Integrating multiple sources of ordinal information in portfolio optimization

Eranda Çela[a], Stephan Hafner[b], Roland Mestel[c], Ulrich Pferschy[b]

[a]*Department of Discrete Mathematics, TU Graz, Steyrergasse 30, 8010 Graz, Austria,* `cela@math.tugraz.at`
[b]*Department of Operations and Information Systems, University of Graz, Universitaetsstrasse 15, 8010 Graz, Austria,* `ulrich.pferschy@uni-graz.at`
[c]*Institute of Banking and Finance, University of Graz, Universitaetsstrasse 15, 8010 Graz, Austria,* `roland.mestel@uni-graz.at`

Abstract

Active portfolio management tries to incorporate any source of meaningful information into the asset selection process. In this contribution we consider qualitative views specified as total orders of the expected asset returns and discuss two different approaches for incorporating this input in a mean-variance portfolio optimization model. In the robust optimization approach we first compute a posterior expectation of asset returns for every given total order by an extension of the Black-Litterman (BL) framework. Then these expected asset returns are considered as possible input scenarios for robust optimization variants of the mean-variance portfolio model (max-min robustness, min regret robustness and soft robustness). In the order aggregation approach rules from social choice theory (Borda, Footrule, Copeland, Best-of-k and MC4) are used to aggregate the total order in a single "consensus total order". Then expected asset returns are computed for this "consensus total order" by the extended BL framework mentioned above. Finally, these expectations are used as an input of the classical mean-variance optimization. Using data from EUROSTOXX 50 and S&P 100 we empirically compare the success of the two approaches in the context of portfolio performance analysis and observe that in general aggregating orders by social choice methods outperforms robust optimization based methods for both data sets.

---

*Corresponding author: Ulrich Pferschy

1. Introduction

In this paper, we focus on the Black-Litterman model (hereafter BL) which derives an optimal portfolio based on two sets of information: backward-looking market equilibrium and forward-looking expert views. We extend the original model in three ways. First, instead of requiring investors to estimate the parameters of the distribution of future asset returns, we allow for views that simply rank these assets according to their expected returns. Second, rather than considering the view of only one expert, we allow for multiple experts to input their separate views on (possibly, but not necessarily) different assets. Third, we propose alternative methods from robust optimization and social choice theory to incorporate these different views into portfolio optimization. To the best of our knowledge this is the first time that aggregation methods from social choice are incorporated into the world of asset allocation. In addition to theoretical explanations, we compare the performance of the portfolios built upon the different approaches using historical stock market data. In this way we can determine the best of the approaches studied.

The relevance of incorporating multiple exogenous insights, such as portfolio managers' judgments, into formal models of portfolio management was already highlighted in Harry Markowitz's seminal work on portfolio selection, commonly referred to as mean-variance (MV) optimization or Modern Portfolio Theory (MPT):

> *"Various types of information concerning securities can be used as the raw material of a portfolio analysis. One source of information is the past performance of individual securities. A second source of information is the beliefs of one or more security analysts concerning future performances."* (Markowitz (1959), p. 3)

MPT (Markowitz (1952)) is undoubtedly the most important quantitative framework for portfolio optimization. While the theory is intuitive and coherent, reputation and implementation have suffered among practitioners due to numerous difficulties. In particular, to construct alternative portfolios with different combinations of risk and return, with one portfolio being the optimal (or *efficient*) one, investors must provide estimates of expected

returns, volatilities and correlations for all securities in the investment universe under consideration. However, early theory remains vague on how to specifically estimate these inputs to the investment process. Traditionally, the use of historical market data has been suggested. Unfortunately, using historical distributions to estimate expected returns and the covariance matrix is prone to measurement errors, which can lead to significant discrepancies between *ex-ante* and *ex-post* optimal portfolio weights for a variety of reasons. Empirical evidence of the poor out-of-sample performance of the classical MV optimization model is provided, among others, by Michaud (1989), Best and Grauer (1991), and Chopra and Ziemba (1993).

To mitigate the adverse effects of neglecting uncertainty in the inputs and solving the portfolio optimization problem in the original MV framework as a deterministic problem, several approaches have been taken in the literature (for a concise overview, we refer to Kolm et al. (2014)). Some of these approaches focus on modifying the constraints and/or the goal objective function, e.g. by imposing restrictions on portfolio weights (e.g., Michaud (1989), Levy and Levy (2014)) and incorporating higher moments and tail-risk measures (e.g., Harvey et al. (2010), Lassance and Vrins (2021)). Other approaches target appropriate modifications to the input of the portfolio selection problem including robust optimization methods (e.g., DeMiguel and Nogales (2009), Huang et al. (2010), Bayesian statistics (e.g., Pastor and Stambaugh (2000), Bodnar et al. (2022)) and the Black-Litterman model (Black and Litterman (1991, 1992). The latter provides the framework for our analysis.

The original BL-approach assumes that the "true" expected asset returns $\boldsymbol{\mu}$ are both, unknown and *random*. Allowing for this uncertainty in estimation, the model starts with equilibrium risk premia as the neutral reference point (the *prior distribution*) for expected returns, derived from an asset pricing model such as the CAPM using reverse optimization. The prior can be interpreted as the *market view* on the portfolio optimization process, since all investors who do not have specific views on expected asset returns will hold the same optimal portfolio, the market portfolio. In a second step, Black and Litterman (1991) assume that investors formulate their own individual views on (at least some of) the securities in the investment universe.[1] These can be views on absolute returns of individual assets (e.g. "the expected one-year return on security X is 3.5%") or relative predictions about the difference in returns between two or more assets (e.g. "security X will

---

[1]Note that a view is generally defined as an uncertain statement about the expected return of one or more securities

3

outperform security Y by 1% over the next 6 months"). As with the prior distribution for $\boldsymbol{\mu}$, investor individual views are assumed to be stochastic, with uncertainty inversely proportional to an investor's confidence in his/her views.

After separately specifying the prior distribution of expected returns implied by the market equilibrium and an investor's views on expected asset returns, the BL-model applies Bayes' theorem to combine these two sources of information to derive the posterior distribution of $\boldsymbol{\mu}$. The lower an investor's confidence in his/her views, the closer his/her posterior distribution of expected returns will be to that implied by the market equilibrium. On the other hand, the more (s)he trusts her own views the more expected returns will tilt away from market equilibrium in the direction of the views. The vector of expected returns derived from the BL-model ($\boldsymbol{\mu}_{\mathrm{BL}}$) is obtained as the conditional expectation of the prior distribution given the investor's views. These estimates finally enter the MV portfolio optimization process.

With respect to investor's views, the traditional BL-model makes two assumptions. First, views, whether formulated in absolute or relative terms (see above), are assumed to be metrically scaled variables (we will refer to such views as *quantitative* views hereafter). Although this assumption simplifies the calculation of $\boldsymbol{\mu}_{\mathrm{BL}}$, it is unlikely that this kind of views will prevail in practice. Portfolio managers and financial analysts are likely to feel more comfortable expressing *qualitative* views, e.g., by simply ranking their expectations concerning the future performance of individual stocks, different asset classes, several sectors of the economy, or individual or multiple stock markets in different geographic regions. In this sense Fabozzi et al. (2007), p. 233, state that "...trading strategies ... rather just provide *relative* rankings of securities that are predicted to outperform/underperform other securities." Further (p. 234):"Clearly, it is not an easy task to translate any of these relative views into the inputs required for the modern portfolio theoretical framework."

Second, the BL-model and most of its extensions assume that it is individual views being processed as input. This is in contrast to team-based portfolio management, which has become very popular in recent years, especially among institutional investors such as mutual funds (e.g., Patel and Sarkissian (2017)). In a group such as an investment committee, group members usually have different opinions and perspectives on the possible future performance of individual assets, sectors, or even entire markets. In such a situation, the team's decisions inevitably will be compromises between these different opinions (e.g., Sah and Stiglitz (1986)). Multiple views may also matter to individual investors. Consider an investor who incorporates the

views of several financial analysts into his/her portfolio selection process in addition to his/her own views. In both cases, the assumption that only individual views are processed to derive the posterior distribution of expected asset returns is not appropriate. However, considering multiple views immediately raises the question of how to combine these views to derive an optimal portfolio.

We contribute to the literature dealing with the integration of qualitative views into portfolio optimization using the parametric BL-framework. Following the argument of Fabozzi et al. (2007) above, we apply an approach recently proposed by Çela et al. (2021) that translates qualitative views, expressed as ordering relations between asset returns, into quantitative estimates of expected returns. Moreover, we extend Çela et al. (2021) by allowing multiple qualitative views and applying two different approaches to incorporate these views into portfolio optimization: the first one (I) is based on robust optimization and the second one (II) on social choice theory. We are not aware of any work that has applied methods from social choice to portfolio optimization before.[2]

In (I), we consider K different ranking orders (e.g., views of K financial analysts) and apply the approach discussed in Çela et al. (2021) to compute a vector of estimated returns for each of them. These K (different) vectors of expected returns are considered as the possible scenarios of expected returns in the MV portfolio optimization model (MVO). We solve three robust variants of the MVO which are obtained by applying the following concepts of discrete robustness: *max-min robustness*, *min-max regret robustness*, and *soft robustness*. Depending on the concept of discrete robustness involved, we obtain different robust optimization-based methods as representatives of this "first estimate, then aggregate" approach.

In (II) we follow a completely different approach. Given K views, we first aggregate them into a single "consensus" ordering by applying some method from social choice theory (varying among the *Borda rule*, the *Footrule aggregation*, the *Best-of-k-Algorithm*, the *Copeland method* and the *MC4-Algorithm*). We then apply the approach discussed in Çela et al. (2021)

---

[2]It should be noted that Simonian (2014) contains a first pointer into the direction of social choice theory. The paper addresses a specific situation in which each view includes a set of stochastically related components. The goal is to aggregate the views such that the components of the aggregated view fulfill the stochastic relation to the greatest extent possible. The components of the aggregated view are seen in analogy to an allocation vector in cooperative game theory and a Shapley-value-based heuristic is proposed to compute them.

to compute a vector of estimated returns from this single total order. Finally, we use this vector as input to the classical MVO and compute the respective optimal portfolio. In summary, in (II), we first aggregate the orderings and then generate the expected return estimator. Depending on the aggregation rule from social choice theory we obtain different social choice-based methods as representatives of this "first aggregate, then estimate" approach.

