

Robust optimisation by constructing near-optimal portfolios

Martin van der Schans,¹ Tanita de Graaf,² and Wessel van Eeghen¹

Abstract

Many investors use optimisation to determine their optimal investment portfolio. Unfortunately, optimal portfolios are sensitive to the optimisation's required input, ie they are not robust. Traditional robust optimisation approaches seek to provide an optimal and robust portfolio and, in doing so, replace the investor's investment decision process. In practice, however, portfolio optimisation supports but seldomly replaces the investment decision process. In this paper, we present an approach that both solves the robustness problem and aims to support rather than replace the investment decision process. We determine a region with near-optimal portfolios that, especially in light of the robustness problem, are all satisfactory allocation decisions. Then, as is already common practice, the investor can bring in expert opinions and additional information to select a preferred near-optimal portfolio. We will show that the region of near-optimal portfolios is more robust than the optimal portfolio itself.

JEL classification: C61, G11.

¹ Ortec Finance, PO Box 4074, 3006 AB Rotterdam, The Netherlands, www.ortec-finance.com. This version: 5 June 2019.

² ORTEC BV.

1. Robust optimisation

Many investors use portfolio optimisation to determine their optimal investment portfolio. Unfortunately, optimal portfolios are sensitive to the optimisation's input parameters. Although this sensitivity is often studied in the context of mean-variance optimisation, where optimal portfolios are sensitive to the estimated mean and covariance matrix (Frankfurter et al (1971), Michaud (1989), Chopra and Ziemba (1993)), sensitivity is a generic problem in portfolio optimisation (Kondor et al (2007), Ciliberti et al (2007)). As discussed in Hurley and Brimberg (2015), the sensitivity is caused by an interaction of an estimation error in the input and the optimisation objective.

The literature proposes several robust optimisation approaches to deal with the sensitivity problem. Although they differ methodologically, approaches such as shrinkage (Ledoit and Wolf 2004), robust statistics (Reyna et al (2005)), Black-Litterman inverse optimisation (Bertsimas et al (2012)) and Bayesian optimisation (Schöttle et al (2010)) reduce the sensitivity by diminishing the role of the data on which the optimisation parameters are estimated, ie they reduce estimation errors in the optimisation input or their effect. Other methods, such as regularisation, change the optimisation objective to make it less sensitive to estimation errors in the optimisation parameters (DeMiguel et al (2009), Brodie et al (2009)). Also, there are hybrid methods that do both. For example, the optimisation community proposes a general robust optimisation framework. Given a convex optimisation problem, the framework proposes a robust counterpart that lets the input vary within a specified range and selects the worst case outcome (Ben-Tal and Nemirovski (1998)). Finally, there is the resampled frontier (Michaud (1998)), which is constructed by resampling the input from a distribution and averaging over the resamplings' optimisation results.

Robust optimisation approaches work well when all market information is quantified and incorporated in the optimisation problem. However, despite efforts to incorporate information such as transaction costs, expert opinion and liquidity into portfolio optimisation problems, they remain a simplification of reality. In practice, investors often combine the optimal portfolio with additional information that was not or could not be incorporated. Thus, for investors, portfolio optimisation is a tool that supports but does not replace their decision process. In this paper, we take this as the starting point for developing a robust optimisation approach.

2. Near-optimal portfolios

Generally, the result of a portfolio optimisation problem is an efficient frontier with optimal portfolios. Now, given an optimal portfolio w_0 on the efficient frontier, we construct a near-optimal region just below the efficient frontier as indicated by the shaded region in Figure 1. The portfolios in this region are referred to as near-optimal portfolios. As shown in Chopra (1993) and Section 4, the near-optimal portfolios' weights can differ completely from those of the optimal portfolio. The idea is that all near-optimal portfolios still have a satisfactory risk-return trade-off so that investors can, as is already common practice, bring in additional arguments and select their preferred near-optimal portfolio.

