An Introduction to Mean-Variance Portfolio Allocation

Jay Sivadas, Clara Ee, Ronak Datta

Introduction

As investors, the primary objective is to maximize returns while minimizing potential downside risk. Many investors spend ample amounts of time trying to select an ideal asset with the above characteristics, without providing the same level of thought to the entire portfolio allocation process. The process of optimizing the return-to-risk profile of a portfolio involves both selecting a proper set of assets and allocating those assets proportionally so that the maximal risk-return profile among those assets can be achieved.

Mean-variance analysis was invented by nobel prize winner Harry Markowitz in his development of Modern Portfolio Theory. The goal of the analysis is to provide a set of portfolio allocation weightings given any combination of assets that optimizes the Sharpe ratio, or mean return to volatility, of that portfolio. This invention is one of the most widely used formulas for portfolio allocation and is implemented by both retail and institutional investors.

In this paper, we aim to provide an introduction to mean-variance analysis, so that readers can understand both its theoretical derivation and practical applications.

Return Notation

This section outlines the return notation that will be used throughout the paper.

<u>Excess Returns</u>

We describe return as the percentage in which the price of an asset changes or the interest rate on a security.

notation	description	formula	example
r^i	return rate of asset i		
r^f	risk-free return rate		
$\tilde{r}^i$	excess return rate of asset i	$r^i - r^f$	

Figure 1: Return Notation

A risk-free asset is an asset with a certain future return and has no possibility of depreciating or becoming worthless. In reality, a true risk-free asset does not exist, and we instead use the United States 10-year Treasury Bond as a proxy as it is relatively safe to assume that the United States will not default on its bonds in the next 5, 10, or more years.

The excess return given any asset is thus the return rate minus the risk-free rate or $r^i - r^f$. This is important to note because excess returns, not absolute returns, are what investors are typically concerned with.

Allocation Weight

If we consider a portfolio of stocks and bonds, the return rate of the bonds is denoted as r^b and the return rate of stocks is denoted as r^s (Figure 2).

	return	allocation weight
bonds	r^b	w
stocks	r^s	$1 - w$

Figure 2: Return Notation for a Hypothetical Portfolio

The portfolio has a total allocation weight of 1—or 100 percent of the whole portfolio. If the allocation weight of bonds is w, the allocation weight of stocks is equal to 1 – w as illustrated in Figure 2 above.

Portfolio Return Statistics

mean	variance	correlation
μ	σ^2	ρ

Figure 3: Return Statistics Notation

The return statistics of a portfolio or asset could be described by its mean, variance and correlation.

The mean, denoted by μ, gives the expected return of a portfolio of assets. This is the weighted average return considering the holdings, weights, and individual assets in a portfolio.

The variance, denoted by σ^2, shows the extent to which a portfolio may vary from its expected return or mean. This is also described as volatility or risk in portfolio theory. (Note: σ denotes standard deviation.)

The correlation, denoted by ρ, is a measure of the degree to which two variables move in relation to each other, where $-1 \leq \rho \leq 1$. For example, if variable x moves by +2 and has a correlation of +0.5 with respect to variable y, variable y will move by +1. In finance, we typically take variable x to be the performance of a company's stock, and variable y to be the performance of some basis, which is typically the overall stock market or a proxy of it such as the S&P 500 Index.

Investment portfolios have means and variances given by the following equations:

Expected return of a portfolio, $\mu_p = w\mu_b + (1 - w)\mu_s$
Variance of a portfolio, $\sigma_p^2 = w^2\sigma_b^2 + (1 - w)^2\sigma_s^2 + 2w(1 - w)\rho\sigma_s\sigma_b$
where subscripts p, b and s represent the investment portfolio, bonds, and stocks respectively

From the equations above, it is evident that the expected return of a portfolio is simply a weighted average of the expected returns of bonds and stocks—thus exhibiting a linear relationship with respect to *w*, the allocation weight of bonds. On the other hand, there is a non-linear relationship between the variance of the portfolio and the allocation weight of bonds, *w*.

