Liesiö, Juuso; Kallio, Markku; Argyris, Nikolaos

Incomplete risk-preference information in portfolio decision analysis

Published in:
European Journal of Operational Research

DOI:
10.1016/j.ejor.2022.04.043

Published: 01/02/2023

Document Version
Publisher's PDF, also known as Version of record

European Journal of Operational Research 304 (2023) 1084–1098

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Incomplete risk-preference information in portfolio decision analysis

Juuso Liesiö[a,*], Markku Kallio[a], Nikolaos Argyris[b]

[a] Department of Information and Service Management, Aalto University School of Business, Finland
[b] School of Business and Economics, Loughborough University, Leicestershire, UK

ARTICLE INFO

Article history:
Received 16 December 2020
Accepted 28 April 2022
Available online 4 May 2022

Keywords:
Decision support systems
Portfolio decision analysis
Stochastic dominance
Multi-objective programming/optimization

ABSTRACT

Portfolio decision analysis models support decisions on the allocation of resources among assets with uncertain outcomes (e.g., investments, projects or stocks). These models require information on the decision maker's risk-preferences which can be difficult to obtain in practice. Stochastic dominance criteria show promise in this regard as they can compare portfolios without exact specification of risk-preferences, but the current literature lacks practical approaches for generating the efficient frontier, i.e., the set of those portfolios that are not stochastically dominated by any other portfolio. We address this gap by developing models to identify sets of portfolios that are efficient in the sense of second- or third-order stochastic dominance (SSD, TSD). These models provide novel insights into the composition of portfolios belonging to the efficient frontier by, e.g., identifying those assets that are included in all efficient portfolios. Moreover, the identification of the efficient frontier makes it possible to utilize additional information on the decision maker's risk-preferences to further reduce the set of admissible portfolio alternatives, and to analyze the implications this information has on the amount of capital that should be allocated to each individual asset. We illustrate the usefulness of these models with applications in project portfolio selection and financial portfolio diversification.

1. Introduction

Organizations need to decide how to diversify capital and other resources among assets with uncertain outcomes such as R&D projects, products or financial instruments. These decisions can be supported with portfolio decision analysis (PDA; Salo, Keisler, & Morton, 2011) models that harness utility theory to capture the decision makers' risk-preferences and optimization techniques to identify the optimal portfolio. PDA models usually encode risk preference with a single utility function that is consistent with the preference statements the decision maker provides on pairs of uncertain outcomes. Providing such complete preference information can be time-consuming and difficult for the decision maker (see, e.g., Hazen, 1986, Moskowitz, Preckel, & Yang, 1993) and hence models that allow for incomplete information especially on multiobjective preferences have been widely deployed to support real-life portfolio decision making in, for instance, infrastructure asset management (Liesiö, Mild, & Salo, 2007; Mild, Liesiö, & Salo, 2015), strategy formation (Vilkkumaa, Liesiö, Salo, & Ilmola-Sheppard, 2018) and military applications (Kangaspunta, Liesiö, & Salo, 2012; Harju, Liesiö, & Virtanen, 2019).

Indeed, confidence in decision recommendations can be hindered if they are based on unrealistic assumptions about the accuracy of the process used to elicit risk-preferences. This opens up the important question of how robust the decision recommendations are relative to changes in the elicited risk-preference information utilized in the analysis. Motivated by this questions, we consider the problem of providing support for portfolio decisions without requiring complete specification of risk-preference through a single utility function. In particular, we address the general questions of (i) which portfolios can be optimal if only incomplete information of risk-preferences is available, and (ii) what decision recommendation can be given on the level of individual assets based on incomplete preference information, e.g., are there projects whose selection/rejection is not contingent on exact shape of the utility function?

Stochastic dominance (SD) offers an appealing tool to tackle these questions. SD can be used to compare portfolios without exact specification of risk-preferences and it has strong decision theoretic foundations in Expected Utility Theory (EUT; Neumann & Morgenstern, 1947). In particular, second-order stochastic dominance (SSD, Hadar & Russell, 1969) identifies if a portfolio is preferred to another by any decision maker with a non-decreasing

* Corresponding author.
E-mail addresses: juuso.liesio@aalto.fi (J. Liesiö), markku.kallio@aalto.fi (M. Kallio), n.argyris@lboro.ac.uk (N. Argyris).

https://doi.org/10.1016/j.ejor.2022.04.043

concave utility function. Third-order stochastic dominance (TSD, Whitmore, 1970) further assumes that the decision maker prefers higher skewness of the outcome distribution, i.e., given two distributions with equal expectations and variances, the DM would choose the one more skewed to the right. Plenty of research has been devoted to utilize SD in portfolio selection, diversification and resource allocation problems (see, e.g., Karsu, Morton, & Argyris, 2017; Liesiö & Salo, 2012; Ringuest, Graves, & Case, 2000). Perhaps the most widely used approach for incorporating SD into portfolio selection problems has been the deployment of SD constraints to ensure that the optimal portfolio stochastically dominates (weakly) some predefined benchmark (see, e.g., Dentcheva & Ruszczyński, 2003; Kallio & Deghan Hardoroudi, 2017; Kuosmanen, 2004; Liesiö, Xu, & Kuosmanen, 2020b; Post, 2003; Post & Kopa, 2013; Post & Kopa, 2017), although important work has also been carried-out to develop statistical tests for SD-efficiency (see, e.g., Arvanitis, Hallam, Post, & Topaloglou, 2019; Linton, Post, & Whang, 2014; Scaillet & Topaloglou, 2010)

However, these existing approaches cannot fully address the two questions (i) and (ii) stated above. Firstly, even when completely specifying a utility function is not possible, it may be still be possible — even desirable — to utilize limited preference information in providing portfolio and project recommendations. Between the complete specification of a utility function and mere specification that the preferences are risk-averse lies a broad range of possibilities on information about preferences; it may be desirable to arrive at recommendations by utilizing more information than only specifying a benchmark portfolio. Secondly, arriving at project level recommendations requires examination of the entire SD-efficient set of portfolios, so existing approaches cannot be used to that effect.

Both of the above shortcomings can be addressed through the identification of the entire SD-efficient set of portfolios. Once this has been generated, it immediately becomes possible to provide project-level recommendations by examining the composition of all portfolios in the efficient set. In addition, it also becomes possible — in a way that we specify later in the paper — to incorporate additional preference information to the mere specification of risk-attitude, and examine the effect of this on portfolio and project recommendations. In fact, as we will illustrate below, it even becomes possible to arrive at recommendations through interactive decision support.

In view of the above, the lack of tools for identifying the SD-efficient set might be seen as a shortcoming. Indeed, within the financial portfolio selection literature, the need for models to identify the SD-efficient frontier was recognized explicitly by Levy (1992): *"Ironically, the main drawback of the SD framework is found in the area of finance where it is most intensively used, namely in choosing the efficient diversification strategies. This is because as yet there is no way to find the SD efficient set of diversification strategies... Therefore, the next important contribution in this area will probably be in this direction."* Subsequent work has made important advances in this area, mainly focused on the SSD-efficient set (for a summary, see, e.g. Levy, 2016). First, Roman, Darby-Dowman, & Mitra (2006) establish a seminal result linking SSD-efficient portfolio to the Pareto optimal solutions of a multi-objective optimization problem in case of equally likely states, which can be viewed, as we will show, as a special case of the general equivalence considered here. Second, Longarela (2016) showed that any portfolio that is optimal for some non-decreasing concave utility function must satisfy a specific system of linear constraints involving both continuous and integer-valued variables. Although this formulation is theoretically interesting we do not utilize it in this paper because the number of binary variables it requires grows polynomially in the number of states and as a result the formulation quickly becomes computationally infeasible for larger problems. Furthermore, there is no clear path to extending the formulation by Longarela (2016) from concave utility functions to those representing preferences consistent with TSD. Third, Arvanitis et al. (2019) invented the concept of stochastic spanning: an investment opportunity set is stochastically spanned by its subset if each portfolio in the set is weakly SSD dominated by some portfolio in the subset. Indeed, the smallest subset spanning the entire investment opportunity set would correspond to the set of SSD efficient portfolios. However, Arvanitis et al. (2019) do not provide computational approaches for generating the minimal spanning subset. While Longarela (2016) and Arvanitis et al. (2019) make important advances in the SD-literature, they do not consider approaches for identification of the efficient portfolios, because the focus of these articles is on the plethora of SD-applications in which the generation of the entire efficient set is not needed. Thus, the current literature does not offer approaches for identification of the efficient set that are general in the sense that they can (i) incorporate preferences consistent with both SSD and TSD, (ii) be utilized for continuous and integer stochastic optimization problems, and (iii) can handle unequal state probabilities.

This gap in the literature appears stark in contrast to the attention given to frontier identification in related fields. For example, within the multi-objective optimization literature the problem of identifying the efficient frontier of multi-objective optimization problems has received considerable attention (see, e.g. Ehrgott, 2005). Moreover, several approaches have been developed for portfolio selection and resource allocation under multiple objectives that make use of the identification and analysis of the efficient frontier (see, e.g., Stummer & Heidenberger, 2003, Liesiö et al., 2007; Liesiö, Mild, & Salo, 2008, Kangaspunta et al., 2012, Mild et al., 2015, Mancuso, Compare, Salo, Zio, & Laakso, 2016, Liesiö, Andelmin, & Salo, 2020a).

This paper addresses this gap in the literature by developing methods to identify the frontiers of SSD- and TSD-efficient portfolios. These methods build on the equivalence between the SD-efficient portfolios and the Pareto optimal solutions to multi-objective programming problems. We utilize this general equivalence to develop models for solving the SSD- and TSD-efficient portfolios when the uncertain returns are captured by a finite state-space. In particular, we develop linear and quadratic multi-objective optimization models and show how existing multi-objective methods that can be used to generate the SSD- and TSD-efficient frontiers. We demonstrate that in some applications it is possible to utilize the Benson algorithm (see, e.g., Löhne, 2011) to solve the finite number of polyhedral sets that form the SSD-efficient set. For other applications we utilize weighted sum and Tchebychef-norm methods to generate a finite number of SSD- and TSD-efficient portfolios by solving a series of single objective optimization problems. Overall this paper contributes to the growing literature on the interfaces between stochastic and multi-objective optimization (for a survey, see, e.g., Gutjahr & Pichler, 2016).

The developed models enable novel ways to analyze risk-averse portfolio diversification problems and to provide decision support for portfolio selection. First, identification of the efficient frontier makes it possible to address question concerning the common properties of efficient (or inefficient) portfolios: For instance, are there assets which are allocated resources in all or none of the efficient portfolios? Identification of such assets enables to determine which decisions are actually contingent on decision maker's risk preferences, and to highlight the decisions that are not affected by risk-preferences. Second, our models enable the identification of potentially optimal portfolios, i.e. those that may be optimal for some utility function compatible with incomplete preference information elicited from the decision maker. Third, identification of the SD-efficient frontier makes it possible to utilize additional information on the decision maker's risk-preferences to further reduce the number of potentially optimal portfolios and thus offer more conclusive decision recommendations on the level of individual assets. Specifically, identification of those portfolios that are potentially optimal in view of the additional preference information can be carried-out using linear programming or convex optimization. We demonstrate the developed models and their usage in project portfolio selection and financial portfolio diversification.

The rest of the paper is structured as follows. Section 2 introduces the mathematical framework, and formalizes relevant SD concepts. Section 3 develops multi-objective optimization models for identifying the SSD- and TSD-efficient portfolios. Sections 4 and 5 presents applications in scenario-based project portfolio selection and financial portfolio diversification, respectively. Section 6 concludes.

2. Portfolio selection under uncertainty

Although our notation is in line with that generally used in financial portfolio selection, our assumptions do not rule out applications that require the use of integer-valued decision variables (e.g. project portfolio selection) or applications in which portfolios are evaluated based on other measures than total returns. In particular, the total returns of the m assets (projects) are modelled with real-valued random variables $R_1, \dots, R_m$ on the set of mutually exclusive and collectively exhaustive states $S = \{s_1, \dots, s_n\}$ with state probabilities $\mathbb{P}(\{s_i\}) = p_i$. The total return of the jth asset in the ith state is denoted by $r_{ij} = R_j(s_i)$, and we use $r = (r_{ij}) \in \mathbb{R}_+^{n\times m}$ to denote a matrix consisting of the assets' state-specific returns. The indices of the states and the assets belong to sets $N = \{1, \dots, n\}$ and $M = \{1, \dots, m\}$, respectively, and $e = (1, \dots, 1)^T$ denotes a sum vector.

