

# Strategic asset allocation from theory to practice: new decision-support tools\*

Golan Benita and David Hoffman\*\*

## Abstract

The mean-variance (MV) model is widely used in portfolio management, particularly when choosing a strategic asset allocation (SAA) benchmark. But it suffers from several drawbacks, such as sensitivity to small changes in the estimates of mean returns and concentration in a small number of assets. Therefore, when portfolio managers use the MV model as part of their SAA process, they usually do not accept its results “as-is”, but rather make adjustments to them – adjustments which are often based on arbitrary rules or intuition. We describe two decision-support tools that were developed at the Bank of Israel which can help portfolio managers decide whether such adjustments are needed, and examine how they would affect the SAA benchmark’s return in various scenarios. Both tools use a forecast of financial market developments, referred to as the base scenario, to calculate the expected returns provided to the MV model. The first tool identifies a few portfolios which are very similar to the MV optimal portfolio in risk and expected return under the base scenario, but which have quite different exposures to financial market risk factors, including duration, credit spreads and equity returns. The performance of these alternative portfolios can then be compared to that of the MV optimal portfolio under market scenarios different from the base scenario. The second tool allows the portfolio manager to see whether, and by how much, small changes in the base scenario’s financial market forecasts would cause large changes in the makeup of the MV optimal portfolio, and also to see how such differences would affect the portfolio’s return. The information provided by these two decision-support tools gives portfolio managers a more solid quantitative foundation upon which to base their decisions about how to deviate from the MV optimal portfolio in setting their SAA benchmark.

JEL classification: G110, G170.

Keywords: asset allocation, investment decisions, portfolio choice, risk analysis.

\* The views expressed in this paper are those of the authors, and do not necessarily reflect the position of the Bank of Israel. This version: October 2019.

\*\* David Hoffman, Markets Department, Bank of Israel, POB 780, Jerusalem 91007, Israel. Phone: 972-2-655-2385. Fax: 972-2-666-9385. Email: [david.hoffman@boi.org.il](mailto:david.hoffman@boi.org.il).

## 1. Introduction

One of the pillars of modern finance is the mean-variance (MV) model developed by Markowitz (1952, 1956). The MV model is widely used in academic research and among practitioners, *inter alia* for selecting an investment portfolio's strategic asset allocation.<sup>1</sup> However, the model has several drawbacks, of which the two most well known are the following: First, the composition of the portfolios located on the efficient frontier is very sensitive to the assumptions made about the distributions of return-on-assets, particularly about their means. Second, imposing a "no short sell" constraint – ie requiring the weight of every asset in the portfolio to be positive – often causes the model to generate an efficient frontier comprised of portfolios that include a very small number of assets.

If the future distribution of return-on-assets were known with certainty, these drawbacks would be less severe. However, in practice, there is usually considerable uncertainty regarding the nature of this distribution – particularly about the means of return-on-assets. The uncertainty about the true mean of the distribution, coupled with the tendency of small changes in means to lead to large changes in portfolio composition, often results in asset allocations concentrated in a small number of assets. This magnifies the potential impact of any error, which means the results of an MV optimisation cannot be implemented as the asset allocation of an investment portfolio without first being subject to critical examination.

There may also be uncertainty about the target level of risk appropriate for a given investor – which often is best expressed as a narrow range, rather than as a precise value – and about the choice of risk measure. Although portfolio risk is often summarised by a single risk measure (such as volatility, the risk measure used by the MV model), its level is usually a function of the portfolio's exposure to a variety of risk factors. For example, central bank foreign exchange reserves portfolios may be exposed to four main risk factors: duration, credit, equity and currency. In such portfolios, very similar combinations of risk and return can sometimes be obtained with different combinations of exposures to these risk factors. For instance, there may be two portfolios with similar levels of expected return and volatility over a one-year horizon, where the volatility of one is mainly due to its allocation to equities, while that of the other is mainly due to duration exposure (ie a higher interest rate risk).

Faced with a choice between the two, an investment manager may prefer the portfolio that is exposed to factors that they believe have *less uncertainty*. In addition, if the manager found these or any other portfolios on the frontier to be concentrated in very few assets, they may go as far as finding an additional *portfolio* that lies *slightly below* the efficient frontier, but has greater diversification benefits. The latter allocation would have a level of risk that is only slightly different from the target, but would place weights on a greater number of assets.

<sup>1</sup> According to the capital asset pricing model (Sharpe (1964), Lintner (1965, 1969)), all investors should select the market portfolio and mix it with a riskless asset according to their risk preferences. However, in practice there is no pure investable riskless asset available to large-scale investors with relatively long investment horizons. For instance, the appropriate riskless asset for an investor who has a one-year investment horizon should be a 12-month Treasury bill. However, large-scale investors cannot practically invest such a high proportion of their portfolio in a single security. Given that there is, in practice, no riskless asset, the optimal portfolio should be selected from the efficient frontier according to the investor's risk preferences.