The main advantage of our proposed aggregation methods from (I) and (II) is that they are purely deterministic, in the sense that they do not use probabilistic assumptions in the aggregation process itself. This is in contrast to well known approaches used to aggregate multiple views in portfolio optimization, such as the classical Black-Litterman approach (BL) in Black and Litterman (1992), the copula-opinion pooling technique in Meucci (2006), or the generalized Black-Litterman model in Chen and Lim (2020)[3], all of which consider quantitative views, and the entropy pooling approach in Meucci (2008), which also treats qualitative views. All these approaches (implicitly) assume a stochastic dependence structure among the different views. Essentially, this is required because in these approaches aggregation is an integral part of the estimation process and the latter is based on distributional assumptions. The proposed approaches (I) and (II) decouple the aggregation from the estimation process, so that no assumptions about the stochastic dependencies between the different views need to be made. This is an obvious advantage since it is not necessary to figure out suitable distributional assumptions and then calibrate the distribution's parameters.

Another approach to generating the inputs of the MVO (or a more general portfolio optimization model) by "blending" equilibrium arguments with (more general) investor's views, without resorting to specific distributional assumptions, is discussed in Bertsimas et al. (2012). This is a fairly general inverse optimization approach that can incorporate any view that can be expressed as a linear matrix inequality. Thus, it can incorporate views addressing risk measures or risk factors in addition to views addressing expected returns. This approach is conceptually and methodologically quite different from the other approaches mentioned in this paper and we refer to Bertsimas et al. (2012) for further details.

Finally, we analyze the performance of our different approaches from ro-

---

[3]Chen and Lim (2020) consider multiple quantitative views and generalize the classical BL approach to account for biases in both the prior distribution and the expert's views. The biased views and the biased prior are aggregated into a posterior distribution by using historical forecast data in a Bayesian hierarchical model.

bust optimization and social choice theory, respectively, in a MV framework using historical financial market data. To this end, we compare two performance measures of the in-sample optimal portfolios, the Sharpe Ratio and the Certainty-Equivalent return. The analysis uses synthetic views which are orderings of the expected asset returns generated uniformly at random at some distance d from the realized ordering (which is obviously known in an in-sample setting). This distance value can be seen as a control parameter for the quality of information represented in the views. The second parameter in our computational experiments is the confidence level in the views, denoted by c. We evaluate the performance of approaches (I) and (II) for different combinations of values of parameters c and d. In a first step we compare the methods within groups (I) and (II) separately, and in a second step we compare the best performing methods from each group with each other.

Our findings can be summarized as follows. In general, the best methods of group (I) are the min-max regret robustness and the soft robustness. In group (II), the methods based on the Borda rule and on the Footrule aggregation clearly perform better than the others. The comparison across the two groups of methods (I) and (II) reveals that the best social choice-based methods outperform the best robust optimization-based methods for almost all parameter settings, albeit the absolute difference between the performance measures is not very accentuated. Thus, our results show that an investor considering incorporating multiple qualitative views into the portfolio allocation decision should use a method from social choice and aggregate first, and then estimate, rather than the other way around.

The remainder of the paper is organized as follows. Section 2 prepares the formal framework for our analysis. Section 3 presents different methods from robust optimization for deriving an optimal portfolio given multiple expert views. Section 4 relates to selected methods from social choice theory that can be applied in a first step to compute a single consensus total order of expected asset returns from different orders. In a second step the aggregated order feeds into the portfolio selection process. Section 5 presents the results of our computational experiments based on historical stock market data. We conclude and discuss future work in section 6.

2. Formal framework for ordinal information

In this section we introduce the formal framework and briefly review the extension of the BL model derived in Çela et al. (2021). We consider a

7

universe of assets $1, \ldots, n$ from which an investor has to choose a portfolio represented by a weight vector $\mathbf{w} = (w_1, \ldots, w_n)^\mathsf{T}$ with $\sum_{i=1}^{n} w_i = 1$, where w_i represents the fraction of capital invested in asset i. We assume that $w_i \geq 0$ $(i = 1, \ldots, n)$, i.e. no short selling is allowed. Let $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)^\mathsf{T} = \mathrm{E}(\mathbf{R})$ be the expected value of the random vector of asset returns $\mathbf{R} = (R_1, \ldots, R_n)^\mathsf{T}$ with an $n \times n$ covariance matrix $\boldsymbol{\Sigma} = \mathrm{Cov}(\mathbf{R})$. We assume that $\boldsymbol{\mu}$ is a normally distributed random vector

$$\boldsymbol{\mu} \sim \mathrm{N}_n(\boldsymbol{\pi}, \tau\boldsymbol{\Sigma}) \tag{1}$$

with mean vector $\boldsymbol{\pi}$ and covariance matrix $\tau\boldsymbol{\Sigma}$. $\boldsymbol{\Sigma}$ is estimated from historical data and $\boldsymbol{\pi}$ is determined by invoking CAPM equilibrium arguments as described for example in Çela et al. (2021). The scale parameter τ indicates the uncertainty of the prior, the smaller the value of τ, the higher the investor's confidence in the estimation of the implied equilibrium return $\boldsymbol{\pi}$.

The classical BL model considers quantitative views which express statements about the expected return of linear combinations of the considered assets.[4] These statements are represented as equalities of the form $\mathbf{P}\boldsymbol{\mu} = \mathbf{v}$, where $\mathbf{P}$ is a so-called pick matrix and $\mathbf{v}$ is a declared outcome. For example, the view that asset 1 will outperform asset 3 by 2% and asset 2 will outperform asset 4 by 1% can be modeled by setting $\mathbf{v} = (0.02, 0, 01)^\mathsf{T}$ and the pick matrix

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & -1 & 0 & \ldots & 0 \\ 0 & 1 & 0 & -1 & \ldots & 0 \end{pmatrix}.$$

Uncertainty is captured by assuming that $\mathbf{v}$ is subject to normally distributed perturbations $\boldsymbol{\varepsilon}$. Thus $\mathbf{v}$ is a random variable $\mathbf{V}$ such that $\mathbf{V} = \mathbf{P}\boldsymbol{\mu} + \boldsymbol{\varepsilon}$ with $\boldsymbol{\varepsilon} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Omega})$ and $\boldsymbol{\mu}, \boldsymbol{\varepsilon}$ independent.

Çela et al. (2021) consider an extended BL-model denoted by *EBL*. Instead of using quantitative views, the authors process qualitative views which simply express ordering relations between asset returns. These are represented as inequalities $\mathbf{P}\boldsymbol{\mu} \geq \mathbf{0}$, where $\mathbf{P}$ is again a pick matrix. For example, the qualitative view that asset 1 will outperform asset 3 and asset 2 will outperform asset 4 can be modeled by the pick matrix

$$\mathbf{P} = \begin{pmatrix} 1 & 0 & -1 & 0 & \ldots & 0 \\ 0 & 1 & 0 & -1 & \ldots & 0 \end{pmatrix}.$$

---

[4]For a more detailed review of the BL model we refer to Meucci (2010) or Walters (2014).

The uncertainty is captured by the random variable $\mathbf{V}$ fulfilling $\mathbf{V} = \mathbf{P}\boldsymbol{\mu} + \boldsymbol{\varepsilon}$ and $\mathbb{V} \geq 0$ with $\boldsymbol{\varepsilon} \sim \mathrm{N}(\mathbf{0}, \boldsymbol{\Omega})$ and $\boldsymbol{\mu}, \boldsymbol{\varepsilon}$ independent. The conditional expectation of $\boldsymbol{\mu}$ given the event $\mathbf{V} \geq \mathbf{0}$ will be denoted by $\boldsymbol{\mu}_{\mathrm{EBL}}$ and serves as the input to a mean-variance portfolio optimization problem.

A frequent choice for the uncertainty $\boldsymbol{\Omega}$ is a matrix proportional to the covariance matrix of $\mathbf{P}\boldsymbol{\mu}$ (see e.g. Walters (2014)), i.e.

$$\boldsymbol{\Omega} = c\,\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^{\mathsf{T}}, \tag{2}$$

where the parameter c represents an overall confidence level in the views with $c = 0$ corresponding to full confidence in the views. Under this setting Çela et al. (2021) derive the following expression for $\boldsymbol{\mu}_{\mathrm{EBL}}$:

$$\boldsymbol{\mu}_{\mathrm{EBL}} = \boldsymbol{\pi} + \frac{\tau}{\tau + c}\boldsymbol{\Sigma}\mathbf{P}^{\mathsf{T}}\left(\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^{\mathsf{T}}\right)^{-1}\left(\mathrm{E}(\mathbf{V}|\mathbf{V} \geq \mathbf{0}) - \mathbf{P}\boldsymbol{\pi}\right) \tag{3}$$

For the numerical computation of $\boldsymbol{\mu}_{\mathrm{EBL}}$ according to (3) the authors apply an an *Importance Sampling* approach.

Çela et al. (2021) provide us with an algorithmic framework which outputs a vector $\boldsymbol{\mu}_{\mathrm{EBL}}$ of expected return estimators combining information from historical data with a view on the future development represented by order relations on the expected future asset returns. The outcome $\boldsymbol{\mu}_{\mathrm{EBL}}$ can finally be used as an input for the following classical mean-variance portfolio optimization model (Markowitz (1952)) with a risk-aversion parameter $\delta > 0$:

$$\max_{\mathbf{w}} \quad \mathbf{w}^{\mathsf{T}}\boldsymbol{\mu} - \frac{\delta}{2}\mathbf{w}^{\mathsf{T}}\boldsymbol{\Sigma}\mathbf{w} \tag{4}$$

$$\text{s.t.} \quad \sum_{i=1}^{n} w_i = 1 \text{ and } w_i \geq 0 \quad (i = 1, \ldots, n)$$

While the approach of Çela et al. (2021) allows to account for qualitative information in portfolio optimization, the restriction of views to one single total order is a major loss of potentially relevant information. Rather, it is reasonable to assume that investors consider multiple views (e.g., from different analysts) when constructing their optimal portfolios. In the following, we consider *multiple* views represented by K different total orders of expected asset returns. Formally we use K permutations σ_k, $k = 1, \ldots, K$, each of them describing a total order of expected asset returns.