A portfolio is characterized by a vector of asset weights $\lambda = (\lambda_1, \dots, \lambda_m)^T$. In financial portfolio selection these weights capture the share of initial capital allocated to each asset and hence sum up to one ($e^T\lambda = 1$). In project portfolio selection these weights take binary values $\lambda \in \{0, 1\}^m$ indicating whether the jth project is selected ($\lambda_j = 1$) or not ($\lambda_j = 0$), and limited resources (e.g. budget) can be modelled by introducing linear constraints on these weights. Hence, in order to preserve the generality of the model we only assume that the bounded set of feasible weight vectors is given by

$$\Lambda = \{\lambda \in \underline{\Lambda} \mid A\lambda \le b\}, \tag{1}$$

where $\underline{\Lambda} = \mathbb{R}^m$ or $\underline{\Lambda} = \{0, 1\}^m$, and matrix $A \in \mathbb{R}^{h\times m}$ and vector $b \in \mathbb{R}^h$ contain the coefficients of the h linear constrains. The return of a portfolio λ is captured by random variable $R^\lambda = \sum_{j\in M} \lambda_j R_j$, whose state-specific return vector is denoted by $r^\lambda = (r_i^\lambda) = r\lambda \in \mathbb{R}^n$. The lowest and highest state-specific returns across all portfolios are denoted by $\underline{r}$ and $\bar{r}$, respectively.

Establishing stochastic dominance between two portfolios is based on comparing the cumulative distribution functions (CDFs) of the portfolios' returns. In particular, let $F_\lambda^1(t) = \mathbb{P}(R^\lambda \le t)$ denote the CDF of portfolio return R^λ, and F_λ^d the function that is obtained when the CDF is integrated $d-1$ times, i.e., $F_\lambda^d(t) = \int_{-\infty}^t F_\lambda^{d-1}(y)dy$. With this notation dth-order stochastic dominance can be defined as follows (see, e.g., Gotoh & Konno, 2010).

Definition 1. Portfolio λ weakly dominates portfolio τ in the sense of dth-order stochastic dominance, denoted by $\lambda \succeq^d \tau$, if

$$F_\lambda^d(t) \le F_\tau^d(t) \ \forall\ t \in \mathbb{R}.$$

Strict dominance $\lambda \succ^d \tau$ holds if $\lambda \succeq^d \tau$ and $\neg(\tau \succeq^d \lambda)$.

We focus on second- and third-order stochastic dominance (SSD, TSD) as their economic interpretations are particularly appealing from the view-point of portfolio diversification. Specifically, if a portfolio weakly dominates another in the sense of SSD ($d = 2$), then any expected utility maximizing risk-averse or -neutral decision maker would weakly prefer the former portfolio over the latter (Hanoch & Levy, 1969). Formally, this result can be stated as

$$\lambda \succeq^2 \tau \Leftrightarrow \mathbb{E}[u(R^\lambda)] \ge \mathbb{E}[u(R^\tau)] \text{ for all } u \in U^2,$$

where U^2 is the set of all non-decreasing concave utility functions that are non-constant on the interval $[\underline{r}, \bar{r}]$.

TSD ($d = 3$) imposes a further restriction on the set of utility functions (Whitmore, 1970). In particular, it requires a higher or equal expected utility across those concave utility functions that exhibit preference for higher skewness of the return distribution.These restrictions result in utility functions $u \in U^3 \subset U^2$ such that the marginal utility $u'(t)$ is a convex function in $\mathbb{R}$. The set U^3 of utility functions satisfies the equivalence

$$\lambda \succeq^3 \tau \Leftrightarrow \mathbb{E}[u(R^\lambda)] \ge \mathbb{E}[u(R^\tau)] \text{ for all } u \in U^3.$$

Stochastic dominance relations can be used to identify those portfolios which are not strictly dominated by any other portfolio. The resulting set of SD-efficient portfolios is defined as follows.

Definition 2. The set of dth-order SD-efficient portfolios is

$$\Lambda_E^d = \{\tau \in \Lambda \mid \nexists \lambda \in \Lambda \text{ s.t. } \lambda \succ^d \tau\}.$$

If a portfolio is efficient for some order of SD, it is also efficient for any lower order of SD. In particular, the TSD-efficient portfolios form a subset of the SSD-efficient portfolios. This well-known result is formally stated by the following lemma.

Lemma 1. $\Lambda_E^3 \subseteq \Lambda_E^2$.

The set of efficient portfolios provides a frontier of defensible risk diversification strategies: If a portfolio is selected outside SSD-efficient set Λ_E^2, then there exists a portfolio in this set that yields a greater or equal expected utility for any utility function $u \in U^2$, and hence the latter portfolio would be (weakly) preferred by any risk-averse or -neutral decision maker. Similarity, if a portfolio is selected outside the TSD-efficient set Λ_E^3, then some portfolio in this set would be (weakly) preferred by all risk-averse 'skewness loving' decision makers. Moreover, for any decision maker with a utility function belonging to set U^2 (U^3), the highest expected utility is offered by one of the SSD-efficient (TSD-efficient) portfolios.

Instead of using efficient portfolios (Definition 2) as the fundamental solution concept, an alternative approach would be to focus on the potentially optimal portfolios, i.e., those portfolios that maximize expected utility for some utility function in the set U^d. The sets of efficient and potentially optimal portfolios are not necessarily the same. For instance, consider two portfolios such that the return distribution of one portfolio is a mean preserving spread of the return distribution of the other portfolio. Then it is possible that both are potentially optimal as they yield the same expected utility for the linear utility function. However, it is also possible that one dominates the other in sense of SSD and hence both cannot be SSD-efficient. Indeed, our choice of focusing on efficient portfolios on the outset is motivated by the concepts' robustness against such pathological cases. Note that this choice does not involve the risk of discarding some portfolios that could be seen as defensible choices. In particular, any potentially optimal portfolio outside the efficient set would be a poor selection as the dominating portfolio (i) is also potentially optimal, (ii) yields a greater or equal value for all utility functions in U^d, and (iii) has a strictly greater expected utility for some utility function in U^d. In addition, the identification of the SSD-efficient portfolios that are also potentially optimal can be pursued after the SD-efficient set has been solved.

We utilize potential optimality in analyzing how additional information preference information beyond that encoded in the SSD-criterion reduces the set admissible portfolio alternatives. In particular, we propose an approach in which the decision maker states her preferences between pairs of uncertain returns or outcomes and identify the set of utility functions that are consistent with these preferences. Decision recommendations can then be based on identifying those efficient portfolios that yield maximal expected utility for some utility function that is consistent with the given preference statements. Since the approach is based on using the efficient frontier as the set decision alternatives, any portfolio labeled as potentially optimal will also be efficient.

Formally, suppose (X^k, Y^k), $k \in \{1, \ldots, K\}$, are K pairs of real-valued random variables on the state-space $S = \{s_1, \ldots, s_n\}$ about which the DM has stated her (pairwise) preferences. These random variables can correspond to the returns of actual portfolios (i.e., $X^k = R^\lambda$ for some $\lambda \in \Lambda$) or they can be lotteries structured only for the purposes of preference elicitation. Without loss of generality, suppose that the decision maker weakly prefers X^k over Y^k for each $k \in \{1, \ldots, K\}$. The set of risk-averse utility functions that are compatible with these preferences is

$$\tilde{U}^K = \{u \in U^2 \mid \mathbb{E}[u(X^k)] \geq \mathbb{E}[u(Y^k)] \ \forall\ k \in \{1, \ldots, K\}\}. \tag{2}$$

In the ensuing we will assume that DM's elicited pairwise preferences are 'rationalizable' i.e. compatible with at least one utility function in U^2, implying that set $\tilde{U}^K$ is non-empty. In practice this can be ensured via incremental preference elicitation, where a pair X^{k+1}, Y^{k+1} is chosen such that either response from the DM is consistent with all previously elicited statements (i.e. there exists $u, u' \in \tilde{U}^K$ such that $\mathbb{E}[u(X^{K+1})] \geq \mathbb{E}[u(Y^{K+1})]$ and $\mathbb{E}[u'(Y^{K+1})] \geq \mathbb{E}[u'(X^{K+1})]$). One way to generate such a pair is to identify portfolios that maximize expected utility for some utility function $u \in \tilde{U}^K$, as we will discuss subsequently in more detail.

The SSD-efficient portfolio λ is *potentially optimal* with regard to this set of utility functions if

$$\lambda \in \arg\max_{\tau \in \Lambda_E^2} \mathbb{E}[u(R^\tau)] \text{ for some } u \in \tilde{U}^K. \tag{3}$$

When the set of utility functions is of the form (2), linear programming can be used to establish which of the SSD-efficient portfolios are potentially optimal (see, e.g., Argyris, Morton, & Figueira, 2015; Armbruster & Delage, 2015; Liesiö & Salo, 2012; Post & Khanjani Shiraz, 2019; Post & Kopa, 2013).

Although in this paper we focus on analysing which SSD-efficient portfolios remain admissible decision alternatives when additional preference information is introduced, a similar analysis can be carried out for the TSD-efficient portfolios too. In this case, non-linear convex optimization (Kallio & Dehghan Hardoroudi, 2019) can be used to identify those TSD-efficient portfolios that are potentially optimal with regard to the subset of utility functions in U^3 that are compatible with preferences.

3. Identification of the SD-efficient frontiers

Our approaches for identifying the set of SD-efficient portfolios are based on the observation that they have a one-to-one correspondence with the *Pareto optimal solutions* of a particular multiple objective optimization problem with an infinite number of objective functions. In multi-objective optimization literature the terms 'Pareto optimal' and 'efficient' are often used as synonyms, but to avoid confusion, this paper uses the term 'efficient' only to refer to portfolios that are not stochastically dominated (cf. Definition 2). A feasible solution to a multi-objective optimization problem is said to be Pareto optimal, if no other feasible solution provides (i) a better or equal value in each objective function, and (ii) a strictly better value in at least one objective function.[1] Hence, constructing the objective functions by evaluating the integrated CDF of portfolio returns $F_\lambda^d(t)$ for each return level $t \in \mathbb{R}$ directly results in the following proposition, which can thus be presented without a proof.

Proposition 1. *Portfolio $\lambda \in \Lambda$ is efficient in the sense of d th-order SD, i.e., $\lambda \in \Lambda_E^d$, if and only if it is a Pareto optimal solution to the multi-objective programming (MOP) problem*

$$\text{v–min}_{\lambda \in \Lambda} \left(F_\lambda^d(t)\right)_{t \in \mathbb{R}}. \tag{4}$$

In the following sections we develop practical approaches for solving problem (4) when $d = 2$ and $d = 3$ to identify the SSD and TSD efficient frontiers.

3.1. Identification of the SSD-efficient frontier

In case of SSD, the objective function of problem (4) evaluated at t can be formulated as the LP problem

$$\begin{aligned} F_\lambda^2(t) &= \int_{-\infty}^{t} \overbrace{\Bigg(\sum_{\substack{i \in N \\ r_i^\lambda \leq t}} p_i\Bigg)}^{=F_\lambda^1(t)} dy = \sum_{\substack{i \in N \\ r_i^\lambda \leq t}} p_i (t - r_i^\lambda) \\ &= \sum_{i \in N} p_i \max\{t - r_i^\lambda, 0\} = \min_{\phi_i \geq 0}\{\sum_{i \in N} p_i \phi_i \mid \phi_i \geq t - r_i^\lambda\}, \end{aligned} \tag{5}$$

where the last equality holds because at optimum the value of each ϕ_i will be as low as possible, i.e., equal to maximum of the two lower bounds $t - r_i^\lambda$ and zero.

For any portfolio the integrated CDF F^2 is zero below the lowest state-specific return $\underline{r}$. Moreover, the integrated CDF increases at a unit rate for all return levels above the highest state-specific return $\bar{r}$. Hence, it is sufficient to evaluate the objective function of problem (4) only for the return levels $t \in [\underline{r}, \bar{r}]$. This observation together with Eq. (5) implies that portfolio $\lambda \in \Lambda$ is SSD-efficient if and only if (λ, Φ) is a Pareto optimal solution to the MOP problem

$$\begin{aligned} \underset{\lambda,\Phi}{\text{v–min}} \quad & \Big(\sum_{i \in N} p_i \Phi_i(t)\Big)_{t \in [\underline{r}, \bar{r}]} \\ & \Phi_i(t) \geq t - r_i^\lambda \ \forall\ i \in N, t \in [\underline{r}, \bar{r}] \\ & \Phi_i(t) \geq 0 \ \forall\ i \in N, t \in [\underline{r}, \bar{r}] \\ & \lambda \in \Lambda. \end{aligned} \tag{6}$$

The decision variables Φ_i are functions on the interval $[\underline{r}, \bar{r}]$. Hence, implementing and solving this problem in practice requires choosing some discretization points $\{t_1, \ldots, t_l\} \subset [\underline{r}, \bar{r}]$ in which these functions are evaluated. By denoting $\phi_{i,k} = \Phi_i(t_k)$ problem (6) becomes

$$\begin{aligned} \underset{\lambda,\Phi}{\text{v–min}} \quad & \Big(\sum_{i \in N} p_i \phi_{i,k}\Big)_{k \in \{1,\ldots,l\}} \\ & \phi_{i,k} \geq t_k - r_i^\lambda \ \forall\ i \in N, k \in \{1, \ldots, l\} \\ & \phi_{i,k} \geq 0 \ \forall\ i \in N, k \in \{1, \ldots, l\} \\ & \lambda \in \Lambda, \end{aligned} \tag{7}$$

which is a multi-objective linear programming problem with l objective functions.