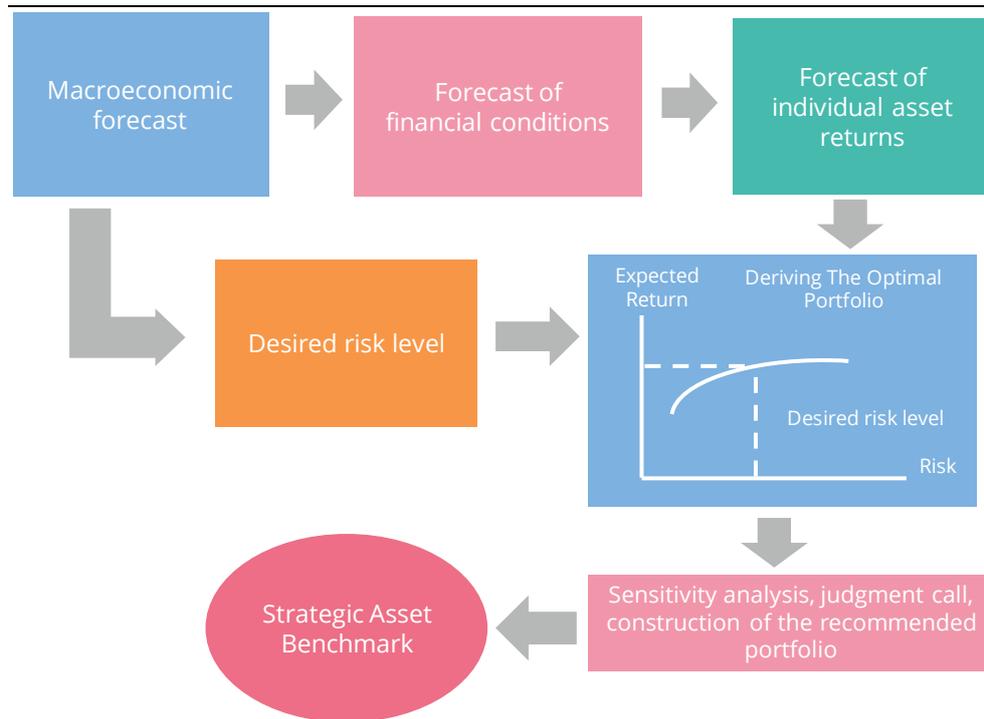
By identifying portfolios using these notions, the investment manager could reduce their exposure to uncertainty about the future distribution of return-on-assets, and may be able to select a strategic asset allocation (SAA) that appears more appropriate than the one obtained by simply and mechanically applying the MV model to solve the problem.

As understood in this paper, the SAA decision is a structured decision-making process, conducted at regular intervals (eg annually or semi-annually). During the process, inputs such as economic and financial forecasts are chosen by the portfolio manager, which are used to select an allocation of funds among asset classes available for investment. This allocation is kept as a *benchmark* portfolio until the next SAA decision is made, and serves as a comparison for the *actual* portfolio of assets held. The portfolio manager may manage against this benchmark by either passively tracking it, or with some latitude allowed to deviate from its composition (in other words, to engage in active management).

A schematic view of one approach to SAA is given in Figure 1. In this example, the process starts with a macroeconomic forecast, which then is mapped both to a forecast of how conditions in the financial markets, eg bond yields, will change over the coming year, and to the choice of a target risk level that the stewards or stakeholders of the portfolio feel is appropriate in the forecasted macroeconomic conditions. The issues involved in formulating such a macroeconomic forecast, and in mapping it into a forecast of financial conditions and an appropriate choice of risk level, are outside the scope of this paper. However, the uncertainty about whether such a forecast will be realised as anticipated, and the difficulty in expressing such uncertainty quantitatively, underlines the importance of taking into account the uncertainty about the means of return-on-assets distributions and about the appropriate level at which to set the target of the risk measure.

Example of a strategic asset allocation process

Figure 1



Source: Bank of Israel.

For a given distribution of asset returns and a target risk level, the MV model can be used to derive the efficient frontier, as shown in the blue rectangle of Figure 1. The composition of the portfolio on the efficient frontier for the target risk level can then be identified as theoretically (ex-ante) optimal. However, due to the aforementioned limitations of the MV model, the SAA decision-making process does not normally end there. In other words, the MV model serves as a decision-support tool, not a decision-making tool. In most cases, the portfolio obtained from it will be modified in ways that reduce the impact of its shortcomings on the result – for example, by introducing additional diversification in an ad hoc manner (ie personalising the weight of each asset on the optimal solution). Such modifications may be supported by analyses such as stress testing, but the final call is often made on the basis of the portfolio manager’s intuition or knowledge, which may sometimes be perceived as arbitrary.

To improve the quality of the decisions made in this final step, we have developed two new decision-support tools, which are described in this paper.

1. **The near-optimal analysis (NOA) tool**, which allows the investment manager to compare the MV optimal portfolio with other portfolios that have similar risk and return characteristics but different exposures to three of the four risk factors listed above – ie duration, credit and equity.
2. **The forecast sensitivity analysis (FSA) tool**, which allows the investment manager to analyse the sensitivity of the composition of the ex-ante *optimal* portfolio against changes in the forecasts of risk factors over reasonable ranges. Mapping the optimal allocation over a range of values allows the investment manager to examine how outcomes of the financial variables that are different from those in their forecast would affect the asset allocation chosen for the portfolio.

By allowing the investment manager to explore the trade-off between the various asset classes and different risk-return assumptions, these tools can be of help in selecting an SAA that is close to optimal under given assumptions, while having less exposure than the theoretically optimal portfolio to uncertainty about how accurate those assumptions are.

The approach suggested is not intended to provide a single, unequivocal answer to the question of how to invest, but rather to give the portfolio manager additional information, which will allow the final call to be based less on intuition and more on a solid quantitative foundation.

The remainder of the paper is organised as follows. In Section 2, we describe a simple case study to help illustrate how these two tools are used. In Sections 3 and 4, respectively, we describe how each tool works, and give an example of its results, based on the case study. Section 5 discusses how the results of the two tools together can be used to gain insights that may influence the SAA decision in the case study. Section 6 concludes the paper.

## 2. The case study

We consider the decision of a portfolio manager who has nine assets available to invest: three US Treasury indices with maturity ranges of 0–1, 1–5 and 1–10 years, respectively; three European domestic government bond indices of the same maturities; an index of investment grade (IG) corporate bonds issued in the United States; an index of IG euro-denominated corporate bonds; and a multicurrency, developed market equity index. The portfolio manager has no desire to expose the portfolio to currencies other than the US dollar and the euro, which entails that a proportion of the equity's currency risk will be hedged as part of the asset allocation decision.

