

THE SWOP ALGORITHM

USING ADVANCED NUMERICAL METHODS
TO RE-BALANCE CLIENT PORTFOLIOS

EXECUTIVE SUMMARY

The current case study presents why Unigestion, a financial institution with headquarters in Geneva, has decided to use the Shortfall and Windfall Optimized Portfolios (SWOP) asset allocation algorithm and how it was able to improve their legacy clients' portfolio asset allocation method by opting for the SWOP module. It shows that the SWOP module can be used as a portfolio strategic allocation and benchmarking tool as well as a marketing tool for the front office, while at the same time benefiting from a sound and easy to explain underlying numerical algorithm, which makes it a perfect tool for a portfolio manager.

The case study will start by presenting why SWOP has fit Unigestion's needs so well, what department is using this module and how it is used. The old asset allocation module will be briefly described as well as SWOP's main principles and the main SWOP differentiating factors will be highlighted.

The theoretical underpinnings of the SWOP model can be found in the following papers by the University of Lausanne professors Michael Rockinger and Eric Jondeau:

- > *"The Impact of Shocks on Higher Moments"*
Journal of Financial Econometrics, 2008, 1-29.
- > *"Optimal Portfolio Allocation Under Higher Moments"*
2006, European Financial Management.
- > *"The Economic Value of Distributional Timing"*
2006.

The owner of the SWOP module is FinQuant, a subsidiary of SAGE S.A.

TABLE OF CONTENT

1	INTRODUCTION	4
2	UNIGESTION	5
2.1	BACKGROUND	5
2.2	OLD ASSET ALLOCATION ALGORITHM	6
2.3	THE NEED FOR A MORE ADVANCED MODEL	6
3	SWOP IMPLEMENTATION	8
4	THE BASIC PRINCIPLES OF SWOP	9
5	EF VERSUS SWOP	11
6	CONCLUSIONS	13

1 INTRODUCTION

SWOP, the Shortfall and Windfall Optimized Portfolios, is a module that uses a sound mathematical model for finding the allocation of a portfolio of financial assets which maximizes the investor's expected return, while at the same time minimizing the maximum loss that this portfolio could incur in extreme cases like crashes. The model goes beyond the standard Mean-Variance model in defining risks by concentrating on the extreme positive and negative asset returns (happening at the tail of the return distribution).

One of the main benefits of the SWOP model is the fact that its principles, mechanism and results can be easily explained to anyone, even a non-professional, despite the fact that there are serious calculations going on behind the algorithm. The tractability of the model, along with its friendly graphical user interface, quick response time and the ease with which it can be understood and explained to the customer, are among the reasons why Unigestion has taken part in the development of the SWOP model.

One of the many lessons the latest financial crisis has taught the companies in the financial sector that survived the crisis is that using hyper-complicated models for quantifying risk – models that even their creators cannot explain in simple words, models that use variables that are not easily measurable or are provided by other companies – often produce results that are wrong.

Before the crisis, banks and other financial institutions were massively using models that had as an input the credit ratings that were computed by third-party firms like Standard & Poor's, Fitch, Moody's etc. These ratings proved to be wrong or lagging far behind the real evolution of the companies' creditworthiness, which rendered the models (even good ones) that were using these ratings also wrong. The conclusion that has to be drawn here is that all the input parameters of a good and reliable model also have to be reliable and understandable.

Another example would be the failure of extremely complicated (be it asset allocation or risk measurement) models which principles cannot be explained in simple words. Because of their lack of tractability, it is very hard to explain what exactly in these models went wrong. The conclusion to be drawn here is that a good financial model is one which main mechanism can be clearly understood by those who use it and the influence of each of the input parameters on the final output of the model are very clear. In this sense, SWOP has a clear advantage over more sophisticated, but less tractable, and thus less reliable, models. As it will be discussed further on, although the SWOP optimization algorithm is computationally intensive, its principles are simple and easily understandable, which makes it a reliable tool for portfolio managers.

Another important lesson we have learned from the recent crisis is that numerical algorithms alone cannot be used as the ultimate tool for asset allocation purposes. All the multifaceted aspects of risk management should be taken into consideration when deciding which assets should be invested in and which ones can prove dangerous for the portfolio. Due diligence, counterparty risk, and systemic risk are examples of the different components that financial institutions should use when creating value for their customers.