In the case where the set of feasible asset weights Λ (1) is convex, Pareto optimal solutions to problem (7) can be obtained

[1] Formally, denoting the index set of the objectives functions by T, solution $\omega^* \in \Omega$ is Pareto optimal to the multi-objective problem $\text{v–min}_{\omega \in \Omega}\big(h_\omega(t)\big)_{t \in T}$ if there does not exist $\omega' \in \Omega$ such that $h_{\omega'}(t) \leq h_{\omega^*}(t)$ for all $t \in T$ and $h_{\omega'}(t) < h_\omega^*(t)$ for some $t \in T$.

with the multi-objective simplex (see, e.g., Ehrgott, 2005), the Benson algorithm (Benson, 1998, Löhne, 2011), or the weighted sum method. In case Λ is non-convex, Pareto optimal solutions can be obtained with the reference point method (Wierzbicki, 1980), the weighted Tchebychef-norm method (Bowman, 1976) or multi-objective mixed-integer linear programming algorithms (see, e.g., Przybylski & Gandibleux, 2017).

3.2. SSD-efficient frontier for sample-generated state-space

In practical applications the states often correspond to a sample of returns obtained from historical observations, or from a joint distribution of forecasted future returns, and thus it is reasonable to assume that the states are equally likely. With equal state probabilities there exists a one-to-one correspondence between the SSD-efficient portfolios and the Pareto optimal solutions to a MOP problem with a finite number of objective functions. Thus, we assume throughout the rest of this subsection that $p_i = \frac{1}{n}$ for each $i \in N$.

Under equal state probabilities the integrated CDF (5) of portfolio λ can be written as

$$F_\lambda^2(t) = \sum_{i\in N} \frac{1}{n}\max\{t - r_i^\lambda, 0\}$$
$$= \frac{1}{n}\max\left\{0, t-g_1(\lambda), 2t-g_2(\lambda), \ldots, nt-g_n(\lambda)\right\}, \quad (8)$$

where $g_k(\lambda)$ denotes the sum of the k smallest state-specific returns $r_1^\lambda, \ldots, r_n^\lambda$. Suppose that portfolio λ does not dominate portfolio τ in SSD, i.e., $F_\tau^2(t) < F_\lambda^2(t)$ for some return level t. Then, according to Eq. (8), there exists $k \in N$ such that

$$\frac{1}{n}(kt - g_k(\lambda)) = F_\lambda^2(t) > F_\tau^2(t) \geq \frac{1}{n}(kt - g_k(\tau)) \Rightarrow g_k(\lambda) < g_k(\tau).$$

In turn, if $g_k(\lambda) < g_k(\tau)$ for some $k \in N$, then choosing t such that $F_\tau^2(t) = \frac{1}{n}(kt - g_k(\tau))$ gives

$$F_\tau^2(t) = \frac{1}{n}(kt - g_k(\tau)) < \frac{1}{n}(kt - g_k(\lambda)) \leq F_\lambda^2(t) \Rightarrow F_\tau^2(t) < F_\lambda^2(t).$$

Together these observations imply that $\neg(\lambda \succeq^2 \tau)$ holds if and only if $g_k(\lambda) < g_k(\tau)$ for some $k \in N$. In fact, the negation of this statement is the well-known result that $g_k(\lambda) \geq g_k(\tau)$ for all $k \in N$ if and only if $\lambda \succeq^2 \tau$ (see, e.g., Hardy, Littlewood, & Polya, 1934, Kuosmanen, 2004). Thus, under equal state probabilities the SSD-efficient portfolios can be identified by maximizing the sums of the k smallest state-specific portfolio returns $g_k(\cdot)$ as stated by the following theorem.

Theorem 1. *Portfolio $\lambda \in \Lambda$ is SSD-efficient if and only if it is a Pareto optimal solution to the MOP problem*

$$\operatorname*{v-max}_{\lambda\in\Lambda} \left(g_k(\lambda)\right)_{k\in N}, \quad (9)$$

where $g_k(\lambda)$ is the sum of the k smallest state specific returns of portfolio λ.

Roman et al. (2006) have established a multi-objective linear programming formulation whose Pareto optimal solutions have a similar one-to-one correspondence to the SSD-efficient portfolios. Deriving this linear problem from the MOP problem (9) requires formulating each $g_k(\lambda)$ as an LP problem in which the objective function is maximized, since a minimization LP would yield a max-min problem if substituted into problem (9). Fortunately, LP duality can be harnessed to obtain

$$g_k(\lambda) = \min_{\nu_k\in\mathbb{R}^{1\times n}} \{\nu_k r^\lambda \mid \nu_k e = k,\ 0 \leq \nu_k \leq e^T\}$$
$$= \max_{\beta_k\in\mathbb{R},\ z_k\in\mathbb{R}^n} \left\{k\beta_k - e^T z_k \mid r\lambda - \beta_k e + z_k \geq 0,\ z_k \geq 0\right\}.$$

Substituting these maximization problems into (9) for each $k \in N$ gives multi-objective linear programming formulation by Roman et al. (2006). This is formalised by the following corollary, which introduces the auxiliary decision variables $\beta = (\beta_k) \in \mathbb{R}^n$ and $z = (z_{ik}) \in \mathbb{R}_+^{n\times n}$.

Corollary 1. *Portfolio $\lambda \in \Lambda$ is SSD-efficient if and only if (λ, z, β) is a Pareto optimal solution to the multi-objective linear programming problem*

$$\operatorname*{v-max}_{\lambda,z,\beta} \ \left(\ k\beta_k - e^T z_k\ \right)_{k\in N} \quad (10)$$

$$r\lambda - \beta_k e + z_k \geq 0\ \forall\ k \in N$$
$$z_k \geq 0\ \forall\ k \in N$$
$$\lambda \in \Lambda.$$

Roman et al. (2006) also note that there exists an intutive interpretation for MOP problem (9) based on Conditional Value-at-Risk (CVaR_π) risk measures, which we will present here for the sake of completeness. Specifically, this interpretation can be derived by using the generally known fact that the objective functions of a MOP problem can be subjected to strictly increasing transformations without affecting the set of Pareto optimal solutions. Choosing a transformation in which the kth objective of the MOP problem (9) is multiplied by k^{-1} results in the objective function values $g_1(\lambda)/1, \ldots, g_n(\lambda)/n$ for each efficient portfolio $\lambda \in \Lambda_E^2$. In particular, the kth objective function is equal to

$$\frac{g_k(\lambda)}{k} = \frac{\sum_{i=1}^k r_{\sigma(i)}^\lambda}{k} = \mathbb{E}[R^\lambda | R^\lambda \leq r_{\sigma(k)}^\lambda] \triangleq -\text{CVaR}_{\frac{k}{n}}[R^\lambda],$$

where the permutation σ arranges the components of r^λ in a non-decreasing order and, again, $R^\lambda = \sum_{j\in M}\lambda_j R_j$ denotes the random variable representing the returns of portfolio λ. Moreover, multiplying the nth objective function with n^{-1} yields the expected portfolio return $g_n(\lambda)/n = \sum_{i=1}^n r_i^\lambda/n = \mathbb{E}[R^\lambda]$. Hence, the choice among the SSD-efficient portfolios is contingent on how the decision maker values minimizing the CVaR_π-measures with different confidence levels $\pi \in \{\frac{1}{n}, \frac{2}{n}, \ldots, \frac{n-1}{n}\}$ on one hand, and maximizing the expected portfolio return on the other. Thus, portfolio $\lambda \in \Lambda$ is SSD-efficient if and only if it is a Pareto optimal solution to the multi-objective linear programming problem

$$\operatorname*{v-min}_{\lambda\in\Lambda} \left[\left(\text{CVaR}_{\frac{k}{n}}[R^\lambda]\right)_{k\in N\setminus\{n\}}, -\mathbb{E}[R^\lambda]\right]. \quad (11)$$

This result also helps establish a geometric intuition of the set of SSD-efficient portfolios and the corresponding efficient frontier in the objective function space. Since the set of SSD-efficient portfolios corresponds to the set of Pareto-optimal solutions to problem (11), the image of the SSD-efficient set is a connected, possibly non-convex set consisting of a finite union of polyhedra. These polyhedra lie in the n-dimensional space whose coordinates correspond to different CVaR-measures and expected return. As a result, also the set of SSD-efficient portfolios Λ_E^2 is a connected, generally non-convex set consisting of a union of polyhedra in the $(m-1)$-dimensional space $\{\lambda \in \mathbb{R}^m | \sum_{j=1}^m \lambda_j = 1\}$.

In the case where the set of feasible asset weights (1) is convex, a portfolio is SSD-efficient if and only if it maximizes the weighted sum of all CVaR-measures for some positive weights. This is implied by the equivalence between the SSD-efficient portfolios and the Pareto optimal solutions established by Theorem 1, and the fact that a solution to a continuous multi-objective linear programming is Pareto optimal if and only it maximizes the weighted sum of the objective functions for some positive weights. This result is formally established by the following corollary.

Corollary 2. *Assume Λ is convex. Then portfolio λ is SSD-efficient if and only if there exists $\alpha \in \mathbb{R}_+^n$, $\alpha_i > 0$, such that (λ, z, β) is an*

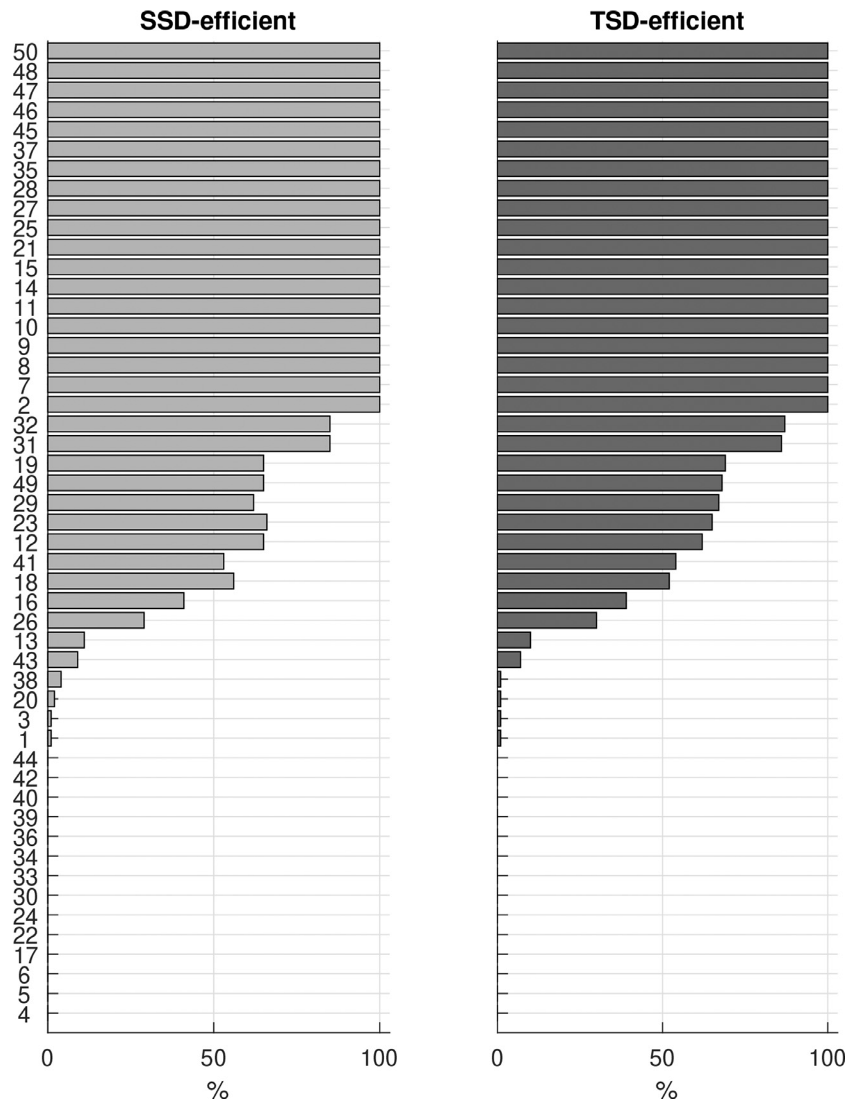

Fig. 1. Share of SSD- (light gray) and TSD-efficient portfolios (dark gray) that include each project.

optimal solution to the linear programming (LP) problem

$$\begin{aligned} \max_{\lambda,\ z,\ \beta} \sum_{k\in N}\ & \alpha_k\left[k\beta_k - e^T z_k\right] \\ & r\lambda - \beta_k e + z_k \geq 0 \quad \forall\ k \in N \\ & z_k \geq 0 \qquad\qquad\quad\ \forall\ k \in N \\ & \lambda \in \Lambda. \end{aligned} \tag{12}$$

3.3. SSD-efficient frontier for convex sets of asset weights under unequal state probabilities

An extension of linear programming problem (12) can be used to generate SSD-efficient portfolios also in situations where the state probabilities $p = (p_1, \dots, p_n)^T$ are unequal. This requires that the set of feasible asset weights Λ (1) is convex. We also assume each scenario probability p_i is a rational number. Note that this assumption is not necessary for the developed mathematical results, but it makes the proofs simpler. Moreover, this assumption does not limit the applicability of the developed theory as in practice any numerical estimates for the state probabilities are rational numbers.