Portfolio Diversification

Consider different correlation values and their implications on returns and variation.

Case 1: ρ = 1 (Perfect Correlation)
We assume two securities, *b* and *s*, to be perfectly correlated.

The volatility (or standard deviation) of a portfolio, σ_p, then becomes directly proportional to the weight of asset b, *w*, such that both the mean and volatility are linear with respect to *w* as illustrated in the two equations below:

Expected return of a portfolio, $\mu_p = w\mu_b + (1 - w)\mu_s$
Volatility of a portfolio, $\sigma_p = w\sigma_b + (1 - w)\sigma_s$
where subscripts p, b and s represent the portfolio, security b, and security s respectively

The second equation for volatility is simply a mathematical derivation from the formula for the variance of a portfolio and represents a simple, weighted average of the volatilities of its constituent securities.

Case 2: ρ < 1 (Imperfect Correlation)
We assume two securities, *b* and *s*, to be imperfectly correlated.

The resulting volatility is convex relative to *w*, but the expected return of a portfolio remains linear relative to *w* as shown below:

Expected return of a portfolio, $\mu_p = w\mu_b + (1 - w)\mu_s$
Volatility of a portfolio, $\sigma_p < w\sigma_b + (1 - w)\sigma_s$
where subscripts p, b and s represent the portfolio, security b, and security s respectively

The term *convex* is used to describe this scenario where the volatility of a portfolio is less than the weighted average volatility of its constituent securities.

Case 3: ρ = -1 (Perfect Inverse Correlation)
We assume two securities, *b* and *s*, to be perfectly, inversely correlated.

This relationship entails that one can expect security *b* to increase by 5 basis points if security *a* decreases by 5 basis points, where a basis point is equivalent to 0.01%. However, it is important to note that this relation is *not* causal in nature.

This presents an opportunity for the "perfect hedge". Theoretically, the risk or volatility of a portfolio, σ_p, could be 0 as long as the volatility of security s, σ_s, as a fraction of the sum of the volatilities of securities *b* and *s* is equal to w:

$$\text{weighting of security b, } w = \frac{\sigma_s}{\sigma_b + \sigma_s}$$

This is simply a rearrangement of the formula $\sigma_p = w\sigma_b + (1 - w)\sigma_s$, where $\sigma_p = 0$.

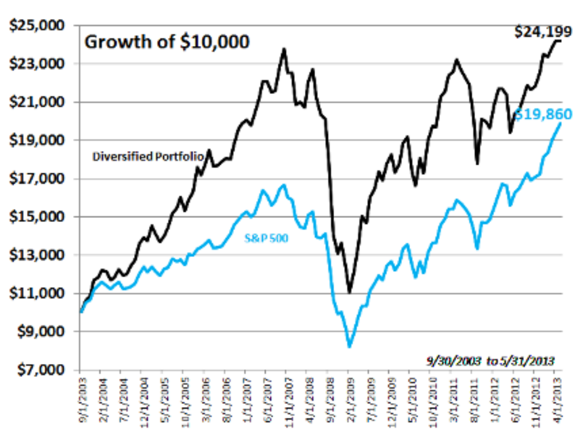

Figure 4: Performance of a Diversified Portfolio relative to a Market Benchmark (S&P 500)

This concept of correlation will serve as the basis for our discussion on portfolio diversification. Portfolio diversification balances risk and reward through exposure across non-perfectly correlated assets. Mathematically, the expected return of a portfolio, μ_p, must be a simple, weighted average of the expected returns of its constituent assets while the volatility of the portfolio, σ_p, must be less than the weighted average of the volatilities of these assets. This implies that correlation between the constituent assets must be imperfect or $\rho < 1$. If $\rho = 1$, there is effectively no diversification.

