asset Pricing Models: Implications for Expected Returns and Portfolio Selection. *Review of Financial Studies*, *13*(4), 883–916. https://doi.org/10.1093/rfs/13.4.883

Markowitz, H. M. (1952). Portfolio Selection. *Journal of Finance*, *7*(1), 77. https://doi.org/10.2307/2975974

Michaud, R. O. (1989). The Markowitz Optimization Enigma: Is 'Optimized' Optimal? *Financial Analysts Journal*, *45*(1), 31–42. https://doi.org/10.2469/faj.v45.n1.31

Nawrocki, D. N. (1992). The characteristics of portfolios selected by n-degree Lower Partial Moment. *International Review of Financial Analysis*, *1*(3), 195–209. https://doi.org/10.1016/1057-5219(92)90004-n

Park, J. D., Lim, B. C., Lee, Y., & Young, M. R. (1998). A Minimax Portfolio Selection Rule with Linear Programming Solution. *Management Science*, *44*(5), 673–683. https://doi.org/10.1287/mnsc.44.5.673

Pástor, Ľ., & Stambaugh, R. F. (2000). Comparing asset pricing models: an investment perspective. *Journal of Financial Economics*, *56*(3), 335–381. https://doi.org/10.1016/s0304-405x(00)00044-1

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *The Journal of Risk*, *2*(3), 21–41. https://doi.org/10.21314/jor.2000.038

Roy, A. (1952). Safety First and the Holding of Assets. *Econometrica*, *20*(3), 431. https://doi.org/10.2307/1907413

Serraino, G., & Uryasev, S. (2013). Conditional Value-at-Risk (CVaR). Springer US EBooks, 258–266. https://doi.org/10.1007/978-1-4419-1153-7_1232

Snoek, J., Larochelle, H., & Adams, R. P. (2012). Practical Bayesian Optimization of Machine Learning Algorithms. *Neural Information Processing Systems*, *25*, 2951–2959. http://books.nips.cc/papers/files/nips25/NIPS2012_1338.pdf

Stambaugh, R. F. (1997). Analyzing investments whose histories differ in length. *Journal of Financial Economics*, *45*(3), 285–331. https://doi.org/10.1016/s0304-405x(97)00020-2

Tu, J., & Zhou, G. (2010). Incorporating Economic Objectives into Bayesian Priors: Portfolio Choice under Parameter Uncertainty. *Journal of Financial and Quantitative Analysis*, *45*(4), 959–986. https://doi.org/10.1017/s0022109010000335

Tütüncü, R., & Koenig, M. (2004). Robust Asset Allocation. *Annals of Operations Research*, *132*(1–4), 157–187. https://doi.org/10.1023/b:anor.0000045281.41041.ed

Vose, D. (2008). *Risk Analysis: A Quantitative Guide*. John Wiley & Sons.

Xidonas, P., Steuer, R. E., & Hassapis, C. (2020). Robust portfolio optimization: a categorized bibliographic review. *Annals of Operations Research*, *292*(1), 533–552. https://doi.org/10.1007/s10479-020-03630-8

Yam, S. C. P., Yang, H., & Yuen, F. L. (2016). Optimal asset allocation: Risk and information uncertainty. *European Journal of Operational Research*, *251*(2), 554–561. https://doi.org/10.1016/j.ejor.2015.11.011

Zhang, Y., Li, X., & Guo, S. (2018). Portfolio selection problems with Markowitz's mean–variance framework: a review of literature. *Fuzzy Optimization and Decision Making*, *17*(2), 125–158. https://doi.org/10.1007/s10700-017-9266-z

Zitzler, E. (1999). *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*.