In particular, let $Q = (Q_1, \dots, Q_n) \in \mathbb{R}^n$ be a parameter vector and consider the extended version of the linear programming problem (12)

$$\begin{aligned} \max_{\lambda,\ z,\ \beta} \sum_{k\in N}\ & \alpha_k\left[Q_k\beta_k - p^T z_k\right] \\ & r\lambda - \beta_k e + z_k \geq 0 \quad \forall\ k \in N \\ & z_k \geq 0 \qquad\qquad\quad\ \forall\ k \in N \\ & \lambda \in \Lambda, \end{aligned} \tag{13}$$

where $\alpha_1, \dots, \alpha_n$ are positive parameters as in Corollary 2. Note that in case of equal probabilities $p_k = \frac{1}{n}$, choosing $Q_k = \frac{k}{n}$, for all k, the problems (12) and (13) are equivalent up to multiplication with the positive constant n.

The optimal solutions (λ, β, z) for problem (13) with different parameter values $\alpha, Q \in \mathbb{R}$ are closely linked to the SSD-efficient portfolios. However, this requires that the parameter values $Q = (Q_1, \dots, Q_n)$ satisfy a special consistency condition that roughly speaking requires higher values for states in which the optimal portfolio has a higher return. The precise formulation of this condition is given by the following definition.

Definition 3. Parameter values $Q = (Q_1, \dots, Q_n)$ are consistent with portfolio λ if there exists a permutation $\sigma : N \to N$ such that

$r^{\lambda}_{\sigma(1)} \leq r^{\lambda}_{\sigma(2)} \leq \cdots \leq r^{\lambda}_{\sigma(n)}$ and

$$\sum_{i=1}^{k-1} p_{\sigma(i)} < Q_k \leq \sum_{i=1}^{k} p_{\sigma(i)} \ \forall\ k \in N. \tag{14}$$

If the components of vector $r^{\lambda} = r\lambda$ are distinct (i.e., $r^{\lambda}_i \neq r^{\lambda}_k$ for all $i \neq k$), then the permutation σ in Definition 3 is unique, and hence the bounds for parameters $Q = (Q_1, \ldots, Q_n)$ correspond to the values of the CDF evaluated at the state-specific returns, i.e., $\sum_{i=1}^{k} p_{\sigma(k)} = F^1_{\lambda}(r^{\lambda}_{\sigma(k)})$. With this consistency condition we can formulate the main result as follows.

Theorem 2. *Assume Λ is convex and state probabilities $p = (p_1, \ldots, p_n)^T$ are rational numbers.*

(i) *If portfolio $\lambda \in \Lambda$ is SSD-efficient, then there exists parameter values Q that are consistent with λ and parameter values $\alpha \in \mathbb{R}^n_+$ such that λ is an optimal solution to problem (13).*
(ii) *Let portfolio $\lambda \in \Lambda$ be an optimal solution to problem (13) for some parameter values Q that are consistent with λ and for some parameter values $\alpha \in \mathbb{R}^n_+$, then either λ is SSD-efficient or there exists an alternative optimal solution λ' which is SSD-efficient.*

A detailed proof of the theorem is presented in Appendix A, but we summarize it here. In particular, note that if (λ, z, β) is an optimal solution to problem (13), then $p^T z_k = F^2_{\lambda}(\beta_k)$ holds for all $k \in N$. Note also that $F^2_{\lambda}(t)$ is a piece-wise linear and increasing convex function for all return levels t exceeding the smallest state-specific return of portfolio λ and zero otherwise. Hence, the hypo-graph H^{λ} of F^2_{λ} is a convex polyhedral set, and for all $k \in N$, the optimal β_k in (13) is obtained as an optimal value of t in a 2-dimensional linear program $\max_{t,v}\{Q_k t - v \mid (t, v) \in H^{\lambda}\}$. Consequently, at the optimum we have for all $k \in N$, $\beta_k = r^{\lambda}_i$ for some $i \in N$. Furthermore, if Q satisfies the condition of Definition 3, then $\beta_k = r^{\lambda}_{\sigma(k)}$, for all $k \in N$, and at the optimum each term $Q_k\beta_k - p^T z_k$ of the objective function of problem (13) corresponds to the expected return of portfolio λ over the Q_k-quantile. The optimal portfolio thus maximizes a weighted sum of expected returns across different quantiles, which is effectively the same as minimizing the weighted sum of Conditional Value-at-Risk measures $\text{CVaR}_{Q_k}[R^{\lambda}] = -(1/Q_k)(Q_k\beta_k - p^T z_k)$, $k \in N$.

An important practical implication of Theorem 2 is that any efficient solution can be obtained as an optimal solution for (13) with suitably chosen parameter vectors α and Q. In Appendix B, we introduce a simple procedure searching for such vectors Q. We tested this procedure for 120 randomly generated problem instances with the number of scenarios n ranging from 10 to 1000 and the number of stocks m from 5 to 100 (for details see Appendix B). In these tests the procedure never failed and the number of iterations needed for finding parameter values Q that meet the condition of Definition 3 varied from 1 to 9 with an average of 3.9 iterations.

3.4. Identification of the TSD-efficient Frontier

In the case of TSD ($d = 3$), the MOP problem (4) is somewhat more challenging. In particular, TSD cannot be established by evaluating the twice integrated CDF (F^3_{λ}) for all return levels between the lowest $\underline{r}$ and highest $\bar{r}$ state-specific returns. This is since the twice integrated CDFs of two portfolios may intersect for some return level above $\bar{r}$, even though one is strictly higher for all return levels between $\underline{r}$ and $\bar{r}$. However, it is sufficient to consider the values of $F^3_{\lambda}(t)$ only for return levels t on the interval $[\underline{r}, \bar{r}]$ if the expected portfolio return is included as an additional objective function to be maximized (see Appendix A for a detailed proof).

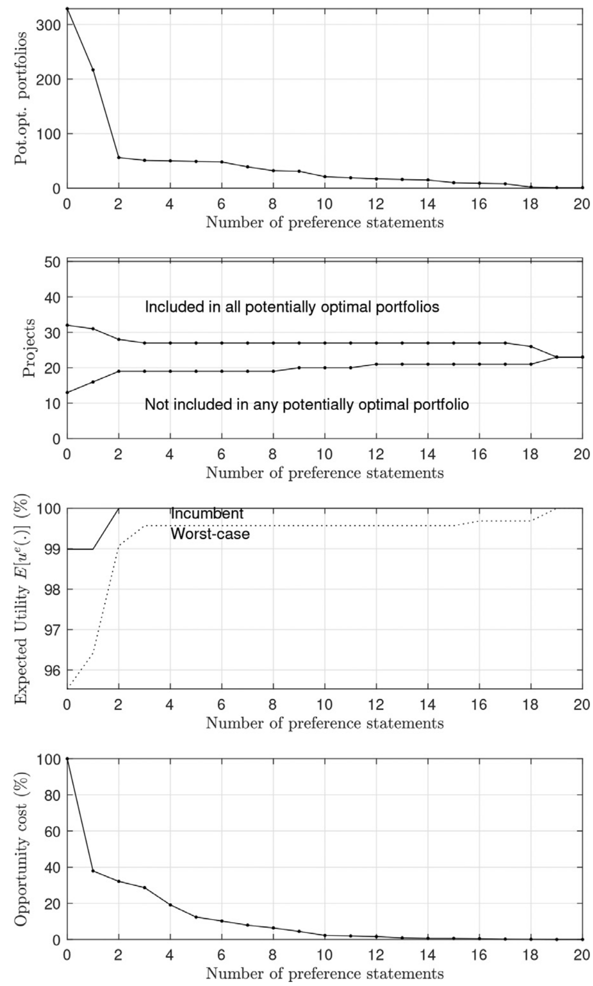

Fig. 2. Impact of preference information.

Lemma 2. *Portfolio $\lambda \in \Lambda$ is TSD-efficient if and only if it is a Pareto optimal solution to the MOP problem*

$$\operatorname*{v-min}_{\lambda\in\Lambda} \left[\left(F^3_{\lambda}(t)\right)_{t\in[\underline{r},\bar{r}]}, -\mathbb{E}[R^{\lambda}] \right].$$

The second challenge is that the twice integrated CDF is piecewise quadratic rather than piecewise linear. Fortunately, for a fixed t, $F^3_{\lambda}(t)$ can be evaluated by solving a convex quadratic minimization problem (see, e.g., Post & Kopa, 2017). Combining this observation with Lemma 2 yields the following theorem (see Appendix A for a detailed proof).

Theorem 3. *Portfolio $\lambda \in \Lambda$ is TSD-efficient if and only if (λ, Φ) is a Pareto optimal solution to the MOP problem*

$$\operatorname*{v-min}_{\lambda,\Phi} \left[\left(\sum_{i\in N} p_i \Phi^2_i(t) \right)_{t\in[\underline{r},\bar{r}]}, -\sum_{i\in N} p_i r^{\lambda}_i \right] \tag{15}$$
$$\Phi_i(t) \geq t - r^{\lambda}_i \ \forall\ i \in N, t \in [\underline{r}, \bar{r}]$$
$$\Phi_i(t) \geq 0 \ \forall\ t \in [\underline{r}, \bar{r}]$$
$$\lambda \in \Lambda.$$

Practical implementation of problem (15) requires evaluating the functions Φ_i only at a finite set of points $\{t_1, \ldots, t_l\} \subset [\underline{r}, \bar{r}]$. By denoting $L = \{1, \ldots, l\}$ and $\phi_{k,i} = \Phi_i(t_k)$ for all $k \in L$, this discretization yields the multi-objective quadratic programming (MOQP) problem

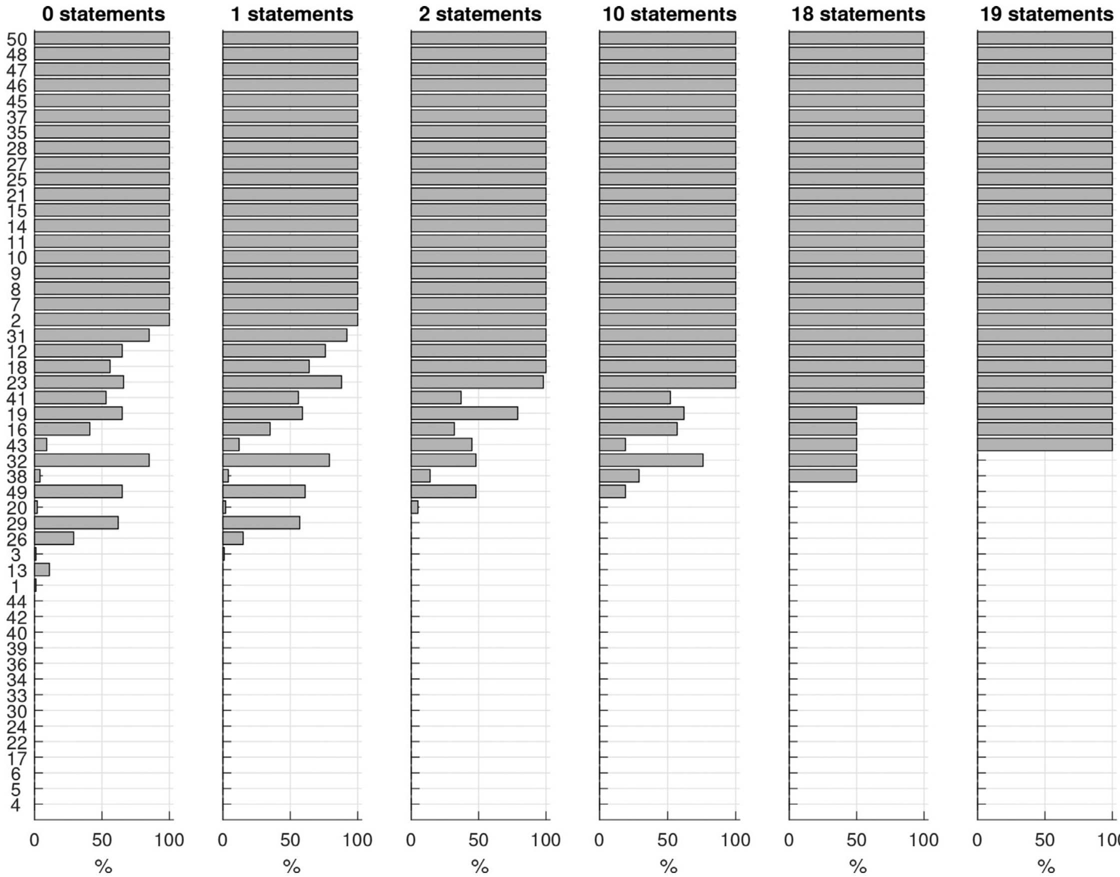

Fig. 3. Share of potentially optimal portfolios that include each of the $m = 50$ projects.

$$\begin{aligned} \text{v-}\min_{\lambda,\phi} \quad & \left[\Big(\sum_{i\in N} p_i \phi_{i,k}^2 \Big)_{k\in L}, -\sum_{i\in N} p_i r_i^{\lambda} \right] \qquad (16) \\ & \phi_{i,k} \geq t_k - r_i^{\lambda}, \;\; \forall\; i \in N, \;\; k \in L \\ & \phi_{i,k} \geq 0, \forall\; i \in N, \;\; k \in L \\ & \lambda \in \Lambda, \end{aligned}$$

which has $l+1$ objective functions.