To find the near-optimal region represented by the shaded region in Figure 1, we construct near-optimal portfolios w_1, \dots, w_n far away from each other and show that any weighted average of these portfolios, ie any portfolio in their convex hull,

$$\text{Conv}(w_1, \dots, w_n) := \{ \sum_{i=1}^n \theta_i w_i \mid \theta_i \geq 0, \sum_{i=1}^n \theta_i = 1 \}, \quad (1)$$

Figure 1: Efficient frontier and near-optimal region




Table 1: Statistics and optimal allocation

	Stocks	Bonds	T-bills
Mean	1.32%	1.03%	0.73%
Std. dev.	4.79%	3.98%	0.22%
Correlations:			
Stocks	1.000		
Bonds	0.341	1.000	
T-bills	-0.081	0.050	1.000
Optimal allocation	58.10%	22.80%	19.10%

Figure 1 shows a mean-variance efficient frontier based on the statistics in Table 1, an optimal allocation (orange dot) and a shaded region with near-optimal portfolios. Table 1 shows statistics of monthly returns from January 1980 to December 1990 as reported by Chopra (1993) and the portfolio weights of the orange dot in Figure 1.

is near-optimal. We continue constructing near-optimal portfolios until their convex hull sufficiently covers the near-optimal region. We will show that the near-optimal region is more robust than the optimal portfolio w_0 itself. The intuitive understanding is that, with slightly different input parameters, the near-optimal region slightly changes in shape. For example, the old optimum becomes near-optimal and one of the near-optimal portfolios becomes optimal. In particular, we will show that most near-optimal portfolios remain near-optimal. And, because near-optimal portfolios have, by construction, satisfactory risk-return trade-offs, the investor's allocation does not need to be revised and becomes more robust.

3. Methodology

3.1 Constructing near-optimal portfolios

The construction of near-optimal portfolios consists of the following steps:

1. start with an efficient frontier and an optimal portfolio w_0 on the frontier as in Figure 1;
2. specify the near-optimal region as indicated by the shaded region in Figure 1;
3. find the portfolio w_1 in the near-optimal region with the highest risk and return;
4. find the portfolio w_2 in the near-optimal region that differs most in allocation from w_1 ;
5. find the portfolio w_3 in the near-optimal region that differs most in allocation from the weighted averages of w_1 and w_2 , ie from all portfolios in $\text{Conv}(w_1, w_2)$; and
6. continue until $\text{Conv}(w_1, \dots, w_n)$ covers the near-optimal region up to a required precision ϵ .

In this section, we specify these steps in more detail. Although near-optimal portfolios can be found below any efficient frontier, we assume for simplicity that the investor is interested in near-optimal mean-variance portfolios. Thus, the efficient frontier is constructed by solving

$$\min_w \lambda w^T \Sigma w - (1 - \lambda) w^T \mu, \quad (2a)$$

$$Aw = b, \quad (2b)$$

$$Gw \leq h, \quad (2c)$$

where the vector μ contains the asset's mean returns, Σ is their covariance matrix, A is a matrix representing, together with the vector b , the equality constraints, G is a matrix representing, together with the vector h , the inequality constraints and $0 \leq \lambda \leq 1$ represents the investor's risk aversion. Also, equality constraints (2b) should at least enforce that the sum of all portfolio weights w equals one. Apart from this restriction, constraints (2b) and (2c) can be chosen freely and, for example, be used to prevent short selling or fix the allocation to certain asset classes. In Step 1, the investor solves mean-variance optimisation problem (2) for a number of risk aversion parameters λ and obtains the efficient frontier as denoted in Figure 1 by the solid blue line. From the frontier, the investor selects an optimal portfolio w_0 with an appropriate risk-return trade-off. In Step 2, the investor specifies a region R of near-optimal portfolios around the optimal portfolio w_0 by specifying a minimum return μ_{\min} so that

$$w^T \mu \geq \mu_{\min}, \quad (3)$$

and a maximum variance σ_{\max}^2 so that

$$w^T \Sigma w \leq \sigma_{\max}^2. \quad (4)$$

The near-optimal region $R(\mu_{\min}, \sigma_{\max}^2)$ thus consists of all portfolios that satisfy (2b), (2c), (3) and (4), and, it is represented by the shaded region in Figure 1. It can be shown that the near-optimal region is convex, which means that any weighted average of the near-optimal portfolios is near-optimal (see Appendix A).