All of these lessons were taken into account at Unigestion, which does not directly apply the results of the SWOP optimized portfolios, but rather applies them as one of the tools for client portfolio asset allocation. Portfolio managers at Unigestion also add their strategic views to the SWOP results given their perception of the forecasted performance of different asset classes. Used cautiously in this context, SWOP has proven a very good and reliable model that performed well during the crisis, suggesting, for example, a much lower investment in the equities asset class than the Mean-Variance (a.k.a. Efficient Frontier) model did.

2 UNIGESTION

2.1 BACKGROUND

In this section we will quickly summarize the position and functions of the Unigestion department that uses the SWOP module in its day-to-day work and describe how SWOP fits the needs of this department.

The Family Investment Office (called FIO) department at Unigestion manages assets of private clients. It uses a bespoke questionnaire to determine the client risk tolerance level that will then determine the most appropriate investment strategy to be pursued for the client portfolio. Unigestion assigns to each investment strategy a so-called risk profile colour (or simply “risk colour”), which in turn has a level of portfolio volatility that the customer is willing to bear. Unigestion also attaches a set of constraints to the portfolio allocation based on the client risk profile colour and client's personal wishes. For example, a risk averse client might be quantified by the results of the questionnaires as belonging to the Green risk profile. Such profile has certain constraints attached to it, which could then be combined with client specific constraints.

For example, the client might want to have a weight of not less than 80% of his portfolio invested in domestic bonds and stocks. Of course, these are theoretical constraints - each Unigestion client does, however, have a certain level of acceptable portfolio volatility and a set of constraints attached to the detected risk profile.

And this is where the SWOP algorithm fits in – after the client's risk profile is determined and the set of constraints is defined, the SWOP algorithm is used by Unigestion to determine the strategic asset allocation that will maximize the investor's wealth (the expected return) while minimizing the risk of incurring large losses, taking into account all the single, group or relative constraints defined earlier. But what SWOP does differently from the standard Markovitz' mean-variance framework is that it goes beyond the framework where risk means volatility. The volatility of positive returns (good volatility) should not penalize a portfolio or an asset that exhibits this kind of large positive returns. The SWOP module will determine a portfolio that has a minimum expected shortfall (new risk benchmark instead of the volatility) for a given level of return, thus redefining the risk vs. return space. Furthermore, the module allows the optimization results to be visualized and compares different portfolios (optimized and not optimized) with respect to various statistics such as maximum and minimum rolling performances for different time lengths, different financial ratios, and so on: this is, of course, invaluable when explaining the asset allocation that can be applied to the portfolio and discussing this with the client.

2.2 OLD ASSET ALLOCATION ALGORITHM

The first model used by Unigestion for this purpose was the Efficient Frontier (EF) optimization algorithm, which served the same purpose. It has been used for several years by the Family Investment Office department at Unigestion for determining the allocation at the asset class level for clients of different risk appetites based on the well-known expected return versus volatility Markovitz model. The implementation of this model with enriched graphical user interface capabilities was also implemented by SAGE for Unigestion in the late 90's. Being aware of the shortcomings of the model that have long been discussed in the financial research literature, the results of the EF module optimizations were not applied directly to the client portfolio allocation but were used for benchmarking and strategic allocation. The benefits of the EF optimization for Unigestion have been:

- > results made logical sense and could be easily explained to customers;
- > the module has a very good response time and thus optimizations could be run in front of clients, making it not only a middle office but also a front office tool;
- > there was very good integration with the existing information system (IS) and the optimizations could be run from inside the IS and the results of the optimization could be further used inside the system;
- > the module was very flexible in terms of number and complexity of constraints, which was very important to Unigestion.

Overall, the Efficient Frontier module has gained a good reputation at Unigestion due to its stability, ease of use and flexibility, despite its obvious shortcomings that are discussed in numerous academic papers.

2.3 THE NEED FOR A MORE ADVANCED MODEL

Unigestion is a flexible organisation, which is able to quickly adapt its practices to new needs or market conditions. A decision to change the optimization algorithm used for asset allocation that could take years in a big company, was taken in a matter of weeks at Unigestion. The FIO department was conscious of the fact that using volatility as the only measure of risk has proven to be too simple over the years, thus turning to a more sophisticated portfolio allocation algorithm has become imperative for Unigestion, which has always been very innovative and flexible.