Using standard multi-objective programming methods to generate Pareto optimal solutions to problem (16) is relatively straightforward because the objective functions are convex and quadratic and the constraints are linear. Thus, deploying, for instance, the reference point or the weighted Tchebychef-norm method to problem (16) requires only solving a series of single-objective quadratic programming problems. Clearly, if the set of feasible asset weights Λ in (1) is non-convex, then these problems will include binary decision variables. The weighted sum method can be used if Λ is convex because in this case problem (16) is convex, i.e., it has convex objective functions and a convex feasible region. The weighted sum method is particularly attractive choice as it does not require the introduction of additional decision variables or constraints. It should be noted that even for convex problems there can exist border points of the Pareto optimal set that the method cannot identify (see, e.g., Ehrgott, 2005). Nevertheless, all interior points of a convex Pareto optimal set can be found with the weighted sum method, which seems sufficient for practical applications.

4. Efficient frontiers in scenario-based project portfolio selection

In this section we demonstrate the use of the developed methods in project portfolio selection. We utilize the real-world application context reported in Liesiö & Salo (2012) in which scenario-based portfolio models were used to identify combinations of projects expected to prove most valuable in the future. However, the real data set contains only 24 projects, which makes it possible to identify SSD- and TSD-efficient portfolio through complete enumeration of all project combinations. Hence, we randomly generate a larger data set consisting $m = 50$ projects (see Appendix C). Otherwise the problem attributes are kept intact. In particular, there are $n = 8$ scenarios, binary decision variables λ_j indicating whether or not each project is included in the portfolio (i.e., $\underline{\Lambda} = \{0, 1\}^{50}$), and a single budget constraint that limits the total portfolio cost to be no greater than one third of the sum of all project candidates' costs.

The SSD- and TSD-efficient frontiers were generated by solving the multi-objective programming problems (7) and (16), respectively, with the Tchebychef-norm method (Bowman, 1976). For problem (7) this approach results in a series of single objective MILP problems, the average solution time of which was some 0.14 s when using Gurobi running on a standard laptop. For problem (16) this approach requires solving single objective quadratic mixed-integer programming problems and the average solution time with Gurobi was some 0.22 s.

Fig. 1 shows the share of the 341 SSD-efficient and the 290 TSD-efficient portfolios that include each of the 50 projects. Somewhat surprisingly there are 19 projects included in all of the SSD-efficient portfolios. These projects should be selected by any risk-averse or neutral decision maker. Moreover, there are 14 projects that are not included in any of the SSD-efficient portfolios and should not be selected by any risk-averse or -neutral decision maker.

To demonstrate how preference information affects which efficient portfolios remain potentially optimal (cf. Eq. (3)), we utilize

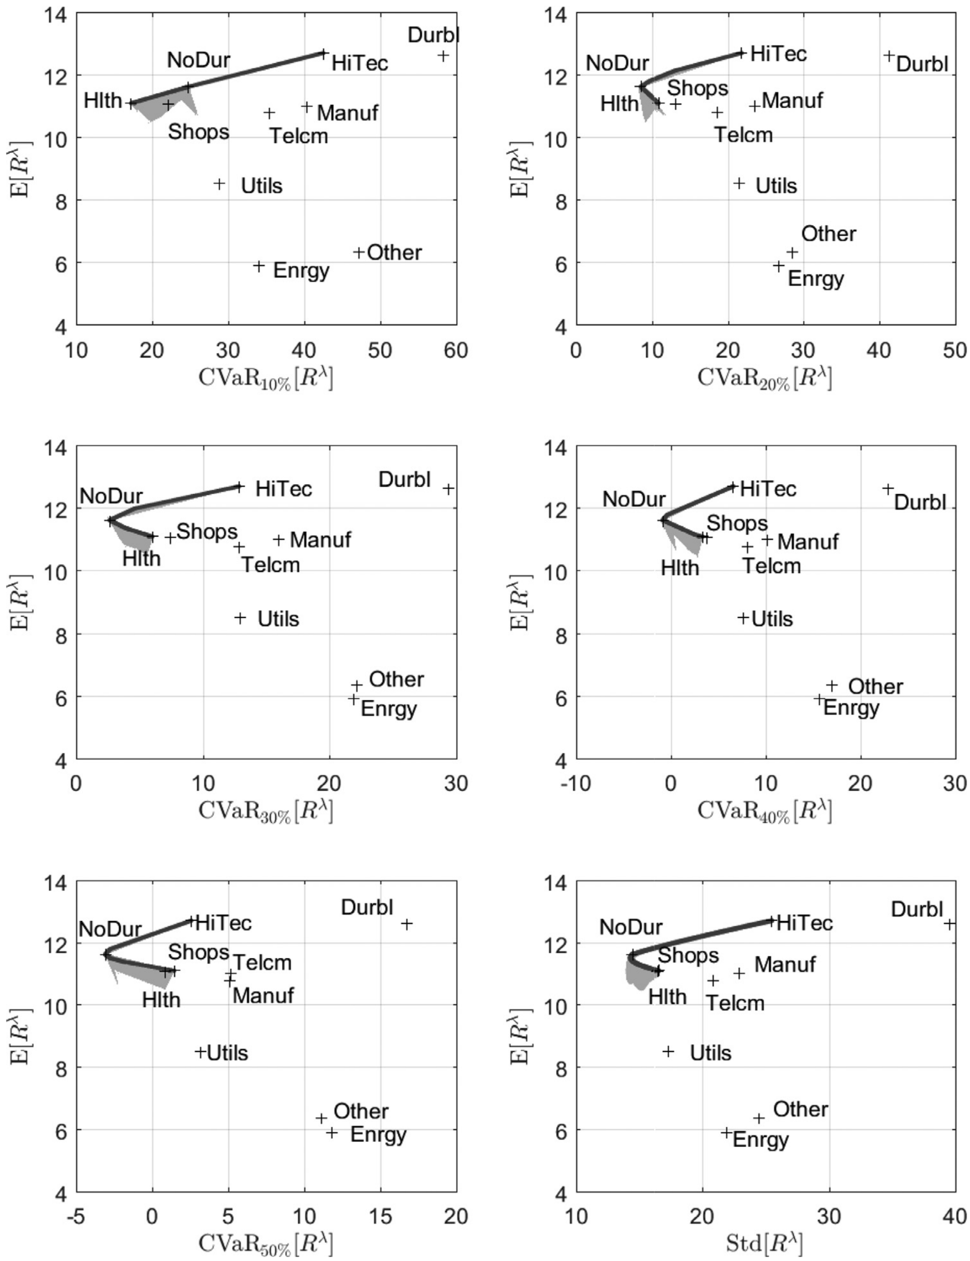

Fig. 4. SSD-efficient portfolios (light gray) and TSD-efficient portfolios (dark gray) projected onto six mean-risk planes. The crosses correspond to the underlying $m = 10$ industry portfolios serving (cf. base assets).

preference statements given by a simulated decision maker. This decision makers is asked repeatedly to state her preferences between return distributions of two SSD-efficient portfolios named the incumbent portfolio and the challenger portfolio. Decision maker's preference for the incumbent portfolio $\hat{\lambda}^k$ over the challenger $\tilde{\lambda}^k$ on the kth elicitation question is modelled by setting $X^k = R^{\hat{\lambda}^k}$ and $Y^k = R^{\tilde{\lambda}^k}$ in Eq. (2), while preference for the challenger yields $X^k = R^{\tilde{\lambda}^k}$ and $Y^k = R^{\hat{\lambda}^k}$. For the first elicitation question the incumbent portfolio is the one with the median expected value across the SSD-efficient frontier. If in the kth elicitation question the challenger is preferred to the incumbent then the challenger becomes new incumbent ($\hat{\lambda}^{k+1} = \tilde{\lambda}^k$) and otherwise the incumbent remains unchanged ($\hat{\lambda}^{k+1} = \hat{\lambda}^k$). The challenger is chosen as the portfolio that maximizes the expected utility difference to the current incumbent portfolio across all compatible utility function, i.e.,

$$\tilde{\lambda}^k \in \underset{\lambda \in \Lambda_E^2, u \in \tilde{U}^k}{\arg\max} \left(\mathbb{E}[u(R^{\lambda})] - \mathbb{E}[u(R^{\hat{\lambda}^k})] \right). \tag{17}$$

Fig. 2 shows the how preference information affects the set of potentially optimal portfolios when the decision maker's true risk-preferences are captured by the exponential utility function $u^e(t) = 1 - e^{-0.018t}$. Checking for potential optimality of a portfolio (cf. Eq. (3)) took approximately 0.06 s on Gurobi.

The first panel shows the number of potentially optimal portfolios as a function of the number of preference statements given by the decision maker, while the second panel shows the share of the $m = 50$ projects that are included in all, some or none of the potentially optimal portfolios. Already with two preference statements we observe a substantial reduction in the number of

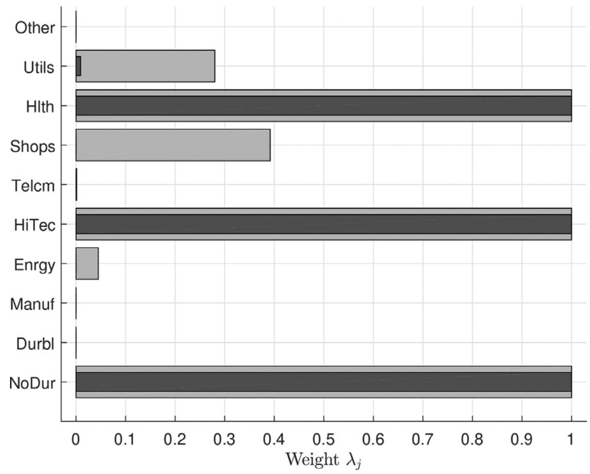

Fig. 5. Share of capital invested in each industry across SSD-efficient portfolios (light gray) and TSD-efficient portfolios (dark gray).

potentially optimal portfolios, and these portfolios differ with regard to decision concerning only 8 projects. However, it takes 19 preference statements to reach a situation with a single potentially optimal portfolio. Together these results suggest that there are several portfolios with very similar outcome distributions, although their project composition is slightly different. This is exemplified by Fig. 3 which visualizes the project level decision recommendations obtained at specific points of the preference elicitation process.

The third panel of Fig. 2 visualises the expected utilities of the potentially optimal portfolios evaluated using the exponential utility function u^e. The dashed line shows the minimum expected utility across all potentially optimal portfolios and the solid line corresponds to the expected utility of the incumbent portfolio. These are shown relative to the potentially optimal portfolio with the highest expected utility ($=100\%$). After only two preference statements selecting any one of the potentially optimal portfolios would yield over 99% of the expected utility of the true optimal portfolio. Moreover, at this point the incumbent portfolio is in fact an optimal portfolio, i.e., it maximizes expected utility $\mathbb{E}[u^e(\cdot)]$.

In practice the information conveyed in the third panel of Fig. 2 is unknown, since the true utility function is not known. However, decision support can be provided on the basis of computing the decision maker's opportunity cost. For any portfolio that the decision maker might be considering to select, we can compute an upper bound on the 'lost' potential expected utility associated with that particular portfolio choice. In the context of the interactive preference elicitation discussed above, this opportunity cost is obtained as the maximum expected utility difference across all feasible utility functions.