Next, in Step 3, we find a portfolio w_1 with maximum risk and return which is located in the region's upper right corner. Because, apart from degenerate cases, the efficient frontier is concave, increasing and consists of unique portfolios, this portfolio cannot be written as a weighted average of other near-optimal portfolios. Therefore, it is a natural starting point for the construction of the near-optimal region. In Step 4, we find a portfolio w_2 in the region R that differs most from w_1 , ie w_2 satisfies:

$$w_2 = \arg \max_{w \in R(\mu_{\min}, \sigma_{\max}^2)} \|w_1 - w\|, \quad (5)$$

where $\|\cdot\|$ indicates the Euclidean norm, ie the root of the sum of the components squared. Note that maximising the Euclidean norm favours large deviations in one component over small deviations in several components. More generally, as follows from Lemma A1, once we find $i-1$ near-optimal portfolios w_1, \dots, w_{i-1} , all portfolios in the convex hull $\text{Conv}(w_1, \dots, w_{i-1})$, ie all weighted averages, are near-optimal. Therefore, in Step 5, we find w_i by finding the portfolio that differs most from all the portfolios in the convex hull of w_1, \dots, w_{i-1} :

$$w_i = \arg \max_{w \in R(\mu_{\min}, \sigma_{\max}^2)} d(w, \text{Conv}(w_1, \dots, w_{i-1})), \quad (6)$$

where the function d indicates the distance to the convex hull:

$$d(w_i, \text{Conv}(w_1, \dots, w_{i-1})) = \min_{w \in \text{Conv}(w_1, \dots, w_{i-1})} \|w_i - w\|. \quad (7)$$


As indicated in Step 6, we continue constructing near-optimal portfolios until the constructed convex hull $\text{Conv}(w_1, \dots, w_n)$ covers the near-optimal region $R(\mu_{\min}, \sigma_{\max}^2)$ up to a required precision $\varepsilon > 0$, ie until:

$$d(w_n, \text{Conv}(w_1, \dots, w_{n-1})) < \varepsilon. \quad (8)$$

Although we cannot say a priori to what extent a certain choice of ε will capture the entire near-optimal region, criterion (8) enforces that there are no near-optimal portfolios that have allocation weights that differ more than ε with the nearest portfolio in the convex hull of w_1, \dots, w_n . Therefore, in practice, we recommend choosing ε equal to the absolute difference in allocation that is considered insignificant.

3.2 Support vector machines

Although Section 3.1 specifies how to construct the near-optimal region, optimisation problem (6) is difficult to solve since evaluating its objective requires calculating the distance of a portfolio to a convex hull which in turn requires solving optimisation problem (7). Optimisation problem (6) can be simplified by applying theory on support vector machine (SVM) classification developed since Vapnik (1963).



The blue line separates the green portfolios from the orange portfolio and has maximal separation margin $2/\|x\|$.

In its easiest form, SVM classification is a machine learning method that tries to separate two classes of points in space by a plane (Boser et al (1992)). Figure 2 shows a two-dimensional example where the green points, representing portfolios w_1 to w_{i-1} , are separated in space from the orange point, representing the portfolio w_i , by the line $w^T x + z = 0$; here x is a vector and z is a scalar.

In SVM classification, the separating line $w^T x + z = 0$ maximises its distance to the two classes, which, as indicated in Figure 2, can be shown to equal $1/\|x\|$. As shown in Boser et al (1992), the separating line of an SVM classification problem can, when it exists, be found by solving a quadratic programming problem:

$$\min_{x,z} \|x\|^2, \quad (9a)$$

$$x^T w_j + z \leq -1 \text{ for } j = 1, \dots, i-1, \quad (9b)$$

$$x^T w_i + z \geq 1. \quad (9c)$$

As shown in Bennett and Bredensteiner (2000), Bennett and Campbell (2000), and Mavroforakis and Theodoridis (2006), the separation margin $2/\|x\|$ equals the distance between the convex hull $\text{Conv}(w_1, \dots, w_{i-1})$ and the point w_i as defined by (7). Optimisation problem (6) can therefore be written as:

$$\min_{x,z,w_i} \|x\|^2, \quad (10a)$$

$$x^T w_j + z \leq -1 \text{ for } j = 1, \dots, i-1 \quad (10b)$$

$$x^T w_i + z \geq 1. \quad (10c)$$

$$Aw_i = b, \quad (10d)$$

$$Gw_i \leq h, \quad (10e)$$

$$w_i^T \mu \geq \mu_{\min}, \quad (10f)$$

$$w_i^T \Sigma w_i \leq \sigma_{\max}^2, \quad (10g)$$

where constraints (10d) to (10g) enforce that w_i is near-optimal, ie w_i is in the region $R(\mu_{\min}, \sigma_{\max}^2)$, and constraints (10b) and (10c) together with objective (10a) enforce that w_i is at maximum distance of the convex hull $\text{Conv}(w_1, \dots, w_{i-1})$.

Figure 3: Mean-variance efficient frontier with a near-optimal region

Table 2: Near-optimal portfolios and near-optimal region's centre of mass

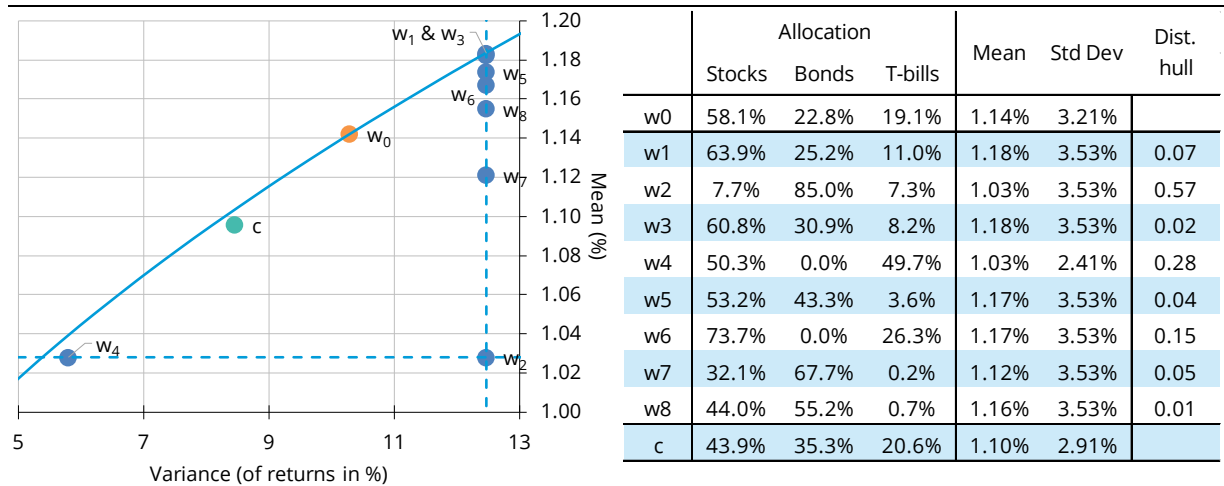


Figure 3 shows, as in Figure 1, the mean-variance efficient frontier and optimal portfolio (orange dot) based on the statistics in Table 1. In addition, the figure shows a near-optimal region (shaded) covered by the convex hull of several near-optimal portfolios (blue dots). The green dot represents the convex hull's centre of mass. Table 2 lists the allocation weights and statistics of the optimal portfolio w_0 , the near-optimal portfolios and centre of mass portfolio c . The last column indicates the distance to the convex hull of the preceding near-optimal portfolios.

Although optimisation problem (10) is non-linear and not even convex due to inequality constraint (10c), it is solvable with standard optimisation software such as the NLOPT's SLSQP solver (Johnson 2010) once a good starting point for the algorithm is found. Finding a good starting point is not straightforward. Imagine, for example, the near-optimal region as a circle and the convex hull approximating it as a triangle connecting three points on the boundary of the circle. The main difficulty is that the complement of the triangle consists of three disjoint parts of the circle. To find a good starting point, we use the so-called hit-and-run (Smith 1984) and shake-and-bake (Boender et al (1991)) algorithms to sample points uniformly from the boundary of the near-optimal region. In practice, a large enough sample will generate a starting point for the optimisation that is far enough away from the convex hull if it exists.