The optimal portfolio must conform to the following constraints: its currency composition must be 65% in US dollars and 35% in euros; the allocation to non-government securities (equity and corporate bonds) cannot exceed 25%; short positions and leverage are not allowed; and finally, the risk of the portfolio, as measured by the volatility (standard deviation) of returns, must fall in the range between 1.95 to 2.05% – this is the portfolio manager's target risk level range. The goal, then, is to find the portfolio with the highest level of return for this level of risk. This is an optimisation problem, which in theory could be solved by applying the traditional Markowitz MV model with restrictions.

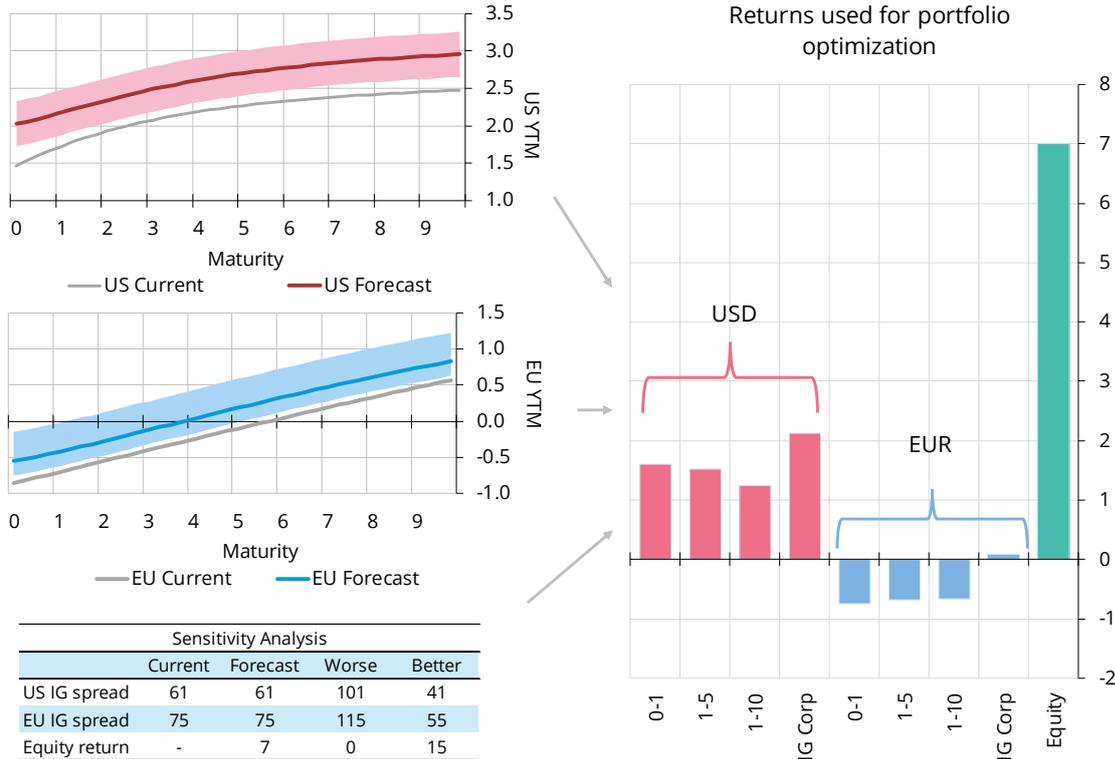
In order to use the Markowitz MV model to select an optimal portfolio, assumptions are required about the parameters of the distribution. For example, to use a multivariate normal distribution to model the returns of the aforementioned assets, assumptions on the vector of expected returns and their variance-covariance matrix are required.<sup>2</sup>

<sup>2</sup> The tools we have developed do not address the issues involved in macroeconomic forecasting or the conversion of a macroeconomic forecast into a forecast of financial market outcomes, as represented by the first two boxes in the top row of Figure 1, which, as previously stated, are outside the scope of this paper. Nor do they address the problem of defining the variance-covariance index, which is an exogenous input. However, they do include a process by which a financial forecast can

These elements may be, in practice, derived from other inputs such as (i) forecasts for the government yield curves throughout the projection period specified for the SAA process (in this example, one year ahead), together with their current values; (ii) forecasts for the average credit spreads of the two IG corporate bond indexes that are included in the available assets, together with their current values; and (iii) a forecast for the return of the equity index. Discussing the generation of these inputs, and how these are translated into asset returns is out of the scope of this paper, but the process is illustrated in Figure 2.<sup>3,4</sup>

Financial market and asset return forecasts

Figure 2



Source: Bank of Israel.

be directly translated into expected returns on specific assets, reflecting the move from the second (purple) box to the third (green) box. For the FSA tool, this ability is not just a convenience, but is integral to its functioning, as will be explained.

- <sup>3</sup> The FSA tool also requires a range of uncertainty around each of the forecasts as part of its input. This is shown by the shaded area around the yield curves in Figure 2, and by the two rightmost columns in the table at the bottom left corner of the figure. As with the forecasts themselves, setting these ranges belongs to the stage of formulating a macroeconomic forecast and translating it into a forecast of financial conditions, which has already been done before the tools we are presenting come into play.
- <sup>4</sup> The process of translating financial forecasts into asset return forecasts is computationally intensive for the bond indexes, but simple in its conception. The transition from the initial to the final values of the yield curves and credit spreads is assumed to occur at a constant rate over the course of the year. We also assume that the makeup of the index is the same at the beginning of each month throughout the year. We can then use the forecast yield-to-maturity of each bond, its weight in the index and its duration to calculate monthly holding-period returns of the index, the geometric sum of which gives its annual holding-period return.

By using the returns shown on the right-hand panel of Figure 2 as the expected-returns vector, and supplying a variance-covariance matrix for returns on these assets, it is straightforward to find the efficient frontier of portfolios and identify the MV optimal portfolio at our target level of risk. In what follows, we describe how the NOA and the FSA may aid the decision-maker by improving this process.