While there are different ways to diversify, the underlying idea is to increase the number of assets such that each is imperfectly correlated to the other. One could, for instance, purchase securities from different sectors such that exposure to the risk of each asset or sector is largely limited by offsetting characteristics of other assets. This protects a diversified portfolio from significant downside in any of its assets. While some investors recommend a minimum of 32 assets to completely offset individual asset risk, also known as idiosyncratic risk, others, like Warren Buffet, choose to incorporate only 10 assets across different industries in his portfolio.

Portfolio Variance Decomposition

The statistics underscoring the relationship between two different securities can be generalized to illustrate the risk-return characteristics of a whole portfolio.

Covariance

Covariance, like correlation, is a measure of the relationship between two variables. Correlation can be thought of as the scaled version of covariance where the data is standardized, such that the range of values are limited to between –1 and 1, inclusive.

$corr(X, Y) = \frac{cov(X, Y)}{\sigma_x \sigma_y}$, *where* σ_x *and* σ_y *represent the standard deviations of x and y resp.*

A notable difference between correlation and covariance is that the former encapsulates the strength of the relationship between x and y. A greater absolute correlation value would hence imply a stronger relationship, but a greater covariance value does not necessarily suggest such a relation.

Pairwise Covariances refer to covariances of all possible pairings of values. Given two lists (x, y, z) and (i, j, k), a list of pairwise covariances would look like this: (cov(x, i), cov(x, j), cov(x, k), cov(y, i), cov(y, j), cov(y, k), …).

Deriving Portfolio Variance

Consider a portfolio of *n* equally weighted assets, each with identical volatilities and correlations. Variance of this portfolio could be given by the following equation

$$\sigma_p^2 = \frac{1}{n}\, avg[\sigma_i^2] \; + \frac{n-1}{n} avg[\sigma_{i,j}]$$

where the average of the variances of these assets, $avg[\sigma_i^2] = \frac{1}{n}\sum_{i=1}^{n} \sigma_i^2$

and the average of the pairwise covariances of these assets, $avg[\sigma_{i,j}] = \frac{1}{n\,(n-1)}\sum_{j\neq i}\sum_{i=1}^{n} \sigma_{i,j}$

Note that the average variance and covariance values are simply weighted averages of the constituent assets' variances and covariances.

Observe that the first term of the resulting equation, $\frac{1}{n}\, avg[\sigma_i^2]$, approaches 0 as n approaches infinity and the second term of the resulting equation, $\frac{n-1}{n} avg[\sigma_{i,j}]$ approaches $avg[\sigma_{i,j}]$ as n approaches infinity. We describe them as diversifiable and non-diversifiable risk or the idiosyncratic and systematic risks, respectively.

Consider the equation, where instead $avg[\sigma_{i,j}]$ is substituted for $\rho\sigma^2$, given $\sigma_{i,j} = \rho\sigma^2$.

$$\sigma_p^2 = \frac{1}{n}\sigma^2 + \frac{n-1}{n}\rho\sigma^2$$

Because $\lim_{n\to\infty} \sigma_p^2 \to \rho\sigma^2$, no amount of diversification—by increasing the number of assets, *n*—can lower the variance of the portfolio below $\rho\sigma^2$. As such, we find that the total portfolio risk is always greater than or equal to the systematic risk, which is given by $\sigma_p^2 \geq \rho\sigma^2$.

The idiosyncratic risk in the above scenario is 0 in the above situation, as $\lim_{n\to\infty} \frac{1}{n}\sigma^2 \to 0$.

Because the idiosyncratic risk can easily be diversified away, the variances of each asset within the portfolio becomes a trivial consideration. This is because investors can easily limit their idiosyncratic risk simply by holding a larger number of assets in their portfolio. Systematic risk, however, is a factor that can not be easily removed and must be a considerable factor in the decision of portfolio allocation. Investors will demand more returns based purely on the systematic. Given these considerations, we consider only the following equation:

$$\lim_{n\to\infty} \sigma_p^2 = avg[\sigma_{i,j}]$$

Correlation and Risk

Recall from the previous section that $\sigma_p^2 = \frac{1}{n}\sigma^2 + \frac{n-1}{n}\rho\sigma^2$. If ρ is taken to be the correlation, we could derive a range of risk profiles of a portfolio based on differing asset correlations.