We compare two different approaches to incorporate these views into portfolio optimization. In Section 3 we first apply the extended BL approach

described above to obtain a return estimator $\boldsymbol{\mu}_k$ for each permutation σ_k. Thus, $\boldsymbol{\mu}_k$ is obtained by applying (3), where $\mathbf{P}$ is replaced by the pick matrix $\mathbf{P}_k$ corresponding to permutation σ_k. Then, we take the vectors $\boldsymbol{\mu}_k$, $1 \leq k \leq K$ as possible scenarios for $\boldsymbol{\mu}$ in a robust version of (4) with the discrete uncertainty set $\{\boldsymbol{\mu}_k \colon 1 \leq k \leq K\}$. A completely different approach is pursued in Section 4, where we apply aggregation rules from social choice theory to reach a single total order representing a consensus over the given permutations σ_k, $1 \leq k \leq K$. Then, according to the extended BL approach, we use (3) to compute a return estimator $\bar{\boldsymbol{\mu}}$ for the consensus total order. Finally, the MVO (4) is solved with $\bar{\boldsymbol{\mu}}$ as an input.

3. Robust optimization with discrete scenarios

Robust optimization (RO), introduced by Ben-Tal and Nemirovski in the late 1990s in the context of linear and convex optimization problems, is a methodology for dealing with optimization problems whose input has some uncertainty in the form of deterministic variability on the data (for a comprehensive survey, we refer to Ben-Tal and Nemirovski (2002)). The possible values of a parameter or a set of parameters are called scenarios. The scenarios represent possible realizations of data and take their values from a discrete uncertainty set or an uncertainty continuum. The goal is to find solutions that are of high quality for each scenario or realization of the data. There are different ways to define the quality of a solution with respect to all scenarios, and different definitions of solution quality lead to different robust optimization models.

In this paper we consider the classical mean-variance portfolio optimization model (4), where the input parameter $\boldsymbol{\mu}$ is uncertain and its realizations build a discrete uncertainty set $\mathcal{U} := \{\boldsymbol{\mu}_k \colon k \in \{1, 2, \ldots, K\}\}$. We assume that there are K qualitative views, each of which is specified as a total order σ_k of the assets in the portfolio according to the magnitude of their expected returns, $k \in \{1, 2, \ldots, K\}$. For any k we apply the algorithmic framework of Çela et al. (2021) to incorporate the k-th total order σ_k in the vector $\boldsymbol{\mu}_k$ of expected asset returns. In particular $\boldsymbol{\mu}_k$ is obtained by applying (3) with a pick matrix $\mathbf{P}_k$ which represents the total order of expected asset returns specified by the permutation σ_k, for all $k \in \{1, 2, \ldots, K\}$:

$$\boldsymbol{\mu}_k := \boldsymbol{\pi} + \frac{\tau}{\tau + c} \boldsymbol{\Sigma} \mathbf{P}_k^{\mathsf{T}} \left(\mathbf{P}_k \boldsymbol{\Sigma} \mathbf{P}_k^{\mathsf{T}} \right)^{-1} \left(\mathrm{E}(\mathbf{V} | \mathbf{V} \geq \mathbf{0}) - \mathbf{P}_k \boldsymbol{\pi} \right). \qquad (5)$$

In the following we briefly discuss the three robust optimization variants of (4) addressed in this paper: max-min robustness, min-max regret robustness

10

and soft robustness.

3.1. Max-min robustness

In max-min robustness, also called the strictly robust optimization approach, see Ben-Tal and Nemirovski (2002), the goal is to find a solution which maximizes the worst value of the objective function over all scenarios in $\mathcal{U}$.

$$\max_{\mathbf{w}} \min_{\boldsymbol{\mu}} \mathbf{w}^\mathsf{T} \boldsymbol{\mu} - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w} \tag{6}$$
$$\text{s.t. } \boldsymbol{\mu} \in \{\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_K\}$$
$$\sum_{i=1}^n w_i = 1 \text{ and } w_i \geq 0 \;\; (i = 1, \ldots, n).$$

It is obvious that this robust version of the mean-variance formulation is very conservative; a candidate portfolio $\mathbf{w}$ is considered to be optimal if it maximizes the worst case, i.e. the smallest value of the risk adjusted expected return that this portfolio achieves over all possible realizations of expected asset returns in $\mathcal{U}$. The above problem can be reformulated equivalently as a concave quadratic maximization problem.

$$\max_{\mathbf{w}, y} \; y - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w} \tag{7}$$
$$\text{s.t. } y \leq \boldsymbol{\mu}_k^\mathsf{T} \mathbf{w} \;\; (k = 1, \ldots, K)$$
$$\sum_{i=1}^n w_i = 1 \text{ and } w_i \geq 0 \;\; (i = 1, \ldots, n).$$

It is well known that concave quadratic programming, i.e. maximizing a concave quadratic objective function subject to linear and affine constraints is (weakly) polynomially solvable (see e.g. Vavasis (2001)) and efficient optimization algorithms for solving this type of problems are available in every standard optimization software. Thus, in this case the max-min robust problem (7) (or equivalently (6)) remains polynomially solvable and is computationally not harder than the non-robust counterpart (4). Note, however, that the time complexity of the algorithms for solving (7) depends on the number K of scenarios (or views).

3.2. Min-max regret robustness

The min-max regret robustness model is less conservative than the max-min robustness model. In this model the quality of a solution is measured

11

in terms of the *maximum regret* defined as follows.

Definition 1. *The maximum regret $MaxReg(\boldsymbol{w})$ of a portfolio $\boldsymbol{w}$ with respect to the mean-variance portfolio optimization model (4) and the discrete uncertainty set $\mathcal{U} := \{\boldsymbol{\mu}_k \colon k \in \{1, 2, \ldots, K\}\}$ is defined as*

$$MaxReg(\boldsymbol{w}) := \max_{\boldsymbol{\mu} \in \{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots, \boldsymbol{\mu}_K\}} [f(\boldsymbol{w}_{\boldsymbol{\mu}}, \boldsymbol{\mu}) - f(\boldsymbol{w}, \boldsymbol{\mu})], \tag{8}$$

where $f(\boldsymbol{w}, \boldsymbol{\mu})$ is the objective function of (4), i.e. $f(\boldsymbol{w}, \boldsymbol{\mu}) := \boldsymbol{w}^\mathsf{T} \boldsymbol{\mu} - \frac{\delta}{2} \boldsymbol{w}^\mathsf{T} \boldsymbol{\Sigma} \boldsymbol{w}$, and $\boldsymbol{w}_{\boldsymbol{\mu}}$ maximizes $f(\boldsymbol{w}, \boldsymbol{\mu})$ for fixed $\boldsymbol{\mu}$, i.e.

$$f(\boldsymbol{w}_{\boldsymbol{\mu}}, \boldsymbol{\mu}) = \max \left\{ f(\boldsymbol{w}, \boldsymbol{\mu}) \colon \boldsymbol{w} \in \mathbb{R}^n_+, \sum_{i=1}^{n} w_i = 1 \right\}.$$

The expression in the square parentheses in (8) is called the *regret* of the portfolio $\mathbf{w}$ for a given $\boldsymbol{\mu}$ and is equal to the portion of the objective function value lost when the optimal portfolio $\mathbf{w}_{\boldsymbol{\mu}}$ with respect to $\boldsymbol{\mu}$ is replaced by the portfolio $\mathbf{w}$.

The min-max regret robustness problem is then given as

$$\min\{MaxReg(\mathbf{w}) \colon \mathbf{w} \in \mathbb{R}^n_+, \sum_{i=1}^{n} w_i = 1\}. \tag{9}$$

For $k \in \{1, 2, \ldots, K\}$ the portfolios $\mathbf{w}_{\boldsymbol{\mu}_k}$ and the values $f_k := f(\mathbf{w}_{\boldsymbol{\mu}_k}, \boldsymbol{\mu}_k)$ of the risk-adjusted return can be computed by setting $\boldsymbol{\mu} := \boldsymbol{\mu}_k$ for $k \in \{1, 2, \ldots, K\}$ and solving K mean-variance portfolio optimization problems as in (4). Observe that

$$MaxReg(\mathbf{w}) := \max_{k \in \{1,2,\ldots,K\}} \left[f_k - \mathbf{w}^\mathsf{T} \boldsymbol{\mu}_k + \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w}\right] = \max_{k \in \{1,2,\ldots,K\}} \left[f_k - \mathbf{w}^\mathsf{T} \boldsymbol{\mu}_k\right] + \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w}.$$

Thus, (9) can be equivalently written as

$$- \max \left\{ - \max_{k \in \{1,2,\ldots,K\}} [f_k - \mathbf{w}^\mathsf{T} \boldsymbol{\mu}_k] - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w} \colon \mathbf{w} \in \mathbb{R}^n_+, \sum_{i=1}^{n} w_i = 1 \right\}.$$

By setting $-y_{\mathbf{w}} := \max_{k \in \{1,2,\ldots,K\}} \left[f_k - \mathbf{w}^\mathsf{T} \boldsymbol{\mu}_k\right]$ we get the next equivalent

formulation of (9) as

$$- \max \left\{ y_{\mathbf{w}} - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \mathbf{\Sigma} \mathbf{w} : -y_{\mathbf{w}} \geq f_k - \mathbf{w}^\mathsf{T} \boldsymbol{\mu}_k, k \in \{1, 2, \ldots, K\}, \mathbf{w} \in \mathbb{R}^n_+, \sum_{i=1}^{n} w_i = 1 \right\}.$$

Finally, by observing that in the last formulation the variable $y_{\mathbf{w}}$ can be replaced by a single variable y (and by neglecting the minus in front of the problem) we get the following concave quadratic maximization formulation of (9)

$$\max_{\mathbf{w}, y} \quad y - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \mathbf{\Sigma} \mathbf{w} \qquad (10)$$
$$\text{s.t. } y \leq \boldsymbol{\mu}_k^\mathsf{T} \mathbf{w} - f_k \,, \ k = 1, \ldots, K$$
$$\sum_{i=1}^{n} w_i = 1 \text{ and } w_i \geq 0 \,, \ i = 1, \ldots, n.$$

Thus, similarly to the case of the max-min robustness approach, also the min-max regret robustness problem (9) (or equivalently (10)) remains polynomially solvable and is not harder than its non-robust counterpart (4). Note, however, that the time complexity of algorithms for solving (10) heavily depends on the number K of scenarios (or views). Other than in the case of the max-min robust version of the mean-variance formulation, here we also need to solve K concave quadratic minimization problems to determine f_k, $k \in \{1, 2, \ldots, K\}$, prior to solving the concave quadratic maximization problem (10) which analogously to (7) involves $K + 1$ linear constraints besides the nonnegativity constraints.