### 3. The near-optimal analysis (NOA) tool

To address the issues that arise when using the traditional MV model, we will identify other portfolios that have similar risk and return to the “efficient” or “optimal” one, but, insofar as possible, different exposures to duration risk, credit risk and equity risk. This is done in four steps:

- (1) Identify the MV optimal portfolio at our target level of risk, and take note of its expected return.
- (2) Identify a large number of portfolios that are below the efficient frontier, but close to it.
- (3) From among the portfolios found in Step 2, select those that have risk and return approximately equal to that of the MV optimal portfolio identified in Step 1.
- (4) From among the subset of nearly optimal portfolios found in Step 3, identify a few that are as different as possible, both from the MV optimal portfolio and from one another, in terms of the three risk factors that are of interest.

Step 1 requires the application of the Markowitz MV model to derive the efficient frontier using the expected-returns vector, the variance-covariance matrix and subject to the constraints imposed. The portfolio that has the required standard deviation – the target risk level – is then identified. In the case study, this was defined as a range between 1.95 and 2.05%. For this first step, we choose a standard deviation at the middle of this range, that is, 2.0%. The expected return of this portfolio is found to be 1.72%.

To carry out Step 2, we repeat Step 1 many times, introducing random fluctuations into the expected return vector each time. To generate such random fluctuations, we define a range where the return of each asset may fall, modelled through a uniform distribution. The range where this uniform distribution is defined (upper and lower bound) can be derived in many ways.<sup>5</sup> For this case study, we calculate them as 75% and 125% of the base expected return for the lower and upper bound, respectively.<sup>6</sup> Using a different expected return vector every time allows us to derive a quasi-efficient frontier in every iteration. However, these new portfolios’ expected returns are then calculated using the base expected return for every asset, which, by construction, places them on or below the efficient frontier. This is repeated

<sup>5</sup> Some other possibilities to consider include (i) ranges based on each asset’s standard deviation of return, of some fixed number of standard deviations above and below the base mean; and (ii) ranges above and below the base mean set in absolute terms by the analyst as a judgement call, either per asset or by groups of assets.

<sup>6</sup> These base expected returns were those derived from our financial forecasts, which were used to find the original efficient frontier itself.

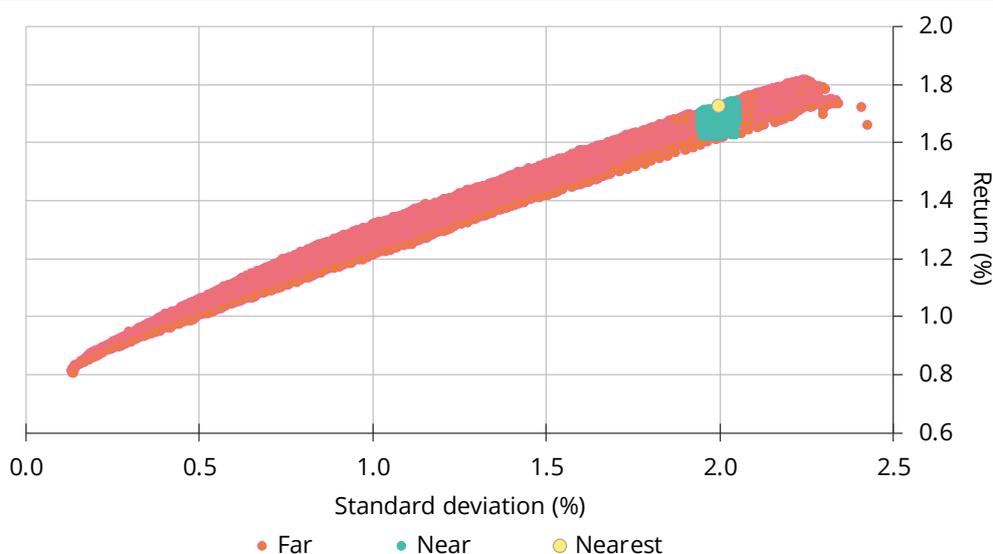
until we have accumulated a number of simulations sufficient for empirical analysis. In this case study, 25,000 sub-optimal portfolios were used.<sup>7</sup>

In Step 3, we define the maximum allowed differences between a portfolio's expected return and standard deviation and those of the MV optimal portfolio. We have earlier specified a standard deviation target in the range of 1.95 to 2.05%. So we look for portfolios with plus/minus five basis points of volatility around the standard deviation of 2.0. For the return, in this case, we choose a lower bound of eight basis points below the return of the MV optimal portfolio found in Step 1, which is assumed to be acceptable by the decision-makers. The 25,000 portfolios found in Step 2 are then examined, and only those meeting the bounds defined around the return and the volatility of the MV optimal portfolio are retained. In our case, this set includes about 1,000 portfolios.

Figure 3 illustrates Steps 1 to 3 graphically. The pink region represents the outcome of Step 2, and the green one, the outcome of Step 3. Note that the actual number of portfolios identified in Step 2 is larger than those plotted in the graph; for greater readability, and due to the limits of the software graphics package that we used, those beyond a certain distance from the efficient frontier are omitted.

Portfolios near the efficient frontier

Figure 3



Source: Bank of Israel.