Given its ties with the time-proven reliable partner SAGE, Unigestion decided to assist with the implementation of the new asset allocation solution SAGE proposed after providing a proof of concept. The Shortfall and Windfall Optimized Portfolios (SWOP) proposed by SAGE back in 2006 went beyond the risk/return framework, beyond the much disputed Value-at-Risk and was focused on minimizing the average loss a portfolio incurs when it is in the n% worst-case scenario (called Shortfall) or maximizing the average gain of a portfolio in an n% best-case scenario (Windfall). Unigestion asked to incorporate the results of the mean-variance asset allocation algorithm into the graphical user interface of SWOP, and the new tool would allow to:

- > minimize extreme losses during crashes (Shortfall Optimized Portfolios algorithm) – left tail of the return distribution;
- > maximize the expected return while using the volatility as a proxy for risk (Efficient Frontier algorithm) – an algorithm that is concentrated on the center of return distribution;
- > maximize the extreme gains during rallies (Windfall Optimized Portfolios) – the right tail of the return distribution.

The original SWOP model was developed by well-known professors in finance, Michael Rockinger and Eric Jondeau, who participated substantially in the implementation stage of the project.

The idea of Unigestion was for its FIO department to use the SWOP module after its implementation in parallel with the old Efficient Frontier module, compare the results of the two algorithms and, if the results made logical and financial sense, migrate to use the SWOP module only and drop the Efficient Frontier module. Currently SWOP is the main numerical algorithm used for modeling and re-balancing customer portfolio asset allocation. The SWOP product is now mature and has proved to be a stable, reliable and very useful tool.

3 SWOP IMPLEMENTATION

The negotiations between SAGE and Unigestion on implementing the SWOP module and integrating it into the GAMS information system (also a product of SAGE S.A.) started in early 2006. After presenting the proof of concept and holding some theoretical seminars with senior managers, the contract to implement the SWOP module at Unigestion was signed in February 2006.

The scope of the project was to create a module that implemented the theoretical SWOP model with a graphical user interface (GUI) that would allow up to 100 optimized portfolios given different investment constraints to be computed and graphed in less than a minute. The requirement for such a short response time was dictated by the need to run the optimization in front of clients. Another requirement was that the resulting GUI should include the Efficient Frontier optimized portfolios in order to compare the SWOP optimized portfolios with the EF allocation. The GUI should also contain certain statistics pre-computed for each optimized portfolio and the possibility to export the optimization results back to the GAMS information system, where they would be further used for portfolio allocation.

Due to the high computational burden that would be imposed by the SWOP module and the tight implementation schedule, it was decided that it would be implemented in Matlab, a high-level programming language with lots of existing mathematical and statistical functions present in different Matlab toolboxes and well as a powerful optimizer.