3.3. Soft robustness

The soft robustness model allows to control the degree of conservativeness of the robust solution in the vein of the so called Γ-robustness (see Bertsimas and Sim (2003)). This model includes a parameter $\Gamma \in (0, 1]$ which allows to adjust the extent at which the possible estimators of the expected asset returns $\boldsymbol{\mu}_k$, $k \in \{1, 2, \ldots, K\}$, are taken into account. More precisely, we want to choose a portfolio $\mathbf{w}^*$ which maximizes a threshold of risk adjusted return which would be exceeded by at least ΓK estimators $\boldsymbol{\mu}_k$, $k \in \{1, 2, \ldots, K\}$. In other words $\mathbf{w}^*$ maximizes the empirical $(1 - \Gamma)$-quantile $y_\Gamma(\mathbf{w})$ of the risk-adjusted expected returns over all feasible portfolios $\mathbf{w}$.

Definition 2. *The empirical $(1 - \Gamma)$-quantile $y_\Gamma(\boldsymbol{w})$ of the set*

13

$\{f(\boldsymbol{w}, \boldsymbol{\mu}_1), f(\boldsymbol{w}, \boldsymbol{\mu}_2), \ldots, f(\boldsymbol{w}, \boldsymbol{\mu}_K)\}$ *is defined as follows*

$$y_\Gamma(\boldsymbol{w}) := \sup \left\{ z \in \mathbb{R} : \left| \{ k : k \in \{1, 2, \ldots, K\}, f(\boldsymbol{w}, \boldsymbol{\mu}_k) \geq z \} \right| \geq \Gamma K \right\}.$$

The corresponding optimization problem is then given as

$$\max \left\{ y_\Gamma(\mathbf{w}) : \mathbf{w} \in \mathbb{R}^n_+, \sum_{i=1}^{n} w_i = 1 \right\}. \tag{11}$$

By introducing a new variable y to represent the maximum of $y_\Gamma(\mathbf{w})$ we obtain the following equivalent formulation of (11).

$$\max_{\mathbf{w}, y} \quad y \tag{12}$$
$$\text{s.t.} \ \left| \{ k : k \in \{1, 2, \ldots, K\}, f(\mathbf{w}, \boldsymbol{\mu}_k) \geq y \} \right| \geq \Gamma K$$
$$\sum_{i=1}^{n} w_i = 1 \text{ and } w_i \geq 0, \ i = 1, \ldots, n.$$

Obviously, the larger the value of Γ, the more conservative is the model. In the limit case $\Gamma = 1$ we get $y_\Gamma(\mathbf{w}) = \min\{f(\mathbf{w}, \boldsymbol{\mu}_1), f(\mathbf{w}, \boldsymbol{\mu}_2), \ldots, f(\mathbf{w}, \boldsymbol{\mu}_K)\}$ and the soft robustness model coincides with max-min robustness model. On the other hand, if Γ becomes very small, e.g. $\Gamma = \frac{1}{K}$, we get $y_\Gamma(\mathbf{w}) = \max\{f(\mathbf{w}, \boldsymbol{\mu}_1), f(\mathbf{w}, \boldsymbol{\mu}_2), \ldots, f(\mathbf{w}, \boldsymbol{\mu}_K)\}$ and the soft robustness model becomes

$$\max_{\mathbf{w}} \max_{\boldsymbol{\mu}} \quad \mathbf{w}^\mathsf{T} \boldsymbol{\mu} - \frac{\delta}{2} \mathbf{w}^\mathsf{T} \boldsymbol{\Sigma} \mathbf{w}$$
$$\text{s.t.} \ \boldsymbol{\mu} \in \{ \boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \ldots, \boldsymbol{\mu}_K \}$$
$$\sum_{i=1}^{n} w_i = 1 \text{ and } w_i \geq 0, \ i = 1, \ldots, n.$$

Observe now that (12) can be formulated as a convex mixed integer nonlinear problem (MINLP) whith a linear objective function involving just continuous variables as well as linear and convex quadratic constraints. Indeed, we introduce the binary variables $v_k \in \{0, 1\}$, $k \in \{1, 2, \ldots, K\}$, such that $v_k = 0$ iff $y \leq f(\mathbf{w}, \boldsymbol{\mu}_k)$ and $v_k = 1$, otherwise. In the latter case $y \leq f(\mathbf{w}, \boldsymbol{\mu}_k) + v_k M$ is fulfilled for a large enough constant M. The constraint $\sum_{k=1}^{K} v_k \leq (1 - \Gamma)K$ guarantees that at least ΓK of the variables

v_k equal zero, i.e. that $f(\mathbf{w}, \boldsymbol{\mu}_k) \geq y$ holds for at least ΓK estimators $\boldsymbol{\mu}_k$. Finally, by introducing an additional auxiliary variable z as an upper bound on $\mathbf{w}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w}$ the inequality $y \leq f(\mathbf{w}, \boldsymbol{\mu}_k) + v_k M$ can be formulated as a linear constraint $y \leq \mathbf{w}^\mathsf{T}\boldsymbol{\mu}_k - \frac{\delta}{2}z + v_k M$. Summarizing we obtain the following equivalent formulation of (12):

$$\max_{\mathbf{w}} \quad y \tag{13}$$

$$\text{s.t. } y \leq \mathbf{w}^\mathsf{T}\boldsymbol{\mu}_k - \frac{\delta}{2}z + v_k M \,, \, k = 1, 2, \ldots, K$$

$$\mathbf{w}^\mathsf{T}\boldsymbol{\Sigma}\mathbf{w} \leq z$$

$$\sum_{k=1}^{n} v_k \leq (1 - \Gamma)K$$

$$\sum_{i=1}^{n} w_i = 1 \text{ and } w_i \geq 0 \,, \, i = 1, \ldots, n$$

$$v_k \in \{0, 1\} \,, \, k = 1, \ldots, K.$$

In general the convex MINLP is a NP-hard problem because it includes the mixed integer linear programming (MILP), see e.g. Belotti et al. (2013). There is a wealth of literature on exact and heuristic solution approaches for convex MINLP, see e.g. the recent review by Kronqvist et al. (2019). Note, however, that (13) has a particular structure: the constraints involving the binary variables are linear and the objective function just involves continuous variables. The outer approximation method of Duran and Grossmann (1986) is an iterative approach which exploits the particular structure mentioned above and consists in solving an alternate finite sequence of convex quadratic subproblems and relaxations of a MILP as a master problem. A rigorous analysis of the convergence properties of the algorithm can also be found in Belotti et al. (2013).

4. Aggregation by social choice methods

In this section we pursue the idea of computing a single consensus total order of expected asset returns which is a sort of "center" of all K given total orders. Formally, we are looking for a total order σ^* which minimizes the sum of distances to the total orders $\sigma_1, \ldots, \sigma_K$ (see Marden (1996) for an overview). The distance between two total orders σ' and σ'' will be measured by the classical Kendall-Tau distance which counts the number of pairwise disagreements between two total orders, i.e. the number of pairs

(i, j) with $\sigma'(i) < \sigma'(j)$ and $\sigma''(i) > \sigma''(j)$. For our purposes we use the normalized version of the Kendall-Tau distance $KT(\sigma', \sigma'')$ given as the number of pairwise disagreements between σ' and σ'' divided by the overall number of pairs (i, j) with $i \neq j$, see e.g. Dwork et al. (2001). Then we are looking for an optimal total order σ^* such that

$$\sigma^* = \arg \min_{\sigma \in S_n} \sum_{k=1}^{K} KT(\sigma, \sigma_k), \qquad (14)$$

where S_n is the set of all permutations of numbers $1, \ldots, n$. This optimization problem, which is called the Kemeny-Young problem, is the basis for the well-known Kemeny-Young method for group decisions. However, the problem is NP-hard for $K \geq 4$ (Dwork et al. (2001)) and becomes practically intractable already for a moderate number of assets. Thus, we will focus on approximate solutions to the Kemeny-Young problem, in particular on methods developed in the area of social choice.

Social choice theory (see e.g. Arrow et al. (2002)) studies the aggregation of individual preferences or opinions into a group decision considering issues such as fairness, social welfare or acceptability. Typical areas of application are voting, selection of committees and allocation of resources. A central concept in this field is a *social welfare function* (cf. the seminal work of Nobel price winner Sen (1970)), which receives as input K total orders of the objects $1, \ldots, n$ representing the preferences of K individuals and determines as output a single total order. This is exactly the task we are facing in this section. The mapping of the total orders σ_k, $1 \leq k \leq K$, to σ^* in (14) can be seen as one particular social welfare function. For a total order σ we call the sum in the righthand side of (14) the Kendall-Tau score of σ (with respect to the input total orders σ_k, $1 \leq k \leq K$).

In social choice theory a large number of social welfare functions were introduced and their theoretical properties investigated. A good overview can be found in Brams and Fishburn (2002). In the following we will describe some of the more prominent examples which will then be applied to compute a total order of expected asset returns. For a broader discussion and computational aspects we refer the reader to Brandt et al. (2016).

4.1. Borda rule

The Borda rule was formulated by Jean-Charles Chevalier de Borda in 1770 and is one of the most prominent and natural aggregation rules. It belongs to the family of scoring rules, where an element at a certain position

of a total order is assigned a certain number of points. In case of the Borda rule, the asset i at position $\sigma(i)$ in a total order gets a Borda score $Bor_\sigma(i) := n - \sigma(i)$. The aggregation works by simply summing up for every asset i the scores over all K total orders, i.e. $\sum_{k=1}^K Bor_{\sigma_k}(i)$, and sorting the assets in decreasing order of these total scores.

4.2. Footrule Aggregation

This aggregation is based on an optimization problem similar to (14), but it applies the *Spearman footrule* $SF(\sigma', \sigma'')$ for measuring the distance between two total orders σ' and σ'' instead of the Kendall-Tau distance. It is based on measuring for every i the absolute difference between the positions $\sigma'(i)$ and $\sigma''(i)$ and summing up these values. Formally, we are looking for a total order σ^* such that

$$\sigma^* = \arg\min_{\sigma \in S_n} \sum_{k=1}^K SF(\sigma, \sigma_k), \quad \text{with } SF(\sigma, \sigma_k) = \sum_{i=1}^n |\sigma(i) - \sigma_k(i)|. \quad (15)$$

It was shown by Diaconis and Graham (1977) that the difference between the Spearman footrule distance and the Kendall-Tau distance can be bounded by

$$KT(\sigma', \sigma'') \leq SF(\sigma', \sigma'') \leq 2\, KT(\sigma', \sigma''). \quad (16)$$

Moreover, $SF(\sigma', \sigma'')$ can be computed efficiently in polynomial time by solving a linear assignment problem as shown e.g. in Dwork et al. (2001).

