Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Nonadditive Multiattribute Utility Functions for Portfolio Decision Analysis

Juuso Liesiö, Eeva Vilkkumaa

informs
http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 69, No. 6, November–December 2021, pp. 1886–1908
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Methods

Nonadditive Multiattribute Utility Functions for Portfolio Decision Analysis

Juuso Liesiö,[a] Eeva Vilkkumaa[a]

[a] Department of Information and Service Management, Aalto University School of Business, 00076 Aalto, Finland

Contact: juuso.liesio@aalto.fi, https://orcid.org/0000-0002-8178-3179 (JL); eeva.vilkkumaa@aalto.fi, https://orcid.org/0000-0002-4217-8697 (EV)

Received: January 15, 2020
Revised: June 16, 2020
Accepted: June 26, 2020
Published Online in Articles in Advance: January 18, 2021

Subject Classifications: decision analysis: multiple criteria; utility/preference: multiattribute
Area of Review: Decision Analysis

https://doi.org/10.1287/opre.2020.2046

Abstract. Portfolio decision analysis models support selecting a portfolio of projects in view of multiple objectives and limited resources. In applications, portfolio utility is commonly modeled as the sum of the projects' multiattribute utilities, although such approaches lack rigorous decision-theoretic justification. This paper establishes the axiomatic foundations of a more general class of multilinear portfolio utility functions, which includes additive and multiplicative portfolio utility functions as special cases. Furthermore, we develop preference elicitation techniques to assess these portfolio utility functions as well as optimization models to identify the most preferred portfolio in view of resource and other constraints. We also examine how the functional form of the portfolio utility function affects decision recommendations by using randomly generated and real problem instances.

Funding: This research has been supported by the Academy of Finland [grant 323800].

Keywords: multiattribute utility theory • portfolio decision analysis • resource allocation

1. Introduction

Selecting which portfolio of project candidates (e.g., products, infrastructure investments, research programs, policy options) to implement with the available resources is an important decision problem faced frequently by both businesses and public organizations (e.g., Parnell et al. 2002, Ewing Jr. et al. 2006, Phillips and Bana e Costa 2007, Grushka-Cockayne et al. 2008, Montibeller et al. 2009, Lourenço et al. 2012, Lopes and de Almeida 2015, Mild et al. 2015, Barbati et al. 2018). Often, such decisions are complicated by the presence of several objectives, uncertainties, multiple resource types, and project interdependencies. Consequently, there has been a great deal of research on developing and deploying methods to support such decisions within the fields of, for example, resource allocation decisions (Kleinmuntz 2007) and portfolio decision analysis (PDA; Salo et al. 2011). These methods combine decision analysis and optimization techniques to capture preferences on risk and multiple objectives and to identify the most preferred combination (i.e., portfolio) of projects to implement.

Multiattribute PDA applications commonly utilize the additive preference model (e.g., Golabi et al. 1981, Parnell et al. 2002, Ewing Jr. et al. 2006, Kleinmuntz 2007, Phillips and Bana e Costa 2007, Gurgur and Morley 2008, Lourenço et al. 2012, Lopes and de Almeida 2015, Mild et al. 2015). In this model each project is first evaluated with regard to a set of attributes, after which a suitable multiattribute value/utility function is used to obtain the overall value/utility of each project. The overall value/utility of a portfolio is then obtained as the sum of the multiattribute values/utilities of projects included in the portfolio. The widespread use of additive preference models in PDA applications can be mainly attributed to two factors. First, in terms of preference elicitation, the models require only the assessment of a multiattribute value/utility function over the measurement scales of the attributes used to evaluate individual projects. Second, the computational requirements of additive preference models are modest, because the most preferred portfolio can be identified using standard integer linear programming tools.

A concern with the additive preference model is that it is not particularly well-suited for handling uncertain project outcomes. In particular, the axiomatization of the additive preference model in the seminal paper by Golabi et al. (1981) builds on measurable value theory (Dyer and Sarin 1979), which characterizes preferences between certain outcomes and changes thereof (cf. strength of preference).

Essentially, this choice implies that the model can only be deployed in applications where the projects' outcomes are deterministic. It is worth highlighting that, even though Golabi et al. (1981) utilize utility functions to capture multiattribute preferences at the project level, these utility functions are not directly used by the additive portfolio preference model. Instead, results of Dyer and Sarin (1979) are used to identify a value function that represents the same preferences as the utility function among *deterministic* project outcomes, and this value function is used in the additive model capturing preferences among portfolios. Hence, although outcomes in practical applications are often uncertain, the additive preference model lacks a solid axiomatic foundation to accommodate these uncertainties.

Conceivably, this lack of axiomatic foundation could be ignored and the additive preference model could be applied in a setting with uncertain project outcomes by using expectations of portfolio values to provide decision recommendations. This approach does not eliminate the fact that the additive portfolio value function entails a fixed attitude toward risk and does not include any parameters to control this attitude through preference elicitation questions. For instance, the additive preference model assumes that there are no benefits from diversification, but that a single project is always equally preferred to two projects the sum of whose expected values is equal to that of the single project.

We address this gap in the existing literature by developing an axiomatic foundation for PDA under uncertainty. Specifically, we show that under reasonably mild assumptions about preferences among portfolios with uncertain outcomes, these preferences can be represented with a symmetric multilinear utility function. Importantly, this portfolio utility function makes use of a multiattribute utility function constructed over the project evaluation scale, and as a result the preference elicitation effort grows only linearly in the number of projects when compared with additive portfolio preference models. Moreover, we establish results that identify the types of preferences under which the multilinear portfolio utility function is reduced to either a multiplicative or an additive portfolio utility function. In the multiplicative utility function, portfolio utility is obtained as a product of the scaled multiattribute utilities of the projects in the portfolio. In the additive portfolio utility function, portfolio utility is obtained as a sum of the projects' multiattribute utilities. This result formally establishes those preference assumptions that are implicitly made in practical applications, where project utilities are added up to obtain portfolio utility.

Our work provides continuation for the growing interest in developing multilinear preference models for various decision analytic contexts. For instance, Keller and Simon (2019) develop multilinear value/utility functions for spatial decision analysis, where the alternatives' outcomes vary across a geographical region (Simon et al. 2014). Bordley and Kirkwood (2004) utilize multilinear utility functions to represent preferences for achieving a certain target for each attribute in general multiattribute decision problems. Abbas (2009) shows that multilinear utility functions can be viewed as a special case of multiattribute utility copulas.

In addition to the theoretical contribution, this paper also develops and examines approaches that are required to apply multilinear portfolio utility functions in practice. First, we show how multilinear portfolio utility functions can be constructed with preference assessment questions involving uncertain portfolio outcomes. Second, we develop mixed-integer linear programming (MILP) models to identify the feasible portfolio that maximizes expected multilinear utility. This contribution is especially important from the viewpoint of practical applicability of these new preferences models, as it makes it possible to utilize standard off-the-shelf optimization software to produce decision recommendations. Finally, multilinear portfolio utility functions are applied to real-life and randomly generated data sets to examine how the decision recommendations they provide differ from those resulting from the standard additive portfolio utility function.

The portfolio utility functions developed here can be viewed as symmetric versions of the general nonsymmetric multiattribute utility functions developed in standard multiattribute utility theory (MAUT; Keeney and Raiffa 1976; see also Fishburn 1973, Farquhar 1975). However, there are important theoretical and practical differences between these functions that we wish to highlight. First, the preference assumptions required to establish a specific form for the portfolio utility function can be quite different from those commonly applied in MAUT. In particular, the preference assumption implying the symmetric portfolio utility function is fairly strong in the sense that only relatively weak additional preference assumptions are required to obtain its multilinear, multiplicative, and additive forms. Second, preference elicitation for a symmetric portfolio utility function requires different techniques and involves fewer preference parameters. For instance, the number of preference parameters in the symmetric multilinear portfolio utility function grows only linearly in the number of projects, whereas for a nonsymmetric multilinear utility function this number grows exponentially in the number of attributes. Finally, in general MAUT, the question of how to identify the expected utility maximizing decision alternative is not usually addressed, as it is (implicitly) assumed that it is possible to evaluate the expected utility of each alternative. In PDA applications the number of alternatives (i.e., feasible project portfolios) is often so high that such a complete enumeration approach is not possible. This necessitates the development of portfolio optimization approaches to accompany each of the developed functional forms of portfolio utility.

Besides MAUT, our work intersects with other active strands of current research as well. For instance, preference models for multiattribute portfolio selection with deterministic project outcomes have been discussed and developed by Clemen and Smith (2009), Argyris et al. (2011, 2014), Liesiö (2014), Liesiö and Punkka (2014), Morton (2015), and Morton et al. (2016). In particular, Argyris et al. (2014) develop nonadditive preference models for deterministic PDA using a more conventional multiobjective programming framework. This framework differs from the one used here in that it assumes the attributes measure performance at the portfolio level. Hence, the framework is suitable for cases in which all attributes have natural measurement scales, such as revenues: For such attributes it is meaningful to say that a portfolio containing two projects with revenues \$200 and \$500 has a composite revenue of \$700. In turn, our framework evaluates multiattribute outcomes already at the project level, after which the resulting project utilities are aggregated to obtain the portfolio utility. Such a framework is best suited for settings with attributes that do not have a natural measurement scale, such as innovativeness or strategic fit, but are evaluated on, for example, a 1–5 Likert scale or even a verbal scale. In such a setting a portfolio consisting of, for instance, two projects each with Likert score 2 would not necessarily be equally preferred to a portfolio containing a single project with score 4.

A framework similar to the one developed here is used by Liesiö (2014), but the resulting multilinear preference models are different. First, Liesiö (2014) assumes deterministic project outcomes and, therefore, builds on measurable value theory (Dyer and Sarin 1979). In contrast, we use MAUT, which allows for uncertain project outcomes to be modeled through probability distributions. Consequently, the models by Liesiö (2014) can be used to identify the portfolio with the highest measurable value, whereas the models proposed in this paper seek to optimize expected portfolio utility. In general, the optimal portfolios resulting from these two models are not the same. In particular, using the projects' expected outcomes as the deterministic outcomes required by the portfolio value model by Liesiö (2014) would generally result in a different decision recommendation than what would be obtained by maximizing expected portfolio utility across all possible outcomes. Second, the different theoretical foundations of these models imply different approaches to preference elicitation. In particular, the preference assessment procedure in Liesiö (2014) relies on the decision maker (DM) comparing changes in deterministic portfolio outcomes (i.e., strength of preference), whereas the elicitation methods proposed in this paper seek to capture the DM's preferences with regard to uncertain portfolio outcomes. Finally, the structure of the axiomatic framework in this paper is fundamentally different from that in Liesiö (2014). In particular, the models developed here do not build on restrictive assumptions on how the preferences are modeled within projects, but allow for an arbitrary functional form to aggregate the projects' multiattribute outcomes into project utilities. The focus is on preference assumptions at the portfolio level and on the functional forms that these assumptions imply for the portfolio utility function.

The rest of this paper is structured as follows. Section 2 introduces the basic framework for multiattribute portfolio selection under uncertainty. Section 3 presents the preference assumptions under which preferences among uncertain multiattribute portfolio outcomes can be represented with a multilinear utility function. Section 4 discusses preference assessment techniques to construct multilinear portfolio utility functions. Section 5 provides results on two special types of preferences that lead to additive and multiplicative forms for portfolio utility. Section 6 develops optimization models to identify the most preferred portfolio under resource and other portfolio constraints. Section 7 analyzes the effects that the use of a multilinear portfolio utility function has on decision recommendations compared with standard approaches. Section 8 concludes.

2. Preliminaries on Multiattribute Project Portfolio Selection

Consider selecting a portfolio from a set of project candidates. The quality of each project candidate is evaluated with regard to n attributes with measurement scales $Y_1, \ldots, Y_n$. These attributes can represent both quantitative measures (e.g., present value, production volume) and qualitative evaluations (e.g., fit to company strategy, quality considerations; see, e.g., Kleinmuntz 2007, Clemen and Smith 2009, Lopes and de Almeida 2015). The outcome of each project corresponds to a vector $y = (y_1, \ldots, y_n)$ belonging to the set

$$Y = Y_1 \times \cdots \times Y_n. \tag{1}$$

Moreover, we use $y^0 = (y_1^0, \ldots, y_n^0)$ and $y^* = (y_1^*, \ldots, y_n^*)$ to denote the least and most preferred outcomes in set Y, respectively.

We denote the uncertain outcome of the jth project by $\tilde{x}_j$, which is technically a vector-valued random variable whose realizations take values in the set Y given by (1). The possible outcomes and distribution of this random variable depend on whether the jth project is selected or not. To formalize this dependence, we denote the outcome of project j by $\tilde{x}_j^F$ if it is selected (cf. funded) and by $\tilde{x}_j^B$ if not (cf. baseline; Clemen and Smith 2009, Liesiö and Punkka 2014, Morton 2015). The realizations of both of these random variables belong to set Y. Thus, the uncertain outcome of the jth project is formally given by

$$\tilde{x}_j = \begin{cases} \tilde{x}_j^B & \text{if } z_j = 0, \\ \tilde{x}_j^F & \text{if } z_j = 1, \end{cases} \tag{2}$$

where the binary decision variable $z_j = 1$ if the project is included in the portfolio. If the baseline outcome is deterministic and the same for all projects—which is a common assumption in reported applications—we will denote it by $y^B \in Y$.

The expected utility of the jth project is thus $\mathbb{E}[u(\tilde{x}_j)]$, where the multiattribute utility function $u : Y \to \mathbb{R}$ maps the multiattribute project outcomes in Y to a single-dimensional utility. Perhaps the most widely used is the additive utility function $u(y) = \sum_{i=1}^n w_i u_i(y_i)$, where $u_i : Y_i \to \mathbb{R}$ is the marginal utility function for the ith attribute and w_i is the importance weight of attribute i.