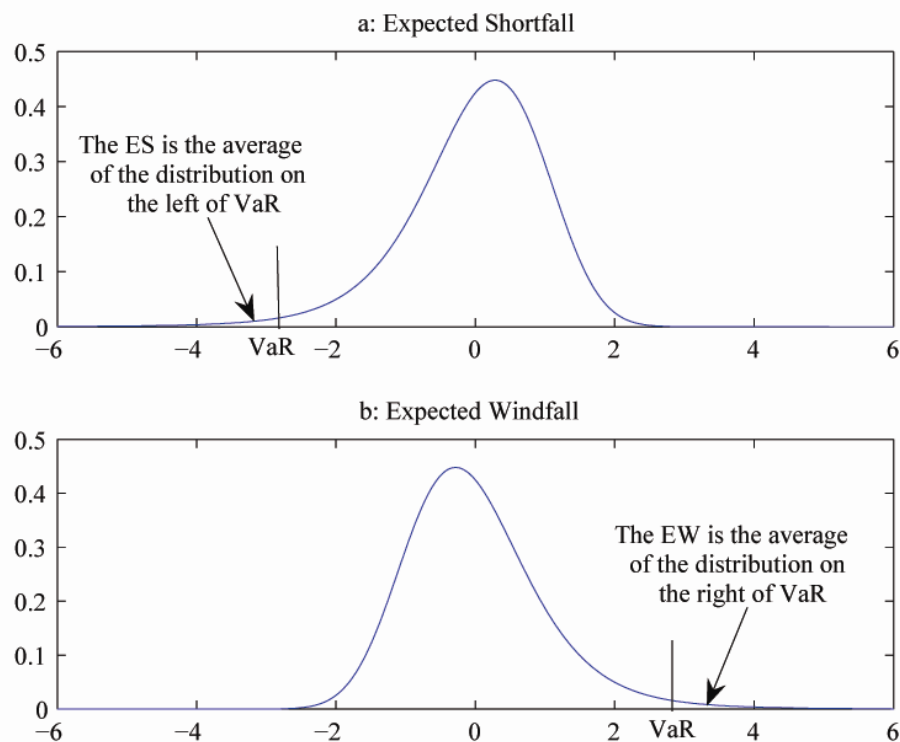
Generally, the implementation of the SWOP algorithm was over after some 3 months of work with the continuous support from professors Rockinger and Jondeau. The rest of the implementation time was dedicated to adding user interface options that allow users to add another portfolio for comparison purposes onto the graph, change the weight of existing portfolios, change the size or colour of existing portfolios on the graph, save the weights (allocation) of existing optimized portfolio, compute statistics of all optimized portfolios (like month to date, year to date, best and worst 1, 3, 6, 12 months rolling performance, Sharpe and Sortino ratios, asset class and currency distribution, etc) and so on.

At the end of November 2006, the final version of the module was integrated into the existing information system and went live without major issues. Stage 2 of the project – maintenance and enrichment of the SWOP module interface – went on for another year and there have been no bug reports on the module since February 2008: thus the project is currently in a mature stage with all the infant problems already solved.

Currently, the SWOP module is a mature module with proved efficiency and reliability. It can be installed at any client that has a financial assets data feed such as Bloomberg, Reuters or Telex, as a stand-alone application, or integrated into any legacy information system currently in use.

4 THE BASIC PRINCIPLES OF SWOP

The SWOP model starts by answering the following question: what is the risk of a portfolio of financial assets? As mentioned earlier, volatility is not a good enough answer because a portfolio that has encountered large positive returns (right-hand side of the portfolio distribution) in the past, will have a high volatility even if the negative returns (left-hand side of the portfolio distribution) had a much smaller magnitude. SWOP defines the portfolio risk as the average loss the portfolio incurs in the $n\%$ worst case scenario (n is normally chosen to be 1% or 5%), called Expected Shortfall. This means looking at the left side of the portfolio distribution and recognizing that a portfolio of financial asset returns is not normally distributed, and thus the higher moments of asset return distribution – the skewness (prevalence of positive returns over negative ones or vice versa) and kurtosis (a high probability of extreme positive or negative returns) – should be taken into consideration.



Finally what the SWOP algorithm does is to search for a portfolio that:

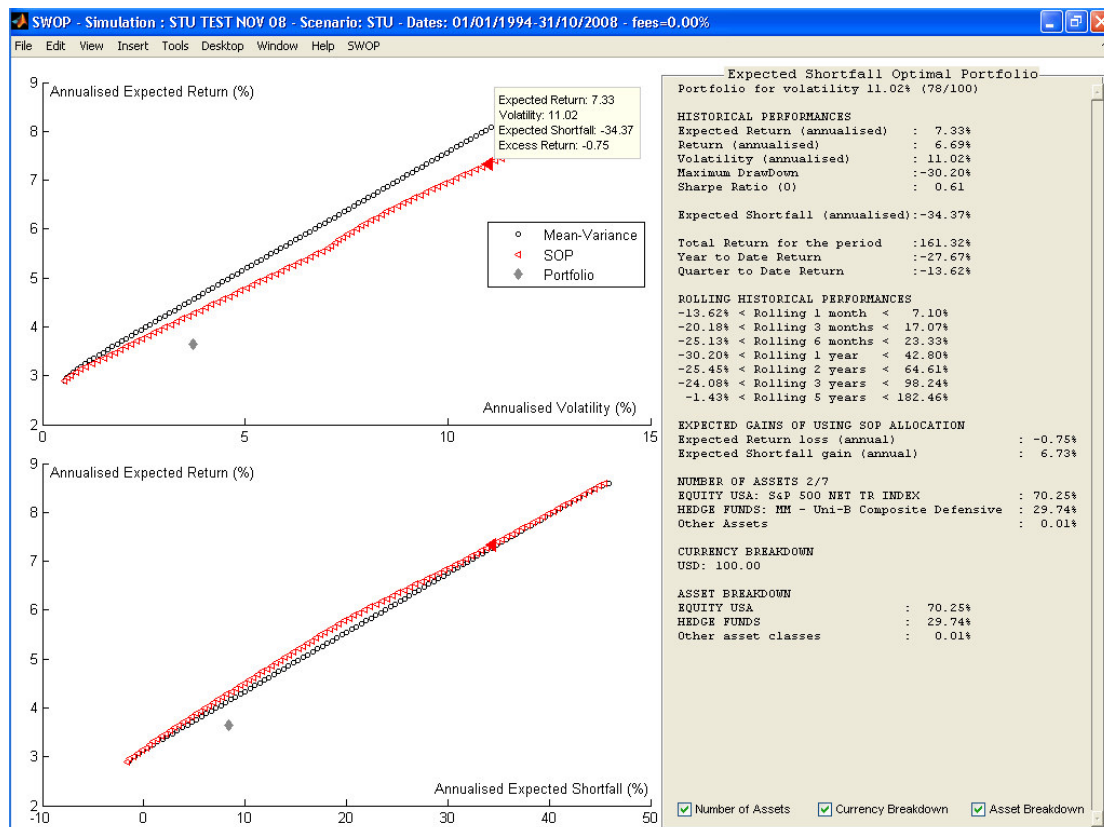
- > for a given level of expected return has the lowest level of Expected Shortfall (the Shortfall Optimized Portfolio – SOP algorithm), or
- > for a given level of volatility has the highest level of Expected Windfall (the Windfall Optimized Portfolio – WOP algorithm).

The SWOP module starts from the definition of the so-called investor's investment universe – the totality of the financial assets where the investor is willing to invest. Then, using the historical assets' returns to infer information about the distribution of each asset's returns, the SWOP algorithm will try combining these assets with different weights (the weights of all the assets should be of course be equal to 1, short-sale restriction is applied) until it finds the weights that make a portfolio that has the minimum expected shortfall of all the possible portfolios made of these securities. The result of the optimization is thus a vector of asset weights that define the optimized portfolio for different levels of expected returns. The WOP (windfall optimized portfolios) algorithm, as opposed to the SOP (shortfall optimized portfolios) algorithm, will produce vectors of weights that for different levels of volatility have the highest level of expected windfall – the expected level of gain during economic or financial booms.

As mentioned at the beginning of this paper, the basic principles of SWOP are easy to understand. If we change one of the input parameters (say the expected return for one of the assets), the change in the output parameters can be explained and will make logical sense, which makes it very easy to deal with, even though there are serious mathematical computations that run during the SWOP optimization. This means that this makes SWOP a good tool in the hands of a portfolio manager, as compared to overly complicated models that the manager cannot understand and thus explain: if something goes wrong, he does not know exactly what went wrong with those models.

5 EF VERSUS SWOP

This section will provide a short overview of the difference in allocations suggested by the legacy Mean-Variance algorithm previously used at Unigestion and the SWOP allocation that Unigestion currently uses. The graph is a typical result screen of a SOP optimization (just as a reminder, SWOP is composed of two separate components – SOP and WOP – that can be run separately).



The SWOP GUI screen is divided into three regions:

- > the region on the right side of the screen with the properties and relevant statistics of the optimized portfolio on which the user has parked;
- > the top left region with the mean-variance criteria optimized portfolios in black and the shortfall optimized portfolios in red in the volatility vs expected return space, and
- > the bottom left region with the same mean-variance and SOP curves but this time in the more usual SWOP space – expected shortfall (the measure of risk) vs expected return (the return measure). This also explains why the SOP curve is smoother on the bottom curve.

