Full-scale optimisation: applications for public investors[1]

Alejandro Barrios Heras[2], Diana Karolyn Acevedo Becerra[3], Carlos Antonio Cano Cordova[4] and Mike McMorrow[5]

Abstract

Central banks and other public investors typically employ portfolio construction techniques as part of their strategic asset allocation processes, usually relying on one or more summary risk measures to identify and select optimal portfolios. Full-scale optimisation (FSO) is an alternative that allows the entire distribution of expected returns to determine an optimal portfolio by maximising expected utility. This requires the investor to select a utility function and set the necessary parameters based on their preferences. We conduct portfolio construction exercises for a stylised reserve manager using FSO and compare the outcomes with those derived from more conventional approaches, such as expected shortfall and variance minimisation. When the utility function is solely dependent on returns and asset return distributions are elliptical, FSO produces results comparable with traditional risk-minimisation methods. However, the advantages of FSO become more pronounced when broader utility dimensions are considered or when return distributions deviate significantly from ellipticity. Furthermore, we discuss how the performance of the differential evolution algorithm – employed in the FSO process – is sensitive to the initial population configuration, and we explore some strategies to enhance its robustness and efficiency.

JEL classifications: C61, G11.

1 The authors are grateful to Murad Tovmasyan, Simone Letta, Tommaso Perez and Patrizia Ferrauto for helpful comments. We would like to express our gratitude to the institutions that participated in the survey on utility functions and investment preferences. In addition, we wish to acknowledge our colleagues from the reserve and asset management teams at the Bank of Mexico, the Central Reserve Bank of Peru and the Bank for International Settlements, whose valuable contributions to their respective strategic asset allocation processes fostered an environment in which this work could flourish. Any remaining errors are the responsibility of the authors. The views expressed are those of the authors and do not necessarily reflect those of their respective institutions.

2 Bank of Mexico, abarrios@banxico.org.mx.

3 Central Reserve Bank of Peru, diana.acevedo@bcrp.gob.pe.

4 Central Reserve Bank of Peru, carlos.cano@bcrp.gob.pe.

5 Bank for International Settlements, mike.mcmorrow@bis.org.

1. Introduction

Over recent decades, financial professionals – including public investors such as central bank reserve managers – have employed quantitative models for portfolio construction. Among the first – and most prominent – is the Markowitz model, which represents a mean-variance optimisation framework. Other approaches have been considered, either to focus on a risk metric that better represents the tail risk of the return distribution to integrate views on the financial and economic environment in the distribution of returns, or to have a more balanced portfolio in terms of risk contributions. These methodologies have a common feature in that they seek to identify an optimal allocation based on a single, summary risk metric in the objective function. Therefore, the result can be considered only an approximation of the optimal solution for the complete distribution of returns. With this idea in mind, and with the objective of fostering flexibility and precision when specifying the risk and return preferences, we introduce the application of full-scale optimisation (FSO) into the investment management framework.

Full-scale optimisation, introduced by Cremers et al (2004), is an alternative optimisation framework to maximise the investor's expected utility. This approach boasts several advantages, including its flexibility in directly incorporating the investor's full set of preferences into the objective function. Notably, FSO allows for the representation of an investor's risk profile, with various parametrisations of utility functions reflecting different levels of risk appetite based on the monotonicity of their risk aversion coefficient. An additional advantage of FSO is that, by definition, this technique takes into account the complete probability distribution of returns rather than solely relying on a set of metrics such as expected return, volatility and other risk measures like the conditional value at risk (CVaR). This becomes particularly advantageous when handling asymmetric return distributions with heavy tails.

Applications of FSO for investment decision-making purposes have been explored in the literature, though with varying objectives. For example, Adler and Kritzman (2006) presented evidence that the in-sample superiority of full-scale optimisation over the mean-variance framework prevails out of sample. Hagströmer and Binner (2009) then extended this conclusion for the selection of stocks. Thereafter, Seymour et al (2014) elaborated on the application for the construction of an optimal hedging strategy for equity and foreign exchange exposure.

This paper seeks to add to the existing literature by exploring FSO in a central bank and public investment management context. Central bank experience with FSO appears limited. A survey of selected central bank reserve managers in support of this paper suggests few have attempted to incorporate FSO into the reserve management process, though some indicate an interest in doing so. In principle, this method may be useful for establishing a connection between the three traditional reserve management objectives – liquidity, safety (sometimes termed capital preservation) and return – and the outcome of portfolio optimisation exercises, insofar as these objectives can be represented in a utility function.

2. Utility functions in a reserve management context

The first challenge that public investors such as central bank reserve managers will face when considering FSO is to define and parametrise a utility function. This may prove challenging from an institutional perspective, given the need to find common agreement among stakeholders on a technical subject potentially requiring multiple specific decisions. It should be noted that reserve managers employing common techniques have implicitly selected a utility function in one form or another, for example mean-variance optimisation gives rise to maximising quadratic utility – which may not have fully desirable features.

A multitude of utility functions have been introduced, debated and used in the literature (eg Pratt (1964), Arrow (1971)). We briefly discuss the features of five – quadratic, exponential, isoelastic (also known as a power function), kinked log and s-shaped – before homing in on those that may be particularly useful for central bank reserve portfolios or public investors more generally.