As an example of a real-life portfolio selection problem with uncertain project outcomes, we consider an application in healthcare resource allocation by Airoldi et al. (2011). In this application, public healthcare officials decide on a portfolio of interventions (i.e., projects) that seek to improve the quality of life and reduce health inequalities in a particular geographical area. Specifically, $m = 23$ project candidates are evaluated with regard to $n = 2$ attributes, namely, health benefits and health inequality reduction. Moreover, the projects' outcomes with regard to these attributes are uncertain due to the possibility that an intervention fails and does not deliver the planned outcomes. The expected utility for a funded project candidate j is $\mathbb{E}[u(\tilde{x}_j^F)] = p_j u(x_j^S) + (1 - p_j)\, u(x_j^U)$, where p_j is the success probability of the intervention, and x_j^S and x_j^U represent its outcome in case the intervention is successful or unsuccessful, respectively (see Airoldi et al. 2011 for details). If a project is not funded, it yields a baseline utility of zero (i.e., $\mathbb{E}[u(\tilde{x}_j^B)] = 0$). The expected utilities (scaled here between 0 and 1), implementation costs, and benefit-to-cost ratios of the projects are shown in Table 1.

In applications, it is often assumed that the most preferred portfolio is the one that maximizes the sum of the (expected) utilities of the projects included in the portfolio. If, for instance, there is a single budget constraint and the baseline utility for each project is zero, then the most preferred portfolio can be identified by solving the knapsack problem

$$\max_{z_j \in \{0,1\}} \left\{ \sum_j z_j \mathbb{E}\left[u\left(\tilde{x}_j^F\right)\right] \,\middle|\, \sum_j z_j c_j \le b \right\}, \tag{3}$$

Table 1. Expected Multiattribute Utilities (Airoldi et al. 2011) and Costs of the Intervention Projects

Project	j	$\mathbb{E}[u(\tilde{x}_j^F)]$	c_j (k£)	$\mathbb{E}[u(\tilde{x}_j^F)]/c_j$
Pneumonia	1	0.7	75	0.1579
Dementia services	2	0.31	50	0.1036
TIA and secondary prevention	3	0.32	130	0.0415
Prison MH	4	0.27	150	0.0301
Obesity training	5	0.1	60	0.0288
Workforce development	6	0.16	100	0.0278
Psych therapies	7	0.18	120	0.0254
Early detection and diagnostics	8	0.34	300	0.0191
CAMHS school	9	0.16	160	0.0172
Prevention	10	0.62	650	0.0161
CAMHS 1:1	11	0.07	80	0.0158
Cardiac rehab	12	0.08	100	0.0129
Alcohol misuse svc	13	0.22	300	0.0126
Social inclusion	14	0.22	300	0.0125
Palliative and EOL	15	0.54	760	0.0119
Obesity 1:1	16	0.07	140	0.0087
Primary prevention	17	0.27	600	0.0077
Access to dental	18	0.19	480	0.0068
Active treatment	19	0.02	50	0.0062
Stroke emergency	20	0.2	600	0.0056
CHD acute	21	0.05	300	0.0026

Note. The projects are listed in a decreasing order of utility-to-cost ratios $\mathbb{E}[u(\tilde{x}_j^F)]/c_j$.

where c_j is the cost of implementing the jth project and b is the budget. In case of the healthcare example of Table 1, the portfolio that maximizes the sum of the expected project utilities given a budget of £1.6 million (approximately 30% of the cost of implementing all projects) consists of projects $j = 1$ (Pneumonia) through $j = 7$ (Psych therapies), and projects $j = 9$ (CAMHS school), $j = 10$ (Prevention), and $j = 12$ (Cardiac rehab).

This portfolio recommendation assumes that the DM's preferences among portfolios are correctly captured by the sums of the expected utilities of projects included in these portfolios. For instance, project $j = 10$ (Prevention) with a cost of £650,000 is recommended over the equally expensive combination of projects $j = 8$ (Early detection and diagnostics), $j = 13$ (Alcohol misuse svc), and $j = 19$ (Active treatment), because its expected utility (0.62) is higher than the sum of the expected utilities of the three projects (0.58). Nevertheless, the DM might actually prefer the three-project combination, because then the health benefits obtained with this substantial amount of resources would not be entirely contingent on the success of a single project. What type of a utility function would represent such preferences, and what kinds of assumptions would these preferences need to satisfy? To address these questions, the following sections develop multilinear and multiplicative utility functions whose underlying preference assumptions are less restrictive than those required by the additive portfolio utility function. In Section 7.2, we revisit the above example on healthcare resource allocation to analyze the differences in portfolios that are optimal in view of these different utility functions.

3. Multilinear Portfolio Utility Functions

In PDA, decision alternatives correspond to project portfolios. Since each project is modeled as a vector-valued random variable $\tilde{x}_j$ with realizations in set Y, each portfolio corresponds to a random variable $\tilde{x} = (\tilde{x}_1, \ldots, \tilde{x}_m)$ with realizations in the set

$$X = \underbrace{Y \times \cdots \times Y}_{m \text{ sets}}, \tag{4}$$

where Y is given by (1) and m is the number of project candidates. The set of all portfolios is denoted by $\mathcal{X}$. With a slight abuse of notation we use x to denote the degenerate random variable in $\mathcal{X}$ whose outcome is $x = (x_1, \ldots, x_m) \in X$ with probability one (i.e., essentially a deterministic variable). Furthermore, we use the notation x_j instead of $\tilde{x}_j$ to highlight that the outcome of the jth project is deterministic.

The DM's preferences among portfolios in $\mathcal{X}$ are captured by a complete and transitive relation $\succeq$. Specifically, $\tilde{x} \succeq \tilde{x}'$ denotes that portfolio $\tilde{x}$ is (weakly) preferred to portfolio $\tilde{x}'$. Strict preference $\succ$ and indifference $\sim$ are defined in the usual manner. Assuming the relation $\succeq$ satisfies certain additional assumptions (i.e., independence and Archimedean axioms; von Neumann and Morgenstern 1947), there exists a utility function $U : X \to \mathbb{R}$ that represents the DM's preferences $\succeq$ in the sense that

$$\mathbb{E}[U(\tilde{x})] \geq \mathbb{E}[U(\tilde{x}')] \Leftrightarrow \tilde{x} \succeq \tilde{x}' \; \forall \tilde{x}, \tilde{x}' \in \mathcal{X}. \tag{5}$$

This utility function is unique up to positive affine transformations, and hence, it can be scaled by fixing the utility of any two portfolios. We choose the following scaling: zero utility is assigned to a portfolio in which all projects have the worst outcome ($U(y^0, y^0, \ldots, y^0) = 0$), and unit utility is assigned to a portfolio in which the outcome of the first project is changed to its most preferred level ($U(y^*, y^0, \ldots, y^0) = 1$).

We make two assumptions on the preferences among portfolios. The first assumption is that if two portfolios have degenerate outcomes that are equal up to the indexing of projects, then these portfolios are equally preferred.

Assumption 1. *Preferences are independent from project indexing:*

$$(x_1, x_2, \ldots, x_{j-1}, x_j, x_{j+1}, \ldots) \sim (x_j, x_2, \ldots, x_{j-1}, x_1, x_{j+1}, \ldots).$$

Note that Assumption 1 only requires that a portfolio remains equally preferred if the outcomes of the first and jth project are interchanged. However, this assumption implies that the outcomes of any two projects with indices j and k can be interchanged without affecting preferences, because the repeated use of the assumption yields $(x_1,\ldots,x_j,\ldots,x_k,\ldots) \sim (x_j,\ldots,x_1,\ldots,x_k,\ldots) \sim (x_k,\ldots,x_1,\ldots, x_j,\ldots) \sim (x_1,\ldots,x_k,\ldots,x_j,\ldots)$. Thus, the permutation of project indexing does not affect portfolio preferences. This implies that any portfolio utility function that represents preferences satisfying Assumption 1 has to be symmetric.

The second assumption addresses preference between two portfolios in which only the outcome of the first project is uncertain, and the deterministic outcomes of the remaining projects are equal in both portfolios. In particular, the assumption requires that preference between these two portfolios is not affected if the deterministic outcomes are changed. Hence, Assumption 2 can be viewed as a requirement that the first project is utility independent of the other projects (cf. utility independent attributes; Keeney and Raiffa 1976).

Assumption 2. *Preferences for uncertain project outcomes do not depend on the deterministic outcomes of the other projects:*

$$(\tilde{x}_1, x_2, x_3, \ldots) \succeq (\tilde{x}_1', x_2, x_3 \ldots) \Rightarrow (\tilde{x}_1, x_2', x_3', \ldots) \succeq (\tilde{x}_1', x_2', x_3', \ldots).$$

Note that, even when combined with Assumption 1, Assumption 2 only restricts preferences between pairs of portfolios both of which have exactly the same deterministic outcomes for all projects except for one. Thereby, Assumption 2 does not rule out the possibility that preferences may be affected by stochastic dependencies (e.g., correlation) between project outcomes. Moreover, Assumption 2 does not prohibit introducing resource synergies (or cannibalization effects) to the portfolio model using the same modeling techniques as with additive portfolio utility or value functions (for details, see, e.g., Liesiö et al. 2008). These techniques include adding dummy projects that capture synergy effects (e.g., reduction in cost) and linear constraints that ensure these dummy projects can only be included in the portfolio if specific synergy conditions are met (e.g., a specific combination of projects is selected).

Together, Assumptions 1 and 2 ensure that the most and least preferred project outcomes can be defined based on preferences among portfolios: If the assumptions did not hold, these outcomes might be different for each project, and they could also depend on the outcomes of other projects in the portfolio. Moreover, the two assumptions make it possible to define a project utility function u over the set of possible multiattribute project outcomes Y (Equation (1)). If the assumptions did not hold, then assessing a utility function over project outcomes would not make sense, because the shape of the function could depend on the outcomes of other projects in the portfolio. Specifically, the two assumptions together imply that the portfolio utility function is a symmetric multilinear function of the project utility functions. This result is formally stated by the following theorem.

Theorem 1. *Assumptions 1 and 2 hold if and only if the portfolio utility function $U : X \to \mathbb{R}$ is multilinear:*

$$U(x_1, \ldots, x_m) = \sum_{J \subseteq \{1,\ldots,m\}} \lambda(|J|) \prod_{j \in J} u(x_j) \prod_{j \notin J} (1 - u(x_j)), \tag{6}$$

where $u : Y \to \mathbb{R}$ is the project utility function

$$u(y) = U\left(y, y^0, \ldots, y^0\right) \tag{7}$$

and $\lambda : \mathbb{N}^+ \cup \{0\} \to \mathbb{R}$ is a strictly increasing function

$$\lambda(k) = U\left(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\right). \tag{8}$$

Furthermore, $u(y^0) = \lambda(0) = U(y^0, \ldots, y^0) = 0$ and $u(y^*) = \lambda(1) = U(y^*, y^0, \ldots, y^0) = 1$.

Detailed proofs are presented in Appendix A, but the approach for establishing Theorem 1 is relatively intuitive. Assumption 2 requires project x_1 to be utility independent from the rest of the projects. Moreover, Assumption 1 states that the preferences do not depend on project indexing, which extends this utility independence assumption to hold for the other projects $x_2, \ldots, x_m$ as well. Keeney and Raiffa (1976) (theorem 6.3) establish that if each attribute is utility independent, then preferences are represented by a (in general, nonsymmetric) multilinear function (see also Fishburn 1973, Farquhar 1975). Assumption 1 then implies additional constraints on the parameters of this function that ensure preferences do not depend on project indexing, which makes the function symmetric.

The resulting symmetric multilinear portfolio utility function U utilizes only two objects to capture the preference structure among portfolios: (i) the project utility function u that maps the multiattribute project outcomes $y \in Y$ to real values, and (ii) the real-valued parameters $\lambda(k)$ that represent the utilities of portfolios in which each project has either the most or least preferred outcome. Figure 1 illustrates the effect parameter λ has on the shape of the multilinear portfolio utility function. In the middle panel, the utility $\lambda(2) = U(y^*, y^*) = 2.0$ of a portfolio of two projects that both attain the most preferred outcome y^* is equal to the sum of the projects' utilities $u(y^*) + u(y^*) = 2.0$. As a result, the utility of any portfolio is always equal to the sum of the projects' utilities. In the leftmost (rightmost) panel, $\lambda(2)$ is strictly below (above) 2.0, which results in portfolio utilities that are always strictly lower (higher) than the sum of the utilities of the projects included in the portfolio.

Note that the choice of defining the project utility function as the utility of a portfolio with all other projects having the least preferred outcome (Equation (7)) is a convenient scaling choice rather than a fundamental property of the model. In particular, if the outcomes of the other projects were fixed to some other levels $x_2, \ldots, x_m$, then the resulting project utility function u' would be a positive affine transformation of u. This is because

$$\begin{aligned} u'(y) &= U(y, x_2, \ldots, x_m) \\ &= u(y)\Big(\underbrace{U(y^*, x_2, \ldots, x_m) - U(y^0, x_2, \ldots, x_m)}_{=\alpha}\Big) \\ &\quad + \underbrace{U(y^0, x_2, \ldots, x_m)}_{=\beta}, \end{aligned}$$

where constants $\alpha > 0$ and β do not depend on y.