What is worthwhile observing in the graph above is that for a certain level of volatility (on this graph there are a total of 100 different levels) we have two optimized portfolios: a mean-variance one and a shortfall optimized one. Particular attention should be paid to the information area called 'Expected gains of using SOP allocation' on the left middle part of the screen. The numbers here say that, for the same level of variance, the SOP algorithm proposes an asset allocation that will result in:

- > 0.75% loss of expected return as compared to the mean-variance (MV) criteria, and
- > 6.73% gain in expected shortfall.

For the portfolio manager, it means that if he/she chooses a SOP portfolio allocation over a MV allocation, the difference between the return of the MV and SOP portfolio is expected to be 0.75%. But conversely, in the 5% worst-case scenario (usually happening during crises) the SOP portfolio will lose 6.73% less than the MV portfolio. It is perfectly understandable that a gain in expected shortfall will come at the cost of a loss in expected return, but the magnitude of the gain versus the cost (6.73% gain vs. -0.75% loss) is nevertheless impressive.

Although the benefit in monetary terms to Unigestion in using the SWOP module cannot be divulged, it is worth noting that before the 2008 financial crisis the SWOP algorithm was proposing to invest a much lower portion of assets in equities. This advice enabled the clients' portfolios at Unigestion to perform better during the crisis than if they had followed the mean-variance criteria alone.

SWOP's objective is to maximize the investors' wealth while keeping the portfolio risk under control. As the example illustrated in the graph shows, there are real, measurable benefits to using SWOP over the basic mean-variance optimization.

The risk control and investment strategies employed at Unigestion have enabled to generate relative good performance, as almost all client portfolios managed by the FIO department have outperformed the market indexes.

6 CONCLUSIONS

As can be seen from this case study, the SWOP algorithm is a much more advanced tool for asset allocation than the standard mean-variance method that brought real added value to Unigestion in increasing their clients' wealth. Among factors that make SWOP a very practical and useful tool for portfolio managers are:

- > it is based on a sound numerical algorithm of minimizing the risk of losing a large amount of money during crisis;
- > it has the possibility to obtain an optimized allocation for any level of investor risk aversion and with any investment constraints;
- > it is possible to compare different portfolio characteristics on-line before deciding which one should be used as a model for the client portfolio allocation.
- > the optimized allocations proposed by SWOP have a smaller loss in expected return than a gain in expected shortfall as compared to the allocations obtained by the simple mean-variance optimization;
- > the results of the optimization are ready for use in some 10-20 seconds (for an investment universe of 10-20 asset classes) and can be thus run in front of the client while changing different constraints in order to see how they influence the results;
- > it is a mature solution that has a proven track record for a couple of years in a financial institution.

SWOP is certainly not the only solution for quantitative asset allocation currently present on the market. There are also thousands of theoretical papers that recognize the importance of extreme events and provide mathematical models for minimizing the risk and increasing return of a portfolio. With respect to the implemented asset allocation solutions already present in the market, many can be distinguished either by overly simplified models that do not take into account all four moments of the asset return distribution (expected return, volatility, skewness and kurtosis) or by a very long response time (usually tens of minutes to hours) and thus computations generally have to be run overnight in order for the results to be used during the day (and as such are designated only to quantitative professionals).

The majority of the theoretical papers (many of them are very advanced) either can be used only for a very limited number of assets, and are thus not applicable in a real financial institution, or are hard to understand for a non-quantitative professional, which renders them intractable to portfolio managers as the influence of input parameters on the results is hard to understand and explain.

The SWOP module, on the other hand, uses the risk, return, skewness and kurtosis distributional properties that are relevant for minimizing the potential loss (standard and quicker models use only the first 2 or 3 of them). It can take a large number of assets (at the expense of the response time, however). The computational algorithm has been optimized to give a response in up to 20 seconds (for around 10-20 asset classes) with a graphical interface representation of results with many additional options for interacting with the program and testing the results. In addition, the program code is written in a high level programming language (Matlab) and is thus easily maintainable and is very flexible to changes, add-ons and improvements.

The above-mentioned factors allowed the SWOP module implementation in Unigestion to become a success story. After seeing the benefits of SWOP as compared to the MV criteria, Unigestion has migrated its benchmarking and portfolio asset allocation tool to SWOP and is extending the use of SWOP to its offices in other countries. Further improvements to SWOP are envisaged to be implemented in 2009, which will make it an even more successful tool for financial institutions of any size can benefit from.