2.1. Common utility functions in the portfolio optimisation literature

As presented below, many of the commonly used utility functions serve as simplifications to describe specific types of risk aversion. In this work, we define utility functions in terms of returns instead of the alternative definition based on wealth, in order to directly relate the utility level to an investment portfolio selection. This is straightforward, considering the increasing relationship between wealth (w) and return (r), $w = 1 + r$.

a. Quadratic

The assumption of this form of utility function is embedded in different portfolio optimisation frameworks. In particular, the similarity between a framework where the quadratic utility function is maximised and the mean-variance optimisation framework, has been highlighted in different research. This feature is derived from the representation of risk preferences by only two metrics, the mean and the variance of the distribution of returns. In this case, the λ parameter reflects the investor's degree of risk aversion.

$$U(r) = (1 + r) - \lambda\,(1 + r)^2$$

A quadratic utility function exhibits some features that an investor may see as potential disadvantages. In particular, it reflects increasing absolute risk aversion,[6] ie a reduction of the exposure to risky assets as the investor's wealth increases. It can also be shown that the quadratic utility function exhibits increasing relative risk aversion.[7] An additional disadvantage of the quadratic utility function is that it implies the existence of a level of satiation.

6 Absolute risk aversion measures the magnitude of risk aversion for a specific level of wealth. It is defined as the negative ratio between the sensitivity of marginal utility relative to wealth and the marginal utility: $ARA(w) = -U''(w)/U'(w)$.

7 Relative risk aversion measures the degree of change in the investor's willingness to take on risk in response to changes in wealth. Mathematically, it is defined as the negative percentage change in the utility in response to a percentage change in wealth: $RRA(w) = -w \cdot U''(w)/U'(w)$.

b. Exponential

The exponential utility function is an alternative to represent constant absolute risk aversion, meaning that the investor will have the same exposure to risky assets for any given level of wealth. In this case, the parameter A corresponds to the level of absolute risk aversion.

$$U(r) = 1 - e^{-A(1+r)}$$

Having constant absolute risk aversion may be considered unrealistic. As a result, other alternatives, such as the power (isoelastic) utility function, have appeared as well.

c. Power (isoelastic)

This utility function is commonly used to represent the preferences of a risk-averse investor, since it presents a decreasing absolute risk aversion and it reflects a decreasing marginal utility with respect to wealth, ie the marginal utility gained from an additional unit of wealth decreases. It is worth noting that, under this function, the utility level is non-negative for positive levels of wealth.

$$U(r) = \begin{cases} \dfrac{(1+r)^{1-\gamma} - 1}{1-\gamma}, & \gamma \geq 0, \gamma \neq 1 \\ \ln(1+r), & \gamma = 1 \end{cases}$$

Finally, the power utility function has constant relative risk aversion. In fact, under this specification, the parameter γ corresponds precisely to the level of relative risk aversion.

d. Kinked log

The kinked log utility function is a natural choice to represent changes in the investor's risk profile for different ranges of wealth. This feature has been incorporated to provide more flexibility in the definition of the utility function based on the factors that shape an investor's preferences. It is also closely related to the idea of loss aversion observed in the behaviour of different types of investors.

Mathematically, the shift in the investor's risk profile is incorporated as a point of non-differentiability that distinguishes the levels of wealth or return around a reference level, as shown in the following formula.

$$U(r) = \begin{cases} \mathrm{P}(\mathrm{r} - \theta) + \ln(1+\theta), & r < \theta, P > 0 \\ \ln(1+r), & r > \theta \end{cases}$$

This kinked log utility function captures preferences that are strongly influenced by a critical threshold θ, which acts as a pivotal point for the agent's utility. The function exhibits distinct behaviours on either side of this threshold, reflecting the agent's sensitivity to falling below or above it.

For values of r below θ, the utility function incorporates a linear penalty with slope P. This penalty reflects a strong aversion to outcomes that fail to meet the critical threshold, making the utility decline sharply as r moves further below θ. The parameter $P > 0$ determines the steepness of this penalty, highlighting the importance of maintaining outcomes above the threshold.

Above θ, the utility transitions to a logarithmic form, which is characteristic of diminishing marginal utility. This logarithmic utility reflects a standard risk-averse behaviour, where the agent derives progressively smaller increases in utility as r grows larger. This smooth, concave shape captures the agent's reduced sensitivity to additional gains as outcomes improve further.

A key feature of this utility function is the discontinuity in its slope at $r = \theta$. At this critical point, the sharp change in slope reflects the agent's heightened sensitivity to whether outcomes fall above or below the threshold. This creates a "kink" in the utility curve, marking a significant behavioural shift around $r = \theta$.

Overall, this kinked log utility function is particularly well suited for modelling scenarios in which there is a minimum acceptable level of performance or return, such as financial thresholds, consumption minimums or investment goals. It effectively captures the penalties or sharp losses in utility that occur when outcomes fall short of the critical threshold while maintaining standard risk-averse behaviour for higher outcomes.

e. S-shaped

The s-shaped utility function is another alternative to represent a shift in the risk attitude of an investor for different ranges of wealth or return, in particular, for gains or losses. In this case, instead of showing risk neutrality, the investor is risk-averse when making gains, and risk-seeking when making losses. This framework for the risk profile of an investor is consistent with the behavioural finance anomaly known as the disposition effect, which describes the tendency of investors to sell assets that have increased in value, and keep assets that have presented losses.