Figure 1. (Color online) The Multilinear Portfolio Utility Function (6) for $m = 2$ Projects, When Parameter $\lambda(2)$ Is Equal to 1.5, 2, or 2.5

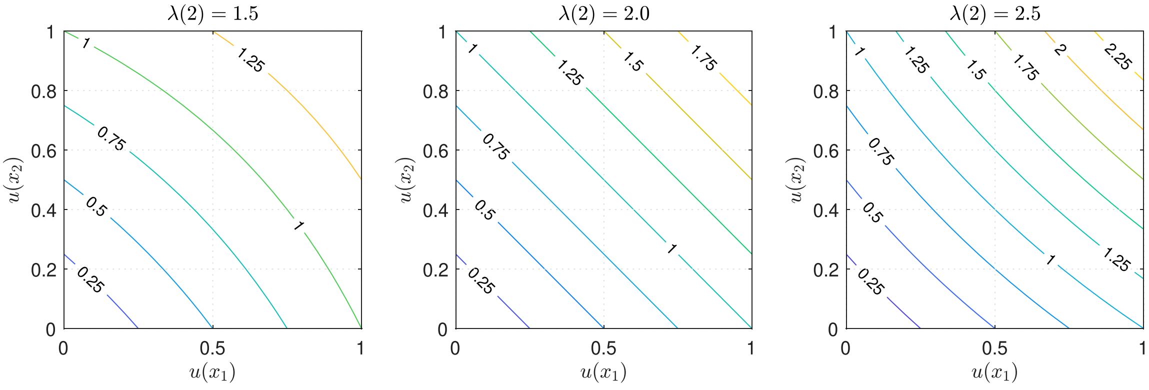

Compared with the standard approach where portfolio utility is computed as the sum of project utilities, the preference elicitation effort is only increased by the need to specify the additional parameters $\lambda(2), \ldots, \lambda(m)$. Importantly, the number of these parameters increases only linearly in the number of project candidates m, whereas applying general nonsymmetric multilinear utility functions (Keeney and Raiffa 1976, theorem 6.3) in the portfolio context would result in the number of preference parameters growing exponentially in m.

4. Preference Elicitation

We develop two alternative approaches for the assessment of the parameters λ. The first approach can be used without having assessed the project utility function u, but requires the DM to compare only portfolios in which each project yields either the most (y^*) or the least preferred (y^0) outcome. The second approach is more flexible in terms of the outcomes of the compared portfolios, but requires that the project utility function has been assessed.

In the first approach, the DM is asked to repeatedly compare two portfolios of different types. One portfolio has a deterministic outcome such that k projects have the most preferred outcome y^*, while the remaining $m-k$ projects have the least preferred outcome y^0. The outcome of the second portfolio is uncertain with two possible outcomes: Either $k+1$ projects have the most preferred outcome with probability p, or $k-1$ projects have the most preferred outcome with probability $1-p$. In both cases, the rest of the projects have the least preferred outcome. The DM is asked to adjust the value of p until the portfolio with the deterministic outcome and the portfolio with the uncertain outcome are equally preferred. Formally, for each $k \in \{1, \ldots, m-1\}$ the DM is asked to determine a probability $p \in (0,1)$ such that portfolios

$$x' = \Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big) \text{ and}$$

$$\tilde{x} = \begin{cases} \Big(\underbrace{y^*, \ldots, y^*}_{k-1 \text{ elements}}, y^0, \ldots, y^0\Big) & \text{with probability } 1-p, \\ \Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big) & \text{with probability } p \end{cases} \tag{9}$$

are equally preferred. Evaluating the expected utilities of these portfolios gives

$$\begin{aligned} \mathbb{E}[U(x')] &= \mathbb{E}[U(\tilde{x})] \\ U\Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big) &= (1-p) \\ &\quad \times U\Big(\underbrace{y^*, \ldots, y^*}_{k-1 \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\quad + pU\Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big) \\ \Leftrightarrow \lambda(k) &= (1-p)\lambda(k-1) + p\lambda(k+1) \\ \Leftrightarrow \lambda(k+1) - \lambda(k) &= \left(\frac{1}{p} - 1\right)(\lambda(k) - \lambda(k-1)), \end{aligned} \tag{10}$$

which is a linear equation with three variables $\lambda(k-1)$, $\lambda(k)$, and $\lambda(k+1)$. Repeating this process for all $k = 1, \ldots, m-1$ results in $m-1$ linear equations with $m-1$ variables $\lambda(2), \ldots, \lambda(m)$, since $\lambda(0) = 0$ and $\lambda(1) = 1$ are fixed.

To illustrate the elicitation approach, consider again the case on healthcare resource allocation presented in Table 1. The DM is first asked to consider a portfolio that contains one project that is a perfect success in that it would yield the highest total health benefits and largest decrease in healthcare inequalities attainable for a single intervention; and 20 projects that would have no impact on health or health inequality. Then, the DM is asked to consider a portfolio with 19 projects that would have no health-related impacts, and two projects with uncertain outcomes in that with probability p they would both be perfect successes and with probability $1-p$ they would both fail to deliver any health-related impacts. Finally, the DM is asked to adjust p such that the DM would be indifferent between the two portfolios. If, for instance, the DM would state that the success probability of the two uncertain projects would need to be 80% for such indifference, then Equation (10) together with $\lambda(0) = 0$ and $\lambda(1) = 1$ would yield

$$\lambda(2) = 1 + \left(\frac{1}{0.8} - 1\right)(1 - 0) = \frac{1}{0.8} = 1.25.$$

Next, the DM is asked to consider a portfolio that contains two projects that would be perfect successes and 19 projects that would have no health-related impacts, and compare this to a portfolio which would contain either (i) three perfectly successful projects and 18 projects with no health-related impacts with probability p or (ii) one perfectly successful project and 20 projects with no health-related impacts with probability $1-p$. Again, then DM is asked to adjust p such that the DM would be indifferent between the two portfolios. Given, for instance, $p = 70\%$, we would have $\lambda(3) = 1.25 + (1/0.7 - 1)(1.25 - 1) = 1 + 0.25/0.7 \approx 1.36$. Continuing in this vein, all values of $\lambda(k)$ can be determined.

The above approach relies on the most and least preferred outcomes (y^0, y^*). These outcomes should be defined in such a way that they are attainable for real-life projects and meaningful in the sense that the DM can state preferences between portfolios that include projects with these outcomes. However, in some applications—like in our example on healthcare resource allocation—it can be the case that actual portfolios do not contain many projects with such extreme outcomes, whereby imagining such portfolios for the purposes of preference statements can be cognitively too demanding.

To address this issue, we propose a second approach for assessing parameters λ in a way that allows the DM to compare portfolios with arbitrary outcomes $y^-, y^+ \in Y$ such that $y^0 \prec y^- \preceq y^+ \prec y^*$. An advantage of this approach is that it enables the analyst to choose whichever y^-, y^+ that would make the most sense to the DM in the given decision situation. For instance, y^- could be the outcome resulting from not selecting a project (i.e., baseline utility y^B) and y^+ could be the outcome of an "average" project. A second advantage of this approach is its apparent simplicity from the DM's perspective: although the choice of assessment questions and the underlying calculations can be somewhat demanding, the effort required from the DM to respond to these questions is moderate. A third advantage is related to the issue of DMs providing inconsistent preference statements, which is not unique to our context. In this approach, the range of feasible responses to each elicitation question can be readily computed a priori, thereby mitigating the problem of inconsistency.

The key challenge in modeling preferences between arbitrary portfolios is that even indifference between two deterministic outcomes x and x' will lead to constraint $U(x) = U(x')$, where both sides of the equality are multilinear functions. A set of such preference statements leads to a system of nonlinear constraints, the solutions to which can be difficult to find—especially in an interactive preference elicitation process that requires fast computation. The key insight to overcome this challenge is that, even though the portfolio utility function U is a nonlinear function on X (4), it is linear in parameters λ given a fixed portfolio outcome $x \in X$. Specifically, combining the terms in the multilinear portfolio utility function (6) that correspond to subsets of equal sizes gives

$$\begin{aligned} U(x) &= \sum_{J \subseteq \{1,\ldots,m\}} \lambda(|J|) \prod_{j \in J} u(x_j) \prod_{j \notin J} \big(1 - u(x_j)\big) \\ &= \sum_{k=0}^{m} \lambda(k) \underbrace{\sum_{\substack{J \subseteq \{1,\ldots,m\} \\ \text{s.t. } |J|=k}} \prod_{j \in J} u(x_j) \prod_{j \notin J} \big(1 - u(x_j)\big)}_{=\nu_k(x)} \\ &= \sum_{k=0}^{m} \lambda(k) \nu_k(x), \end{aligned}$$

which is a linear mapping of the vector $(\lambda(0), \lambda(1), \ldots, \lambda(m))$ when portfolio outcomes x and the project utility function u are fixed.

Suppose that the DM compares two portfolios x and $\tilde{x}$, where x has a deterministic outcome and $\tilde{x}$ has two possible outcomes x' and x'' whose probabilities are p and $(1-p)$, respectively. Suppose the DM is asked to adjust p until the DM is indifferent between the two portfolios. In this case we obtain

$$\begin{aligned} &\mathbb{E}[U(x)] = \mathbb{E}[U(\tilde{x})] \\ &\quad \Leftrightarrow U(x) = pU(x') + (1-p)U(x'') \\ &\quad \Leftrightarrow \sum_{k=0}^{m} \lambda(k)\nu_k(x) = p \sum_{k=0}^{m} \lambda(k)\nu_k(x') \\ &\qquad + (1-p) \sum_{k=0}^{m} \lambda(k)\nu_k(x'') \\ &\quad \Leftrightarrow \sum_{k=0}^{m} \lambda(k) \underbrace{\big(\nu_k(x) + p(\nu_k(x'') - \nu_k(x')) - \nu_k(x'')\big)}_{=\Psi_k} = 0, \end{aligned} \tag{11}$$

which is a linear constraint on the parameters $\lambda(0)$, $\lambda(1), \ldots, \lambda(m)$. An important implication of this result is that it suffices for the DM to provide $m-1$ preference statements of the form (11) to completely specify λ as long as the resulting vectors $[\Psi_0, \ldots, \Psi_m]$ are linearly independent. Thus, this result opens up several possibilities to assess preferences, as finding parameters λ that are consistent with the given preference statements requires only solving a system of linear equations.

The second elicitation approach builds on the above result by allowing the DM to specify the utility functions over portfolios

$$\dot{x}^k = \Big(\underbrace{y^+, \ldots, y^+}_{k \text{ elements}}, y^-, \ldots, y^-\Big), \; k \in \{0, \ldots, m\}, \tag{12}$$

where $y^-, y^+ \in Y$ such that $y^0 \prec y^- \preceq y^+ \prec y^*$. In particular, we propose a bisection-type technique which resembles those commonly used to assess a utility

Operations Research, 2021, vol. 69, no. 6, pp. 1886–1908, © 2021 The Author(s)

function over a monetary scale. In this technique, the DM is repeatedly asked to adjust probability p so that portfolio $\dot{x}^k$ is equally preferred to a portfolio whose outcome is $\dot{x}^{\overline{k}}$ with probability p and $\dot{x}^{\underline{k}}$ with probability $(1-p)$, where $k = \lfloor(\underline{k}+\overline{k})/2\rfloor$. The process begins by determining the utilities of portfolios $\dot{x}^0$ and $\dot{x}^m$ by comparing each of them to a portfolio the outcome of which is $(y^*, \dots, y^*)$ with probability p and $(y^0, \dots, y^0)$ with probability $(1-p)$, and adjusting p in both cases in such a way that the DM would be indifferent between the two portfolios. In the third assessment question we set $\underline{k}=0, \overline{k}=m$, in the fourth question $\underline{k}=0, \overline{k}=\lfloor m/2\rfloor$, and in the fifth one $\underline{k}=\lfloor m/2\rfloor, \overline{k}=m$. Questions after that would further specify the utilities between the portfolios $\dot{x}^0$, $\dot{x}^{\lfloor m/4\rfloor}$, $\dot{x}^{\lfloor m/2\rfloor}$, $\dot{x}^{\lfloor 3m/4\rfloor}$, and $\dot{x}^m$. This process can be continued until it suffices to use simple interpolation to obtain the utilities that remain unspecified.

Note that each of these preference statements in general results in a linear constraint involving all the parameters $\lambda(0), \dots, \lambda(m)$. This is in contrast to preference statements of the form (9), which resulted in a constraint on three values $\lambda(k-1)$, $\lambda(k)$, and $\lambda(k+1)$. The consequence of this is that the DM cannot assign an arbitrary value to probability p in (11), as this could result in a preference statement that is not consistent with the statements the DM has already provided. Fortunately, due to the fact that the statements result in linear constraints on the parameters λ, the interval of consistent values for probability p can be readily identified with linear programming (LP). Suppose that the DM has given L preference statements resulting in constraints $\sum_{k=0}^m \lambda(k)\Psi_k^l = 0$, where $l \in \{1, \dots, L\}$. Then, the minimum utility for portfolio $\dot{x}^k$ is

$$\begin{aligned} \underline{U} = \min_{\lambda} & \sum_{k'=0}^{m} v_{k'}\left(\dot{x}^k\right) \\ & \sum_{k'=0}^{m} \lambda(k')\Psi_{k'}^l = 0 \ \forall l \in \{1, \dots, L\}, \\ & \lambda(k') \le \lambda(k'+1) \quad \forall k' \in \{1, \dots, m-1\}, \end{aligned} \tag{13}$$

and the maximum utility $\overline{U}$ can be obtained by maximizing the objective function subject to the same constraints. Thus, probability p must belong to the interval

$$p \in \left[\frac{\underline{U} - U(\dot{x}^{\underline{k}})}{U\left(\dot{x}^{\overline{k}}\right) - U(\dot{x}^{\underline{k}})}, \frac{\overline{U} - U(\dot{x}^{\underline{k}})}{U\left(\dot{x}^{\overline{k}}\right) - U(\dot{x}^{\underline{k}})}\right], \tag{14}$$

since portfolio utilities $U(\dot{x}^{\overline{k}})$ and $U(\dot{x}^{\underline{k}})$ have been specified by the preceding preference statements.