4. Example

Discussion

In this section, we continue with the Chopra (1993) example as shown schematically in Figure 1. Under a no short selling constraint, we perform a mean-variance optimisation using the statistics in Table 1. This results in the efficient frontier (blue line) and the optimal portfolio w_0 (orange dot) indicated in Figure 3. The near-optimal region, indicated by the shaded region in Figures 1 and 3, consists of portfolios that,

compared to the optimal portfolio (orange dot), have an expected return that is relatively at most 10% lower and a standard deviation is relatively at most 10% higher. To find the near-optimal region, we apply the optimisation procedure presented in Section 3 and obtain portfolios w_1 to w_8 . The procedure continues until no portfolio can be found with an allocation difference larger than 1% to the convex hull found so far (see also the last column in Table 2). So, the convex hull of the portfolios w_1, \dots, w_8 covers the near-optimal region up to a precision $\varepsilon = 0.01$, see equation (8).

Chopra (1993) studies (almost) the same near-optimal region with the purpose to show that near-optimal portfolios can have completely different weights. There are two important differences to note. First, for the purpose in Chopra (1993), a different definition of the near-optimal region is used: portfolios that, compared to the optimal portfolio, provide at least 90% of the expected return for less than 90% of the standard deviation are left out, ie portfolios in the lower-left corner of the near-optimal region indicated in Figure 3 are left out. For our purpose, however, we do consider these portfolios near-optimal because there are portfolios considered near-optimal with the same average return and a higher standard deviation. Regardless of this difference in definition, we verified that, as is implied by the difference, all near-optimal portfolios constructed in Chopra (1993) are contained in the convex hull of w_1, \dots, w_8 .

Second, Chopra (1993) constructs near-optimal portfolios through a grid search, ie tries all possible portfolios, and searches for near-optimal portfolios with the highest upward and downward deviation in one asset class. Because there are three assets, this results in six near-optimal portfolios that, as intended, differ completely in weights. It can, however, be verified that the convex hull of the near-optimal portfolios found in Chopra (1993) does not at all cover the near-optimal region. For example, except for w_4 , all portfolios are near-optimal in the definition of Chopra (1993), but w_1 , w_3 , w_5 and w_8 cannot be written, also not approximately, as a weighted average of the near-optimal portfolios reported in Chopra (1993). This shows that merely constructing portfolios with the highest upward and downward deviation is not sufficient for finding the complete near-optimal region. Additionally, a grid search algorithm is, contrary to the methods presented here, only feasible in low dimensions.

Since optimisation problem (10) is not convex, we performed two consistency checks to ensure that the convex hull of the near-optimal portfolios w_1, \dots, w_8 indeed covers the near-optimal region. First, as noted, we verified that all near-optimal portfolios constructed in Chopra (1993) are contained in the convex hull of w_1, \dots, w_8 . Second, we also verified that all portfolios both on the efficient frontier and in the near-optimal region are contained in the convex hull of w_1, \dots, w_8 . Together, this gives sufficient confidence in the convergence and accuracy of the numerical solvers used.

Finally, note that Figure 3 might wrongly give the impression that the near-optimal portfolios of which the convex hull covers the near-optimal region should lie on the boundary of the region indicated by the shaded region in Figure 3. It can easily be shown that this is not the case. For example, when the near-optimal region is increased to portfolios with a return of at least 1% and a variance of at most 4%, the first near-optimal portfolio found has a 100% allocation to treasury bills and does not lie on the boundary of the region indicated by the shaded region in Figure 3.

Selecting a preferred near-optimal portfolio

Once the near-optimal region is covered by the convex hull of near-optimal portfolio w_1, \dots, w_8 , the investor is free to select a preferred portfolio from this region. Without further information, the convex hull's centre of mass c , indicated by the green dot in Figure 3, can be a good default choice. Geometrically, it can be interpreted as the average portfolio when we would sample portfolios (uniformly) from the convex hull of w_1, \dots, w_8 . To construct the convex hull's centre of mass, we used the Quickhull algorithm (Barber et al (1996)).