The parametrisation of this utility function also includes a critical threshold, θ, and gives more flexibility to represent the attitude towards risk changes by setting parameters A, B>0, γ_1 and γ_2:

$$U(r) = \begin{cases} -A(\theta - r)^{\gamma_1}, & r \leq \theta \\ B(r - \theta)^{\gamma_2}, & r > \theta \end{cases}$$

It is clear that the utility functions presented above are simplistic representations of an investor's preferences, and many of them rely mostly on the concept of risk aversion to determine levels of satisfaction. Moreover, even after a utility function is determined to be a good fit, its parametrisation involves more research.

In this context, a survey[8] was conducted to further analyse the potential application of full-scale optimisation in the public investor community, seeking specifically to identify and incorporate those features that are important to central

8 Survey conducted in the first half of 2024. The conclusions are based on the responses of the 22 central banks that participated.

bank reserve managers when selecting an optimal portfolio. The main findings are presented in the following section.

2.2. Survey on risk preferences and utility functions

The survey was designed to identify the features that are relevant for central banks when performing their strategic asset allocation (SAA), and to spot opportunities for the application of full-scale optimisation in their SAA frameworks. The main survey responses are presented in the following subsections.

a. Relevant attributes and parameters that shape investment preferences

The survey has been useful to highlight the most relevant attributes and parameters for central banks based on their specific objectives, constraints and risk tolerance. In line with the results from other relevant surveys on reserve management, the survey showed that central banks tend to prioritise capital preservation (or safety), liquidity and return, which are typically considered their main objectives. Moreover, in line with the increasing interest in sustainable investment in the public investor community, sustainability factors were listed as another attribute to consider when determining an optimal allocation. Respondents also mentioned more general attributes such as risk tolerance and an inflection point that reflects a shift in their attitude towards risk.

Relevant attributes and parameters when constructing a utility function to represent investment decision-making process[9]

Graph 1

Percentage of respondents

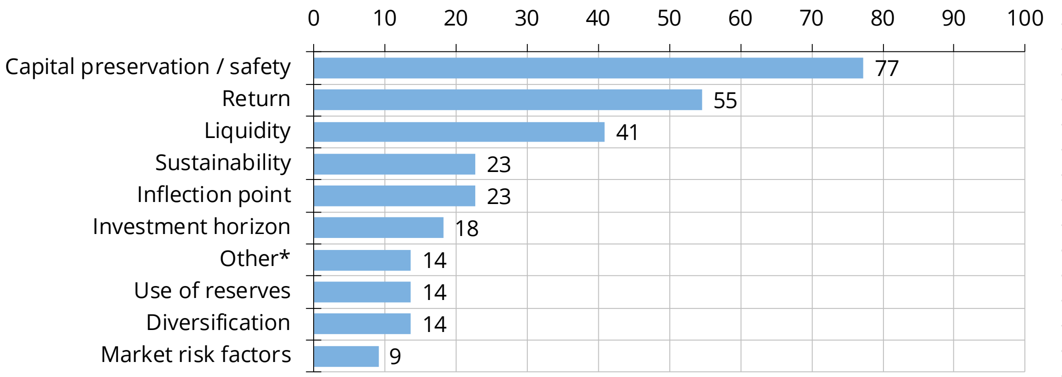

* Macro factors, portfolio size and the portfolio's numeraire.

Source: Survey of 22 central banks on investment preferences and utility functions, conducted by the authors during 2024.

[9] The actual question was: "Given your institution's specific objectives, constraints and risk tolerance, could you specify which attributes and parameters are important when constructing a utility function to represent your investment decision-making process?"

b. Preferred utility functions

Another aspect of interest was the familiarity of respondents with how utility functions can reflect investment preferences, and in particular, whether there are any utility functions that fit their investment profile.

Utility functions that fit a central bank's investment profile[10]

Graph 2

Percentage of respondents

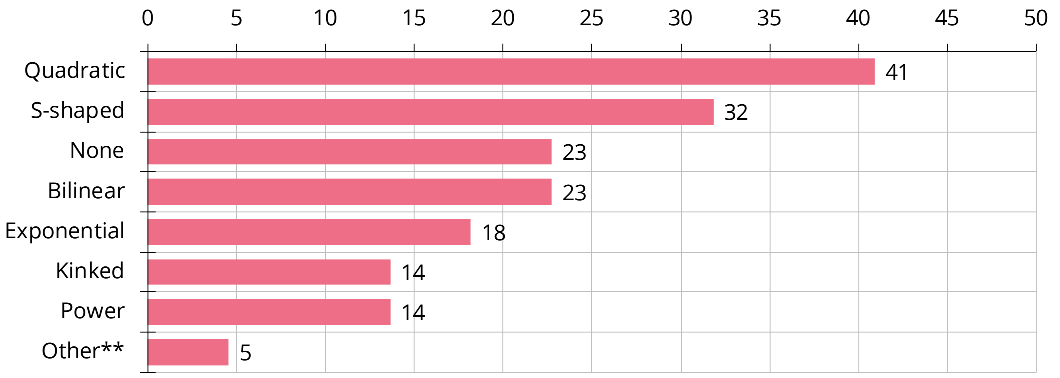

** A function that incorporates the expected shortfall minimisation.

Source: Survey of 22 central banks on investment preferences and utility functions, conducted by the authors during 2024.