The bottom panel illustrates the opportunity cost of selecting the incumbent portfolio. Specifically, for each feasible utility function $u \in \tilde{U}^k$ we compare expected utility of the optimal portfolio (i.e., $\max_{\lambda\in\Lambda_E^2} \mathbb{E}[u(R^\lambda)]$) to that of the incumbent ($\mathbb{E}[u(R^{\hat{\lambda}^k})]$). Opportunity cost is then obtained as the maximum expected utility difference across all feasible utility functions, i.e.,

$$\max_{\lambda\in\Lambda_E^2, u\in\tilde{U}^k} \left(\mathbb{E}[u(R^\lambda)] - \mathbb{E}[u(R^{\hat{\lambda}^k})]\right). \tag{18}$$

Note that the portfolio that optimizes (18) is always potentially optimal. The opportunity costs in Fig. 2 are shown relative to the opportunity cost of selecting the incumbent portfolio without any preference information. Already with 6 preference statements the maximum loss in expected utility from selecting the incumbent portfolio is reduced by 90%.

Overall, these results show that incorporating even a minimal amount of preference information can have a dramatic effect on providing conclusive recommendations and, crucially, on providing associated assurances to the decision maker (via the upper bound on the opportunity cost). As we will demonstrate in the following section, this finding is apparent even when the decision maker's preferences are not simulated via a specific utility function.

5. Efficient frontiers of industry portfolios

In this section we apply the developed models to open data on returns of financial assets. In particular, we present two applications that analyze the risk-return profile and the composition of portfolios that are efficiently diversified across different industries in the sense of SSD and TSD. The assets in the model correspond to the industry portfolios from the Kenneth R. French data library which have been constructed by classifying each NYSE, AMEX, and NASDAQ stock based on its SIC code[2]The first and second application utilize data sets in which the stocks have been divided in to $m = 10$ and $m = 49$ industry portfolios, respectively. In both applications we use the set of asset weights $\Lambda = \{\lambda \in [0,1]^m | \sum_j \lambda_j = 1\}$, which prohibits short selling of industry portfolios. It is perhaps important to highlight that the developed models can handle more complex sets of feasible asset weights, but we seek to demonstrate properties of the SSD- and TSD-efficient frontiers under the standard assumptions.

5.1. SSD- and TSD-efficient diversification among 10 industries

In the first model we use the annual returns of the $m = 10$ industry portfolios from 2007–2016 resulting in a state-space with $n = 10$ states. The MOLP problem (10) thus contains 10 objective functions, 101 constraints and 120 decision variables. The set of all SSD-efficient portfolios was obtained by solving this problem with the Benson algorithm (Löhne & Weißing, 2014) that produces a finite number of polyhedral sets the union of which is equal to the set Λ_E^2. The computation time was approximately 67 s on a standard laptop. The 500 TSD-efficient portfolios were solved by using the weighted sum approach on the MOQP problem (16). We used 200 evenly spaced discretization points $t \in [\underline{r}, \bar{r}]$ between minimum and maximum state-specific assets returns, and thus problem (16) contained 201 objective functions, 2010 decision variables, and 2001 constraints. The resulting quadratic programming (QP) problems took on average 0.07 s to solve on Gurobi.

Fig. 4 illustrates the frontiers of efficient diversification strategies among the 10 industries. In particular, it shows the projections of the SSD- and TSD-efficient frontiers onto five 2-dimensional planes, in which the vertical axis corresponds to the expected return and the horizontal axis corresponds to a specific CVaR-measure or the standard deviation of returns. Recall that the $\mathrm{CVaR}_{\frac{k}{10}}$-measure is equal to the kth objective function value of the MOLP problem (10), and the expected return is equal to the nth objective function up to multiplication with a constant (cf. problem (11)). Hence, the first five projections also illustrate the image of the Pareto optimal set of MOLP problem (10), which consists of a union of polyhedral sets. The figure also highlights that SSD-efficient portfolios satisfying some prespecified target-level for a particular risk measure (e.g., $\mathrm{CVaR}_{10\%} \leq 20\%$) can have significantly different expected returns (e.g., 10.6%-11.5%). Similarly, there exists a TSD-efficient portfolio yielding expected returns of 11% with a 17% standard deviation, even though it is possible to obtain 12% expected return with the same standard deviation.

[2] For details see https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

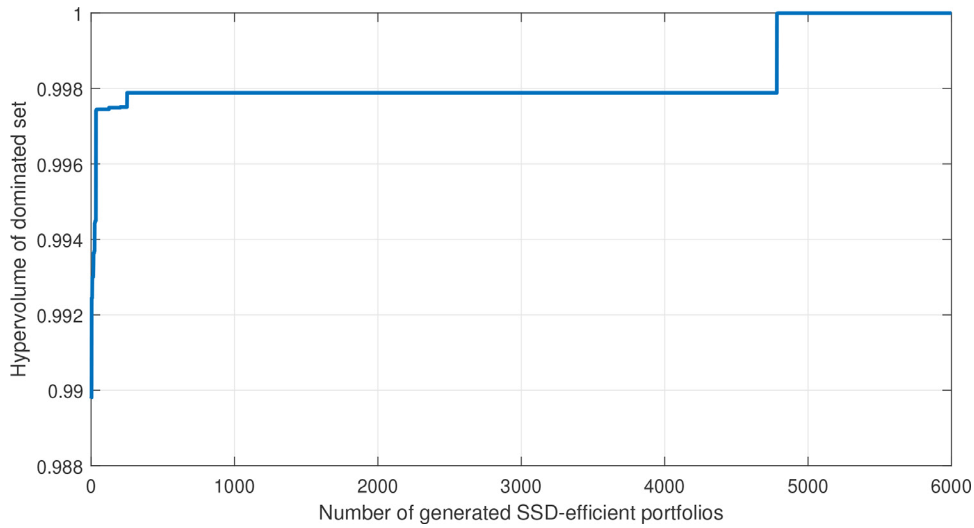

Fig. 6. Hypervolume of the SSD-dominated set as a function of the number of generated SSD-efficient portfolios.

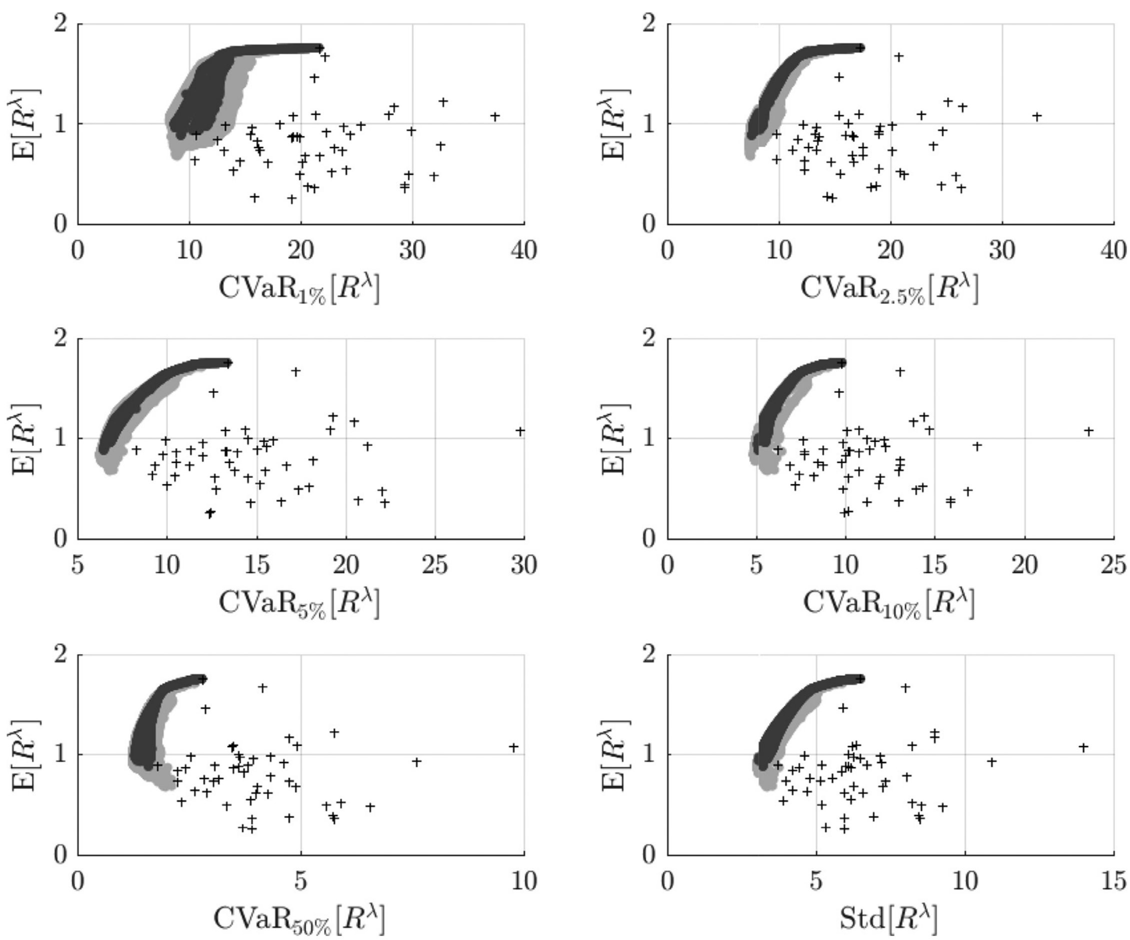

Fig. 7. SSD- (light gray) and TSD-efficient portfolios (dark gray) projected onto six different mean-risk planes. Black crosses correspond to the underlying $m = 49$ industry portfolios (cf. base assets).

Fig. 5 shows the minimum and maximum share of capital allocated to each industry across the SSD- and TSD-efficient frontiers. In particular, for industry $j \in \{1, \ldots, 10\}$, the bars show the interval $[\min_{\lambda \in \Lambda_E^d} \lambda_j, \max_{\lambda \in \Lambda_E^d} \lambda_j]$ for $d \in \{2, 3\}$. In this data set there are no industries to which investments are made in each SSD-efficient or TSD-efficient portfolio. However, there are three industries, namely 'Health', 'High Tech' and 'Consumer Non-Durables', such that investing all capital to any one of these would constitute a TSD-efficient (and hence also a SSD-efficient) diversification. In contrast, the three industries 'Other', 'Manufacturing' and 'Consumer Durables' are not included in any of the SSD- or TSD-efficient portfolios: Investing into any one of these will lead to an inefficient portfolio, meaning that there would exist another portfolio which is preferred by all expected utility maximizing risk-averse or -neutral decision makers.

Between these two extremes there are industries that are included in some but not all SSD-efficient portfolios. For instance, investing more than 39% of initial capital into 'Shops' would lead to a portfolio that is dominated in the sense of SSD, no matter how the remaining 61% of capital is diversified. Moreover, any skewness preferring decision maker with a utility function $u \in U^3$ should invest at most 2% of capital to 'Utils' and diversify the rest of the capital between 'Health', 'High Tech' and 'Consumer Non-Durables'.

5.2. SSD- and TSD-efficient diversification among 49 industries

The second model uses a more detailed stock classification consisting of $m = 49$ industry portfolios, which serve as the base assets. The state-space is constructed using the monthly returns of these portfolios from 2000 to 2016 resulting in $n = 204$ states. The MOLP problem (10) thus contains 204 objective functions, 41,617 constraints and 41,869 decision variables. This problem is far too large to be solved with exact MOLP algorithms and hence we deployed the weighted sum method (Corollary 2). We generated some 40,000 SSD-efficient portfolios, and it took on average 0.26 s to solve a single weighted sum LP problem. For generating the TSD-efficient portfolios we used a 250 point discretization of the return levels resulting in a MOQP problem with 251 objective functions, 51,049 decision variables and 51,001 constraints. Some 1500 TSD-efficient portfolios were generated using the weighted sum method, and solving a single QP problem took approximately 5 s on Gurobi.

Although each of the generated portfolios is efficient, we conducted an additional analysis to examine how representative these portfolios are of the efficient frontier. Specifically, we used Monte Carlo simulation to estimate the hypervolume, which we define as the relative volume of the set of those portfolios that are dominated by at least one of the generated portfolios, i.e., $\mathrm{Vol}(\{\lambda \in \Lambda \mid \exists \lambda' \in \Lambda_E^d \text{ s.t. } \lambda' \succeq^d \lambda\})/\mathrm{Vol}(\Lambda)$, where Λ_E^d is the set of generated SSD- ($d = 2$) or TSD-efficient ($d = 3$) portfolios. Fig. 6 shows the hypervolume of the SSD-dominated set as a function of the number of SSD-efficient portfolios generated. For instance, with 500 SSD-efficient portfolios the hypervolume of the dominated set is almost 99.8%. In case of TSD the convergence is even faster since the TSD-efficient set is a subset of the SSD-efficient set and as a result the hypervolume exceeds 99.99% already with four generated TSD-efficient portfolios.