When $\rho = 1$: $\sigma_p^2 = \sigma^2$, suggesting that diversification is not possible.

When $\rho = 0$: $\sigma_p^2 = \frac{1}{n}\sigma^2$, suggesting that variance depends only on idiosyncratic risks. In this scenario, as $n \to \infty$, the portfolio becomes risk-free as illustrated in the following equation:

$$\lim_{n \to \infty} \sigma_p^2 = 0$$

Mean Variance

Mean-variance analysis, a part of the Nobel Prize-winning Modern Portfolio Theory, was conceptualized by Harry Markowitz. The mean-variance theorem implies that given any combination of *n* assets, an investor could optimize the risk-return profile of the portfolio by selecting a composition of the *n* assets that maximize returns given a certain level of risk.

The Sharpe ratio is a ratio of the return of an investment to its risk or volatility. The mean excess returns of a portfolio could be used to measure return, while risk, or volatility could be derived from the mean variance.

As such, Sharpe ratio $= \frac{\mu^p - r^f}{\sigma^p}$, where μ^p, σ^p and r^f represent the average returns, standard deviation of returns and the risk free rate respectively.

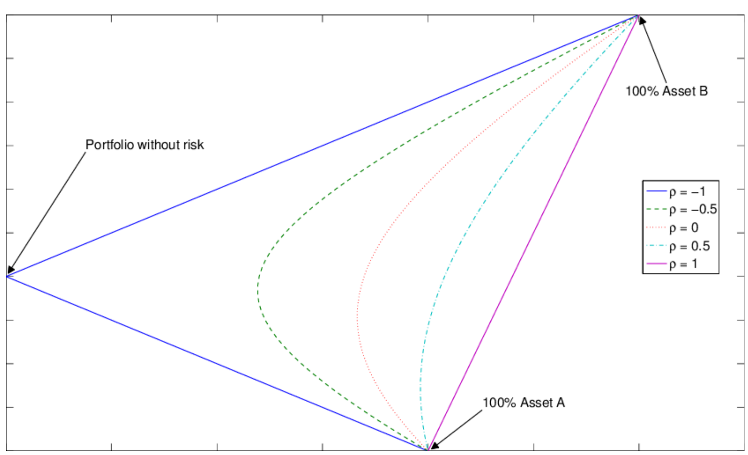

Figure 5: Risk-return trade–off in accordance with modern portfolio theory (Research Gate)

Traditionally, greater reward is always associated with increased risk, and the risk-reward trade-off could be represented by a linear graph with a positive gradient. However, modern portfolio theory contends that the risk-return trade-off follows a hyperbolic graph as shown above (Figure 5). Greater return is thus not necessarily indicative of greater risk, as two different portfolios can simultaneously exhibit similar risk profiles and yield largely different outcomes.

Consider a portfolio of n > 1 assets, each with a portfolio weighting of W_i, where the sum of the constituent weights, $W_1 + W_2 + \ldots + W_n = 1$. Adjusting the combination of individual asset weightings in the portfolio result in different risk-return characteristics for the portfolio.