As shown in Graph 2, the quadratic utility function is the most common representation of the risk preferences of central banks within their optimisation frameworks. Nonetheless, there is notable interest among the sample of respondent central banks for considering utility functions that reflect a change in attitude towards risk for different levels of wealth or return. Utility functions such as the bilinear, s-shaped and kinked are regarded as better representing their investment profiles.

c. Changes in the definition of utility functions

The FSO framework introduces additional considerations into the SAA process. In particular, investors update their strategic asset allocation to reflect new economic and financial developments that may affect their performance. In this sense, the FSO framework could help investors not only to directly update inputs such as return forecasts, but also to update their risk profile by redefining the utility function to be maximised.

[10] The actual question was: "From the well-known utility functions available in the literature, are there any that fit, or you believe would fit, your investment profile particularly well?"

Reasons for changes in the utility function definition[11]

Graph 3

Percentage of respondents

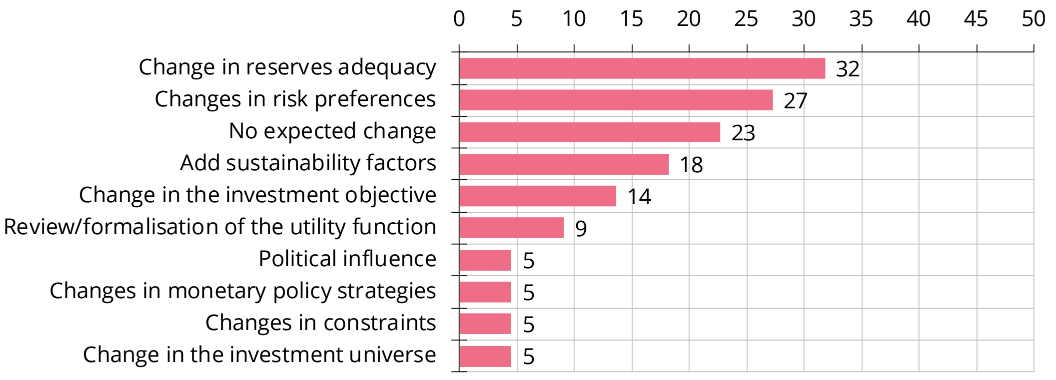

Source: Survey of 22 central banks on investment preferences and utility functions, conducted by the authors during 2024.

When asked to envision any potential reasons for changes in the utility function definition, whether its specification or parametrisation, respondents highlighted changes in the metrics of reserve adequacy and changes in risk preferences. The addition of sustainability factors to the optimisation framework was also among the more frequent responses. Many respondents, however, mentioned that they do not anticipate changes in their preferred utility function.

3. Applying full-scale optimisation to a stylised central bank portfolio

We explore the potential application of FSO in a public investment management framework by running several portfolio construction exercises. We begin by considering a stylised reserve manager. This central bank has a fairly wide range of eligible assets, including US money market instruments, government bonds, SSAs, agency mortgage-backed securities (MBS), corporates, equities and gold, as well as sovereigns issued by several other major economies. Portfolios are subject to a number of constraints.[12] They rely on data randomly sampled from the kernel distributions estimated from historical data during the period January 2004–April 2024. Five thousand simulated annual returns were generated using non-overlapping historical monthly returns and a t-copula model with Kendall's tau correlation[13] and assuming zero autocorrelation.

[11] The actual question was: "Do you have any reason to think your preferred utility function (or its parametrisation) could change over time?"

[12] Minimum 50% in US Treasuries and Treasury Inflation-Protected Securities (TIPS), maximum 35% in supranationals and agencies, maximum 30% in MBS, maximum 30% in corporate bonds and maximum 25% in equities.

[13] The computation of Kendall's tau correlation was derived from the process proposed by Lindskog et al (2003).

This stylised public investor has selected a kinked log utility function, given its simplicity and the need for specifying only two parameters that are fairly easily understood, as presented in the previous section. As mentioned above, this utility function can be associated with loss aversion, given that utility exhibits an accelerated decline when faced with negative returns. For the purposes of this document, a differential evolution (DE) algorithm is implemented to identify an optimal portfolio (see Section 4). Results under the described initial framework suggest a number of conclusions.

First, the parametrisation of the utility function is of crucial importance. Exercises based on the kinked log utility were run for varying levels of P, thereby imposing different degrees of penalty on losses versus similar gains.

Additionally, the optimal portfolios derived under the FSO approach exhibit strong alignment with those on the mean-variance and mean-CVaR efficient frontiers within the corresponding risk-return space.

The utility function representing the preferences of this stylised central bank is characterised by asymmetric risk preferences, exhibiting a pronounced aversion to losses relative to an inclination for gains. This asymmetry originates from the penalty term embedded within the utility function, which disproportionately penalises losses, thereby rendering downside risk a more significant determinant of optimal portfolio selection. Consequently, the structure of the utility function, rewarding elevated returns while penalising losses, closely mirrors the mean-CVaR trade-off between expected return and downside risk. This conceptual congruence provides a robust explanation for the observed alignment between the optimal allocation under the FSO approach and the mean-CVaR efficient frontier.