Fig. 7 illustrates the risk-return profiles of the SD-efficient portfolios and the 49 industry portfolios. These results exemplify that SD-efficient portfolios that are equally risky according to a single risk measure can have large differences in their expected returns. For instance, the expected returns of SSD-efficient portfolios with equal $\mathrm{CVaR}_{10\%}$ values can vary between 0.7% and 1.5%. This observation seems to imply that expected returns can be doubled without increasing risk, but this is an artefact resulting from the use of a single risk measure: Since all of these portfolios are SSD-efficient, increasing expected returns from 0.7% to 1.5% increases CVaR_π for at least some confidence level π (see problem (11)). Indeed, the SSD-efficient frontier needs to be understood as a frontier in the multi-dimensional objective space spanned by all CVaR measures and the expected return. Projections of this frontier onto a 2-dimensional mean-risk plane result in an SSD-efficient 'band'.

Fig. 8 shows the share of capital allocated to each industry across the SSD- and TSD-efficient portfolios. This data set includes one industry ('Smoke' consisting of tobacco companies) that alone constitutes an SSD- and a TSD-efficient portfolio. This industry offers the highest expected return, and thus a pure risk-neutral decision maker should invest 100% of her capital into it. Perhaps more

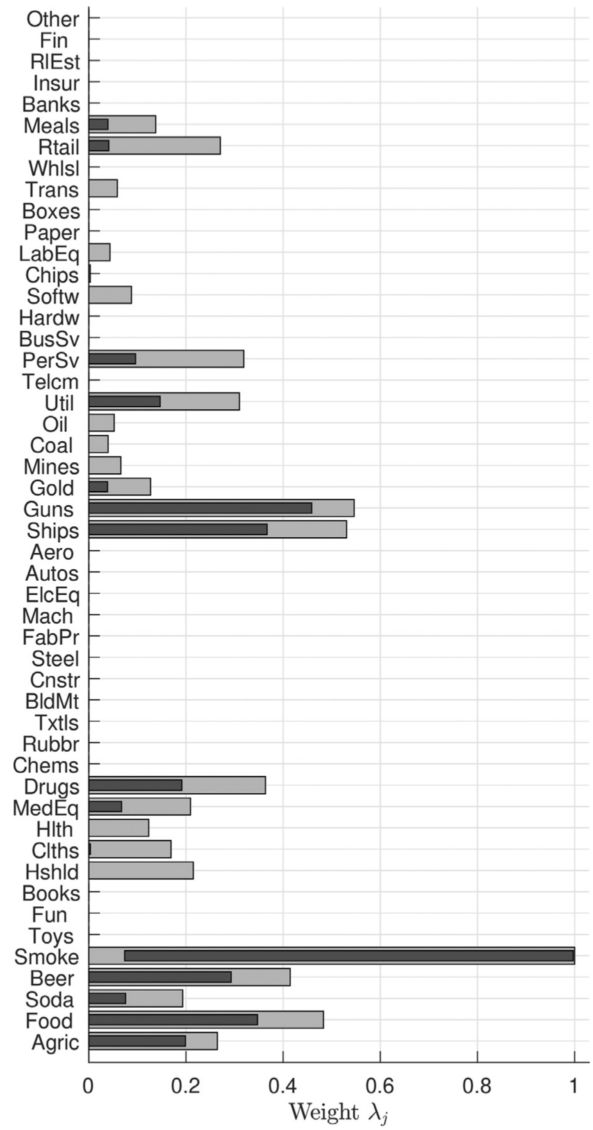

Fig. 8. Share of capital invested in each industry across SSD-efficient portfolios (light gray) and TSD-efficient portfolios (dark gray).

surprising is the observation that over a half of the industries are not included in any of the SSD-efficient portfolios (e.g., 'Toys' and 'Fun' which include stocks of recreation and entertainment companies). This means that maximizing expected utility for *any* concave or linear utility function would result in a portfolio in which no capital is allocated to these industries.

In this data set about two thirds of the industries are not included in any TSD-efficient portfolio. Hence, any risk-averse decision maker with preferences for higher skewness should diversify only among the remaining 15 industries. Moreover, such a decision maker should always allocate more than 5% of her capital to 'Smoke' as otherwise the resulting portfolio will be dominated in the sense of TSD no matter how the remaining capital is allocate.

5.3. Additional preference information

To demonstrate how preference information affects which efficient portfolios remain potentially optimal (cf. Eq. (3)) we will again consider the incremental preference elicitation schema discussed in Section 4. Here, however, we will conduct a more complete investigation, which is not contingent on simulating the re-

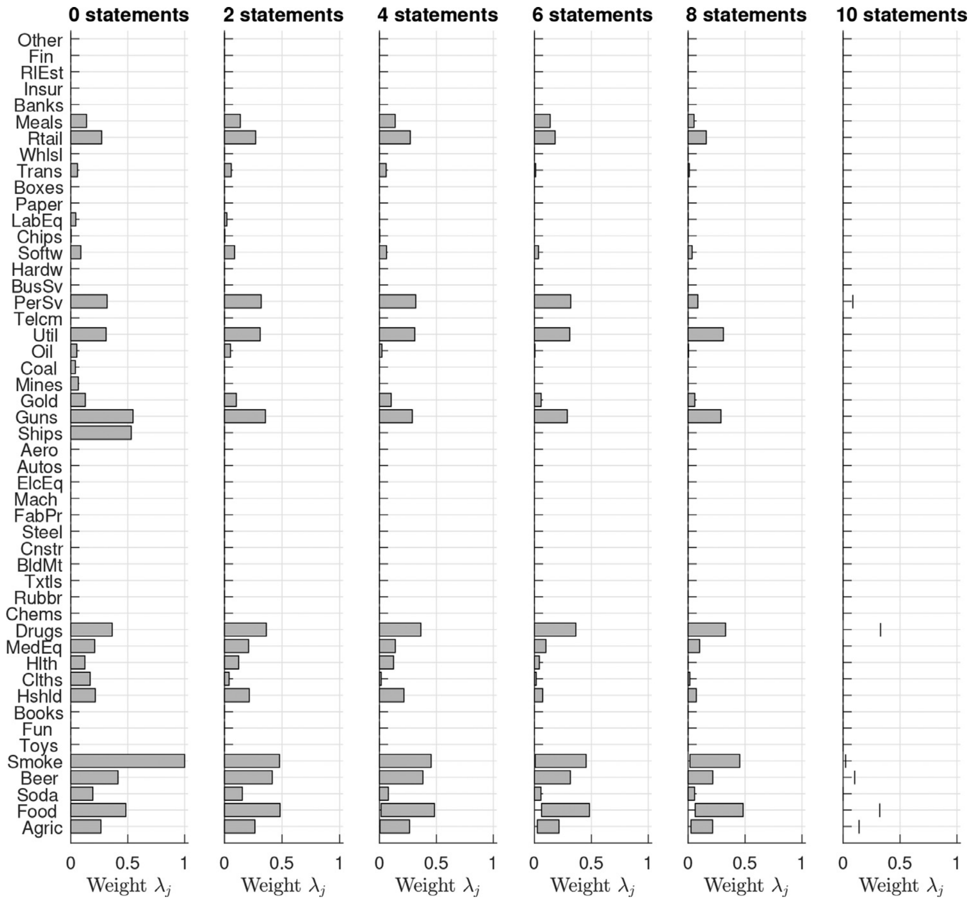

Fig. 9. Share of capital invested in each industry across the potentially optimal portfolios for decision maker who prefers the incumbent over the challenger in the 2nd, 3rd, 5th and 7th elicitation question.

sponses of the DM via the use of a specific utility function. We do this because we want to specifically investigate the effect of utilizing preference information *irrespective* of the specific preferences one DM may provide. In other words we want to consider the effect of utilizing preference information in reducing the set of potentially optimal portfolios across different DMs (i.e. with different preferences).

Our investigation involves generating all potential response sequences that could be observed in the application of the incremental preference elicitation schema of Section 4 after 10 iterations. This generation is a dynamic problem, since the generation of a challenger distribution/portfolio at any iteration is contingent on the previous responses of a DM. We thus construct the complete response tree via exhaustive enumeration. More specifically, when problem (17) is solved for the kth elicitation question, we consider two scenarios for $\tilde{U}^{k+1}$. In one scenario we construct $\tilde{U}^{k+1}$ by modelling a preference of the incumbent $\hat{\lambda}^k$ over the challenger $\tilde{\lambda}^k$, in which case we set $X^k = R^{\hat{\lambda}^k}$ and $Y^k = R^{\tilde{\lambda}^k}$ in Eq. (2) and the incumbent is unchanged ($\hat{\lambda}^{k+1} = \hat{\lambda}^k$). In the other scenario we construct $\tilde{U}^{k+1}$ by modelling a preference of the challenger $\tilde{\lambda}^k$ over the incumbent $\hat{\lambda}^k$, in which case we set $X^k = R^{\tilde{\lambda}^k}$ and $Y^k = R^{\hat{\lambda}^k}$ and the incumbent is updated ($\hat{\lambda}^{k+1} = \tilde{\lambda}^k$). These two scenarios for defining $\tilde{U}^{k+1}$ define two branches of the potential response tree, and each is explored further (leading to further branching). The complete response tree is constructed in this fashion to a depth of $k = 10$ questions. It should be noted that, across 10 questions, the total number of possible answer combinations is $2^{10} = 1024$. Moreover, after k elicitation questions there are 2^k alternative sets of feasible utility functions (Eq. (2)) and corresponding sets of potentially optimal portfolios (Eq. (3)). Due to this exponential number of alternative preferences, this analysis is carried out using 500 portfolios randomly chosen from the total set of 40,000 SSD-efficient portfolios generated.

As an example, Fig. 9 demonstrates the decision recommendations obtained throughout the 10 step preference elicitation process by a decision maker who prefers the challenger to the incumbent in the 2nd, 3rd, 5th and 7th elicitation question. With these preference, there are 31 assets not included in any potentially optimal portfolio after 6 elicitation questions, and the widths of the weight ranges of the remaining assets vary between 0.0007 and 0.4475.

Fig. 10 illustrates the weight ranges' widths across all possible answer combinations. In particular, for each asset we first compute the width of the weight range (i.e., difference between the minimum and maximum weight across all potentially optimal portfolios) for every alternative set of potentially optimal portfolios. Hence, after k elicitation questions, there is a distribution of widths across the $m = 49$ assets and the 2^k alternative sets of potentially optimal portfolios. Fig. 10 illustrates the spread in this distribution by showing the 80th, 90th, 95th and 100th percentiles of this distribution as well as its median. Overall, assets' weight ranges become narrower as more preference information is elicited and for instance after nine preference statements 80% of the assets have a point-estimate weight.

Fig. 11 visualizes the impact additional preference information has on the reduction in the opportunity cost. Specifically, for each alternative set of potentially optimal portfolios the opportunity cost is obtained by maximizing the expected utility differ-

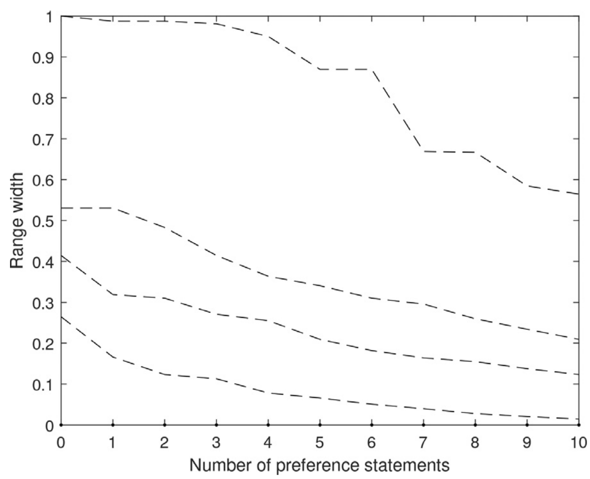

Fig. 10. Width of the assets' weight ranges across potentially optimal portfolio as a function of the number preference statements. The dashed lines show the 100th, 95th, 90th and 80th percentiles. The median width is zero for any number of preference statements.

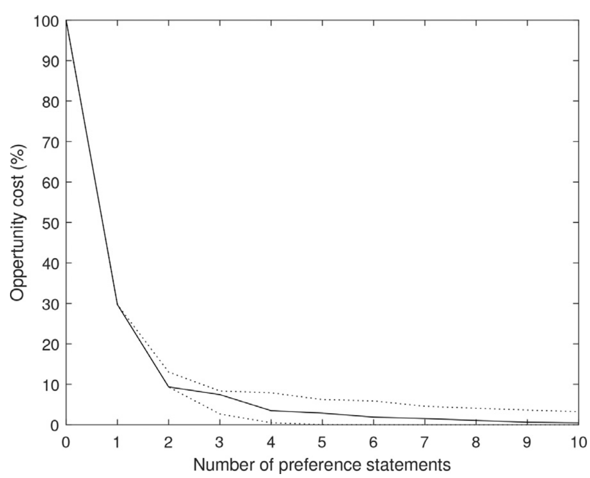

Fig. 11. Maximum expected utility difference to incumbent portfolio across all potentially optimal portfolios and utility functions compatible with the given preference statements. Solid line corresponds to the median expected utility difference across different preferences, and the dashed lines show the minimum and maximum.

ence to the incumbent portfolio across all portfolios in this this set and across all compatible utility functions. Hence, with k elicitation questions, there is a distribution of opportunity costs, one for each of the 2^k alternative sets of potentially optimal portfolios. Fig. 11 shows the minimum, maximum and median of these costs relative to the opportunity cost in case no preference information is given (= 100%). For instance, after three preference elicitation questions the opportunity cost is reduced by over 90% regardless of what the answers to these question are.