To illustrate this preference assessment approach, consider a situation in which the DM is selecting a portfolio from $m = 40$ candidate projects and, for the purposes of preference elicitation, feels comfortable with comparing portfolios in which each project has either outcome y^- or y^+ such that $u(y^-) = 0.1$ and $u(y^+) = 0.6$. Preference elicitation question (11) is first used to specify the utilities of portfolios $\dot{x}^0 = (y^-, \dots, y^-)$ and $\dot{x}^{40} = (y^+, \dots, y^+)$. Assume, for instance, that the DM would be indifferent between portfolio $\dot{x}^0 = (y^-, \dots, y^-)$ and a portfolio whose outcome was $x^* = (y^*, \dots, y^*)$ with probability $p = 0.1$ and $x^0 = (y^0, \dots, y^0)$ with probability $(1-p) = 0.9$; and between portfolio $\dot{x}^{40} = (y^+, \dots, y^+)$ and a portfolio whose outcome was $x^* = (y^*, \dots, y^*)$ with probability $p = 0.7$ and $x^0 = (y^0, \dots, y^0)$ with probability $(1-p) = 0.3$. These assessments would result in preference statements $U(\dot{x}^0) = 0.9U(x^0) + 0.1U(x^*)$ and $U(\dot{x}^{40}) = 0.3U(x^0) + 0.7U(x^*)$, respectively. Figure 2(a) shows the entire range of utilities that portfolios $\dot{x}^k$ can achieve for different values of λ satisfying these statements (cf. LP problem (13)). Next the DM is asked to assess probability p such that portfolio $\dot{x}^{20}$ is equally preferred to a portfolio that has outcome $\dot{x}^0 = (y^-, \dots, y^-)$ with probability $1-p$ and outcome $\dot{x}^{40} = (y^+, \dots, y^+)$ with probability p. Suppose the DM states that $p = 0.7$, which corresponds to

$$U\left(\dot{x}^{20}\right) = 0.3U\left(\dot{x}^0\right) + 0.7U\left(\dot{x}^{40}\right),$$

and results in the minimum and maximum utilities shown in Figure 2(b). Figure 2(c) presents the utility ranges after an additional preference statement $U(\dot{x}^{10}) = 0.5U(\dot{x}^0) + 0.5U(\dot{x}^{20})$, and Figure 2(d) presents the utility ranges after preference statement $U(\dot{x}^{30}) = 0.7U(\dot{x}^{20}) + 0.3U(\dot{x}^{40})$. After the last preference statement, the DM could continue the process by assessing utilities of portfolios $\dot{x}^5, \dot{x}^{15}, \dot{x}^{25}, \dot{x}^{35}$. As the ranges of feasible probabilities in (14) become narrower, however, the DM might find the task of providing further assessments futile. Hence, an alternative approach would be for the analyst to simply choose some feasible solution λ to LP problem (13), which would result in utilities $U(\dot{x}^k)$ that lie in the intervals shown in Figure 2(d).

5. Additive and Multiplicative Portfolio Utility Functions

Certain types of preferences satisfying Assumptions 1 and 2 can be represented with a special case of the multilinear portfolio utility function that has a simpler functional form. To characterize such preferences, we utilize preference elicitation question (9) that involves a portfolio with a certain outcome and a portfolio with two possible outcomes. In particular, we first consider a setting in which the DM states that these two portfolios are always equally preferred as long as the two uncertain outcomes are equally likely. Formally, such preferences satisfy the following assumption.

Figure 2. Ranges of Utilities that Portfolios $\dot{x}^k, k \in \{0,\ldots,40\}$ (Equation (12)) Can Obtain After (a) Two, (b) Three, (c) Four, and (d) Five Preference Statements

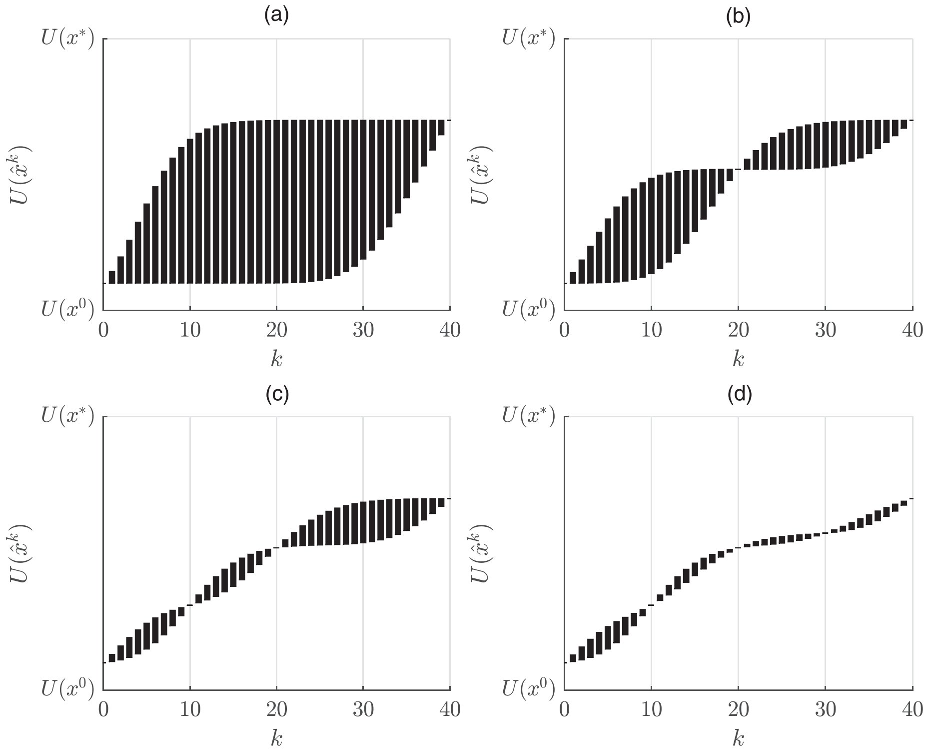

Assumption 3. *The portfolios*

$$x = \Big(\underbrace{y^*,\ldots,y^*}_{k \text{ elements}}, y^0,\ldots,y^0\Big) \text{ and}$$

$$\tilde{x} = \begin{cases} \Big(\underbrace{y^*,\ldots,y^*}_{k-1 \text{ elements}}, y^0,\ldots,y^0\Big) & \text{with probability } 0.5, \\ \Big(\underbrace{y^*,\ldots,y^*}_{k+1 \text{ elements}}, y^0,\ldots,y^0\Big) & \text{with probability } 0.5 \end{cases}$$

are equally preferred for any $k \in \{2,\ldots,m\}$.

Clearly, this is quite a strong assumption. However, this assumption is implicitly made in many multiattribute project portfolio selection applications, because it is necessary for representing the portfolio utility as the sum of project utilities. The following theorem shows that preferences satisfying Assumption 3 in addition to those of Theorem 1 are represented by an additive portfolio utility function.

Theorem 2. *Assumptions 1, 2, and 3 hold if and only if the portfolio utility function U is additive:*

$$U(x) = \sum_{j=1}^{m} u(x_j), \tag{15}$$

where u is the project utility function (7).

The additive portfolio utility function is the symmetric special case of the general nonsymmetric additive multiattribute utility function, which represents preferences if and only if the attributes are additive independent (see, e.g., Keeney and Raiffa 1976, theorem 6.4; Fishburn 1965). In particular, additive independence means that if two uncertain outcomes have equal marginal distributions on each attribute, then the outcomes are equally preferred. Thus, Assumptions 1–3 together imply that projects $x_1,\ldots,x_m$ are additive independent and hence preferences among portfolios depend only on the marginal distribution of each project outcome, not their joint distribution. Indeed, Assumption 1 (symmetric preferences) together with the assumption of additive independent projects would be sufficient for representing preferences with the additive portfolio utility function (15). However, in the portfolio context, Assumptions 2 and 3 are perhaps easier for a DM to verify than the less concrete (but more general) concept of preferences not being affected by the joint distribution of project outcomes.

Theorem 2 raises the question about how the portfolio utility function would change in case the uncertain outcomes in Assumption 3 were not equally likely. In particular, consider a DM who states that the two portfolios are always equally preferred as long as the probabilities of the two outcomes are fixed to some values $p^* \neq \frac{1}{2}$ and $1-p^*$. Such preferences are formally defined by the following assumption.

Assumption 4. *There exists $p^* \in (0, \frac{1}{2}) \cup (\frac{1}{2}, 1)$ such that the portfolios*

$$x = \Big(\underbrace{y^*, \dots, y^*}_{k \text{ elements}}, y^0, \dots, y^0\Big) \text{ and}$$
$$\tilde{x} = \begin{cases} \Big(\underbrace{y^*, \dots, y^*}_{k-1 \text{ elements}}, y^0, \dots, y^0\Big), & \text{with probability } 1-p^*, \\ \Big(\underbrace{y^*, \dots, y^*}_{k+1 \text{ elements}}, y^0, \dots, y^0\Big), & \text{with probability } p^* \end{cases} \tag{16}$$

are equally preferred for any $k \in \{2, \dots, m\}$.

Preferences satisfying this assumption in addition to the assumptions of Theorem 1 are represented by a multiplicative portfolio utility function, as stated by the following theorem.

Theorem 3. *Assumptions 1, 2, and 4 hold if and only if the portfolio utility function U is multiplicative:*

$$U(x) = \frac{1}{\theta}\prod_{j=1}^{m}\big(1 + \theta u(x_j)\big) - \frac{1}{\theta}, \tag{17}$$

where u is the project utility function (7) *and* $\theta = 1/p^* - 2 \in (-1, 0) \cup (0, \infty)$.

The multiplicative portfolio utility function is the symmetric special case of the general nonsymmetric multiplicative multiattribute utility function (Keeney and Raiffa 1976, theorem 6.1), which represents preferences if and only if the attributes are mutually utility independent (i.e., each subset of attributes is utility independent from its complement). Theorem 3 thus shows that Assumptions 1, 2, and 4 together imply that the projects are mutually utility independent. Indeed, Assumption 1 together with the assumption of mutually independent projects would be sufficient to represent preferences with the multiplicative portfolio utility function. However, in the portfolio context, Assumptions 2 and 4 seem more convenient to verify than mutual utility independence, which would require checking the utility independence of subsets of projects.

Compared with the additive portfolio utility function, the multiplicative function contains one additional parameter θ. The value of this parameter depends on the probability p^* in Assumption 4 through formula $\theta = 1/p^* - 2$. Thus, a negative θ corresponds to the case $p^* > 0.5$, and a positive θ corresponds to the case $p^* < 0.5$.

Parameter θ can also be interpreted as a measure that captures the deviation from the additive portfolio utility function. In particular, both additive and multiplicative portfolio utility functions assign the same utility to a portfolio having a single project with the most preferred outcome while the rest of the projects have the least preferred outcome, that is, $U(y^*, y^0, \dots, y^0) = 1$. However, they assign a different utility for a portfolio having two projects with the most preferred outcome: the additive utility of this portfolio is $U(y^*, y^*, y^0, \dots, y^0) = 2$, whereas the multiplicative utility is $U(y^*, y^*, y^0, \dots, y^0) = \frac{1}{\theta}(1+\theta)^2 - \frac{1}{\theta} = 2 + \theta$. Thus, $\theta \in (-1, 0)$ implies a decreasing marginal utility of changing the outcomes of projects from the least to the most preferred level, whereas $\theta \in (0, \infty)$ implies an increasing marginal utility. Moreover, when θ approaches zero, the multiplicative portfolio function approaches the additive portfolio utility function, as the probability $p^* = 1/(\theta + 2)$ approaches 0.5.

6. Optimization Models for Maximizing Expected Portfolio Utility

In standard multiattribute decision problems, the alternative with the highest expected utility can be identified by simply evaluating the expected utility of each of the mutually exclusive decision alternatives. In portfolio problems, however, decision alternatives (i.e., project portfolios) often cannot be explicitly enumerated, but rather they are defined implicitly as subsets of project candidates that satisfy relevant resource and other constraints. Therefore, the problem of identifying the portfolio with the highest expected utility has to be formulated as an optimization model.

In this section, we show that the problem of identifying the portfolio that maximizes the expected multilinear utility under linear constraints can be formulated as a MILP problem. In particular, the section first develops a MILP formulation suitable for a setting in which the project outcomes are stochastically independent, and then extends this formulation to capture stochastic dependencies. Section 6.1 then develops a computationally less demanding formulation for settings in which the portfolio utility function is multiplicative and project outcomes are stochastically dependent. These results contribute to the practical applicability of the developed preference models as they enable the use of standard off-the-shelf optimization software to produce decision recommendations.

To formulate an optimization model for identifying the portfolio with the highest expected utility, we utilize binary decision variables $z = (z_1, \dots, z_m)^T$, where $z_j = 1$ if the jth project candidate is included in the portfolio. A feasible portfolio satisfies $Az \leq B$, where the elements of matrix $A \in \mathbb{R}^{m \times q}$ and vector $B \in \mathbb{R}^q$

include the coefficients of q linear constraints. Such linear constraints can be used to model limited resources (budget), logical dependencies between project decisions (e.g., follow-up projects), and interdependencies between the projects' resource consumption (e.g., synergies), for instance (see, e.g., Liesiö et al. 2008). With this notation, the portfolio that maximizes expected utility is the optimal solution to the nonlinear integer programming (IP) problem

$$\max_{z\in\{0,1\}^m} \mathbb{E}[U(\tilde{x}_1,\ldots,\tilde{x}_m)], \tag{18}$$

$$\tilde{x}_j = \begin{cases} \tilde{x}_j^B \text{ if } z_j = 0, \\ \tilde{x}_j^F \text{ if } z_j = 1, \end{cases} \quad \forall j \in \{1,\ldots,m\}, \tag{19}$$

$$Az \le B, \tag{20}$$

where random variables $\tilde{x}_1^F,\ldots,\tilde{x}_m^F$ capture the uncertain outcomes of the m project candidates in case they are funded and $\tilde{x}_1^B,\ldots,\tilde{x}_m^B$ capture their baseline outcomes (cf. Equation (2)).

Evaluating the objective function of IP problem (18)–(20) by using the definition of the multilinear portfolio utility function (6) requires the enumeration of all 2^m subsets of the set $\{1,\ldots,m\}$. This seems to indicate that there is little hope of building efficient portfolio optimization models, because even the evaluation of the objective function in a deterministic case would require a computational effort that increases exponentially in the number of projects m. Fortunately, $U(x)$ can be evaluated efficiently by using a binomial lattice as described by the following lemma.