The centre of mass's geometric interpretation indicates why it makes a good default choice. Now, there are many sets of portfolios whose convex hull approximates the near-optimal region, and we have merely chosen w_1, \dots, w_8 because they conveniently differ from each other as much as possible. Unless the investor has reasons not to, a preferred portfolio should have the property that it depends on the near-optimal region, but not on the particular choice of portfolios whose convex hull approximates it. Because, loosely speaking, the centre of mass is the average portfolio of a region, similar regions have approximately the same centre of mass. Therefore, the centre of mass has this property. Also, conceptually, it is one of the most simple portfolios with this property. Note that, for example, a simple average of portfolios whose convex hull approximates the near-optimal region does not have this property and does depend on the portfolios chosen.

In Section 4.3, we will show that the centre of mass portfolio is robust in two ways. First, the centre of mass portfolio is less sensitive to changing input parameters than portfolios on the frontier. Second, because the centre of mass portfolio is not on the boundary of the convex hull, it is expected to remain near-optimal with slightly different input parameters.

Although the centre of mass portfolio can be a good default choice, preferably, the investor brings in additional arguments to select a preferred near-optimal portfolio. In light of the sensitivity problem, these arguments should be additional to the risk-return statistics in Table 1, eg selecting the portfolio with the highest sharp ratio would not suffice. As an example, suppose the investor prefers, for whatever reason, not to invest in bonds. It follows from Table 2 that, in that case, any weighted average of portfolio w_4 and w_6 would suffice: such a portfolio is near-optimal and has no allocation to bonds. When the investor currently owns 60% equity and 40% cash, which is to be invested in either stocks or treasury bills, the investor could, for example, choose to leave his current exposure to equity intact and to buy treasury bills with his cash.

In practice, the additional arguments brought in for a selecting a preferred near-optimal portfolio can take many forms. Although some arguments to select a preferred near-optimal portfolio can be quantitative, eg low transaction costs, good performance on another investment horizon or good performance in another risk measure, the methodology especially lends itself to be combined with qualitative arguments such as sustainable investing preferences. The methodology gives the range of possibilities through the near-optimal region, after that, it is up to the investor. Also, the methodology can serve as a decision support tool for investment boards. An advisor can present the range of possibilities, ie the near-optimal portfolios. Then, the investment board can choose from the range of possibilities following a line of reasoning they prefer.

Robustness of the preferred near-optimal portfolio

To investigate the robustness of the near-optimal portfolio method, we view the mean and covariance matrix presented in Table 1, ie the optimisation’s input parameters, as estimated on a sample of size N . To perturb the optimisation’s input, we draw a sample of size N from a normal distribution with mean and covariance as in Table 1 and estimate a perturbed mean and covariance matrix.¹ So, the larger the sample size the closer the perturbed means and covariance matrices are to original ones in Table 1. For each perturbed mean and covariance matrix, we determined the perturbed optimal mean-variance portfolio, near-optimal portfolios, convex hull and centre of mass portfolio.

Figures 4 and 5 show the average robustness using 100 samples for each sample size. In Figure 4, the yellow dots and their trend line show for each sample size the average percentage overlap using 100 samples of size N between the original and perturbed convex hulls. The increase of the yellow trend shows that, with perturbed input parameters, roughly 60% to 90% of the near-optimal portfolios remains near-optimal and, consequently, no new investment advice is required.

Figure 4: Near-optimal region vs mean variance

Figure 5: Near-optimal portfolios




Figure 4 and 5 show, as a function of the sample size N , the average percentage of overlap between the original and perturbed convex hull (yellow), the average turnover between the original and the perturbed mean-variance optimal portfolio (orange) and the average turnover between portfolios in the original convex hull and the closest portfolio in the perturbed convex hull (blue). In addition, Figure 5 shows the original and the perturbed centre of mass (green). And, the average turnovers between the closest portfolio in the perturbed convex hull and: the original convex hull (blue), the original mean-variance optimal portfolio (orange dotted), and the original centre of mass portfolio (green dotted).