Perhaps more intriguing is the observation that the FSO portfolio also lies close to the mean-variance efficient frontier, despite the fact that mean-variance optimisation is typically associated with symmetric risk preferences and does not explicitly penalise downside risk. This finding, however, can be rationalised by noting that the asset return distributions in this study do not exhibit significant departures from normality. In such cases, the assumption of elliptical distributions is justified.

Under elliptical distributions, any linear combination of asset returns – such as a portfolio return – is itself elliptically distributed and fully characterised by its mean and variance. The probability density of these distributions is symmetric around the mean, and higher-order central moments (eg skewness and kurtosis) are either zero or functionally dependent on the variance, rendering them redundant for characterising the return distribution.

Consequently, the expected utility associated with any strictly increasing utility function defined over portfolio returns depends solely on the mean and variance, provided that asset returns are elliptically distributed. This result is formally established by Chamberlain (1983), who demonstrates that under elliptical symmetry, utility rankings are fully determined by the first two moments, regardless of the specific form or smoothness of the utility function. Ingersoll (1987) further supports this conclusion by demonstrating that, under elliptical return distributions, the expected utility of any strictly increasing and integrable utility function – whether quadratic, discontinuous or kinked – can be expressed as a function of the portfolio's mean and standard deviation alone. This conclusion is further substantiated by Meucci (2005), who shows that when the asset returns are elliptically distributed, the investor's satisfaction, irrespective of the specific functional form of the utility function, depends exclusively on the mean and variance of the portfolio return. In this context, the otherwise infinite-dimensional space of higher-order moments collapses onto a two-dimensional manifold fully characterised by these first two moments, justifying the use of the mean-variance framework as an exact representation of utility maximisation under ellipticity assumptions. Accordingly, the proximity of the utility-maximising portfolio – derived using a kinked log utility function – to the mean-variance frontier is consistent with theoretical expectations. In summary, when asset returns conform to an elliptical distribution, the presence of non-linearities or asymmetries in the investor preferences does not necessarily imply a departure from mean-variance efficiency.

An empirical examination of the statistical properties of the asset return data utilised in this portfolio optimisation exercise does not reveal substantial deviations from the normal distribution assumption, thereby reinforcing the validity of the results.

Given the similarity of the FSO solution to the optimal portfolio obtained under the mean-variance and mean-CVaR frameworks, we performed additional exercises to understand the circumstances under which FSO may provide unique value above and beyond more conventional methods.

First, to assess the degree of influence of the choice of eligible assets, a set of alternative assets[14] was included in the optimisation problem. As shown in the table below, even under the extended universe of eligible assets, the FSO portfolio is still close to the portfolio in the efficient frontiers with the same level of return.

Optimal allocation under selected methodologies and asset universes

Table 1

In per cent

Assets and metrics	Traditional assets			Traditional and alternative assets		
	(1) FSO	(2) Min variance with equal mean	(3) Min CVaR with equal mean	(1) FSO	(2) Min variance with equal mean	(3) Min CVaR with equal mean
US Treasuries	50.0%	50.0%	50.0%	50.0%	50.0%	50.0%
Agencies and supranationals	34.0%	31.9%	33.4%	8.0%	9.9%	7.4%
US corporate high yield	0.7%	3.8%	0.0%	0.0%	0.0%	0.0%
Gold	5.0%	4.8%	6.4%	7.3%	5.5%	9.2%
US equities	10.4%	9.6%	10.2%	9.7%	9.6%	8.5%
Alternatives	0.0%	0.0%	0.0%	25.0%	25.0%	25.0%
Expected return	3.54%	3.54%	3.54%	4.59%	4.59%	4.59%
Volatility	2.84%	2.82%	2.85%	3.29%	3.26%	3.41%
CVaR (99%)	4.09%	4.14%	4.03%	4.32%	4.65%	4.15%
Expected utility	2.28%	2.26%	2.27%	3.46%	3.42%	3.38%

Sources: Bloomberg; authors' computations.

The preceding analyses yield significant insights; however, institutional investors frequently encounter scenarios in which the adoption of the FSO approach offers

[14] The alternative assets considered were: private equity, equity strategies (long/short and market neutral), global macro strategies, funds (equal, weighted and diversified), event-driven, arbitrage (convertible, fixed income and merger), commodities, distressed securities, global property and global infrastructure. A constraint for the maximum exposure to alternative assets was set as 25%.

distinct advantages over traditional mean-variance and mean-CVaR methodologies. For instance, the FSO framework demonstrates considerable efficacy in portfolio construction scenarios characterised by non-linear, multi-objective or complex investor preferences, which often surpass the capabilities of traditional mean-variance and mean-CVaR models. By incorporating advanced utility functions, the FSO approach facilitates the simultaneous optimisation of multiple dimensions, enabling investors to balance risk and return while integrating additional objectives, as exemplified in the case study presented here. Drawing upon the methodology proposed by Hayes et al (2015), this study incorporates an illiquidity penalty factor that is contingent upon both the marginal cost of illiquidity and the overall illiquidity level of the portfolio. The primary objective of this approach is to explicitly address liquidity constraints within the portfolio allocation process by integrating the investor's liquidity needs and preferences directly into the utility function.