Lemma 1. *Let $x = (x_1,\ldots,x_m) \in X$, and assume that $\delta_{j,k}$, $j \in J$, $k \in \{0,1,\ldots,j\}$, solve the system of linear equations:*

$$\begin{aligned} \textit{For } j = 1: \quad & \delta_{1,1} = u(x_1), \\ & \delta_{1,0} = 1 - u(x_1). \\ \textit{For } j = 2,\ldots,m: \quad & \delta_{j,0} = (1-u(x_j))\delta_{j-1,0}, \\ & \delta_{j,k} = (1-u(x_j))\delta_{j-1,k} \\ & \qquad + u(x_j)\delta_{j-1,k-1},\ k = 1,\ldots,j-1, \\ & \delta_{j,j} = u(x_j)\delta_{j-1,j-1}. \end{aligned}$$

Then, $U(x_1,\ldots,x_m) = \sum_{k=0}^m \lambda(k)\delta_{m,k}$.

Lemma 1 allows us to formulate the multilinear utility function as a system of linear constraints. If project outcomes are stochastically independent, then the project-specific expectations can be computed prior to optimization, after which these expectations can be mapped through U to obtain the expected portfolio utility. This is because the expectation of the sum (product) of independent random variables is the sum (product) of the random variables' expectations. As a result, the expected utility maximizing portfolio can be obtained by solving a MILP problem, as stated by the following theorem.

Theorem 4. *Assume that $\tilde{x}_1,\ldots,\tilde{x}_m$ are independent random variables. Then, z^* is an optimal solution to IP problem* (18)–(20) *if and only if z^* is an optimal solution to the MILP problem*

$$\max_{z,\delta} \sum_{k=1}^m \lambda(k)\delta_{m,k} \tag{21}$$

$$Az \le B, \tag{22}$$

$$\delta_{j,k} \le \left(1 - \mathbb{E}\left[u\left(\tilde{x}_j^B\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}_j^B\right)\right]\delta_{j-1,k-1} + z_j,\ j = 1,\ldots,m,\ k = 0,\ldots,j, \tag{23}$$

$$\delta_{j,k} \le \left(1 - \mathbb{E}\left[u\left(\tilde{x}_j^F\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}_j^F\right)\right]\delta_{j-1,k-1} + (1-z_j),\ j = 1,\ldots,m, k = 0,\ldots,j, \tag{24}$$

$$\delta_{j,k} \in [0,1],\ j = 1,\ldots,m,\ k = 0,\ldots,j, \tag{25}$$

$$z \in \{0,1\}^m, \tag{26}$$

where $\delta_{0,0} = 1$ and $\delta_{j,-1} = \delta_{j,j+1} = 0$ for all $j = 0,\ldots,m$.

It is important to highlight that Theorem 4 does not rule out the possibility of the attribute-specific outcomes of an individual project being correlated. Specifically, solving the MILP problem gives the optimal portfolio even if the components of each vector-valued random variable $\tilde{x}_j$ are not stochastically independent. However, Theorem 4 does assume that the outcomes are stochastically independent across projects: the probability of a specific project outcome does not depend on the realized outcomes of other projects. Note that, in case of uncertain baseline outcomes, this assumption also rules out the possibility that the baseline outcomes of projects would be stochastically dependent. Moreover, the probability of a given outcome for a funded project cannot depend on the realizations of other projects' baseline outcomes.

These independence assumptions may not be appropriate in applications where some exogenous uncertainties affect the outcomes of all projects. For instance, revenues from research and development projects can depend on the growth of a specific market, which results in correlated project outcomes. Such dependencies among project outcomes can be modeled by extending the MILP model (21)–(26) to capture the projects' outcomes in several states (cf. scenarios). For this purpose we introduce an integer-valued random variable $\tilde{s}$ that indicates which of the d states is realized. Furthermore, we assume that the states are constructed in such a way that the project outcomes are conditionally independent given the state, that is, $\mathbb{P}(\tilde{x}_j = x_j | \tilde{x}_i = x_i, \tilde{s} = s) = \mathbb{P}(\tilde{x}_j = x_j | \tilde{s} = s)$ when $j \ne i$. This allows us to compute the expected portfolio utility conditioned on a specific state with the same approach that is used in Theorem 4 to compute the portfolio utility across projects whose outcomes are stochastically independent (cf. constraints (23)–(25)). The state-specific conditional expected utilities can then be aggregated in the objective function to maximize the expected portfolio utility across the states. This results in the MILP model

$$\begin{aligned}
&\max_{\substack{z\\ \delta^1,\ldots,\delta^d}} \sum_{s=1}^{d}\sum_{k=1}^{m} \mathbb{P}(\tilde{s}=s)\lambda(k)\delta^s_{m,k},\\
&\delta^s_{j,k} \le \left(1-\mathbb{E}\left[u\left(\tilde{x}^B_j\right)|\tilde{s}=s\right]\right)\delta^s_{j-1,k}\\
&\qquad + \mathbb{E}\left[u\left(\tilde{x}^B_j\right)|\tilde{s}=s\right]\delta^s_{j-1,k-1} + z_j,\\
&\qquad\qquad j=1,\ldots,m, k=0,\ldots,j, s=1,\ldots,d,\\
&\delta^s_{j,k} \le \left(1-\mathbb{E}\left[u\left(\tilde{x}^F_j\right)|\tilde{s}=s\right]\right)\delta^s_{j-1,k}\\
&\qquad + \mathbb{E}\left[u\left(\tilde{x}^F_j\right)|\tilde{s}=s\right]\delta^s_{j-1,k-1} + \left(1-z_j\right),\\
&\qquad\qquad j=1,\ldots,m, k=0,\ldots,j, s=1,\ldots,d,\\
&\delta^s_{j,k}\in[0,1],\ j=1,\ldots,m,\ k=0,\ldots,j, s=1,\ldots,d,\\
&Az\le B,\\
&z\in\{0,1\}^m,
\end{aligned} \tag{27}$$

where $\delta^s_{0,0}=1$ and $\delta_{j,-1}=\delta_{j,j+1}=0$ for all $j=0,\ldots,m$, $s=1,\ldots,d$.

6.1. Optimizing Additive and Multiplicative Portfolio Utility Functions

Computationally less demanding optimization models for maximizing the expected multilinear portfolio utility can be formulated if some additional assumptions regarding preferences and the projects' outcomes hold. In particular, if preferences satisfy Assumption 3, the portfolio utility function is additive (Theorem 2), and hence, the objective function of IP problem (18)–(20) can be written as

$$\begin{aligned}
\mathbb{E}[U(\tilde{x}_1,\ldots,\tilde{x}_m)] &= \mathbb{E}\left[\sum_{i=1}^{n} u(\tilde{x}_j)\right]\\
&= \sum_{i=1}^{n}\mathbb{E}\left[z_j u\left(\tilde{x}^F_j\right) + \left(1-z_j\right)u\left(\tilde{x}^B_j\right)\right]
\end{aligned}$$

even in cases where there are stochastic dependencies among the projects' outcomes. In this case, z^* is an optimal solution to the IP problem (18)–(20) if and only if it is an optimal solution to the integer linear programming (ILP) problem

$$\max_{z\in\{0,1\}^m}\left\{\sum_{j=1}^{m}\left(\mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]-\mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)z_j \;\middle|\; Az\le B\right\}. \tag{28}$$

A similar ILP problem for identifying the most preferred portfolio can be formulated if preferences satisfy Assumption 4, which implies that the portfolio utility function takes a multiplicative form (Theorem 3). However, this formulation requires that the projects' outcomes are stochastically independent, as formally stated by the following theorem.

Theorem 5. *Assume that the portfolio utility function U is multiplicative (see* (17)*) and* $\tilde{x}_1,\ldots,\tilde{x}_m$ *are independent random variables. Then,* z^* *is an optimal solution to IP problem* (18)–(20) *if and only if it is an optimal solution to the ILP problem*

$$\max_{z\in\{0,1\}^m}\left\{\sum_{j=1}^{m}\frac{1}{\theta}\log\left(\frac{1+\theta\mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]}{1+\theta\mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]}\right)z_j \;\middle|\; Az\le B\right\}. \tag{29}$$

It should be noted that the objective function value of ILP problem (29) is not equal to the expected portfolio utility, as was the case for (M)ILP problems (21)–(26), (27), and (28). Theorem 5 only states that the optimal solution of ILP problem (29) corresponds to the portfolio that maximizes the multiplicative portfolio utility function (17). Note also that the arguments of the logarithmic functions in (29) are always positive, even though parameter θ can take also negative values in the interval $(-1,0)$. Moreover, for any project whose expected utility $\mathbb{E}[u(\tilde{x}^F_j)]$ exceeds the expected baseline utility $\mathbb{E}[u(\tilde{x}^B)]$, the corresponding objective function coefficient is positive regardless of the value of parameter $\theta\in(-1,0)\cup(0,\infty)$. In turn, if a project has an expected utility below the baseline utility, then its objective function coefficient is always negative. Hence, from a computational point of view, optimization problem (29) is equivalent to (28) in that both of them are (multiconstraint) knapsack problems.

Extending ILP problem (29) to handle dependent project outcomes through the use of multiple states (cf. MILP problem (27)) is not possible without the model becoming nonlinear. This is because the proof of Theorem 5 relies on taking a logarithm of the multiplicative portfolio utility function to enable additive aggregation across the contributions of each project to the overall portfolio utility. With multiple states, this approach could be used to obtain the logarithm of the conditional expectations of portfolio utility in each state. However, the aggregation of these state-specific logarithmic expected utilities into an overall expected utility would not be a linear mapping. Nevertheless, in problems with stochastically dependent project outcomes and a multiplicative portfolio utility function, the optimal portfolio can be identified by solving MILP problem (27), because the multiplicative portfolio utility function is a special case of the multilinear portfolio utility function.

7. Effect of Multilinear Portfolio Utility Function on Decision Recommendations

In general, the multilinear portfolio utility function—and its special case, the multiplicative portfolio utility function—will yield different expected utilities for the feasible portfolios than the widely employed additive portfolio utility function. However, because the most important output of decision analysis models are the decision recommendations, the key question is whether the portfolios that optimize the expected multilinear utility functions differ from those that maximize additive utility, that is, whether these portfolios consist of different projects. To address this question we conducted two studies: The first study is based on generating random instances of project portfolio selection problems and then comparing the optimal portfolios corresponding to additive, multiplicative, and multilinear portfolio utility functions. The second study applies different multilinear utility functions to real data from healthcare resource allocation.

7.1. Simulation Study

We generated 150 instances of project portfolio selection problems, each with $m = 50$ project candidates and a single budget constraint that limits the total portfolio cost to be at most 50% of the sum of the costs of all candidate projects. The randomly generated project costs c_j and expected utilities $\mathbb{E}[u(\tilde{x}_j^F)]$ have uniform marginal distributions, and their joint distribution is a Gaussian copula. Problem instances were generated using values 0.3, 0.6, and 0.9 for the copula's correlation coefficient ρ. The least preferred outcome y^0 is used as the common baseline y^B, which results in $\mathbb{E}[u(\tilde{x}_j^B)] = u(y^B) = u(y^0) = 0$.

For each problem instance, the following computations were carried out. First, Theorem 5 was utilized to find the portfolios z^θ that maximize the expected multiplicative utility (17) for different values of parameter $\theta \in \{-0.9, -0.8, \ldots, 2.9, 3\} \setminus \{0\}$. Second, problem (28) was solved to find a benchmark portfolio z^A that maximizes the expected additive utility (15). Finally, two measures were computed to capture the difference between the benchmark portfolios z^A and z^θ for different values of θ: (i) the share of different project decisions $\Delta(\theta) = \sum_{j=1}^m |z_j^\theta - z_j^A|/m$, and (ii) the share of the optimal expected multiplicative utility achieved by the benchmark portfolio, that is, $\Delta'(\theta) = EU_\theta(z^A)/EU_\theta(z^\theta)$, where $EU_\theta(z)$ denotes the expected multiplicative utility of portfolio z. Recall that the multiplicative portfolio utility function approaches the additive function when parameter θ approaches zero, whereby we can define $\Delta(0) = 0$ and $\Delta'(0) = 1$.

Figure 3 illustrates the differences between decision recommendations resulting from the additive and multiplicative portfolio utility functions for different values of ρ and θ. In particular, the solid black line shows the share of different project decisions $\Delta(\theta)$ across all problem instances, and the dashed black lines show the 5th and 95th percentiles. Notably, even when the multiplicative portfolio utility function is close to the additive one (i.e., θ is close to zero), the decision recommendations given by the two functions can be different. The gray lines show the share of the optimal expected utility achieved by the

Figure 3. Share of Changed Project Decisions (Black) and Relative Loss in Expected Multiplicative Utility (Gray) When the Portfolio Maximizing the Expected Additive Portfolio Utility Is Selected Instead of the Portfolio Maximizing the Expected Multiplicative Utility with Parameter θ.