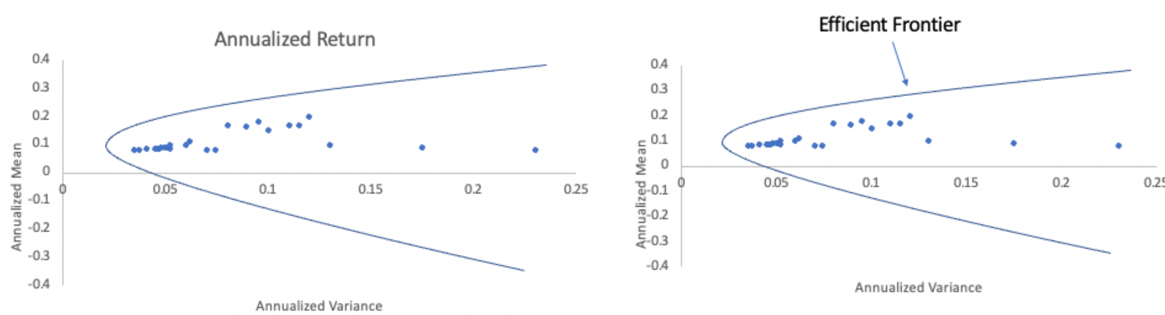

Figure 6: Annualized Returns given a portfolio of n > 1 assets

Figure 6 above illustrates a representation of the risk-return profiles of different portfolios, each comprising the same *n* assets, but with different weights allocated to each asset. The set of all possible compositions of portfolios forms a convex set in the mean-variance space, and the boundary of the set delineates a hyperbola. This boundary, within which every possible combination of risk-return profiles generated from n assets is contained, is known as the mean-variance frontier.

Notice on the boundary that there are two return values (y) given each variance (x), with each value of y representing the largest and smallest possible returns respectively. An investor who constantly looks to maximize his returns will select a portfolio on the upper bound. This upper half of the hyperbola is denoted as the *efficient frontier*, which consists of portfolios yielding the highest possible rates of return for a given level of risk. The objective of adjusting the weights of a portfolio's constituent assets is thus to maximize the return of the portfolio such that the portfolio lies on the efficient frontier.

<u>Tangency Portfolio</u>

The tangency portfolio is the point on the efficient frontier with the maximum Sharpe ratio any portfolio could possibly attain given n assets. The tangency portfolio assumes the existence of a risk-free asset. From this assumption, the tangency portfolio is derived by drawing a line from the point of intercept of the risk-free asset–which has zero variance and a given return–that is tangent to the mean-variance frontier. The point of tangency is the portfolio that has the maximal attainable Sharpe ratio from any combination of those n assets, as is called the tangency portfolio. This line, on the other hand, is known as the capital market line or the capital allocation line.

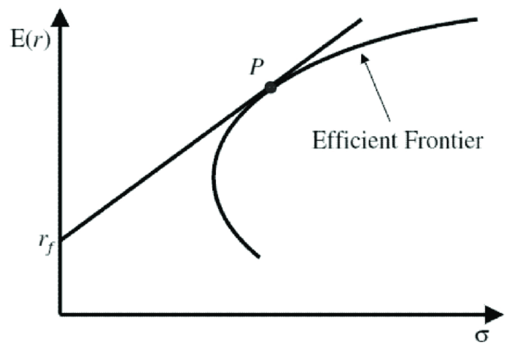

Figure 7: Tangency portfolio (Research Gate)

The capital market line represents a set of portfolios that optimize the risk-return relationship. The slope of the capital market line is equivalent to the Sharpe ratio of the tangency portfolio, meaning that any portfolio constructed on the capital market line holds the same Sharpe ratio as the tangency portfolio. Rational investors can choose to invest in a portfolio anywhere along the capital market line, including the tangency portfolio. Different points on the capital market line differ in that their levels of return and risk vary from each other. Investors can adjust their holdings on the capital market line by holding different combinations of the risk-free and risky assets to achieve targeted levels of risk and return.

The tangency portfolio assumes 100% holdings in risky assets. Holding a larger percentage of risk-free assets moves you left along the line–to a portfolio that has lower return and lower risk–while selling the risk-free asset (i.e. levering your holdings) moves you to the right of the capital market line that has higher risk and higher reward.

<u>Mean-Variance Derivation</u>

Suppose we have a portfolio of n risky assets.

Let *r* be an n-by-1 vector containing the historical return of each of these assets, and μ be another n-by-1 vector containing the mean return of each of the n assets. A vector can simply be understood as a list of values, each corresponding to a risky asset. Here, vectors *r* and μ are related by the following equation:

$$\mu = E[r]$$

Given matrices *r* and μ, the covariance matrix, Σ, could be derived from the following:

$$\Sigma = E[(r - \mu)(r - \mu)']$$

where $(r - \mu)'$ *is simply a vector transposed from* $(r - \mu)$.