Step 4 calls for us to find a *small* number of portfolios from among the roughly 1,000 kept in Step 3 that are quite different from one another in terms of the three

<sup>7</sup> This procedure is similar in spirit to that employed by Michaud and Michaud (2008), but applied somewhat differently. In their model, the goal is to find a near-efficient frontier, in which each portfolio consists of a larger number of assets than there are in the portfolios on the actual efficient frontier. Ours is to identify portfolios that do not necessarily contain more assets than those on the efficient frontier, but do have different risk exposures. For this reason, we introduce a greater degree of randomness, so that a part of the region in the risk-return space bounded above by the efficient frontier is filled with sub-optimal portfolios. Also, we do not average asset allocations across portfolios along the whole length of the efficient frontier, as they do, but rather consider only the final few portfolios found in Step 4 as explicit alternatives, both to the MV optimal portfolio and to each other.

key risk factors that interest us: duration, equity and credit. To do this, we first calculate the exposure of each portfolio to these factors. Each exposure is expressed as follows: portfolio duration in years, exposure to equity in percentage, and exposure to credit in percentage allocated to corporates.<sup>8</sup>

Having thus associated each portfolio with a set of three numbers – which can be thought of as a location vector in three-dimensional space – we next need to calculate the difference, or “distance”, of any one portfolio from another. A natural choice would be the familiar Euclidean distance, found by subtracting one vector from the other, and then squaring the numbers obtained in this manner and adding the squared values together. However, two other distance measures provide results, which we have found more satisfactory.<sup>9</sup>

In the first of these, distance is defined by adding the absolute values of the three values obtained by subtracting one vector from the other. In other words, the distance between points **x** and **y** is calculated as:

$$d(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n |x_i - y_i|$$

This measure is sometimes referred to as the “Manhattan distance”, since it represents the distance travelled between two points when one is constrained to follow a grid of equidistant, perpendicular paths, as in a stylised version of the street map of Manhattan.

In the second measure, distance is defined by the formula:

$$d(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^n \frac{|x_i - y_i|}{|x_i + y_i|}$$

Values of *i* for which the denominator is zero are ignored. This measure is referred to as the “Canberra distance”, for reasons unknown to the authors. Note that the three measures of distance listed above always have positive values.<sup>10</sup>

The distances between all possible pairs of portfolios are calculated using whichever one of the distance measures discussed above is selected. The analysis then proceeds as follows:

- a) First, designate the MV optimal portfolio as portfolio 1.
- b) Then, find the portfolio that is farthest from the MV optimal portfolio; call this portfolio 2.
- c) For every other portfolio, add together its distance to portfolio 1 and its distance to portfolio 2. Call the portfolio with the greatest combined distance portfolio 3.
- d) Continuing in the same manner, portfolio 4 will be that for which the sum of the distances to portfolios 1 to 3 is greatest; 5 will be the one for which the sum of

<sup>8</sup> In principle, these should then be multiplied by factors that equate the profit or loss from a one standard deviation price change for each type of risk. In practice, we have found that using the unmodified values gives satisfactory results.

<sup>9</sup> This is, of course, a subjective impression with which other practitioners are free to differ.

<sup>10</sup> See Bodenhofer et al (2014) pp 37–42 for further discussion of these and other distance measures.

the distances to portfolios 1–4 is greatest, and so on, for as many alternative portfolios as we wish to consider. We choose four portfolios in our case.

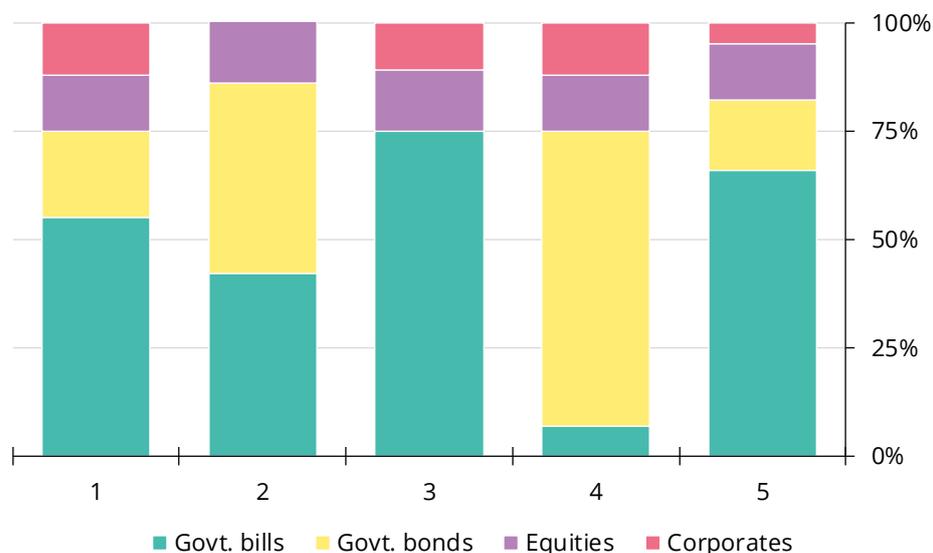
Examining the results of our analysis for the data in our case study (shown in Figure 4) we see that the means and standard deviations of the four alternative portfolios identified are very close to those of the MV optimal portfolio. However, their durations span a broad range, from 0.6 to 2.5 years, and so do their allocations to corporate bonds, which are between 0 and 12%. By contrast, the range of allocations to equity is narrow, from 13 to 15%. From this we can infer that, given the forecasts of asset returns input, there is not any alternative portfolio, having risk and expected return close to those of the MV optimal one, which does not include an allocation to equity within or very near this range.

As an example, it is instructive to compare the two leftmost portfolios (in Figure 4, labeled 1 and 2), which are the MV optimal portfolio and the first alternative identified. If the managers of a portfolio are reasonably comfortable with the quality of their forecasts for equity returns and yield-curve development, but have less confidence in their forecast for corporate credit spreads, they may want to consider a lower allocation to corporates, or perhaps not investing in them at all, while slightly increasing duration and equity exposure to make up the difference in total expected return. The NOA tool's results show that they can do this without deviating greatly from the risk-return profile of the MV optimal portfolio.

It is also possible to examine the performance of the five portfolios that were chosen under alternative scenarios, or to perform stress testing on them. For example, if it turns out that in a reasonably plausible alternative scenario, one of the four alternative portfolios has a slightly positive return, while that of the MV optimal portfolio is significantly negative, an investor who is more sensitive to losses than gains may prefer the alternative portfolio to the original MV optimal portfolio.