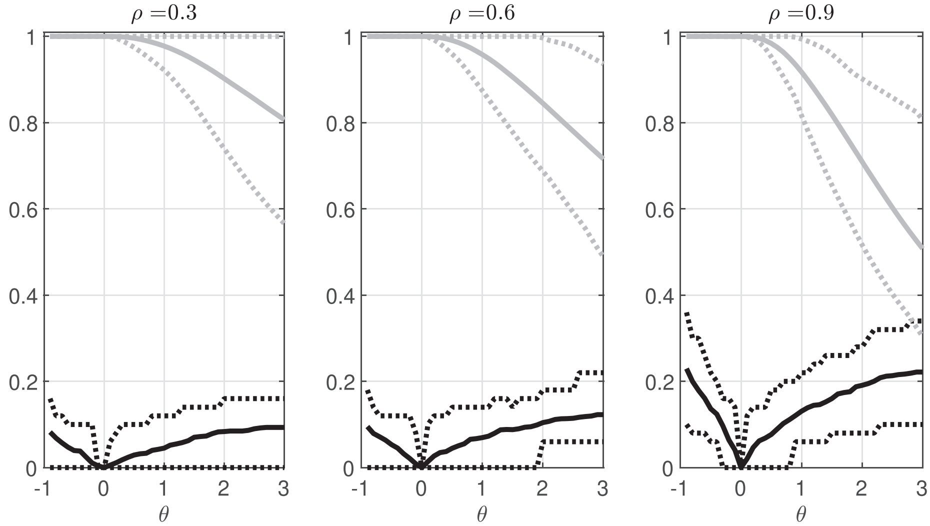

Note. Solid lines show the averages, and dashed lines show the 5th and 95th percentiles.

benchmark portfolio $\Delta'(\theta)$ with the solid line corresponding to the average and the dashed lines corresponding to the 5th and 95th percentiles. For negative values of θ (corresponding to a decreasing marginal utility of changing the outcomes of projects from the least to the most preferred level), the expected multiplicative utility of the benchmark portfolio is very close to that of the optimal portfolio although the composition of the portfolios is different. This indicates that, with negative values of θ, the multiplicative portfolio utility function has a flat optimum in which portfolios close to the optimal portfolio have almost identical expected utilities. For large values of θ and ρ (i.e., when the marginal utility of project outcomes is strongly increasing and the projects' costs and expected utilities are highly correlated), the benchmark portfolio may yield only half of the optimal expected utility. This is because, for higher values of parameter θ, the optimal portfolios contain more projects, that is, projects whose costs are on average lower. The effect becomes more prominent when the projects' costs and expected utilities are highly correlated: With $\rho = 0.3$, the average number of projects in the optimal portfolio increases from approximately 28 to 31 as θ changes from −1 to 3, whereas with $\rho = 0.9$ this number increases from approximately 22 to 32.

Figure 4 shows results from a similar analysis in which decision recommendations resulting from the additive portfolio utility function are compared with those provided by multilinear utility function (6) with λ-parameters specified by the S-shaped sigmoid curve $\lambda(k) = 1/(1+\exp(-\gamma(k-15)))$. Specifically, these results are based on solving 150 problem instances with $m = 30$ projects for different values of parameter $\gamma \in (0, 1]$. When γ approaches zero, then $\lambda(k)$ becomes linear in k, in which case the multilinear portfolio utility function is arbitrarily close to the additive utility function (cf. proof of Theorem 2). In turn, increasing the value of γ results in function $\lambda(k)$ deviating more from the linear function k, thus increasing the difference between the multilinear and additive portfolio utility functions. This difference also affects the composition of the optimal portfolios: the larger the value of γ, the more low-cost projects are included in the optimal portfolio. As was the case with multiplicative portfolio utility functions, this effect becomes stronger as ρ increases.

Results from Figures 3 and 4 suggest that the use of the additive portfolio utility function in cases where preferences do not satisfy Assumption 3 is more likely to yield erroneous decision recommendations and loss in expected portfolio utility when the projects' expected utilities and costs are correlated. This observation is intuitive: If there is little or no correlation between the expected utilities and costs, projects with a low cost and high utility are clear choices to include in the portfolio, and small changes in the portfolio utility function are not likely to change these choices. However, if the utilities and costs are highly correlated, then preference between a high-cost, high-utility project and a project with a low cost and low utility can be highly contingent on the form of the portfolio utility function.

Figure 4. Share of Changed Project Decisions (Black) and Relative Loss in Expected Multilinear Utility (Gray) When the Portfolio Maximizing the Expected Additive Portfolio Utility Is Selected Instead of the Portfolio Maximizing the Expected Multilinear Utility with Parameters $\lambda(k) = 1/(1+\exp(-\gamma(k-15)))$.

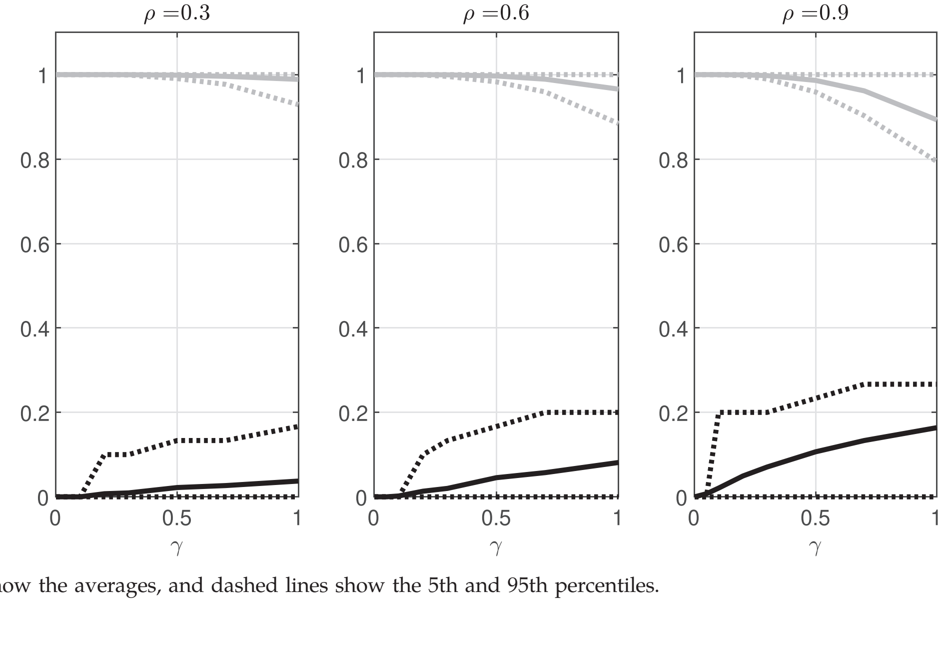

Note. Solid lines show the averages, and dashed lines show the 5th and 95th percentiles.

7.2. Application to Real Data

Besides helping to identify which project proposals to implement, PDA models can also be used to select the appropriate usage levels of resources (e.g., budget). A common approach to this end is to solve the optimal portfolios for a wide range of different budget levels and then to visualize the optimal expected portfolio utility as a function of the budget level. This visualization provides the DMs with clear information about, for instance, the marginal utility of increasing the budget or the amount of utility that would be lost if the budget was cut by a certain amount. Indeed, the user interfaces of many PDA software tools are built around graphs visualizing the efficient budget-utility frontier (see, e.g., Lourenço et al. 2012).

To examine the impact that the choice of the portfolio utility function has on both project-specific decision recommendations and the efficient frontier, we revisit the case on healthcare resource allocation discussed in Sections 2 and 4 (Airoldi et al. 2011; see Table 1). We solved the expected utility maximizing portfolios for all budget levels using three portfolio utility functions: additive, multiplicative, and multilinear. The multiplicative utility function (17) uses parameter value $\theta = -\frac{1}{3}$, which corresponds to the case in which Assumption 4 holds for probability $p^* = 0.6$. For instance, a single project with the most preferred outcome is equally preferred to a portfolio of two projects, both receiving the most preferred outcome with a 60% probability and the least preferred outcome with a 40% probability. For the multilinear utility function (6), the parameters λ are specified using the S-shaped sigmoid function $\lambda(k) = 1/\big(1 + \exp(-(k - 11))\big)$. This corresponds to setting in which the DM assesses the probability p in preference elicitation question (9) to be less than 50% when $k < 11$ and more than 50% when $k > 11$ (cf. risk-seeking for small portfolios, and risk-averse for large portfolios).

Figure 5(a) shows the expected utilities of the optimal portfolios corresponding to the three portfolio utility functions as a function of the budget level. Notably, the choice of the portfolio utility function has a significant impact on the shape of the efficient frontier. For instance, the use of the additive portfolio utility function suggests that a budget of approximately £2.5 million is required to achieve 70% of the benefits that would be obtained by implementing all interventions. However, the use of the multiplicative portfolio utility function suggests that 70% of the maximal benefit can be achieved already with a £1 million budget. This exemplifies that the functional form of the portfolio utility function can have a significant impact on the level of total resource usage if this level is chosen based on the efficient budget-utility frontier.

Figure 5, (b)–(d), shows which projects are included in each of the three optimal portfolios for different budget levels. For instance, project Pneumonia ($j = 1$) with the highest utility-to-cost ratio is included in the optimal portfolio for all three utility functions at all budget levels exceeding its cost of £75K. In turn, project Active treatment ($j = 19$) jumps in and out of the optimal portfolio throughout the budget interval for each of the three utility functions. This is because the project is the least expensive but also has a very low expected utility, whereby it is included in the portfolio only when the budget constraint impedes the selection of another project that would yield higher utility but at a higher cost.

For many budget levels, the composition of the optimal portfolio is independent of which of the three portfolio utility functions is used. The similarity of the optimal portfolios implied by different portfolio utility functions seems intuitive in view of the simulation results of Section 7.1, as the correlation between the projects' costs and expected utilities in these data (Table 1) is approximately 0.37. Furthermore, these differences in project decisions lead to only small differences in the expected portfolio utilities. In particular, suppose that the true preferences are captured by the multilinear utility function, but nevertheless, the portfolio optimizing the expected additive utility function is selected. The worst-case loss in expected multilinear utility of this portfolio choice is some 10% across all budget levels. In turn, if the true preferences are captured by the multiplicative utility function, selecting the additive utility maximizing portfolio results in less than 1% loss in expected multiplicative portfolio utility. Note that the similarity of the optimal portfolios depends on the choices of parameter values θ and γ. If these values were chosen such that the corresponding multilinear/multiplicative utility function deviated more from the additive one, there would generally be larger differences in the optimal portfolios and higher utility losses.

Nevertheless, for some budget levels there are differences between the optimal portfolios. One such instance is at budget level £1.6 million, where each of the three optimal portfolios (marked with asterisks in Figure 5, (b)–(d)) contains projects $j = 1$ (Pneumonia) through $j = 4$ (Prison MH), project $j = 6$ (Workforce development), and project $j = 7$ (Psych therapies) with a combined cost of £625K. Under multiplicative portfolio utility, the remaining budget of £975K is used to fund only two projects $j = 8$ (Early detection and diagnostics) and $j = 10$ (Prevention) with an average cost of £475K and average expected utility of 0.48. On the other hand, using the multilinear portfolio utility function results in allocating the

Figure 5. Expected Utility (a) and Composition (b, c, d) of Optimal Portfolios for Different Budget Levels When Using the Additive (Black), Multiplicative (Dark Gray), and Multilinear (Light Gray) Utility Function

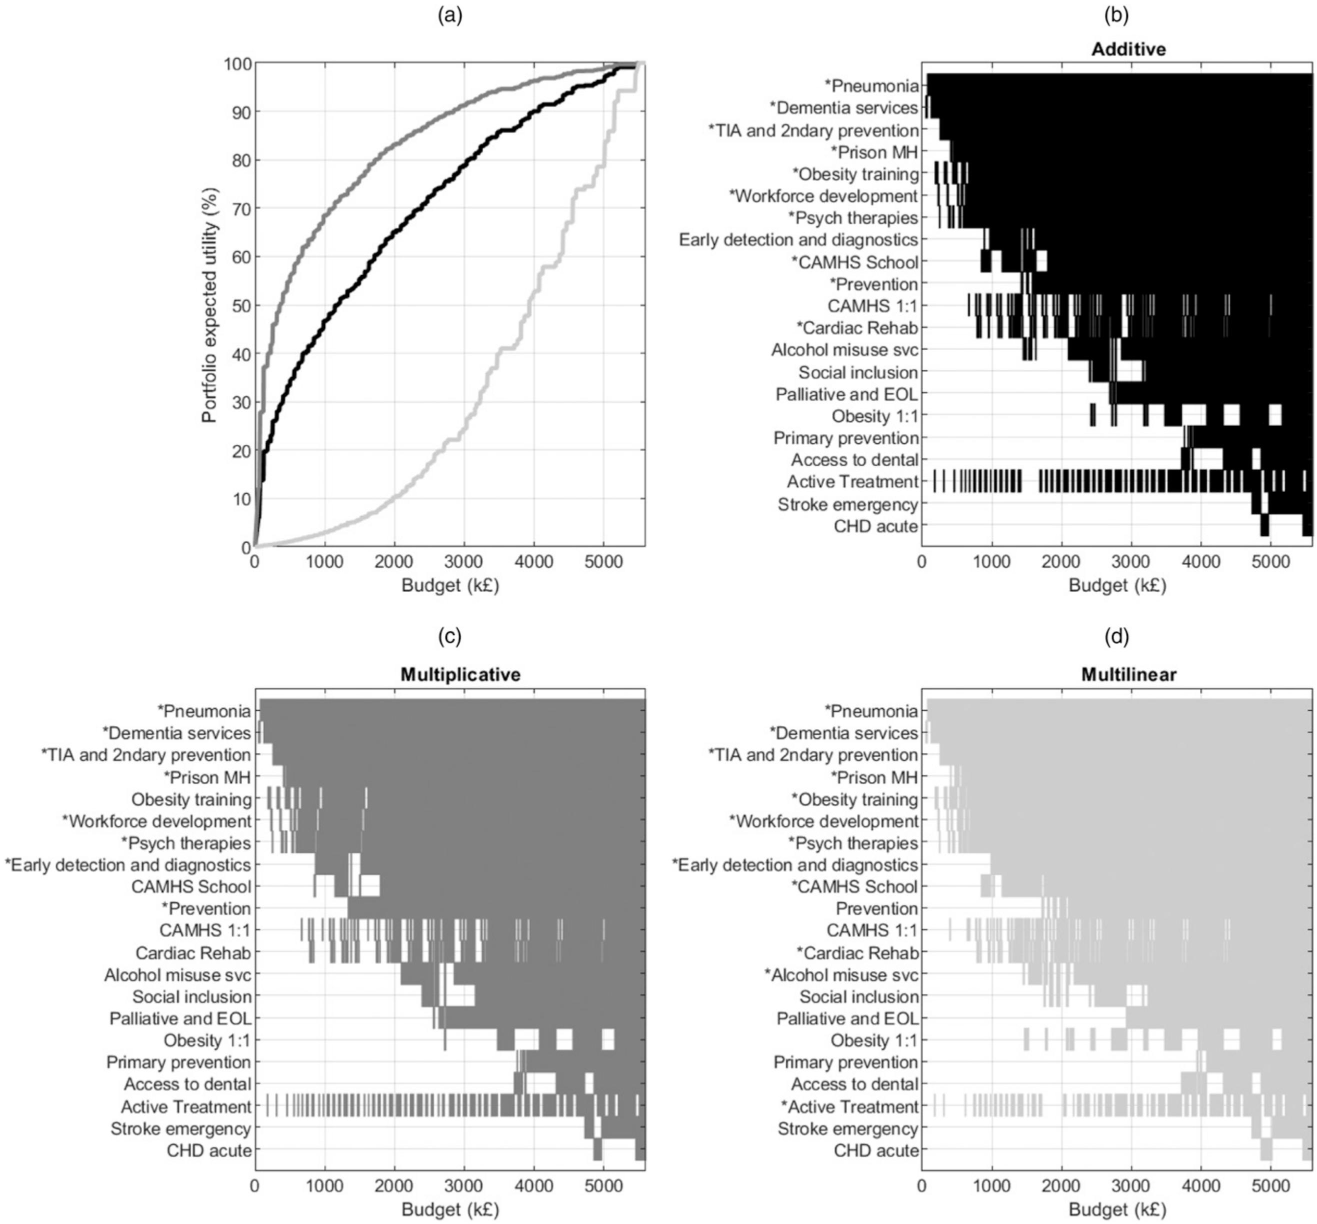

Notes. The budget levels for which a particular project is included in the optimal portfolio are colored. Projects that are included in the optimal portfolio at budget level £1.6 million are indicated with asterisks. The projects are listed in a decreasing order of utility-to-cost ratios $\mathbb{E}[u(\tilde{x}_j^F)]/c_j$.