The resulting covariance matrix, Σ, is simply the pairwise combinations of all possible covariances for the given portfolio of n assets. This could be best visualized using the equation below. Observe that values lying on the diagonal portion of the matrix (from top left down) represent the covariance of each asset and itself.

$$\Sigma =$$

	A	B	C	D
A	0.3	0.2	0.1	0.6
B	0.4	0.5	0.2	0.8
C	0.2	0.8	0.5	0.9
D	0.7	0.3	0.3	0.4

The covariance matrix is thus positive definite, and no asset is a linear function of the others.

Next, consider the existence of a risk-free asset with a return r^f. Given the same constraints above, an investor chooses a portfolio comprising the same n securities with n respective allocation weights, all of which could be transcribed into an n-by-1 vector denoted by *w*. Since the sum of the weights of the portfolio must be added to 1, we allocate to the risk-free asset a weight that is equal to $1 - w'\mathbf{1}$, where w' is the transpose of vector *w* and **1** is an n-by-1 vector with all values comprising 1. Mathematically, $1 - w'\mathbf{1}$ is equivalent to $1 - (w_1 + w_2 + w_3 + ... + w_n)$, with the second term being the result of a matrix multiplication of w' and **1**.

This effectively means that we allocate the remaining amount of the portfolio value to the risk-free asset after investing in all n assets.

Let μ^p denote the mean return on a portfolio given by the following:

$$\mu^p = (1 - w'\mathbf{1})r^f + w'\mu$$

The equation above could be rearranged as $\mu^p = r^f + w'(\mu - \mathbf{1}r^f)$.

Denoting excess returns to be $\tilde{\mu} = \mu - \mathbf{1}r^f$, we get a resulting mean return on a portfolio to be given by $\mu^p = r^f + w'\tilde{\mu}$, and the mean excess return to be given by $\tilde{\mu}^p = w'\tilde{\mu}$.

The return variance of the tangency could then be given by the following:

$$\sigma_p^2 = w'\Sigma w$$

It is important to note that the risk-free asset has a constant and guaranteed return. Moreover, this expected return (referred to as the risk-free rate) has zero variance and thus zero correlation with any security. Consequently, the variance of the tangency portfolio is only dependent on the weights and covariances of the *n* securities.

<u>Mean-Variance Formula</u>

The optimal combination of weightings for a mean-variance portfolio can be given by the following:

$$w^t = \left(\frac{1}{1'\Sigma^{-1}\tilde{\mu}}\right)\Sigma^{-1}\tilde{\mu}$$

where w^t is an n-by-1 vector, and is essentially a "list" of portfolio weightings allocated to each individual risky asset respectively.

When funds are allocated in accordance with these weightings, an investor could expect a portfolio with the optimal in-sample Sharpe ratio.

Two-Fund Separation

The two-fund separation theorem states that an investor can make their allocation decision by first obtaining a tangency portfolio that maximizes their Sharpe Ratio, followed by adjusting the mix of the tangency portfolio and risk-free asset to suit their risk appetite. Essentially, every mean-variance portfolio is a combination of the tangency portfolio and the risk-free asset.