The marginal cost of illiquidity is conceptually defined as the additional return that an incremental unit of an illiquid asset must generate – relative to a theoretically equivalent liquid asset – in order to be deemed preferable. This cost function is dynamic, increasing with the overall illiquidity level of the portfolio. The illiquidity level of the portfolio is computed using the following formulation:

$$\mathrm{L}(\mathrm{w}) = \sum_{\mathrm{i}=1}^{\mathrm{n}} \mathrm{w_i l_i} = \mathrm{w^T l}\,, \quad \mathrm{w}^{\mathrm{T}}\mathbf{1} = 1 \text{ and } \mathrm{w} \geq 0 \text{ for all } i.$$

Where n is the number of assets, $\mathrm{l_i}$ is the illiquidity penalty for asset i, l is a vector of illiquidity penalties, w is the asset weighting vector and $\mathbf{1}$ is the n-vector with all elements equal to one. The values of $\mathrm{l_i}$ range from zero to one, where zero is extremely liquid and one is extremely illiquid. Notably, the assignment of illiquidity penalties remains at the discretion of investors, thereby allowing for customised adjustments based on individual investment objectives.

In this case study, the marginal cost of illiquidity is represented by the power function $\mathrm{c(L)} = \frac{3}{2}\mathrm{L}^{3/2}$. This implies that illiquidity cost increases at an increasing rate with the amount of illiquid assets in the portfolio, which represents an investor with high aversion to illiquidity.

Marginal cost of illiquidity — Graph 4

In per cent

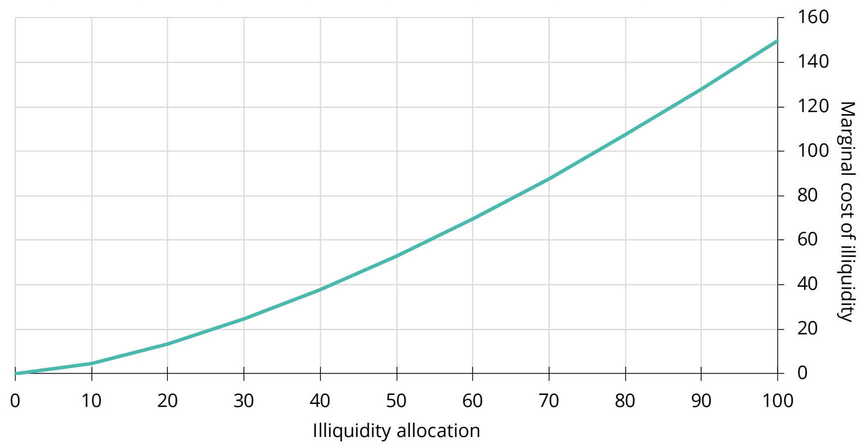

Source: Authors' computations.

The cumulative illiquidity cost function $A(L)$ is derived by integrating the marginal cost function $c(L)$ from 0 to $L(w)$:

$$A(L) = \int_0^{L(w)} c(x)dx$$

For this case study, the accumulated illiquidity cost function is represented as follows:

$$A(L) = \int_0^{L(w)} \frac{3}{2}x^{3/2}dx = \frac{3}{5}L(w)^{5/2}$$

The investor's base utility function U, which represents their fundamental risk-return preferences, is then adjusted to account for the illiquidity penalty:

$$\widehat{U} = U - A(L(w))$$

For the purposes of this study, the base utility function is assumed to be the kinked log. Consequently, the liquidity-adjusted utility function is expressed as:

$$\widehat{U}(x) = \begin{cases} \ln(1+r) - 0.6L^{5/2} & , \quad x \geq \theta \\ \ln(1+\theta) + 10(r-\theta) - 0.6L^{5/2}, & \quad x < \theta \end{cases}$$

where r is the return of the portfolio.

To conduct the empirical analysis, a comprehensive set of asset classes was utilised. As before, 5,000 simulated annual returns were generated, and in this case, a predefined liquidity hierarchy was established, ranking assets from the most liquid to the least liquid as follows: US Treasuries (UST), agency MBS, agencies and supranationals, US equities, US corporate bonds, commodities and alternatives.

Optimal allocation under selected methodologies for the extended asset universe — Table 2

In per cent

Assets and metrics	(1) FSO	(2) Min variance with equal mean	(3) Min CVaR with equal mean
US Treasuries	66.6%	50.0%	50.0%
Agencies and supranationals	0.0%	16.0%	14.8%
US corporate high yield	0.0%	0.0%	0.0%
Gold	4.4%	3.2%	5.2%
US equities	9.3%	5.8%	5.0%
Alternatives	19.7%	25.0%	25.0%
Expected return	3.87%	3.87%	3.87%
Volatility	2.69%	2.36%	2.43%
CVaR (99%)	3.41%	2.86%	2.52%
Expected utility	2.62%	1.95%	2.03%
Portfolio iliquidity	14.01%	21.98%	21.55%

Sources: Bloomberg; authors' computations.

The comparative analysis revealed that the optimal portfolio allocation derived under the traditional mean-variance and mean-CVaR frameworks differed from that obtained under the FSO framework. This divergence arises due to the explicit incorporation of an additional dimension – liquidity – into the allocation decision-making process through the utility function.

The ability to directly integrate relevant investor-specific attributes, such as sustainability factors, into the allocation process underscores the comprehensiveness and flexibility of the FSO framework. This capability positions FSO as a superior methodological approach for addressing complex investment preferences.