4.3. Copeland Method

For a given set of K total orders, we introduce a relation $\succsim$ (majority dominance) on the set of assets such that $i \succsim j$ if and only if i is ranked higher than j by a majority of the total orders, i.e. $|\{k \in \{1, \ldots, K\} : \sigma_k(i) < \sigma_k(j)\}| \geq K/2\}|$, see e.g. Fishburn (1977). The Copeland value $Cop(i)$ of an asset i is defined as the number of assets majority dominated by i minus the number of assets majority dominating i, i.e.

$$Cop(i) := |\{j \in \{1, \ldots, n\} : i \succsim j\}| - |\{j \in \{1, \ldots, n\} : j \succsim i\}|$$

For use as an aggregation rule, assets are ordered in decreasing order of their Copeland values.

4.4. Best-of-k-Algorithm

This simple approach solves the Kemeny–Young problem under the strong restriction that σ^* is chosen only among the given total orders, i.e. in (14)

the domain of the minimization $\sigma \in S_n$ is replaced by $\sigma \in \{\sigma_1, \ldots, \sigma_K\}$. Thus, it suffices to calculate the Kendall-Tau distance between all pairs of total orders and then compare the appropriate sums. This straightforward approach yields a 2-approximate solution to (14) with a Kendall-Tau distance at most twice as large as given by σ^* (see e.g. Ailon et al. (2008)).

4.5. MC4-Algorithm

In their widely cited paper Dwork et al. (2001) present four aggregation algorithms based on Markov chains. Among these, we will consider the so-called MC4-Algorithm, which performed best in the empirical evaluation of Dwork et al. (2001) and Schalekamp and van Zuylen (2009). In this approach, the n assets correspond to the n states of a Markov chain. Starting from a randomly chosen state, the Markov process moves in each iteration from the current state to a new state of the system. From a current state i, the new state j is chosen with a certain transition probability p_{ij}. If the current situation of the system is represented by a probability distribution $\boldsymbol{x}$, where x_i gives the probability that the process is currently in state i, then the probability distribution of the next state is given by $\boldsymbol{x}^\mathsf{T}\boldsymbol{P}$, where the probability matrix $\boldsymbol{P} = (p_{ij})$ consists of the transition probabilities. It is known that under certain conditions, this process reaches a stationary distribution $\boldsymbol{y}$ with $\boldsymbol{y} = \boldsymbol{y}^\mathsf{T}\boldsymbol{P}$, and $\boldsymbol{y}$ can be computed e.g. by a simple power-iteration algorithm.

When applying MC4 in the aggregation of total orders, the probabilities p_{ij} are defined as follows: if i is the current state, a potential new state j is chosen from $\{1, \ldots, n\}$ uniformly at random. If a majority of the total orders ranks j before i, i.e. $j \succsim i$ according to the majority dominance of the Copeland Method[5], then the Markov process moves to state j, otherwise it remains in state i. To ensure that the process has a unique stationary distribution Schalekamp and van Zuylen (2009) add a random "jump": in every step the process moves from i to the randomly chosen j with a tiny probability α, independently from the dominance relation. With probability $1 - \alpha$, the move to j is only chosen if $j \succsim i$. Hence, there is a transition probability $p_{ij} = 1/n + \alpha/n$, if $j \succsim i$, and $p_{ij} = \alpha/n$, otherwise. The non-transition probabilities are then set as $p_{ii} = 1 - \sum_{j \neq i} p_{ij}$, for all i. We use $\alpha = 1/100$ to ensure that for up to $n = 100$ assets the sum of the transition probabilities does not exceed 1.

---

[5]Ties are broken lexicographically, i.e. by the index number of the assets.

After reaching the unique stationary distribution $\boldsymbol{y}$, the aggregation of total orders follows by sorting the assets in decreasing order of y_i.

4.6. Local Improvement

Any total order σ produced by one of the above aggregation methods can be further improved by reducing the corresponding value of the objective function in (14), i.e. the Kendall-Tau score of σ. In our approach this is done by an iterative pairwise exchange operation. We consider the assets in increasing order, i.e. we go through all positions $\ell = 2, 3, \ldots, n$. In the iteration for asset i in position ℓ, $\sigma(i) = \ell$, we compare i with the asset j at the preceding position $\sigma(j) = \ell - 1$. If $j \succsim i$ (according to the majority dominance), then and we proceed to the next iteration for the asset in position $\ell + 1$. Otherwise, if $i \succsim j$, we exchange i and j in σ by setting $\sigma(i) = \ell - 1$ and $\sigma(j) = \ell$, thereby improving the Kendall-Tau score of σ. We continue by comparing i with the asset at position $\ell - 2$ and so on, until either a comparison without improvement is found or i reaches the first position. At this point the iteration for asset i terminates and we move to the next iteration for the immediate successor of asset i at position $\ell + 1$. This procedure is called "Local Kemenization" in Dwork et al. (2001) and "InsertionSort" in Schalekamp and van Zuylen (2009).

In our experiments we apply this process of *Local Improvement* to all total orders produced by the above aggregation procedures.

5. Computational experiments

In this section, we evaluate the performance of the different approaches presented in Sections 3 and 4 for aggregating multiple qualitative views, expressed as total orders of expected returns of assets in the portfolio, in the context of classical MVO using real stock market data. We apply the extended BL-model suggested in Çela et al. (2021), more precisely the equality (3), to compute an estimator of expected asset returns for any given total order of the latter (see Section 2).

We use historical data from two stock indices, one in Europe and one in the U.S. (see Section 5.1). For both samples, we compare the performance of monthly rebalanced MVO portfolios based on (I) the several robust optimization approaches and (II) the various ordering aggregation methods based on social choice theory.

In a first step, we compare the performance of the different approaches within each group separately. To mitigate the effect of estimation errors the analysis is based on in-sample tests. The tests are in-sample in that

(a) the complete time series of realized asset returns are used to estimate the corresponding covariance matrix (which then remains constant over the entire time horizon) and (b) the synthetic views (i.e. total orders) on the expected asset returns in the time interval $(t, t+1)$ are randomly generated (see below) utilizing information of the actual total order of asset returns realized in that time interval. In a second step and after identifying the two best performing approaches from each group (I) and (II), we compare the performance of the two winners of each group to identify the overall "champion".

All experiments were implemented and run in Matlab R2021b on a standard PC.[6]

5.1. Data description

We create two samples with single stock data from the EURO STOXX 50 and the S&P 100. The EURO STOXX 50 is a blue-chip stock index that comprises fifty of the largest and most liquid stocks in the Eurozone. The S&P 100, a subset of the S&P 500, includes the largest U.S. companies across a variety of industries. We collect monthly total return index data from Refinitiv Datastream for the period December 1998 to December 2021 for all stocks included in either index at the end of 2021. We eliminate stocks for which we do not have a continuous price series over the entire time period. In this way we identify 38 (76) stocks from EURO STOXX 50 (S&P 100) that are eligible for our investigation. We calculate monthly returns from the total return index series, with January 1999 being the first observation, giving a total of 276 monthly returns for each stock. Table 1 reports descriptive statistics for cross-sectional averages for our sample of constituents of the EURO STOXX 50 (Panel A) and the S&P 100 (Panel B). For the former, the average monthly holding period return (in EUR) is 1.05% with an average standard deviation of 8.33%. The corresponding values for our S&P 100 sample are 1.14% for the average monthly return (in USD) and 8.21% for the standard deviation. Concerning skewness, we find evidence of slightly positive values, while kurtosis shows the typical pattern of fat tails in both samples.

---

[6]Since the computations required negligible running times and the computational performance was not the focus of this study we refrain from describing further technical details.

Panel A: Selected constituents of EURO STOXX 50

monthly return (%)	mean	q1	q2	q3	min	max
mean	1.05	0.76	0.96	1.22	0.42	2.32
stdev	8.33	6.62	8.57	9.52	4.79	11.79
skew	0.15	-0.13	0.13	0.32	-1.02	1.22
excess kurtosis	3.74	1.10	1.90	4.74	0.34	20.30

Panel B: Selected constituents of S&P 100

monthly return (%)	mean	q1	q2	q3	min	max
mean	1.14	0.81	1.02	1.37	0.22	3.03
stdev	8.21	6.51	7.46	9.12	4.73	20.81
skew	0.31	-0.23	0.08	0.38	-1.38	6.51
excess kurtosis	5.82	1.53	2.64	4.84	0.21	76.35

Table 1: Descriptive statistics for monthly returns (in %) of 38 stocks in the EURO STOXX 50 (Panel A) and 76 stocks in the S&P 100 (Panel B). All stocks are included a) that were constituents of one of the corresponding stock indices on December 31, 2021, and b) for which a complete return time series is available for the period from January 1999 to December 2021. The values for mean and stdev are given in %. q1 (q2) [q3] represents the first (second) [third] quartile. All measures are expressed as equally weighted cross sectional averages. Monthly returns are calculated as holding period returns from the last day of one month to the last day of the following month. Sample period: January 1999 to December 2021. Source: Refinitiv Datastream.

5.2. Setup of the experiments

We choose a large number of different parameter settings for our experiments. For each of the two samples we generate K total orders of all assets for $K \in \{5, 10, 20\}$. The total orders represent different synthetic qualitative views, as they might be expressed by different experts (e.g. financial analysts) under different scenarios. We consider our choice of values for K to be reasonable, but emphasize that our methods can in principle be applied to any number of views. We further differentiate the views according to the validity of the information they contain. We quantify the degree of validity of information in a total order $\mathcal{O}$ in terms of the Kendall-Tau distance d of $\mathcal{O}$ from the total order of the assets returns realized in the time interval $(t, t+1)$ which represents the correct ordering. Note that if the Kendall-Tau distance between $\mathcal{O}$ and the correct ordering is about 0.5, then $\mathcal{O}$ coincides with the correct ordering for about half of the pairs of assets. Such a total order $\mathcal{O}$ would therefore represent a complete lack of knowledge about the correct ordering. On the other hand, a value of d around 0 (and symmetrically a value of d around 1) corresponds to a high degree of useful information, because in this case $\mathcal{O}$ would coincide with the correct ordering for almost all pairs (no pairs) of assets. In order to cover the whole range of the possible values of d we consider $d \in \{0.2, 0.3, 0.4, 0.47\}$. For each value

of d we generate K random total orders of assets, each of them having the same Kendall-Tau distance d from the correct ordering.[7]

Finally, we consider four different levels of overall confidence in the validity of the generated total orders, represented by the parameter c (see equation (2)). Note that the same confidence level applies to each of the K total orders. This should reflect the fact that it is indeed difficult to a priori assess the information content of a (qualitative) view.[8] We span a wide range of relevant values with $c \in \{0.25, 0.5, 0.75, 0.95\}$, where $c = 0.25$ ($c = 0.95$) indicates a very high (low) degree of confidence. As in Çela et al. (2021), we set $\tau = 1 - c$, where the parameter τ represents the confidence in the prior information.