The comparison with mean-variance optimisation can best be made using the turnover measure:

$$T(u, v) = \frac{1}{2} \sum_i |u_i - v_i|,$$

¹ Equivalently, we could have drawn perturbed means and covariance matrices from a normal inverse-Wishart distribution, which is the conjugate prior to the multivariate normal distribution.

where the sum is taken over all the components of the vectors. The turnover measure can be interpreted as the fraction of the portfolio u that has to be sold and reinvested to obtain portfolio v . In Figures 4 and 5, the orange dots and their trend line show that the average turnover between the mean-variance optimal and perturbed mean-variance optimal portfolio ranges from roughly 10% to 25%. The average turnover is significantly decreased with the near-optimal portfolio method: the blue dots and their trend line in Figures 4 and 5 indicate that, on average, less than 5% of the portfolio has to be sold and reinvested to obtain a near-optimal portfolio when input parameters are re-estimated.

Figure 5 shows for two typical near-optimal portfolios, the mean-variance optimal portfolio and centre of mass portfolio, that their robustness significantly increases when their robustness is measured with respect to the near-optimal region. At least for these two near-optimal portfolios, the near-optimal method achieves its increase in robustness, because it measures robustness with respect to a region of near-optimal portfolios instead of with respect to a single optimal portfolio. In other words, the increase in robustness is a free advantage when all near-optimal portfolios are considered valid investment decisions.

5. Conclusion

We have shown how to construct a region of near-optimal portfolios below the efficient frontier. Also, we discussed how this methodology is both robust and is designed to support rather than replace the investor's investment decision process.

There are several directions that can be explored in future research. First, the methodology can be applied to more computationally intensive optimisation objectives such as mean-CVaR optimisation. Second, the robustness with respect to estimates for the mean and covariance matrices can be explored separately. Also, the robustness can be compared with other robust optimisation approaches such as resampling. And finally, the method's integration with the investment decision process or use as a decision support tool can be further explored.

6. Acknowledgments

We would like to thank Dr O.W. van Gaans of Leiden University for contributing to this research and supervising two master thesis projects on this topic. Additionally, we would like to thank our colleagues Marnix Engels, John Kuijt and Loranne van Lieshout for their valuable feedback.

A Convexity of the near-optimal region

Lemma A1. The near-optimal region $R(\mu_{\min}, \sigma_{\max}^2)$ consisting of all portfolios that satisfy (2b), (2c), (3) and (4) is convex.

Proof. When given two near-optimal portfolios $u, v \in R(\mu_{\min}, \sigma_{\max}^2)$, we have to show that any weighted average $w = tu + (1-t)v$ is also near-optimal, ie w satisfies (2b), (2c), (3) and (4) for all $0 \leq t \leq 1$. First, since u and v satisfy (2b), it directly follows that w satisfies (2b). Also, inequality constraints (2c) and (3) follow directly. That w satisfies inequality constraints (4) follows from convexity of the left-hand side of (4) and applying Jensen's inequality:

$$w^T \Sigma w \leq t u^T \Sigma u + (1-t) v^T \Sigma v \leq t(w_0^T \Sigma w_0 + \delta_\Sigma) + (1-t)(w_0^T \Sigma w_0 + \delta_\Sigma) = w_0^T \Sigma w_0 + \delta_\Sigma.$$