remaining budget to six projects with an average cost of approximately £162K and average expected utility of 0.15. The optimal portfolio under additive portfolio utility falls between these two extremes: the remaining budget is allocated to four projects, the average cost and expected utility of which are £243K and 0.24, respectively. Interestingly, the multilinear portfolio utility function recommends selecting the combination of three projects $j = 8$ (Early detection and diagnostics), $j = 13$ (Alcohol misuse svc), and $j = 19$ (Active treatment) instead of the expensive project $j = 10$ (Prevention) recommended by the additive model. In this sense, the multilinear utility function captures the preferences of our hypothetical DM in Section 2, who was reluctant to allocate a large share of the budget to a single expensive project.

8. Discussion and Conclusions

This paper advances both the theory and practice of PDA. In particular, the paper develops an axiomatic theory of utility functions that can be used to support project portfolio selection in view of multiple attributes and uncertain project outcomes. This theory shows that if (i) preferences between portfolios are independent of project indexing and (ii) preference between two uncertain project outcomes is not affected by the outcomes of other projects, then these preferences can be represented with a multilinear portfolio utility function. The paper also presents techniques that can be used to assess these portfolio utility functions based on preference statements given by the DM. Moreover, the paper develops optimization models for identifying the portfolio of projects that yields maximal expected utility, when feasible combinations of projects that can be implemented are defined implicitly through resource and other portfolio constraints.

We also established a solid axiomatic foundation for modeling portfolio utility as the sum of the projects' multiattribute utilities; an approach that is widely used in practical applications. In particular, such an additive portfolio utility function is obtained as a special case of the multilinear portfolio utility function when preferences between certain and uncertain portfolio outcomes have a specific structure. Conversely, multilinear portfolio utility functions extend additive portfolio utility functions by providing more flexibility in that they are able to represent a richer variety of preferences.

The developed theory also includes conditions under which preferences can be modeled by multiplicative portfolio utility functions, in which portfolio utility is obtained as a product of scaled project utilities. Compared with the additive portfolio utility function, the multiplicative utility function contains only one additional real-valued parameter to be assessed based on the DM's preferences, but makes it possible to adjust the level of risk-aversion/nonconstant marginal utility at the portfolio level. Moreover, maximizing the expected multiplicative portfolio utility leads to a simple knapsack problem, as long as the projects' outcomes are stochastically independent. Hence, the multiplicative portfolio utility function offers a readily implementable and computationally straightforward alternative for applications in which the preference assumptions underlying the additive portfolio utility function are not appropriate.

The application of the developed portfolio utility functions to real and randomly generated data sets shows that the choice of the utility function can be expected to influence the decision recommendations obtained from PDA. For instance, our results suggest that the use of an additive portfolio utility function in a setting where preferences are compatible with a multiplicative utility function can lead to selecting a portfolio in which one third of the project-specific recommendations are erroneous. This finding is particularly important, because the additive portfolio utility function is commonly used in applications, but little effort seems to be put into verifying whether the preference assumptions underlying this utility function actually hold. The tools developed in this paper can be used to examine the sensitivity of the decision recommendations when a multilinear or a multiplicative portfolio utility function is used instead of an additive one.

This paper opens up several avenues for future research. First, applied research deploying multilinear and multiplicative portfolio utility functions in real-life cases is needed to create understanding on how readily DMs are able to answer the preference assessment questions that are required to specify nonadditive portfolio utility functions. Second, the preference assessment process could benefit from the development of methods to provide decision recommendations based on incomplete preference information (cf. Montiel and Bickel 2014, Fliedner and Liesiö 2016). Finally, the axiomatic theory developed here could be extended to handle a more diverse family of portfolio preferences by further relaxing some of the underlying preference assumptions.

A. Proofs

Proof of Theorem 1. To make the proof compact we will utilize the results of Keeney and Raiffa (1976) that build on the results by Fishburn (1973) and Farquhar (1975).

Assumption 2 implies that x_1 is utility independent from $x_2, \ldots, x_m$. Together Assumptions 1 and 2 imply that each x_j is utility independent from the other projects $x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_m$. Theorem 6.3 of Keeney and Raiffa (1976) then states that U must be multilinear. In particular, we will here utilize the fact that according to equation (6.28) (proof of theorem 6.3 in Keeney and Raiffa (1976)) any multilinear U must satisfy

$$\begin{aligned} &U(x_1, \ldots, x_{j-1}, y, x_{j+1}, \ldots, x_m) \\ &\quad = U\Big(\underbrace{y^0, \ldots, y^0}_{j-1 \text{ elements}}, y, \underbrace{y^0, \ldots, y^0}_{m-j \text{ elements}}\Big) \\ &\qquad \times U(x_1, \ldots, x_{j-1}, y^*, x_{j+1}, \ldots, x_m) \\ &\qquad + \Big(1 - U\Big(\underbrace{y^0, \ldots, y^0}_{j-1 \text{ elements}}, y, \underbrace{y^0, \ldots, y^0}_{m-j \text{ elements}}\Big)\Big) \\ &\qquad \times U(x_1, \ldots, x_{j-1}, y^0, x_{j+1}, \ldots, x_m). \end{aligned}$$

Moreover, Assumption 1 implies that the value of $U(y^0, \ldots, y^0, y, y^0 \ldots, y^0)$ does not depend on the index of y, and hence, we can denote $u(y) = U(y, y^0, \ldots, y^0) = U(y^0, \ldots, y^0, y, y^0 \ldots, y^0)$ to obtain

$$\begin{aligned} &U(x_1, \ldots, x_{j-1}, y, x_{j+1}, \ldots, x_m) \\ &\quad = u(y)U(x_1, \ldots, x_{j-1}, y^*, x_{j+1}, \ldots, x_m) \\ &\qquad + (1 - u(y))U(x_1, \ldots, x_{j-1}, y^0, x_{j+1}, \ldots, x_m), \end{aligned}$$

which holds for any index of y. This equality can be used repeatedly for $U(x_1, \ldots, x_m)$ to obtain

$$\begin{aligned} &U(x_1, \ldots, x_m) \\ &\quad = u(x_1)U(y^*, x_2, \ldots, x_m) + (1 - u(x_1)) \\ &\qquad \times U(y^0, x_2, \ldots, x_m) \\ &\quad = u(x_1)[u(x_2)U(y^*, y^*, x_3, \ldots, x_m) \\ &\qquad + (1 - u(x_2))U(y^*, y^0, x_3, \ldots, x_m)] \\ &\qquad + (1 - u(x_1))[u(x_2)U(y^0, y^*, x_3, \ldots, x_m) \\ &\qquad + (1 - u(x_2))U(y^0, y^0, x_3, \ldots, x_m)] \\ &\quad = u(x_1)u(x_2)U(y^*, y^*, x_3, \ldots, x_m) \\ &\qquad + u(x_1)(1 - u(x_2))U(y^*, y^0, x_3, \ldots, x_m) \\ &\qquad + (1 - u(x_1))u(x_2)U(y^0, y^*, x_3, \ldots, x_m) \\ &\qquad + (1 - u(x_1))(1 - u(x_2))U(y^0, y^0, x_3, \ldots, x_m) \\ &\qquad \ldots \\ &\quad = \sum_{J \subseteq \{1,\ldots,m\}} \Bigg[\prod_{j \in J} u(x_j) \prod_{j \notin J} (1 - u(x_j)) \\ &\qquad \times U(\hat{y}_1(J), \ldots, \hat{y}_m(J)) \Bigg], \end{aligned}$$

where $\hat{y}_j(J) = y^*$ if $j \in J$ and $\hat{y}_j(J) = y^0$ otherwise. Assumption 1 implies $U(\hat{y}_1(J), \ldots, \hat{y}_m(J)) = U(\hat{y}_1(J'), \ldots, \hat{y}_m(J'))$ whenever $|J| = |J'| = k$. Let us denote this utility by $\lambda(k)$. This yields

$$U(x_1, \ldots, x_m) = \sum_{J \subseteq \{1,\ldots,m\}} \lambda(|J|) \prod_{j \in J} u(x_j) \prod_{j \notin J} (1 - u(x_j)).$$

To show that $\lambda(\cdot)$ is strictly increasing, assume that $\lambda(k+1) \le \lambda(k)$ for some $k \in \{0, \ldots, m-1\}$. Then,

$$\begin{aligned} &U\Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\quad \le U\Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\Leftrightarrow \mathbb{E}\Big[U\Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big)\Big] \\ &\quad \le \mathbb{E}\Big[U\Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big)\Big] \\ &\Leftrightarrow \Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big) \preceq \Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\text{Assumption 2} \Rightarrow (y^*, x_2, \ldots, x_m) \preceq (y^0, x_2, \ldots, x_m), \end{aligned}$$

which violates the assumption that y^* is strictly preferred to y^0.

Multilinear utility function $U(x_1, \ldots, x_m)$ is clearly symmetric, which implies that Assumption 1 holds. As U is a special case of the general nonsymmetric multilinear function (Keeney and Raiffa 1976, theorem 6.3), then each x_j is utility independent, and thus, Assumption 2 holds. □

Proof of Theorem 2. Using the general form of multilinear functions used in theorem 6.3 of Keeney and Raiffa (1976) we can write the symmetric multilinear portfolio utility function as

$$\begin{aligned} U(x) &= \sum_{J \subseteq \{1,\ldots,m\}} \kappa(|J|) \prod_{j \in J} u(x_j), \text{ where} \\ \kappa(k) &= \lambda(k) - \sum_{J' \subset \{1,\ldots,k\}} \kappa(|J'|). \end{aligned} \tag{30}$$

If Assumption 3 holds, then for any $k = 2, \ldots, m-1$,

$$\begin{aligned} &U\Big(\underbrace{y^*, \ldots, y^*}_{k \text{ elements}}, y^0, \ldots, y^0\Big) = \frac{1}{2}U\Big(\underbrace{y^*, \ldots, y^*}_{k-1 \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\qquad + \frac{1}{2}U\Big(\underbrace{y^*, \ldots, y^*}_{k+1 \text{ elements}}, y^0, \ldots, y^0\Big) \\ &\Leftrightarrow \lambda(k) = \frac{1}{2}\lambda(k-1) + \frac{1}{2}\lambda(k+1) \\ &\Leftrightarrow \lambda(k+1) = 2\lambda(k) - \lambda(k-1), \end{aligned}$$

which together with $\lambda(0) = U(y^0, \ldots, y^0) = 0$ and $\lambda(1) = U(y^*, y^0, \ldots, y^0) = 1$ implies $\lambda(k) = k$ for $k = 0, \ldots, m$.

Evaluating κ gives $\kappa(0) = \lambda(0) = 0$, $\kappa(1) = \lambda(1) = 1$, and

$$\begin{aligned} \kappa(k) &= \lambda(k) - \Bigg(\sum_{\substack{J' \subset \{1,\ldots,k\} \\ |J'|=1}} \kappa(|J'|) + \sum_{\substack{J' \subset \{1,\ldots,m\} \\ |J'|>1}} \kappa(|J'|) \Bigg) \\ &= k - k \underbrace{\kappa(1)}_{=1} - \sum_{\substack{J' \subset \{1,\ldots,k\} \\ |J'|>1}} \kappa(|J'|) = - \sum_{\substack{J' \subset \{1,\ldots,k\} \\ |J'|>1}} \kappa(|J'|), \end{aligned}$$

which implies $\kappa(k) = 0$ for all $k \in \{2, \ldots, m\}$. Thus,

$$U(x) = \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=1}} \kappa(|J|) \prod_{j \in J} u(x_j) = \sum_{j=1}^{m} \underbrace{\kappa(1)}_{=1} u(x_j).$$

Since $U(x) = \sum_{j=1}^{m} u(x_j)$ is a special case of the multilinear U (6), Assumptions 1 and 2 hold. Assumption 3 holds since

$$\frac{1}{2}U\Big(\underbrace{y^*,\dots,y^*}_{k-1 \text{ elements}},y^0,\dots,y^0\Big)+\frac{1}{2}U\Big(\underbrace{y^*,\dots,y^*}_{k+1 \text{ elements}},y^0,\dots,y^0\Big)$$
$$=\frac{1}{2}(k+1)u(y^*)+\frac{1}{2}(k-1)u(y^*)$$
$$=ku(y^*)=U\Big(\underbrace{y^*,\dots,y^*}_{k \text{ elements}},y^0,\dots,y^0\Big).$$

Proof of Theorem 3. If Assumption 4 holds, then there exists $p^* \in (0,1), p^* \neq \frac{1}{2}$, such that for any $k = 2,\dots,m-1$ it holds that

$$U\Big(\underbrace{y^*,\dots,y^*}_{k \text{ elements}},y^0,\dots,y^0\Big)=(1-p^*)\times U\Big(\underbrace{y^*,\dots,y^*}_{k-1 \text{ elements}},y^0,\dots,y^0\Big)+p^*U\Big(\underbrace{y^*,\dots,y^*}_{k+1 \text{ elements}},y^0,\dots,y^0\Big) \tag{31}$$
$$\Leftrightarrow \lambda(k)=(1-p^*)\lambda(k-1)+p^*\lambda(k+1)$$
$$\Leftrightarrow \lambda(k+1)=\frac{1}{p^*}\lambda(k)-\Big(\frac{1}{p^*}-1\Big)\lambda(k-1)$$
$$\Leftrightarrow \lambda(k+1)=(\theta+2)\lambda(k)-(\theta+1)\lambda(k-1),$$

where $\theta = (1/p^* - 2) \in (-1,\infty)\setminus\{0\}$.