For each value $K \in \{5, 10, 20\}$, we perform our experiments for all 16 value combinations of the parameters d and c, giving a total number of 48 parameter settings. One might argue that some parameter combinations, such as large distances and high levels of confidence, are less meaningful, but we chose to retain all combinations because it is in general unpredictable where the ordering information will come from and in what context our method will be applied.

5.3. Measuring portfolio performance

Our goal is to find the best out of the aforementioned approaches that allow for integrating multiple input orderings into MVO (cf. (I) and (II) at the beginning of this section). To make a decision in this "horse race" we compare the performance of the optimal portfolios derived from each approach by means of two metrics:[9]

- *Sharpe Ratio (SR)*: The Sharpe Ratio is defined as the excess return of a portfolio divided by its total risk, where the latter is measured by the standard deviation of portfolio returns. The excess return is calculated with respect to the average risk-free rate. Thus, if the return and the

---

[7]Note that we could also handle individual values of d for each total order to account for differences in the quality of the views. This can be useful, for example, when several financial analysts have different track records in terms of their forecasts and stock recommendations in the past. In this case the view of an analyst with a high track record could be assigned a lower value of d, while the view of an analyst with a poor record could be given a higher d.

[8]Similarly to parameter d our model is fully flexible to account for different confidence levels c for each total order.

[9]For a detailed description of the two metrics we refer to Bodie et al. (2021). They are used, among others, in studies by DeMiguel et al. (2009) and Allaj (2020).

volatility of a porfolio P are given by μ_P and σ_P, respectively, then the Sharpe Ratio SR_P of the portfolio P is given as[10]

$$SR_P = \frac{\mu_P - r_f}{\sigma_P},$$

where r_f is the risk-free rate. We use the 3m Euribor for the EURO STOXX 50 sample and the 3m USD Libor for the S&P 100 sample. The higher SR the higher the risk-adjusted performance of a portfolio.

- *Certainty-Equivalent return (CEQ)*: The Certainty-Equivalent return is the zero-risk return that an investor with risk aversion δ is willing to accept rather than pursuing a particular risky investment strategy with a higher but riskier return. Formally, the Certainty-Equivalent return CEQ_P of a portfolio P with return μ_P and volatility σ_P is given as

$$CEQ_P = \mu_P - \frac{\delta}{2}\sigma_P^2.$$

We set δ equal to three which is a commonly used value in the literature (e.g., Çela et al. (2021)). As with SR investors prefer portfolios with higher CEQ values to portfolios with lower values.

5.4. Computational comparison

As mentioned at the beginning of this section, we first evaluate the performance of the robust optimization-based approaches (I) and the performance of the social choice-based approaches (II) separately and try to identify the best performing methods in each of the two groups.

In (I) we consider six robust optimization variants: Max-min robustness (Max-min), Min-max regret robustness (Min regret) and soft robustness (Soft Γ) with $\Gamma \in \{0.25, 0.50, 0.75, 1\}$ as described in Section 3. In (II), we consider five ordering aggregation approaches: Borda rule (Borda), Footrule aggregation (Footrule), Best-of-k algorithm (Best-of-k), Copeland method (Copeland) and MC4-algorithm (Markov). As described in Section 5.2, we conduct experiments for 48 parameter settings.

For each group of approaches (I) or (II) and for each parameter setting, we determine the winners in terms of SR and CEQ. We define as winners all methods with a SR, or CEQ value within 1% of the respective maximum

---

[10]Note that since we are using in-sample mean and variance estimates, we calculate the in-sample Sharpe Ratio with portfolio weights obtained with these estimates.

value. For each method we count the number of wins out of the 48 "races" and report results in Table 2.

In general, we find that there is no particular method in either (I) or (II) that clearly dominates all other methods across all parameter settings. We therefore select the two most successful methods from each of the groups (I) and (II). To this end, we consider the total number of wins over all parameter settings with respect to SR and CEQ, respectively. To illuminate the best methods within each group (I) and (II) we indicate the method with the highest (second highest) number of total wins with a solid (dashed) box in Table 2.

For the methods based on robust optimization, the choice of Soft 0.5 is unambiguous: this method has the largest number of wins in 3 out of 4 subtables. This method is also mostly top-ranked in this group regardless of the number of views K. Other Soft methods also perform well in some races, but Soft 0.5 is clearly the best performing variant from this family. As the second best method within group (I), we identify Min regret with one win and one second place out of 4 subtables. Min regret is closely followed, but not dominated, by Max-min.

Among social choice-based methods, Borda can be seen as the overall winner with 4 wins and Copeland is the runner-up in all 4 subtables (even with some first places in individual settings), although the competitors are quite close. Therefore, and to identify the best method that allows for integrating multiple input orderings into MVO, we focus our further investigations on the four methods Min regret, Soft 0.5, Borda and Copeland.

Table 3 provides more detailed results for the comparison of these "first step winners". For each combination of the Kendall-Tau distance d and the confidence level c, we consider the number of wins across the three portfolios obtained for the three values of K. These numbers are counted separately for SR and CEQ. The full table, while somewhat overwhelming at first glance, allows for relatively clear conclusions.

First, considering the results for different confidence levels c, the two methods based on robust optimization are reasonably competitive with the two social choice methods for higher confidence levels, i.e. $c \leq 0.5$, but are inferior for low confidence levels, e.g. not achieving a single win for $c = 0.95$ in either sample. This can be explained by the fact that these methods follow a "first estimate, then aggregate" philosophy, i.e., they first translate views into posterior return estimators and then consider these estimators as scenarios in the context of the robust optimization model. For low confidence levels (high values of c), the information contained in the views has a minor impact on the posterior estimator of the vector of expected returns, and

Panel A: Selected constituents of EURO STOXX 50

Method	SR				Certainty-Equivalent return CEQ			
	$K=5$	$K=10$	$K=20$	total	$K=5$	$K=10$	$K=20$	total
Max-min	3	3	2	8	3	2	1	6
Min regret	7	2	3	12	3	2	1	6
Soft 0.25	3	4	4	11	3	1	2	6
Soft 0.5	4	6	6	16	9	10	7	26
Soft 0.75	1	1	6	8	1	2	7	10
Soft 1	1	1	0	2	1	2	0	3
Borda	14	10	14	38	15	13	14	42
Footrule	1	5	8	14	3	3	9	15
Copeland	5	6	14	25	6	12	15	33
Best-of-k	3	0	8	11	3	2	9	14
Markov	3	5	12	20	6	8	10	24

Panel B: Selected constituents of S&P 100

Method	Sharpe Ratio SR				Certainty-Equivalent return CEQ			
	$K=5$	$K=10$	$K=20$	total	$K=5$	$K=10$	$K=20$	total
Max-min	4	5	7	16	1	2	1	4
Min regret	9	6	7	22	1	0	0	1
Soft 0.25	3	2	3	8	6	2	1	9
Soft 0.5	5	4	5	14	8	12	11	31
Soft 0.75	0	3	2	5	0	0	5	5
Soft 1	2	1	1	4	4	3	3	10
Borda	10	9	13	32	12	12	14	38
Footrule	2	8	8	18	4	6	9	19
Copeland	9	11	11	31	9	14	13	36
Best-of-k	3	5	7	15	5	3	9	17
Markov	9	6	9	24	6	10	12	28

Table 2: Number of wins of different methods of aggregating qualitative views organized in groups. Methods 'Max-min' to 'Soft 1' denote robust optimization-based approaches (group I), 'Borda' to 'Markov' are social choice-based approaches (group II). $K \in \{5, 10, 20\}$ indicates the number of views in terms of total orders of assets. The input total orders have four different values of Kendall-Tau distance d to the correct order of returns, $d \in \{0.2, 0.3, 0.4, 0.47\}$. The level of confidence in the views c takes four different values: $c \in \{0.95, 0.75, 0.50, 0.25\}$. The maximum number of wins for any K is 16. Results are counted as wins if they deviate from the corresponding optimum within a group (I or II) by less than 1 percent. Solid (dashed) boxes give the highest (second highest) number of total wins within robust estimation and social choice methods, respectively. Panel A (B) shows results for 38 (76) stocks in the EURO STOXX 50 (S&P 100). Sample period: January 1999 to December 2021. Source: Refinitiv Datastream.