As in the previous application, these results show that even small amounts of preference information can substantially improve portfolio recommendations and providing the associated assurances. Crucially, in this second application this finding is independent of the specific form of the underlying utility function (i.e. independent of the underlying decision maker preferences).

6. Implications and conclusions

The models developed in this paper can be readily deployed to support portfolio selection in practical applications. In particular, the equivalence between SSD/TSD-efficient portfolios and Pareto optimal solutions makes it possible to use exact and approximate MOP algorithms to generate the efficient frontier. This frontier can then be used to identify those individual assets that do not belong to any efficient portfolio, and thus investing in them would lead to a sub-optimal diversification regardless of the level of risk aversion. Indeed, in the data sets analyzed in this paper about half of the base assets were not included in any of the SSD-efficient portfolios, which means that no concave (or linear) utility function rationalizes investments into these assets. Moreover, the share of the assets not included in any of the TSD-efficient portfolios was even higher.

Once the SD-efficient frontier has been identified, additional information on the decision maker's risk-preference can be readily utilized to reduce the set of admissible portfolio alternatives. Our results show that this can have a dramatic effect on obtaining conclusive recommendations for the individual assets and on reducing the opportunity cost. For instance, in our application already a few preference statements comparing pairs SSD-efficient portfolios could be enough to conclude that capital should be allocated among only six of the 49 assets.

It is important to highlight some aspects of the developed approaches. First, these approaches identify those portfolios that are SSD- or TSD-efficient for the specified state-space. In financial applications, for instance, if the state-space does not capture the true stochastic process through which future uncertain returns are generated, then there are no guarantees that the identified portfolios are efficient ex-post after observing the empirical return distribution. Second, the computational burden of the developed approaches depends on both the underlying portfolio selection problem as well as the MOP method used to solve the Pareto optimal solutions. For instance, if the portfolio selection problem involves binary decision variables, then the utilization of weighted sum or Tchebychef-norm methods also involves solving optimization problems with binary variables. Moreover, in case of TSD, these problems would be non-linear. Third, in many cases the entire (possibly infinite) efficient frontier cannot be identified, but instead one must rely on MOP methods that generate a finite subset of the efficient portfolios. From a computational perspective these methods are appealing since they often rely on solving several single objective optimization problems that are not linked to each other, which makes it possible to utilise parallel computation with multiples cores.

We believe that this paper opens up several interesting avenues for future research. In terms of methodological development, research efforts are needed to examine if similar links can be established between multi-objective optimization problems and other stochastic dominance criteria such as the first-order stochastic dominance (FSD; Quirk & Saposnik, 1962) or the almost stochastic dominance (Tsetlin, Winkler, Huang, & Tzeng, 2015). For instance, FSD relaxes the assumption of risk aversion, and hence any decision maker preferring more returns to less should make choices that are consistent with FSD. Methods for identifying the FSD-efficient frontier could thus be valuable to behavioral research that does not assume decision makers are fully rational expected utility maximizers (cf. Starmer, 2000). Furthermore, some project portfolio selection models use scenario trees to capture dynamic decision settings (see, e.g., Gustafsson & Salo, 2005). These models often result in mixed-integer linear programming formulations and thus the models developed in this paper could be used to identify the frontiers of SSD- and TSD-efficient project portfolios in such multi-period settings.

In terms of empirical finance applications, the next obvious step is analyzing the sets of SSD- and TSD-efficient portfolios implied by market data to build an understanding on the general structure of SD-efficient frontiers. Such studies should systematically study the effects of factors such as the covariance structure or the numbers of assets and states, and also include analysis of the differences between in-sample and out-of-sample efficient frontiers (cf. Hodder, Jackwerth, & Kolokolova, 2014).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.ejor.2022.04.043.

References

Argyris, N., Morton, A., & Figueira, J. R. (2015). CUT: A multicriteria approach for concavifiable preferences. *Operations Research, 62*(3), 633–642.

Armbruster, B., & Delage, E. (2015). Decision making under uncertainty when preference information is incomplete. *Management Science, 61*(1), 111–128. https://doi.org/10.1287/mnsc.2014.2059.

Arvanitis, S., Hallam, M., Post, T., & Topaloglou, N. (2019). Stochastic spanning. *Journal of Business & Economic Statistics, 37*(4), 573–585. https://doi.org/10.1080/07350015.2017.1391099.

Benson, H. (1998). An outer approximation algorithm for generating all efficient extreme points in the outcome set of a multiple objective linear programming problem. *Journal of Global Optimization, 13*(1), 1–24.

Bowman, V. (1976). On the relationship of the Tchebycheff norm and the efficient frontier of multiple-criteria objectives. In H. Thiriez, & S. Zionts (Eds.), *Multiple criteria decision making: Proceedings of a conference Jouy-en-Josas, France May 21–23, 1975* (pp. 76–86). Berlin, Heidelberg: Springer Berlin Heidelberg.

Dentcheva, D., & Ruszczyński, R. (2003). Optimization with stochastic dominance constraints. *SIAM Journal on Optimization, 14*(2), 548–566.

Ehrgott, M. (2005). *Multicriteria optimization*. Berlin: Springer.

Gotoh, J., & Konno, H. (2010). Third degree stochastic dominance and mean-risk analysis. *Management Science, 46*(2), 289–301.

Gustafsson, J., & Salo, A. (2005). Contingent portfolio programming for the management of risky projects. *Operations Research, 53*(6), 946–956.

Gutjahr, W., & Pichler, A. (2016). Stochastic multi-objective optimization: A survey on non-scalarizing methods. *Annals of Operations Research, 236*(2), 475–499.

Hadar, J., & Russell, W. (1969). Rules for ordering uncertain prospects. *The American Economic Review, 59*(1), 25–34.

Hanoch, G., & Levy, H. (1969). The efficiency analysis of choices involving risk. *The Review of Economic Studies, 36*(3), 335–346.

Harju, M., Liesiö, J., & Virtanen, K. (2019). Spatial multi-attribute decision analysis: Axiomatic foundations and incomplete preference information. *European Journal of Operational Research, 275*(1), 167–181. https://doi.org/10.1016/j.ejor.2018.11.013.

Hardy, G., Littlewood, J., & Polya, G. (1934). *Inequalities*. Cambridge, U.K.: Cambridge University Press.

Hazen, G. (1986). Partial information, dominance, and potential optimality in multiattribute utility theory. *Operations Research, 34*(2), 296–310.

Hodder, J., Jackwerth, J., & Kolokolova, O. (2014). Improved portfolio choice using second-order stochastic dominance. *Review of Finance*, 1–25.

Kallio, M., & Deghan Hardoroudi, N. (2017). Stochastic dominance efficiency tests. *European journal of Operational Research, 264*, 675–685.

Kallio, M., & Dehghan Hardoroudi, N. (2019). Advancements in stochastic dominance efficiency tests. *European Journal of Operational Research, 276*(2), 790–794.

Kangaspunta, J., Liesiö, J., & Salo, A. (2012). Cost-efficiency analysis of weapon system portfolios. *European Journal of Operational Research, 223*(1), 264–275.

Karsu, Ö., Morton, A., & Argyris, N. (2017). Capturing preferences for inequality aversion in decision support. *European Journal of Operational Research, 264*(2), 686–706.

Kuosmanen, T. (2004). Efficient diversification according to stochastic dominance criteria. *Management Science, 50*(10), 1390–1406.

Levy, H. (1992). Stochastic dominance and expected utility: Survey and analysis. *Management Science, 38*(4), 555–593.

Levy, H. (2016). *Stochastic dominance: Investment decision making under uncertainty* (3rd ed.). Springer.

Liesiö, J., Andelmin, J., & Salo, A. (2020a). Efficient allocation of resources to a portfolio of decision making units. *European Journal of Operational Research, 286*(2), 619–636. https://doi.org/10.1016/j.ejor.2020.03.031.

Liesiö, J., Mild, P., & Salo, A. (2007). Preference programming for robust portfolio modeling and project selection. *European Journal of Operational Research, 181*(3), 1488–1505.

Liesiö, J., Mild, P., & Salo, A. (2008). Robust portfolio modeling with incomplete cost information and project interdependencies. *European Journal of Operational Research, 190*(3), 679–695.

Liesiö, J., & Salo, A. (2012). Scenario-based portfolio selection of investment projects with incomplete probability and utility information. *European Journal of Operational Research, 217*(1), 162–172.

Liesiö, J., Xu, P., & Kuosmanen, T. (2020b). Portfolio diversification based on stochastic dominance under incomplete probability information. *European Journal of Operational Research, 286*(2), 755–768.

Linton, O., Post, T., & Whang, Y. (2014). Testing for the stochastic dominance efficiency of a given portfolio. *The Econometrics Journal, 17*(2), S59–S74. https://doi.org/10.1111/ectj.12016.

Löhne, A. (2011). *Vector optimization with infimum and supremum*. Berlin: Springer.

Löhne, A., & Weißing, B. (2014). Bensolve - VLP solver, version 2.0.1. http://www.bensolve.org.

Longarela, I. (2016). A characterization of the SSD-efficient frontier of portfolio weights by means of a set of mixed-integer linear constraints. *Management Science, 62*(12), 3549–3554.

Mancuso, A., Compare, M., Salo, A., Zio, E., & Laakso, T. (2016). Risk-based optimization of pipe inspections in large underground networks with imprecise information. *Reliability Engineering & System Safety, 152*, 228–238.

Mild, P., Liesiö, J., & Salo, A. (2015). Selecting infrastructure maintenance projects with robust portfolio modeling. *Decision Support Systems, 77*, 21–30.

Moskowitz, H., Preckel, P., & Yang, A. (1993). Decision analysis with incomplete utility and probability information. *Operations Research, 41*(5), 864–879.

Neumann, J. v., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton University Press, Princeton.

Post, T. (2003). Empirical tests for stochastic dominance efficiency. *The Journal of Finance, 58*(5), 1905–1932.

Post, T., & Khanjani Shiraz, R. (2019). Uncovering latent stochastic dominance relations using prior rankings. *Working paper*. https://ssrn.com/abstract=3442919

Post, T., & Kopa, M. (2013). General linear formulations of stochastic dominance criteria. *European Journal of Operational Research, 230*(2), 321–332.

Post, T., & Kopa, M. (2017). Portfolio choice based on third-degree stochastic dominance. *Management Science, 63*(10), 3381–3392.

Przybylski, A., & Gandibleux, X. (2017). Multi-objective branch and bound. *European Journal of Operational Research, 260*(3), 856–872.

Quirk, J., & Saposnik, R. (1962). Admissibility and measurable utility functions. *The Review of Economic Studies, 29*(2), 140–146.

Ringuest, J., Graves, S., & Case, R. (2000). Conditional stochastic dominance in R&D portfolio selection. *IEEE Transaction on Engineering Management, 47*(4), 478–484.

Roman, D., Darby-Dowman, K., & Mitra, G. (2006). Portfolio construction based on stochastic dominance and target return distributions. *Mathematical Programming, 108*(2–3), 541–569.

Salo, A., Keisler, J., & Morton, A. (2011). *Advances in portfolio decision analysis: Improved methods for resource allocation*. New York: Springer.

Scaillet, O., & Topaloglou, N. (2010). Testing for stochastic dominance efficiency. *Journal of Business & Economic Statistics, 28*(1), 169–180. https://doi.org/10.1198/jbes.2009.06167.

Starmer, C. (2000). Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal of Economic Literature, 38*(2), 332–382.

Stummer, C., & Heidenberger, K. (2003). Interactive R&D portfolio analysis with project interdependencies and time profiles of multiple objectives. *IEEE Transaction on Engineering Management, 50*(2), 175–183.

Tsetlin, I., Winkler, R., Huang, R., & Tzeng, L. (2015). Generalized almost stochastic dominance. *Operations Research, 63*(2), 363–377.

Vilkkumaa, E., Liesiö, J., Salo, A., & Ilmola-Sheppard, L. (2018). Scenario-based portfolio model for building robust and proactive strategies. *European Journal of Operational Research, 266*(1), 205–220. https://doi.org/10.1016/j.ejor.2017.09.012.

Whitmore, G. A. (1970). Third-degree stochastic dominance. *The American Economic Review, 60*(3), 457–459.

Wierzbicki, A. (1980). The use of reference objectives in multiobjective optimization. In G. Fandel, & T. Gal (Eds.), *Multiple criteria decision making theory and application: Proceedings of the third conference Hagen/Königswinter, West Germany, August 20–24, 1979* (pp. 468–486). Berlin, Heidelberg: Springer Berlin Heidelberg.