4. Technical challenges

The implementation of the FSO approach entails significant computational challenges beyond the fundamental difficulty of defining an investor-representative utility function. These challenges stem from the high dimensionality of the decision space, the presence of non-linear constraints, and the inherent non-convexity of the optimisation landscape, which collectively complicate the identification of globally optimal solutions.

In many portfolio selection problems, real-world constraints further complicate the optimisation process, such as cardinality limits (restricting the number of assets), minimum transaction sizes and risk parity conditions. These constraints disrupt the convexity and differentiability of the problem, reducing the effectiveness of traditional gradient-based optimisation methods, which typically assume smooth and convex search spaces.

To address these limitations, a range of stochastic and heuristic algorithms has been developed across disciplines to explore complex, high-dimensional solution spaces more effectively. In finance, simulated annealing (SA), introduced by Kirkpatrick et al (1983), which uses a probabilistic acceptance rule based on changes in the objective function to transition between candidate solutions, has been applied to portfolio optimisation by Crama and Schyns (2003). A related method, threshold accepting, developed by Dueck and Scheuer (1990), which replaces the probabilistic rule with a deterministic one, accepting deteriorated solutions within a fixed threshold, has been used in econometric optimisation (eg Winker (2000)). Widely used heuristic approaches such as evolutionary and genetic algorithms have also been adopted in portfolio optimisation with complex constraints, offering effective strategies for navigating non-convex and discontinuous search spaces (eg Maringer (2005)).

Among these methods, DE is widely recognised for its robustness in handling non-convex and constrained optimisation problems. DE employs a mutation-based exploration mechanism and an adaptive selection strategy, enabling it to balance global exploration with local exploitation. Its efficacy in financial optimisation has been well documented, particularly in cases where objective functions are non-differentiable or highly irregular. For a comprehensive discussion of this methodology, refer to Maringer (2008), who provides an extensive review of heuristic approaches in financial optimisation.

The performance of the DE algorithm in portfolio optimisation is significantly influenced by the initial population definition and other key input parameters. As observed in previous implementations of evolutionary algorithms, the choice of initialisation strategy plays a crucial role in determining both convergence time and solution quality. Importantly, this impact is highly dependent on the utility function governing the optimisation process, as different utility formulations shape the search landscape in distinct ways.

To systematically evaluate these effects, we investigate multiple initialisation strategies within an evolutionary optimisation framework that employs a stopping criterion based on changes in expected utility per iteration. The empirical results for an investor with risk preferences modelled by a kinked log utility function reveal the following key insights.

1. **Effect of Dirichlet-distributed initial populations**

   Given that numerous studies have indicated that the uniform distribution is not always the most suitable choice for initialisation across different applications (Li et al (2020)), a Dirichlet distribution (a multivariate beta distribution) was employed to generate the initial population. This approach resulted in higher expected utility outcomes compared with uniform or random initialisation. The Dirichlet distribution is particularly well suited for portfolio optimisation because it ensures that asset weights sum to one, thereby satisfying the budget constraint without requiring post-processing (eg Le Courtois and Xu (2019)). Additionally, it introduces correlated proportions among asset allocations, reflecting empirical asset weight distributions observed in real-world portfolios.

   However, while Dirichlet initialisation promotes greater exploration of the solution space, this comes at the cost of a higher number of iterations before convergence. The increased computational effort arises due to the algorithm's extended search process, which systematically refines candidate solutions towards the global optimum. This highlights a trade-off between solution quality and computational efficiency: while Dirichlet-based initialisation leads to better final portfolio allocations, it requires more iterations to reach the stopping criterion.

2. **Projection of transformed population elements onto the feasible region**

   A key modification in the evolutionary optimisation process involves projecting transformed population elements onto the feasible solution space after the mutation and crossover operations of the DE algorithm. Empirical findings indicate that this projection mechanism significantly reduces the number of iterations required for convergence by ensuring that all candidate solutions remain within permissible constraints throughout the search process (eg Alcantara and Lee (2025), Gambella et al (2019)).

   Mathematically, constraint violations introduce computational inefficiencies, as infeasible solutions must be either discarded or repaired, delaying progress towards an optimal portfolio allocation. By enforcing feasibility immediately after genetic operations, the search space is effectively pruned, leading to a more efficient local refinement process.

   However, while feasibility projection accelerates convergence, it also introduces a risk of search bias. By systematically adjusting candidate solutions to fit within feasible regions, the algorithm may overly constrain exploration, leading to premature convergence to locally optimal solutions. This risk is particularly relevant in high-dimensional, non-convex search spaces, where projection-based modifications may force the search into sub-optimal regions.

3. **Incorporating efficient frontier portfolios in the initial population**

   Another proposed enhancement involves incorporating portfolios from an efficient frontier constructed using traditional portfolio optimisation methodologies (eg mean-variance optimisation) into the initial population. This technique significantly reduces the number of iterations required for convergence, as the algorithm can leverage precomputed efficient portfolio sets to either initialise or guide the search process (eg Toğan and Eirgash (2019)).

   The effectiveness of this approach is particularly pronounced when the portfolio selection framework underlying the efficient frontier aligns with the investor's utility function. In such cases, the efficient frontier serves as an informative prior, directing the optimisation process towards high-potential solution regions, thereby enhancing computational efficiency.