Characteristics of case study portfolios identified by the NOA tool

Figure 4



Summary Statistics					
Mean	1.72	1.68	1.71	1.65	1.66
Std Dev	2	2.04	2.04	2.03	1.95
Duration	1.2	1.5	0.6	2.5	1.2
Govt bills	55	42	75	7	66
Govt bonds	20	44	0	68	16
Equities	13	15	14	13	13
Corporates	12	0	11	12	5

Source: Bank of Israel.

#### 4. The forecast sensitivity analysis (FSA) tool

Analysis done in the second step of the NOA tool can provide some insight into how sensitive portfolio allocations using the MV algorithm are to small changes in the vector of expected returns. Nonetheless, the main purpose of the NOA is to offer portfolio alternatives for the decision-maker under the baseline scenario. The FSA tool was developed to more formally explore the sensitivity of the MV algorithm's portfolio allocation to such deviations, within limited ranges. The concern is that small differences in expected returns – which can be interpreted as small deviations of realised financial conditions from the forecast – can lead to a large difference in the composition of the MV optimal portfolio.

Performing the FSA requires that we define bounds of uncertainty for the forecasts of the risk factors. For example, our uncertainty about the direction of yield-curve factors, corporate credit spreads and the return on equity. This is exemplified by the shaded areas above and below each yield curve in Figure 2, and the values labeled "Sensitivity Analysis" in the table at the lower left of the figure.

These bounds reflect estimates about the degree of uncertainty that exists regarding financial market outcomes as implied by the base macroeconomic scenario, rather than the degree of uncertainty on the unconditional realisation of financial market outcomes. In other words, the lack of confidence of the portfolio manager is not on their mapping between macroeconomic and financial variables (which is out of the scope of this paper), but on the realisation of its mean expected path for macroeconomic conditions. As such, the range of possible financial variable outcomes may be narrower than what might be indicated by simply observing the historical (unconditional) distribution of data. It is important to note that it is not necessary for the upper and lower bounds of the outcomes to be symmetric around the mean forecast, nor do they need to be the same range (in percentage points, for example) for similar instruments in different markets.

Intuitively, it is possible to classify the uncertainty about each of the financial market variables that are used to forecast asset returns of this case study as primarily due to one of two risk factors: yield curve risk or business cycle risk. Thus, uncertainties about whether government-bond yield curves will be higher or lower than forecast are classified as yield curve risk. Uncertainties about equity returns and credit spreads are classified as business cycle risk, since business conditions that are better than forecast could lead to a higher equity return and narrower credit spread than expected, while business conditions worse than forecast could lead to the opposite.

The FSA process involves defining a large number of scenarios in a bounded two-dimensional space, in which financial market variables classified as yield curve risk vary in the vertical direction (with higher yields being higher along the vertical axis), while financial market variables classified as business cycle risk vary in the horizontal direction (with higher business conditions being farther to the right along the horizontal axis). Each financial market variable is then mapped to its respective axis, scaled according to the range of its bounds of uncertainty.

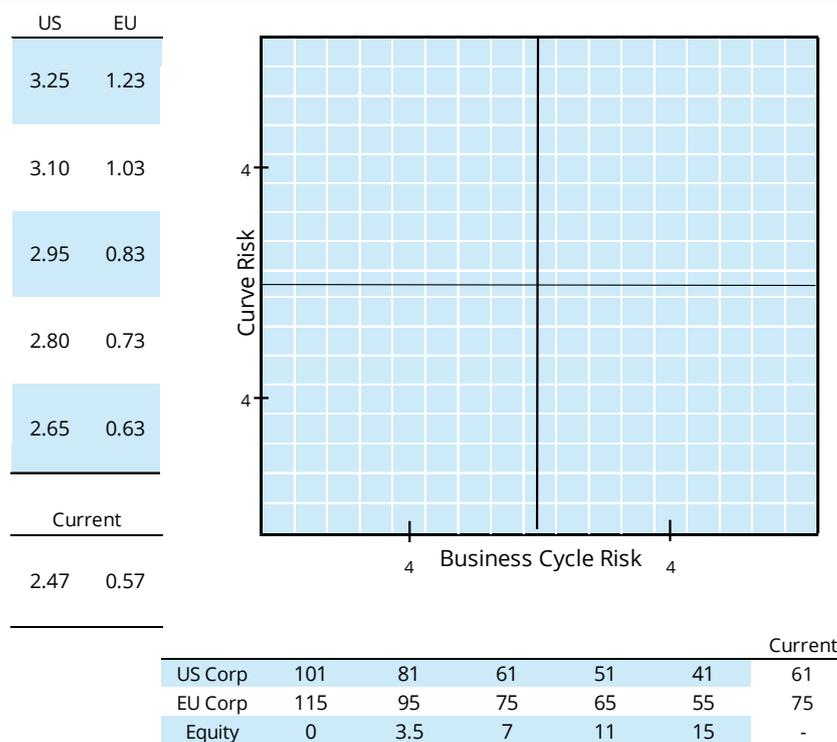
For example, if the range of uncertainty regarding the credit spreads of US IG corporate bonds is defined as 20 basis points (bp) below and 40 bp above the base scenario's forecast, it will be mapped along the horizontal axis from the centre to the left bound as going from 0 to +40 bp, and from the centre to the right bound as going from 0 to -20 bp. In the uppermost of the three horizontal scales in Figure 5, below, this corresponds to the transition from 61 bp (the base forecast) to 101 bp (base + 40) when moving leftwards from the centre line, and from 61 bp to 41 bp (base - 20) when moving rightwards. Implicit in this procedure is the assumption that changes in variables on the same axis (eg US and European credit spreads), move in the same direction and also by the same degree, *relative to the sizes of their bounds of uncertainty in the given direction*.<sup>11</sup> It is important to emphasise that, in each alternative scenario, the stepwise forecasts of financial variables over the uncertainty range refer to their ex ante average ("expected") values, *not* their end-of-period realised values.