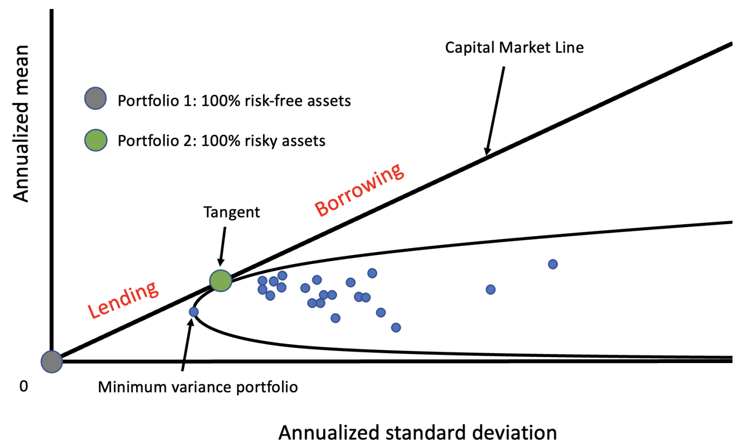

Figure 8: Two-fund Portfolio

Recall from previous sections that the tangential point of intersection between the capital market line and efficient frontier marks the tangency portfolio with the optimal Sharpe ratio, and all portfolios along the capital market line represent portfolios with the same Sharpe ratio (Figure 8). All these portfolios on the capital market line are essentially a different mix of risky (i.e. the tangency portfolio) and risk free assets, such that the tangential point is an example of two-fund separation with 100% risky assets and 0% risk-free assets, while the origin is an example of two-fund separation with 0% risky assets and 100% risk-free assets.

Notice that the capital market line could be split into the “Lending” portion and the “Borrowing” portion. The lending portion is where the investor holds a positive percentage of risk-free assets and is thus extending a “loan” to the US government in exchange for interest payments at the risk-free rate. For portfolios lying on the borrowing portion of the capital market line, their portfolio essentially consists of more than 100% risky assets, and less than 0% risk-free assets (with the sum of both adding up to a 100%). Here, the investor essentially borrows money at the risk free rate to increase his investments in risky assets and magnify his expected returns. Levered portfolios, however, scale up both risk and return in the same proportion.

Two fund separation essentially allows investors to construct a portfolio on the mean-variance frontier, but adjust their holdings to suit their level of risk appetite. This adjustment could simply be made by multiplying the tangency portfolio by a scalar factor. The resulting portfolio is then a point on the capital market line between the origin and tangential point with an adjusted mean return and variation.

This scalar factor could simply be obtained by setting a target mean return for the portfolio and substituting this value into the following equation:

$$\delta = \left(\frac{1'\Sigma^{-1}\tilde{\mu}}{(\tilde{\mu})'\Sigma^{-1}\tilde{\mu}}\right)\tilde{\mu}^{p}$$

As such, adjusted portfolio weights, w*, is given by:

$$w^{*} = \delta\, w^{t}$$

<u>Limitations of the mean-variance theorem</u>
While the mean-variance theorem appears to be a fool-proof method of investing, there are a few limitations to its practicality.

Firstly, mean-variance is extremely sensitized to in-sample data, which limits its ability to make optimal portfolio decisions given new sets of stock market returns. In data science, models are built upon a set of training data, and are tested using testing data sets. When we use stock market data from 2015 to 2020 to build a model and test the same model using stock market data from 2020 to 2022, there is a non-negligible and oftentimes sizable difference between in-sample and out-of-sample performance. This is typically caused by overfitting, where a model and its associated parameters are perfectly optimized and sensitized to the minutiae of its training data set. Performance differences are, as a consequence, pronounced because such hypersensitivity to in-sample data render it incapable of making good generalizations for testing data sets. As such, while counterintuitive, poorer in-sample performance by less-sensitized models often result in better out-of-sample performance. This concept could be applied to mean-variance analyses. While the mathematical formula provides optimized weights for in-sample data used, it may not be able to make accurate, forward projections of optimized portfolio weights to generate desired future returns given the pronounced differences in market environments across different time periods.