Panel A: Selected constituents of EURO STOXX 50

		Kendall tau distance							
		$d = 0.2$		$d = 0.3$		$d = 0.4$		$d = 0.47$	
Method	c	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	0.95	0	0	0	0	0	0	0	0
Soft 0.5		0	0	0	0	0	0	0	0
Borda		3	3	3	3	3	3	2	2
Copeland		2	2	1	1	1	1	3	3
Min regret	0.75	2	0	1	0	0	0	0	0
Soft 0.5		0	0	1	0	1	0	2	0
Borda		0	3	2	3	1	3	3	3
Copeland		1	3	2	3	2	2	0	1
Min regret	0.50	0	0	1	0	0	0	0	0
Soft 0.5		1	0	0	0	1	0	3	1
Borda		3	3	3	3	2	3	1	1
Copeland		2	3	1	2	1	2	0	1
Min regret	0.25	1	0	1	0	0	0	0	0
Soft 0.5		2	0	1	0	1	0	3	2
Borda		1	3	1	2	2	3	0	0
Copeland		0	1	1	3	1	3	0	1

Panel B: Selected constituents of S&P 100

		Kendall tau distance							
		$d = 0.2$		$d = 0.3$		$d = 0.4$		$d = 0.47$	
Method	c	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	0.95	0	0	0	0	0	0	0	0
Soft 0.5		0	0	0	0	0	0	0	0
Borda		3	3	3	3	3	3	1	1
Copeland		2	2	2	1	1	0	2	2
Min regret	0.75	2	0	2	0	0	0	0	0
Soft 0.5		1	0	0	0	1	0	1	0
Borda		0	3	2	3	1	3	0	0
Copeland		1	3	2	3	2	2	2	3
Min regret	0.50	3	0	3	0	2	0	0	0
Soft 0.5		0	0	0	0	1	0	2	1
Borda		0	3	1	3	1	3	1	1
Copeland		0	3	1	3	2	2	0	2
Min regret	0.25	2	0	2	0	1	0	0	0
Soft 0.5		1	0	0	0	1	0	3	2
Borda		0	3	1	3	1	3	0	0
Copeland		0	2	1	3	0	2	1	2

Table 3: Number of wins for selected methods of aggregating qualitative views per different pairs of values of the Kendall tau distance d and the confidence parameter c. The numbers are aggregated across three different levels for the number of views $K \in \{5, 10, 20\}$. Results are counted as wins if they deviate from the optimum within a parameter setting by less than 1 percent. SR (CEQ) denotes the Sharpe Ratio (Certainty-Equivalent return). Solid boxes give the highest number of total wins within a parameter setting. Panel A (B) shows results for 38 (76) stocks in the EURO STOXX 50 (S&P 100). Sample period: January 1999 to December 2021. Source: Refinitiv Datastream.

the latter is not significantly different from the prior estimator. This leads to very similar scenarios for the robust optimization. Consequently, the optimal robust portfolio does not differ much from the optimal portfolio obtained by using only the prior estimator of expected returns and not considering any qualitative views at all. The social choice methods, on the other hand, aggregate the views into one ordering, still taking all information into account. Only when the posterior return estimator is calculated does the dampening effect of low confidence come into play. Therefore, and for both performance measures, the two social choice methods clearly dominate the horse races for higher values of c.

For high confidence levels (i.e., low values of c), the result is not so clear-cut, at least when looking at the Sharpe Ratios. Evaluating all scenarios with $c \leq 0.5$, Min regret shows the highest number of wins for SR in the S&P 100 sample. But even in this setting, the social choice methods remain the predominant methods in the EURO STOXX 50 sample and in general when CEQ is used as the performance measure.

Second, regarding the influence of different degrees of information validity, expressed by different distances d of the views from the correct order, there is almost no effect on the ranking of the different methods. For both samples and for both performance measures, the social choice methods achieve a higher number of wins than group (I) methods almost without exception. The respective overall winner, which in most cases is the Borda rule, achieves first places regardless of the parameter d. Only for a high level of confidence in the views ($c = 0.25$) do the robust optimization methods beat the social choice methods in some settings; however, the degree of information quality turns out to be irrelevant for this result.

Finally, in addition to the influence of the parameters c and d on the rankings of the different methods of aggregating multiple qualitative views in MVO, we are interested in how the number of considered views K (i.e., total orders) affects the outcome of our horse races. To answer this question, we refer to the results in Table 4, where the actual values of SR and CEQ are given for various values of $K = 5, 10, 20$ for each d (averaged over all c) and for each c (averaged over all d). We find that changing the number of views considered has almost no influence on the overall winner of our horse races. Although the levels of our performance measures increase with the number of views considered, as can well be expected, the rankings of the respective methods in the different settings remain largely unaffected. Aggregating the views before estimating the posterior distribution of portfolio returns turns out to be superior to estimating the individual distributions and then aggre-

gating them, mostly regardless of the number of views considered. Within the winning group of social methods, Borda tends to perform slightly better than Copeland for $K = 5$, while for larger K, e.g. $K = 20$, these two social choice approaches are more or less on par.

To emphasize our main experimental result, we present the total number of wins for our selected methods for all 48 parameter settings in Table 5. It shows that the social choice methods dominate the robust optimization methods in our experiments for both panels and for both performance metrics. We therefore recommend that if an investor considers to include several qualitative views into the BL portfolio optimization framework (s)he should apply one of our proposed methods from social choice theory, which commonly "aggregates first and then estimates" rather than the other way around. The approach suggested by Borda should be preferred over the Copeland rule since it is the winner in three out of our four subtables.

Panel A: Selected constituents of EURO STOXX 50

K = 5	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.43	0.11	1.40	0.10	1.23	0.08	0.44	0.02	1.01	0.05	1.17	0.08	1.15	0.09	1.16	0.09
Soft 0.5	1.23	0.10	1.15	0.09	1.06	0.08	0.53	0.03	0.49	0.03	1.14	0.08	1.17	0.10	1.18	0.10
Borda	1.45	0.12	1.42	0.12	1.34	0.10	0.57	0.04	1.12	0.06	1.24	0.10	1.24	0.11	1.17	0.11
Copeland	1.42	0.12	1.41	0.11	1.32	0.10	0.55	0.03	1.10	0.06	1.23	0.10	1.21	0.11	1.16	0.10

K = 10	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.44	0.10	1.40	0.10	1.26	0.08	0.58	0.03	1.05	0.06	1.22	0.08	1.22	0.09	1.19	0.09
Soft 0.5	1.32	0.11	1.28	0.10	1.17	0.09	0.78	0.05	0.57	0.03	1.32	0.09	1.34	0.11	1.33	0.11
Borda	1.51	0.12	1.52	0.12	1.46	0.11	0.84	0.06	1.27	0.07	1.37	0.11	1.38	0.12	1.30	0.12
Copeland	1.52	0.12	1.49	0.12	1.46	0.11	0.85	0.06	1.27	0.07	1.39	0.11	1.34	0.12	1.31	0.12

K = 20	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.46	0.10	1.43	0.10	1.23	0.08	0.65	0.03	1.04	0.06	1.28	0.08	1.24	0.09	1.22	0.09
Soft 0.5	1.37	0.11	1.41	0.11	1.32	0.10	1.04	0.07	0.83	0.05	1.43	0.11	1.42	0.12	1.47	0.12
Borda	1.50	0.12	1.50	0.12	1.48	0.12	1.09	0.08	1.36	0.08	1.43	0.12	1.38	0.12	1.39	0.13
Copeland	1.50	0.12	1.50	0.12	1.48	0.12	1.10	0.08	1.36	0.08	1.43	0.12	1.40	0.13	1.39	0.13

Panel B: Selected constituents of S&P 100

K = 5	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.64	0.13	1.61	0.12	1.54	0.11	0.58	0.03	1.36	0.07	1.39	0.10	1.33	0.11	1.28	0.11
Soft 0.5	1.34	0.12	1.30	0.11	1.27	0.10	0.78	0.06	0.57	0.03	1.46	0.11	1.35	0.13	1.31	0.13
Borda	1.58	0.14	1.63	0.14	1.63	0.13	0.81	0.06	1.53	0.08	1.47	0.12	1.34	0.13	1.32	0.13
Copeland	1.44	0.14	1.62	0.13	1.63	0.13	0.85	0.06	1.50	0.07	1.49	0.12	1.35	0.13	1.32	0.13

K = 10	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.65	0.12	1.64	0.12	1.59	0.11	0.69	0.03	1.44	0.07	1.45	0.10	1.39	0.11	1.30	0.11
Soft 0.5	1.45	0.13	1.35	0.12	1.27	0.11	1.05	0.08	0.75	0.04	1.52	0.12	1.45	0.14	1.39	0.14
Borda	1.64	0.14	1.66	0.14	1.57	0.14	1.20	0.09	1.70	0.09	1.59	0.13	1.46	0.15	1.31	0.15
Copeland	1.63	0.14	1.66	0.14	1.58	0.14	1.23	0.10	1.71	0.09	1.62	0.13	1.46	0.15	1.32	0.15

K = 20	d = 0.2		d = 0.3		d = 0.4		d = 0.47		c = 0.95		c = 0.75		c = 0.50		c = 0.25	
Method	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ	SR	CEQ
Min regret	1.66	0.13	1.64	0.12	1.56	0.11	0.82	0.03	1.45	0.07	1.47	0.10	1.42	0.11	1.33	0.11
Soft 0.5	1.51	0.14	1.47	0.13	1.47	0.13	1.31	0.11	1.08	0.06	1.67	0.14	1.52	0.15	1.48	0.15
Borda	1.66	0.15	1.65	0.15	1.62	0.14	1.50	0.12	1.84	0.09	1.63	0.15	1.53	0.16	1.44	0.16
Copeland	1.65	0.15	1.64	0.14	1.62	0.14	1.52	0.12	1.82	0.09	1.65	0.15	1.52	0.16	1.44	0.16

Table 4: Average values of Sharpe Ratio (SR) and Certainty Equivalent return (CEQ) obtained from portfolio simulation runs computed by two robust optimization-based methods (Min regret; Soft 0.5) and two social choice-based methods (Borda; Copeland). $K \in \{5, 10, 20\}$ indicates the number of views. The average is taken over the values of SR (CEQ) obtained for each fixed value of the Kendall-Tau distance d or the level of confidence in the views c while varying the other parameter in the whole range of values, respectively: $d \in \{0.2, 0.3, 0.4, 0.47\}$, $c \in \{0.95, 0.75, 0.50, 0.25\}$. For each of the 48 combinations of parameters c, d and K, solid boxes give the method with the highest value for SR (CEQ), while dashed boxes mark those methods deviating from the optimum by less than 1 percent. Panel A (B) shows results for 38 (76) stocks in the EURO STOXX 50 (S&P 100). Sample period: January 1999 to December 2021. Source: Refinitiv Datastream.

6. Conclusions

While it is recognized in practice that relative rankings of future asset returns are easier to generate and generally more reliable than estimates of absolute returns, research on portfolio theory concerning this topic is rather modest. Moreover, portfolio allocation decisions are increasingly made by teams requiring compromises among several different views. This effect has not yet been taken into account in the portfolio optimization literature.

In this paper we propose two general approaches that allow the integration of multiple and potentially divergent views on the ranking of expected future stock returns in mean-variance (MV) portfolio optimization using the parametric Black-Litterman-framework (BL). While methods from group (I) first generate an estimator of expected returns for each individual view and then process all these estimators in a robust MV portfolio optimization framework, the methods based on social choice theory (group (II)) first aggregate all orderings into one single consensus view which subsequently enters MV optimization. We are not aware of any previous literature that has combined social choice theory with MV portfolio optimization in the BL framework.