References

- Barber, C, D Dobkin and H Huhdanpaa (1996): "The quickhull algorithm for convex hulls", *ACM Transactions on Mathematical Software*, vol 22, no 4, December, pp 469-483.
- Ben-Tal, A and A Nemirovski (1998): "Robust convex optimization", *Mathematics of Operations Research*, vol 23, no, 4, pp 769–805.
- Bennett, K and E Bredensteiner (2000): "Duality and geometry in SVM classifiers", *Proceedings of the Seventeenth International Conference on Machine Learning (ICML 2000)*, Stanford University, USA, 29 June–2 July, pp 57–64.
- Bennett, K and C Campbell (2000): "Support vector machines: hype or hallelujah?" *SIGKDD Explorations Newsletter*, vol 2, no 2, December, pp 1–13.
- Bertsimas, D, V Gupta and I Paschalidis (2012): "Inverse optimization: a new perspective on the Black-Litterman model", *Operations Research*, vol 60, no 6, pp 1389–403.
- Boender, C, R Caron, J McDonald, A Rinnooy Kan, H Romeijn, R Smith, J Telgen and A Vorst (1991): "Shake-and-bake algorithms for generating uniform points on the boundary of bounded polyhedra", *Operations Research*, vol 39, no 6, November–December, pp 945–54.
- Boser, B, I Guyon and V Vapnik (1992): "A training algorithm for optimal margin classifiers", *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, pp 144–152, COLT '92. Pittsburgh, USA.
- Brodie, J, I Daubechies, C De Mol, D Giannone and I Loris (2009): "Sparse and stable Markowitz portfolios", *Proceedings of the National Academy of Sciences*, vol 106, no 30, pp 12267–12272.
- Chopra, V and W Ziemba (1993): "The effect of errors in means, variances, and covariances on optimal portfolio choice", *The Journal of Portfolio Management*, vol 19, no 2, pp 6–11.
- Chopra, V (1993): "Improving optimization", *The Journal of Investing*, vol 2, no 3, pp 51–59.
- Ciliberti, S, I Kondor and M Mézard (2007): "On the feasibility of portfolio optimization under expected shortfall". *Quantitative Finance*, vol 7, no 4, pp 389–396.
- DeMiguel, V, L Garlappi, F Nogales and R Uppal (2009): "A generalized approach to portfolio optimization: improving performance by constraining portfolio norms", *Management Science*, vol 55, no 5, pp 798–812.
- Frankfurter, G, H Phillips and J Seagle (1971): "Portfolio selection: the effects of uncertain means, variances, and covariances", *Journal of Financial and Quantitative Analysis*, vol 6, no 5, pp 1251– 1262. doi: 10.2307/2329859.
- Hurley, W and J Brimberg (2015): "A note on the sensitivity of the strategic asset allocation problem", *Operations Research Perspectives*, vol 2, pp 133–136.
- Johnson, S (2010): *The NLOpt nonlinear-optimization package*.
- Kondor, I, S Pafka and G Nagy (2007): "Noise sensitivity of portfolio selection under various risk measures", *Journal of Banking & Finance*, vol 31, no 5, pp 1545–1573.

Ledoit, O and M Wolf (2004): "Honey, I shrunk the sample covariance matrix". *The Journal of Portfolio Management*, vol 30, no 4, pp 110–119.

Mavroforakis, M and S Theodoridis (2006): "A geometric approach to Support Vector Machine (SVM) classification". *IEEE Transactions on Neural Networks*, vol 17, no 3, pp 671–682.

Michaud, R (1998): *Efficient asset management: a practical guide to stock portfolio management and asset allocation*, Financial Management Association Survey and Synthesis, Boston, HBS Press.

Michaud, R (1989): "The Markowitz optimization enigma: Is 'optimized' optimal?", *Financial Analysts Journal*, vol 45, no 1, pp 31–42.

Reyna, F, A Duarte Jr, B Mendes and O Porto (2005) : "Optimal portfolio structuring in emerging stock markets using robust statistics", *Brazilian Review of Econometrics*, vol 25, no 2, pp 139.

Schöttle, K, R Werner and R Zagst (2010): "Comparison and robustification of Bayes and Black-Litterman models", *Mathematical Methods of Operations Research*, vol 71, no 3, pp 453–75.

Smith, R (1984): "Efficient Monte Carlo procedures for generating points uniformly distributed over bounded regions", *Operations Research*, vol 32, no 6, pp 1296–1308.

Vapnik, V (1963): "Pattern recognition using generalized portrait method", *Automation and remote control*, vol 24, pp 774–780.

Copyright © 2019 Ortec Finance bv. All rights reserved. No part of this paper may be reproduced, in any form or by any means, without permission from the authors. Short sections may be quoted without explicit permission provided that full credit is given to the source. The views expressed are those of the individual author(s) and do not necessarily reflect the views of Ortec Finance bv.