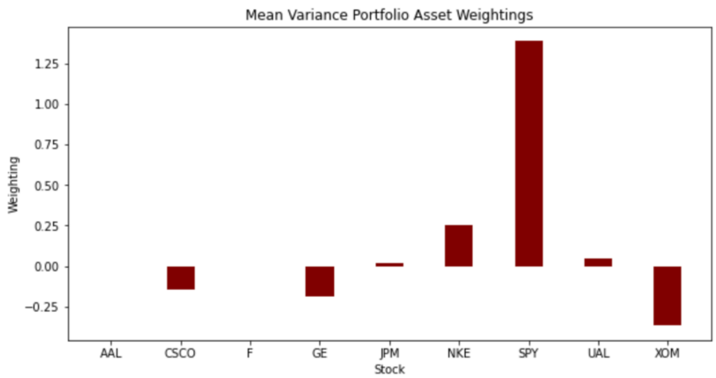

AAL	-0.004421
CSCO	-0.146570
F	-0.000828
GE	-0.189041
JPM	0.018942
NKE	0.252251
SPY	1.390724
UAL	0.048748
XOM	-0.369806

Figure 9: Sample of a mean-variance portfolio

A second limitation is that portfolio weights derived from the mean variance formula are often incompatible with constraints investors face, thus rendering such investment decisions either impractical or outright infeasible. Consider an optimal portfolio that stipulates a 1.3 million investment in security SPY and a 0.6 million short position on another security. If an investor has one million cash on hand, he would be required to lever his portfolio by borrowing 0.3 million to buy SPY. Furthermore, taking a short position would typically entail maintaining a margin account with a value that is 150% of the short position. Given the significance of this short position, unfavorable stock market movement could risk a margin call to his financial detriment.

Portfolio managers thus employ various means of enhancing the practical value of mean-variance analyses. For example, Harvard's investment office applies bounds on its investment holdings to eliminate short positions and limit the size of their long positions.

One method investors use to reduce sensitivity to in-sample data is to diagonalize the covariance matrix (Figure 10).

	A	B	C	D
A	0.3	0.2	0.1	0.6
B	0.4	0.5	0.2	0.8
C	0.2	0.8	0.5	0.9
D	0.7	0.3	0.3	0.4

→

	A	B	C	D
A	0.4	0.0	0.0	0.0
B	0.0	0.5	0.0	0.0
C	0.0	0.0	0.8	0.0
D	0.0	0.0	0.0	0.5

Figure 10: Diagonalization of Covariance Matrices

Previously, mean-variance analyses we adopted rely on the pairwise covariance matrix as the sole measure of portfolio volatility. However, by diagonalizing the matrix, we zero out all the covariances except for those between the asset and itself. This renders the model less sensitive to in-sample data because data points are being limited to a select few. Short positions are also eliminated as well because negative covariance values are effectively removed.

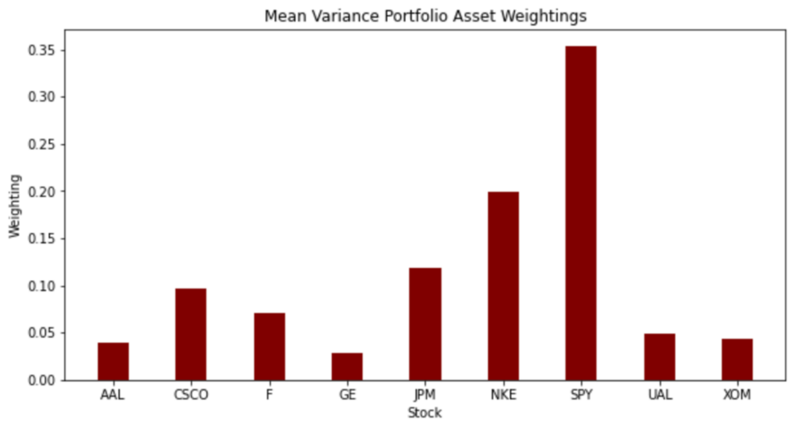

Figure 11: Mean-variance allocation outcome using a diagonalized covariance matrix

As shown in Figure 11 (above), portfolio allocations are typically more feasible and exhibit better out-of-sample performance because of reduced sensitivity to in-sample data.

References

Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance*, *7*(1), 77–91.

Efficient multi-energy generation portfolios for the future - researchgate. (n.d.). Retrieved December 20, 2022, from https://www.researchgate.net/publication/229005093_Efficient_multi-energy_generation_portfolios_for_the_future