   However, this method also introduces an inherent trade-off between search efficiency and global optimality. If the efficient frontier is constructed using a risk-return framework that does not fully align with the investor's specific utility function, the search may become artificially constrained, limiting the diversity of candidate solutions and increasing the likelihood of convergence to a local optimum. To counteract this risk, a hybrid initialisation strategy can be employed, wherein efficient frontier portfolios are supplemented with randomly generated solutions to preserve search diversity.

Concluding remarks

The most important considerations when defining a utility function in the reserve management context include safety, return generation and liquidity. However, given the potentially increased relevance of additional aspects such as sustainability factors, the public investor community may find the FSO approach to be a valuable tool to incorporate such factors directly into their investment decision process.

When investors' preferences reflect loss aversion or asymmetry towards risk, FSO may yield to allocations that are, under certain distributional assumptions, similar to those derived from traditional methodologies such as the mean-variance and mean-CVaR optimisation techniques. In those cases, expected utility might be used as a metric for selecting the optimal portfolio from the efficient frontier.

The FSO approach provides great flexibility in incorporating complex investment preferences, allowing for more sophisticated modelling of investor behaviour. Unlike traditional methods, which often rely on simplified assumptions, FSO enables the use of advanced utility functions, which can lead to significant differences in optimal allocations by more accurately reflecting real-world risk-return trade-offs.

This flexibility extends beyond asset allocation, encompassing broader applications in portfolio optimisation, including the design of optimal currency hedging strategies. By aligning with investor preferences, such as the tendency to prioritise downside protection against substantial currency depreciations over the pursuit of incremental gains, FSO provides a behaviourally consistent approach to risk management. Furthermore, its ability to effectively accommodate non-linear hedging strategies, which involve asymmetric payoffs, path dependency and non-normal return distributions, underscores its superiority over traditional optimisation techniques that fail to fully account for these complex risk dynamics.

References

Adler, T and M Kritzman (2006): "Mean-variance versus full-scale optimisation: in and out of sample", *Journal of Asset Management*, vol 7, no 5, pp 302–11.

Alcantara, J H and C-p Lee (2025): "Global convergence and acceleration of projection methods for feasibility problems involving union convex sets", *Journal of Optimization Theory and Applications*, vol 204, no 2, article 19.

Arrow, K (1971): *Essays in the theory of risk-bearing*, Markham, 1 June.

Chamberlain, G (1983): "A characterization of the distributions that imply mean-variance utility functions", *Journal of Economic Theory*, vol 29, no 1, pp 185–201.

Crama, Y and M Schyns (2003): "Simulated annealing for complex portfolio selection problems", *European Journal of Operational Research*, vol 150, no 3, pp 546–71.

Cremers, J, M Kritzman and S Page (2004): "Optimal hedge fund allocations: do higher moments matter?", *Revere Street Working Paper*, no 272–13.

Dueck, G and T Scheuer (1990): "Threshold accepting: a general purpose optimization algorithm appearing superior to simulated annealing", *Journal of Computational Physics*, vol 90, no 1, pp 161–75.

Gambella, C, J Marecek and M Mevissen (2019): "Projections onto the set of feasible inputs and the set of feasible solutions", 57th Annual Allerton Conference on Communication, Control and Computing.

Hagströmer, B and J Binner (2009): "Stock portfolio selection with full-scale optimization and differential evolution", *Applied Financial Economics*, vol 19, no 19, pp 1559–71.

Hayes, M, J Primbs and B Chiquoine (2015): "A penalty cost approach to strategic asset allocation with illiquid asset classes", *The Journal of Portfolio Management*, vol 41, no 2, pp 33–41.

Ingersoll, J (1987): *Theory of financial decision making*, Rowman & Littlefield, 28 June.

Kirkpatrick, S, C D Gelatt and M P Vecchi (1983): "Optimization by simulated annealing", *Science*, vol 220, no 4589, pp 671–80.

Le Courtois, O and X Xu (2019): "Portfolio optimization in the presence of extreme risks: a Pareto-Dirichlet approach", *SSRN Electronic Journal*.

Li, Q, S-Y Liu and X-S Yang (2020): "Influence of initialization on the performance of metaheuristic optimizers", *Applied Soft Computing*, vol 91.

Lindskog, F, A McNeil and U Schmock (2003): "Kendall's tau for elliptical distributions", in G Bol, G Nakhaeizadeh, S Rachev, T Ridder, K-H Vollmer (eds), *Credit risk: measurement, evaluation and management*, Physica-Verlag, pp 149–56.

Maringer, D (2005): *Portfolio management with heuristic optimization*, Springer-Verlag.

——— (2008): "Risk preferences and loss aversion in portfolio optimization", *Computational Methods in Financial Engineering*, pp 27–45.

Meucci, A (2005): *Risk and asset allocation*, Springer.

Pratt, J (1964): "Risk aversion in the small and in the large", *Econometrica*, vol 32, no 1/2, pp 122–36.

Seymour, A, F Chikurunhe and E Flint (2014): "Full-scale optimization: a flexible framework for hedge design", Peregrine Securities.

Toğan, V and M Eirgash (2019): "Time-cost trade-off optimization with a new initial population approach", *Teknik Dergi*, vol 30, no 6.

Winker, P (2000): *Optimization heuristics in econometrics: applications of threshold accepting*, John Wiley & Sons.