To allow the creation of a finite number of alternative scenarios, each axis is partitioned into a fixed number of steps above and below the centre, which represents the base case value. The number of steps used can be chosen to partition the space into as many scenarios as the manager feels are useful for the analysis. In our work as practitioners, we have found eight steps above and below the base case (for a total of 17, including the base case) to be most convenient, and this is the

<sup>11</sup> This simplifying assumption is necessary for the results of the analysis to be shown in two dimensions.

number used in the case study. Division of each axis into eight steps above and below the base forecast in this manner defines a grid of 288 scenarios, represented by the small squares in Figure 5. The square in the exact centre of Figure 5 represents no deviation from the base forecast for any financial variable, and all others are scenarios in which one or more do deviate from their forecast values.

Grid of scenarios created by the forecast sensitivity analysis tool Figure 5



Source: Bank of Israel.

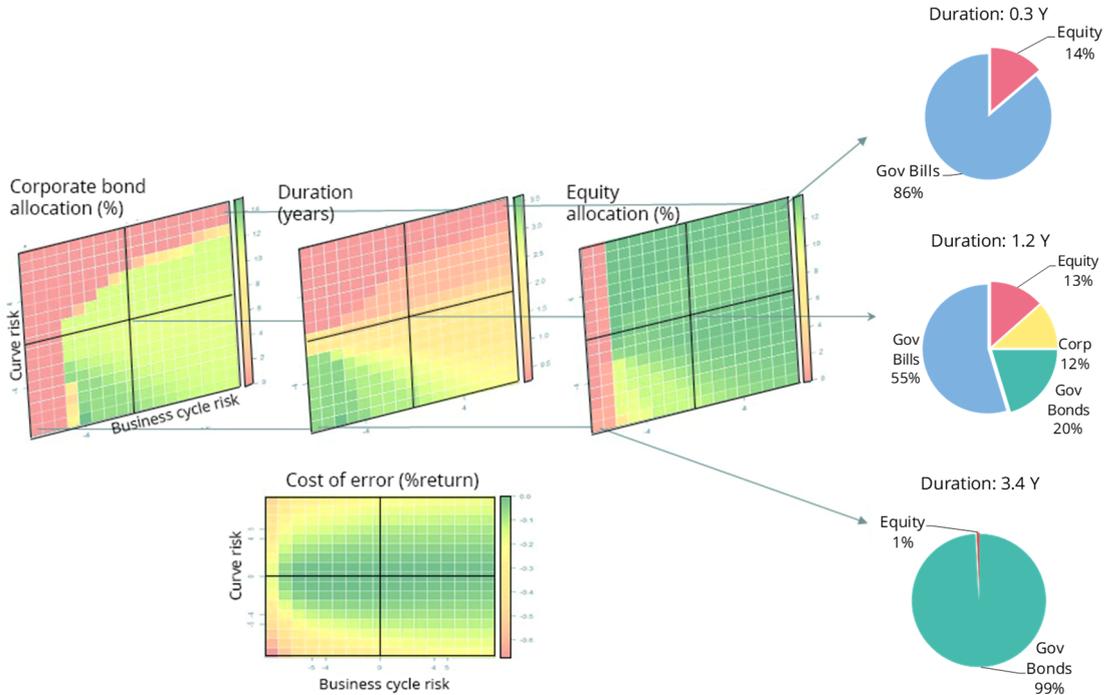
In the scales on the horizontal and vertical axes of the grid in Figure 5, only positions plus 4 and minus 4, out of the range  $[-8, +8]$ , are marked. However, the tables in the horizontal and vertical margins of the grid give the values of selected financial variables at positions  $-8, -4, 0, +4$  and  $+8$ , going from left to right on the horizontal scale and from bottom to top on the vertical scale, as well as their values at the beginning of the projection period (labeled "Current"). For the business-cycle risk factor on the horizontal axis, all three financial variables are given. For the curve risk factor on the vertical axis, the levels of 10-year yields-to-maturity are given as a proxy for the level of the yield curve as a whole.

For each of the 288 scenarios represented by the small squares in the grid, the assumed values of financial market variables are translated into a vector of asset returns, using the method outlined in Section 2, above. The Markowitz MV model is then used to derive the efficient frontier, given the available assets' characteristics, and subject to the constraints imposed, and the MV optimal portfolio at the target risk level is identified. This MV optimal portfolio, representing a possible asset allocation, is associated with the respective grid square. The process is run 288 times. Finally, we summarise each portfolio by the numerical values of the three risk factors we are interested in – its duration, its allocation to corporate bonds and its allocation to equity – just as was done when using the NOA tool. Rather than using these values

as a position vector with which to define the distance between one portfolio and another, we use them to create heat maps, as shown in Figure 6. A brighter shade of red denotes a lower allocation, while a shade of green denotes a higher allocation to a given risk factor.

Forecast sensitivity analysis results for the case study

Figure 6



Source: Bank of Israel.

Looking at the rightmost of the three heat maps in the upper row of Figure 6, we see that the MV optimal portfolio’s allocation to equity is very similar over a broad range of possible scenarios. Only if the expected return on equity approaches zero does it fall off, and then it does so dramatically. This is consistent with the results from the NOA tool, where we also saw that the allocation to equity is very robust.