Applying long-term in-sample analyses to stocks included in one of two broad stock indices and differentiating the views in terms of (i) the level of relevant information they contain, (ii) the level of confidence in the correctness of each generated view, and (iii) the number of views an investor considers in his/her portfolio decision, we find that portfolios built upon our proposed social choice methods (group (II)) clearly outperform portfolios based on robust optimization methods. We therefore recommend that if multiple views are to be considered in the portfolio allocation process, these views should be aggregated into a single view in a first step. Subsequently, this view can be used to derive a quantitative estimator for the posterior distribution of expected stock returns by applying the method proposed in Çela et al. (2021), which allows a qualitative view to be integrated into the BL model. Eventually, MVO can be applied to this extended BL estimator of expected returns.

Since we are (to the best of our knowledge) the first to combine methods from social choice theory with portfolio optimization, a number of questions arises for future research. As a generalization of our ordinal framework, we might consider views represented by partial orders. The extended BL model from Çela et al. (2021) could process partial orders as input, which could then be used as inputs to the robust optimization approaches. However, aggregation of partial orders is typically not studied in the context of social

Panel A: Selected constituents of EURO STOXX 50

Method	Sharpe Ratio SR	Certainty-Equivalent return CEQ
Min regret	6	0
Soft 0.5	16	3
Borda	30	41
Copeland	18	32

Panel B: Selected constituents of S&P 100

Method	Sharpe Ratio SR	Certainty-Equivalent return CEQ
Min regret	17	0
Soft 0.5	11	3
Borda	18	38
Copeland	19	35

Table 5: Number of wins for selected methods of aggregating qualitative views. The numbers are aggregated across three different levels for the number of views $K \in \{5, 10, 20\}$. The input views are total orders of assets with a Kendall tau distance d to the correct order of returns: $d \in \{0.2, 0.3, 0.4, 0.47\}$. $c \in \{0.95, 0.75, 0.50, 0.25\}$ is the parameter representing the confidence in the views. Thus, 48 portfolio simulation runs are performed for each method and the maximum achievable number of wins per method is 48. Results are counted as wins if they deviate from the optimum by less than 1 percent. Solid (dashed) boxes give the highest (second highest) number of total wins within a parameter setting. Panel A (B) shows results for 38 (76) stocks in the EURO STOXX 50 (S&P 100). Sample period: January 1999 to December 2021. Source: Refinitiv Datastream.

choice methods. A further extension could be the combination of qualitative and quantitative information. Clearly, this would be straightforward for robust portfolio optimization methods, after "translating" information of both types into return estimators. Exploiting the potential of social choice methods for such a setting seems to be far less obvious.

As a consequence of our work for practical decision making on asset allocations we would recommend practitioners to take a fresh view at the process of finding consensus between differing opinions (e.g., within an investment committee). It is well-known that personal relations, hierarchies and anticipation of expected behaviour have a notable influence on the outcome of a group decision. We propose to use the tools of robust optimization and social choice theory to formalize and objectify the consensus building process for decisions involving ordered preferences in portfolio optimization.

References

Ailon, N., Charikar, M., Newman, A., 2008. Aggregating inconsistent information: ranking and clustering. Journal of the ACM 55, 1–27. No. 23.

Allaj, E., 2020. The Black–Litterman model and views from a reverse optimization procedure: an out-of-sample performance evaluation. Computational Management Science 17, 465–492. URL: https://doi.org/10.1007/s10287-020-00373-6.

Arrow, K.J., Sen, A.K., Suzumura, K. (Eds.), 2002. Handbook of Social Choice and Welfare. Elsevier.

Belotti, P., Kirches, C., Leyffer, S., Linderoth, J., Luedtke, J., Mahajan, A., 2013. Mixed-integer nonlinear optimization. Acta Numerica 22, 1–131.

Ben-Tal, A., Nemirovski, A., 2002. Robust optimization - methodology and applications. Mathematical Programming, Series B 92, 453–480.

Bertsimas, D., Gupta, V., Paschalidis, I.C., 2012. Inverse optimization: a new perspective on the Black-Litterman model. Operations Research 60, 1389–1403.

Bertsimas, D., Sim, M., 2003. Robust discrete optimization and network flows. Mathematical Programming B 98(1-3), 49–71.

Best, M.J., Grauer, R.R., 1991. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. The Review of Financial Studies 4, 315–342. URL: https://doi.org/10.1023/B:ANOR.0000045281.41041.ed.

Black, F., Litterman, R., 1991. Asset allocation: combining investor views with market equilibrium. The Journal of Fixed Income 1, 7–18. URL: https://doi.org/10.3905/jfi.1991.408013.

Black, F., Litterman, R., 1992. Global portfolio optimization. Financial Analysts Journal 48, 28–43. URL: https://doi.org/10.2469/faj.v48.n5.28.

Bodie, Z., Kane, A., Marcus, A., 2021. Investments. 12th ed., Mc Graw Hill.

Bodnar, T., Lindholm, M., Niklasson, V., Thorsén, E., 2022. Bayesian portfolio selection using VaR and CVaR. Applied Mathematics and Computation 427, 127120. URL: https://doi.org/10.1016/j.amc.2022.127120.

Brams, S.J., Fishburn, P.C., 2002. Voting procedures, in: Arrow, K.J., Sen, A.K., Suzumura, K. (Eds.), Handbook of Social Choice and Welfare. Elsevier. volume 1, pp. 173–236.

Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D., 2016. Handbook of Computational Social Choice. Cambridge University Press.

Çela, E., Hafner, S., Mestel, R., Ulrich, P., 2021. Mean-variance portfolio optimization based on ordinal information. Journal of Banking and Finance 122. Article number 105989.

Chen, S.D., Lim, A.E., 2020. A generalized Black–Litterman model. Operations Research 68, 381–410. URL: https://doi.org/10.1287/opre.2019.1893.

Chopra, V.K., Ziemba, W.T., 1993. The effect of errors in means, variances, and covariances on optimal portfolio choice. Journal of Portfolio Management 19, 6–11. URL: https://doi.org/10.3905/jpm.1993.409440.

DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? Review of Financial Studies 22, 1915–1953. URL: https://doi.org/10.1093/rfs/hhm075.

DeMiguel, V., Nogales, F.J., 2009. Portfolio selection with robust estimation. Operations Research 57, 560–577. URL: https://doi.org/10.1287/opre.1080.0566.

Diaconis, P., Graham, R.L., 1977. Spearman's footrule as a measure of disarray. Journal of the Royal Statistical Society. Series B 39, 262–268.

Duran, M.A., Grossmann, I.E., 1986. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. Mathematical Programming 36, 307–339.

Dwork, C., Kumar, R., Naor, M., Sivakumar, D., 2001. Rank aggregation methods for the Web, in: Proceedings of the 10th International Conference on World Wide Web, ACM. pp. 613–622.

Fabozzi, F.J., Kolm, P.N., Pachamanova, D.A., Focardi, S.M., 2007. Robust Portfolio Optimization and Management. John Wiley & Sons.

Fishburn, P.C., 1977. Condorcet social choice functions. SIAM Journal of Applied Mathematics 33, 469–489.

Harvey, C.R., Liechty, J.C., Liechty, M.W., Müller, P., 2010. Portfolio selection with higher moments. Quantitative Finance 10, 469–485. URL: https://doi.org/10.1080/14697681003756877.

Huang, D., Zhu, S., Fabozzi, F.J., Fukushima, M., 2010. Portfolio selection under distributional uncertainty: A relative robust CVaR approach. European Journal of Operational Research 203, 185–194.

Kolm, P.N., Tütüncü, R.H., Fabozzi, F.J., 2014. 60 years of portfolio optimization: Practical challenges and current trends. European Journal of Operations research 234, 356–371.

Kronqvist, J., Bernal, D.E., Lundell, A., Grossmann, I.E., 2019. A review and comparison of solvers for convex MINLP. Optimization and Engineering 20, 397–455.

Lassance, N., Vrins, F., 2021. Portfolio selection with parsimonious higher comoments estimation. Journal of Banking & Finance 126, 106115. URL: https://doi.org/10.1016/j.jbankfin.2021.106115.

Levy, H., Levy, M., 2014. The benefits of differential variance-based constraints in portfolio optimization. European Journal of Operational Research 234, 372–381. URL: https://doi.org/10.1016/j.ejor.2013.04.019.

Marden, J., 1996. Analyzing and Modeling Rank Data. Chapman and Hall/CRC.

Markowitz, H., 1952. Portfolio selection. Journal of Finance 7, 77–91. URL: https://doi.org/10.1111/j.1540-6261.1952.tb01525.x.

Markowitz, H., 1959. Portfolio Selection. John Wiley & Sons.

Meucci, A., 2006. Beyond Black-Litterman: Views on non-normal markets. Risk Magazine 19, 87–92. URL: https://doi.org/10.1017/S0022109017000655.

Meucci, A., 2008. Fully flexible views: theory and practice. Risk 21, 97–102.

Meucci, A., 2010. Black-Litterman approach, in: Encyclopedia of Quantitative Finance. John Wiley & Sons. URL: https://doi.org/10.1002/9780470061602.eqf14009.

Michaud, R.O., 1989. The Markowitz optimization enigma: is 'optimized' optimal? Financial Analysts Journal 45, 31–42. URL: https://doi.org/10.2469/faj.v45.n1.31.

Pastor, L., Stambaugh, R.F., 2000. Comparing asset pricing models: an investment perspective. Journal of Financial Economics 56, 335–381. URL: https://doi.org/10.1016/S0304-405X(00)00044-1.

Patel, S., Sarkissian, S., 2017. To group or not to group? evidence from mutual fund databases. Journal of Financial and Quantitative Analysis 52, 1989–2021. URL: https://doi.org/10.1017/S0022109017000655.

Sah, R.K., Stiglitz, J.E., 1986. The architecture of economic systems: Hierarchies and polyarchies. American Economic Review 76, 716–727.

Schalekamp, F., van Zuylen, A., 2009. Rank aggregation: Together we're strong, in: Proceedings of the 11th Workshop on Algorithm Engineering and Expermiments (ALENEX 2009), SIAM. pp. 38–51.

Sen, A., 1970. Collective Choice and Social Welfare. Holden-Day.

Simonian, J., 2014. Copula-opinion pooling with complex opinions. Quantitative Finance 14, 941–946.

Vavasis, S.A., 2001. Complexity theory: Quadratic programming, in: Floudas, C.A., Pardalos, P.M. (Eds.), Encyclopedia of Optimization. Springer, Boston, MA. URL: https://doi.org/10.1007/0-306-48332-765.

Walters, J., 2014. The Black-Litterman model in detail URL: http://dx.doi.org/10.2139/ssrn.1314585. working Paper. Available at SSRN.