On the other hand, even a relatively small rise in yields beyond what was forecast will significantly reduce the duration of the MV optimal portfolio. Its allocation to corporate bonds is influenced by changes in both the curves and the credit spreads. Moving vertically from the centre, we see that up to a certain point, a larger-than-forecast rise in yields does not lead to a reduction in the allocation to corporates because total portfolio duration is reduced by lowering the duration of the government bond portfolio. But, beyond the point at which government bond portfolio duration can no longer be reduced, the allocation to corporates falls off sharply. Moving horizontally from the centre, we see that, in a manner similar to the allocation to equity, so long as spread widening does not pass a certain threshold, the allocation to corporates is not reduced, but once spreads widen beyond that point, there is a sharp transition in optimality.

Finally, the bottom heat map in Figure 6 addresses the following question: If our forecast is incorrect, how much of a loss will that cause, relative to how we would have invested given a correct forecast? In other words, how much worse would the

optimal portfolio implied by the base scenario perform in an alternative scenario, as compared with the portfolio that would be, ex ante, optimal for that scenario? Thus, the middle square of the heat map always has a value of zero, since the base portfolio and the alternative portfolio are the same, while other squares have negative values, as by definition the optimal portfolio from the base scenario will be sub-optimal in an alternative scenario, and so will do less well when compared with the optimal portfolio associated with that scenario. For example, we see that, although even a small addition to our forecast for the change in the level of the yield curve causes the MV model to significantly reduce duration, the effect of this reduction on the difference in returns between the two portfolios is not very large. Having this information might allow portfolio managers to accept a higher duration than they would otherwise have felt comfortable with.

## 5. Conclusion

The output of the MV model is typically not used “as-is” to make a final decision about the SAA of a portfolio. One reason is because the model’s results can be very sensitive to the assumptions made about the distribution of the return-on-assets and, particularly, about their means. Furthermore, there may also be some uncertainty about the target level of risk, which often can best be expressed as a narrow range rather than as a precise value. To cope with these problems, adjustments to the MV optimal portfolio found by the model are usually made when choosing a strategic asset allocation, but these adjustments have often been based on arbitrary rules or intuition.

In this paper, we have described two decision-support tools that were developed at the Bank of Israel to address these issues, and that are applied in the consultative process by which SAA decisions are made. The two are based on a common software foundation, which allows forecasts over a planning horizon for certain financial market variables – future government yield curves, future corporate credit spreads and the holding period return on an equity portfolio – to be transformed into values of return-on-assets. These values can then be treated as expected returns, and used to identify the MV optimal portfolio at a given risk level by performing Markowitz MV optimisation to derive an efficient frontier.

The tools discussed can be used by portfolio managers to more clearly see whether adjustments of this type are desirable, and what their consequences would be in various scenarios. The first – the NOA – does so by providing the portfolio manager with a number of alternative portfolios, which are very similar to the MV optimal portfolio in terms of risk and expected return, but significantly different in terms of their exposure to various financial market risk factors, including the general level of yields, credit spreads and equity market returns. These alternative portfolios can then be considered, in light of their performance in alternative scenarios and in light of the relative degree of confidence reposed by the portfolio manager in various aspects of the financial markets forecast.

The second tool – the FSA – allows the portfolio manager to see whether, and to what extent, small changes in the financial markets forecast lead to large changes in the makeup of the MV optimal portfolio, and to what extent such differences in the portfolio composition would impact the portfolio’s expected return. With this information in hand, any decisions made by the portfolio manager about adjustments

to the MV optimal portfolio can be made based on a more solid quantitative foundation.

## Acknowledgements

The development of the decision-support tools we have described was done in the R statistical computing language (R Development Core Team (2008)). A number of R packages are also incorporated into the software implementing the decision-support tools. Of these, the most important for mathematical computations are **apcluster** (Bodenhofer et al (2011)), which is based on work by Frey and Dueck (2007), and **quadprog** (Weingessel (2013)). Two graphics packages that have roles of major importance in the system are **ggplot2** (Wickham (2016)) and **plot3D** (Soetaert (2017)), which were used to produce most of the figures appearing in this paper. Several other R packages play supporting roles. Each of these packages can be accessed through an URL which follows the pattern <https://CRAN.R-project.org/package=quadprog>.

We wish to express our thanks to colleagues in the Markets Department of the Bank of Israel for useful feedback during the development of the tools we have described, and to the editors for useful comments and suggestions on the paper. The usual disclaimer applies.

## References

- Bodenhofer, U, A Kothmeier, and S Hochreiter (2011): "APCluster: an R package for affinity propagation clustering", *Bioinformatics*, vol 27, pp 2463–64.
- Bodenhofer, U, J Palme, C Melkonian, and A Kothmeier (2014): *APCluster: an R package for affinity propagation clustering*, Institute of Bioinformatics, Johannes Kepler University Linz.
- Frey, B and D Dueck (2007): "Clustering by passing messages between data points", *Science*, vol 315, pp 972–7.
- Lintner, J (1965): "The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets", *Review of Economics and Statistics*, vol 47, pp 13–37.
- (1969): "The aggregation of investors' diverse judgements and preferences in purely competitive markets", *Journal of Financial and Quantitative Analysis*, vol 4, pp 346–82.
- Markowitz, H (1952): "Portfolio selection", *Journal of Finance*, vol 7, pp 77–91.
- (1956): "The optimization of a quadratic function subject to linear constraints", *Naval Research Logistics Quarterly*, vol 3, pp 111–33.
- Michaud, R and R Michaud (2008): "Estimation error and portfolio optimization: a resampling solution", *Journal of Investment Management*, vol 6, no 1, pp 8–28.
- R Development Core Team (2008): *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria. <http://www.R-project.org>.
- Sharpe, W (1964): "Capital asset prices: a theory of market equilibrium under conditions of risk", *Journal of Finance*, vol 19, pp 425–42.
- Soetaert, K (2017): *plot3D: plotting multi-dimensional data*, R package version 1.1.1.
- Weingessel, A (2013): *quadprog: functions to solve quadratic programming problems*, R port of original S function by B Turlach, R package version 1.5-5.
- Wickham, H (2016): *ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag, New York.