We again utilize presentation (30) by Keeney and Raiffa (1976) and show by induction that (31) implies $\kappa(k) = \theta^{k-1}$ for $k \geq 1$. First, note that $\kappa(0) = \lambda(0) = 0$ and $\kappa(1) = \lambda(1) = 1 = \theta^0$. Now assume that for any $k' \leq k$ it holds that $\kappa(k') = \theta^{k'-1}$. Solving $\lambda(k)$ from (30) gives

$$\lambda(k)=\kappa(k)+\sum_{J'\subset\{1,\dots,k\}}\kappa(|J'|)=\theta^{k-1}+\frac{(1+\theta)^k-1-\theta^k}{\theta}, \tag{32}$$

where the last equality is obtained by using the binomial formula $\sum_{J'\subseteq\{1,\dots,k\}}\theta^{|J'|} = (1+\theta)^k$. Substituting (32) into (31) gives

$$\kappa(k+1)+\frac{(1+\theta)^{k+1}-1-\theta^{k+1}}{\theta}$$
$$=(2+\theta)\Big(\theta^{k-1}+\frac{(1+\theta)^k-1-\theta^k}{\theta}\Big)$$
$$-(1+\theta)\Big(\theta^{k-2}+\frac{(1+\theta)^{k-1}-1-\theta^{k-1}}{\theta}\Big)$$
$$\Leftrightarrow \kappa(k+1)=\theta^k.$$

Substituting $\kappa(k) = \theta^{k-1}$ into (30) gives

$$U(x)=\sum_{J\subseteq\{1,\dots,m\}}\theta^{|J|-1}\prod_{j\in J}u(x_j)$$
$$=\frac{1}{\theta}\prod_{j=1}^{m}\big(1+\theta u(x_j)\big)-\frac{1}{\theta},$$

which is the symmetric special case of the multiplicative utility function (Keeney and Raiffa 1976, theorem 6.1).

Assumptions 1 and 2 hold, since $U(x) = \frac{1}{\theta}\prod_{j=1}^{m}(1 + \theta u(x_j)) - \frac{1}{\theta}$ is a special case of the multilinear U (6). Furthermore, Assumption 4 holds with

$$p^*=\frac{1}{\theta+2}\Leftrightarrow 1-p^*=\frac{\theta+1}{\theta+2}$$

since

$$(1-p^*)U\Big(\underbrace{y^*,\dots,y^*}_{k-1 \text{ elements}},y^0,\dots,y^0\Big)$$
$$+p^*U\Big(\underbrace{y^*,\dots,y^*}_{k+1 \text{ elements}},y^0,\dots,y^0\Big)$$
$$=\frac{\theta+1}{\theta+2}\Big(\frac{1}{\theta}(1+\theta)^{k-1}-\frac{1}{\theta}\Big)$$
$$+\frac{1}{\theta+2}\Big(\frac{1}{\theta}(1+\theta)^{k+1}-\frac{1}{\theta}\Big)$$
$$=\frac{1}{\theta+2}\Big(\frac{1}{\theta}(1+\theta)^{k}-\frac{1+\theta}{\theta}\Big)+\frac{1}{\theta+2}\Big(\frac{1}{\theta}(1+\theta)^{k+1}-\frac{1}{\theta}\Big)$$
$$=\frac{1}{\theta}\frac{1}{\theta+2}\Big((1+\theta)^k-1-\theta+(1+\theta)^{k+1}-1\Big)$$
$$=\frac{1}{\theta}\frac{1}{\theta+2}\Big((1+\theta)^k(\theta+2)-(\theta+2)\Big)$$
$$=\frac{1}{\theta}\Big((1+\theta)^k-1\Big)=U\Big(\underbrace{y^*,\dots,y^*}_{k \text{ elements}},y^0,\dots,y^0\Big).$$

Proof of Lemma 1. Reordering the summation in the multilinear utility (6) gives

$$U(x)=\sum_{J\subseteq\{1,\dots,m\}}\lambda(|J|)\prod_{j\in J}u(x_j)\prod_{j\notin J}\big(1-u(x_j)\big)$$
$$=\sum_{k=0}^{m}\lambda(k)\sum_{\substack{J\subseteq\{1,\dots,m\}\\|J|=k}}\prod_{j\in J}u(x_j)\prod_{j\notin J}\big(1-u(x_j)\big).$$

Hence, the lemma can be proved by showing that if $\delta^*_{j,k}, j \in \{1,\dots,m\}, k \in \{0,1,\dots,j\}$, is a solution to the system of linear inequalities

$$\delta_{1,0}=1-u(x_1), \tag{33}$$
$$\delta_{1,1}=u(x_1), \tag{34}$$
$$\delta_{j,0}=\big(1-u(x_j)\big)\delta_{j-1,0}, \tag{35}$$

$$\delta_{j,k} = (1 - u(x_j))\delta_{j-1,k} + u(x_j)\delta_{j-1,k-1}, \quad k = 1, \ldots, j-1, \tag{36}$$

$$\delta_{j,j} = u(x_j)\delta_{j-1,j-1}, \tag{37}$$

then

$$\delta^*_{m,k} = \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=k}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)). \tag{38}$$

We show this by using induction. First, let $m = 1$, in which case the system consists of constraints (33)–(34). Then,

$$\delta^*_{1,0} = 1 - u(x_1) = \sum_{\substack{J \subseteq \{1\} \\ |J|=0}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)),$$
$$\delta^*_{1,1} = u(x_1) = \sum_{\substack{J \subseteq \{1\} \\ |J|=1}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)),$$

that is, (38) holds. Now, assume (38) holds for some m and all $k = 0, \ldots, m$. Then, for $m+1$, (35) gives

$$\begin{aligned}\delta_{m+1,0} &= (1 - u(x_{m+1}))\delta_{m,0} = (1 - u(x_{m+1})) \times \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=0}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &= (1 - u(x_{m+1})) \prod_{j\notin \{1,\ldots,m\}} (1 - u(x_j)) \\ &= \sum_{\substack{J \subseteq \{1,\ldots,m+1\} \\ |J|=0}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)).\end{aligned}$$

Moreover, (36) gives for each $k = 1, \ldots, m$

$$\begin{aligned}\delta_{m+1,k} &= (1 - u(x_{m+1}))\delta_{m,k} + u(x_{m+1})\delta_{m,k-1} \\ &= (1 - u(x_{m+1})) \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=k}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &\quad + u(x_{m+1}) \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=k-1}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &= \sum_{\substack{J \subseteq \{1,\ldots,m+1\} \\ |J|=k,\ (m+1)\notin J}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &\quad + \sum_{\substack{J \subseteq \{1,\ldots,m+1\} \\ |J|=k,\ (m+1)\in J}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &= \sum_{\substack{J \subseteq \{1,\ldots,m+1\} \\ |J|=k}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)).\end{aligned}$$

Finally, (37) gives

$$\begin{aligned}\delta_{m+1,m+1} &= u(x_{m+1})\delta_{m,m} = u(x_{m+1}) \times \sum_{\substack{J \subseteq \{1,\ldots,m\} \\ |J|=m}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)) \\ &= u(x_{m+1}) \prod_{j\in \{1,\ldots,m\}} u(x_j) \\ &= \sum_{\substack{J \subseteq \{1,\ldots,m+1\} \\ |J|=m+1}} \prod_{j\in J} u(x_j) \prod_{j\notin J}(1 - u(x_j)).\end{aligned}$$

Hence, (38) holds for $m+1$ and any $k = 1, \ldots, m+1$ and, therefore, by induction for any m. □

Proof of Theorem 4. For any $z \in \{z \in \{0,1\}^m \mid Az \le B\}$, the objective function value for IP problem (18)–(20) is

$$\mathbb{E}[U(\tilde{x}_1, \ldots, \tilde{x}_m)] \text{ where } \tilde{x}_j = \begin{cases} \tilde{x}^B_j & \text{if } z_j = 0, \\ \tilde{x}^F_j & \text{if } z_j = 1. \end{cases}$$

We show that optimizing the MILP problem (21)–(26) with respect to δ provides an equal objective function value for z. Since the objective function coefficient of $\delta_{m,k}$ is $\lambda(k) > 0$, variables $\delta_{m,1}, \ldots, \delta_{m,m}$ are to be maximized. Furthermore, variables $\delta_{j,1}, \ldots, \delta_{j,j}$, $j = 1, \ldots, m$, are bounded from above by the values of variables $\delta_{j-1,1}, \ldots, \delta_{j-1,j-1}$ through constraints

$$\delta_{j,k} \le \left(1 - \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\delta_{j-1,k-1} + z_j,$$
$$\delta_{j,k} \le \left(1 - \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\delta_{j-1,k-1} + (1 - z_j).$$

Hence, the optimal values of these variables are given by

$$\begin{aligned}\delta^*_{j,k} &= \min\Big\{\left(1 - \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\delta_{j-1,k-1} + z_j, \left(1 - \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\delta_{j-1,k} \\ &\qquad + \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\delta_{j-1,k-1} + (1 - z_j)\Big\} \\ &= \begin{cases} \left(1 - \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\delta_{j-1,k-1} & \text{if } z_j = 0, \\ \left(1 - \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\delta_{j-1,k-1} & \text{if } z_j = 1. \end{cases}\end{aligned}$$

For $j = 1$, this results in linear constraints

$$\delta^*_{1,0} = 1 - \mathbb{E}[u(\tilde{x}^B_1)],$$
$$\delta^*_{1,1} = \mathbb{E}[u(\tilde{x}^B_1)]$$

if $z_1 = 0$ and linear constraints

$$\delta^*_{1,0} = 1 - \mathbb{E}[u(\tilde{x}^F_1)],$$
$$\delta^*_{1,1} = \mathbb{E}[u(\tilde{x}^F_1)]$$

if $z_1 = 1$ (since $\delta_{0,0} = 1, \delta_{0,-1} = \delta_{0,1} = 0$). Similarly, for any $j \in \{2, \ldots, m\}$, the optimal variables δ^* satisfy linear constraints

$$\delta^*_{j,0} = \left(1 - \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\delta_{j-1,0},$$
$$\delta^*_{j,k} = \left(1 - \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\delta_{j-1,k-1}, \quad k = 1, \ldots, j-1,$$
$$\delta^*_{j,j} = \mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\delta_{j-1,j-1}$$

if $z_j = 0$ and linear constraints

$$\delta^*_{j,0} = \left(1 - \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\delta_{j-1,0},$$
$$\delta^*_{j,k} = \left(1 - \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\delta_{j-1,k} + \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\delta_{j-1,k-1},$$
$$k = 1, \ldots, j-1,$$
$$\delta^*_{j,j} = \mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\delta_{j-1,j-1}$$

if $z_j = 1$ (since $\delta_{j-1,-1} = 0, \delta_{j-1,j} = 0$). Thus, δ^* is a solution to the system of linear inequalities of Lemma 1, whereby

$$\begin{aligned}\sum_{k=1}^{m}\lambda(k)\delta^*_{m,k} &= \sum_{k=0}^{m}\lambda(k)\delta^*_{m,k}\\ &= \sum_{J\subseteq\{1,\ldots,m\}}\lambda(|J|)\prod_{j\in J}\mathbb{E}[u(\tilde{x}_j)]\prod_{j\notin J}(1-\mathbb{E}[u(\tilde{x}_j)])\\ &= \mathbb{E}[U(\tilde{x}_1,\ldots,\tilde{x}_m)],\end{aligned}$$

where the last equality holds since U is multilinear. Specifically, $U(\tilde{x}_1, \ldots, \tilde{x}_m)$ is obtained by summing products of independent random variables $u(\tilde{x}_1), \ldots, u(\tilde{x}_m)$, and the expected value of the product of independent random variables is equal to the product of the expected values of these random variables. □

Proof of Theorem 5. Let us denote $Z = \{z \in \{0,1\}^m \mid Az \leq B\}$ and

$$\tilde{x}_j = \begin{cases}\tilde{x}^B_j \text{ if } z_j = 0,\\ \tilde{x}^F_j \text{ if } z_j = 1.\end{cases}$$

The optimal solution to IP problem (18)–(20) with a multiplicative utility function is

$$\begin{aligned}z^* &\in \arg\max_{z\in Z}\mathbb{E}\left[\frac{1}{\theta}\prod_{j=1}^{m}(1+\theta u(\tilde{x}_j)) - \frac{1}{\theta}\right]\\ &= \arg\max_{z\in Z}\frac{1}{\theta}\prod_{j=1}^{m}(1+\theta\mathbb{E}[u(\tilde{x}_j)]) - \frac{1}{\theta},\end{aligned}$$

since random variables $\tilde{x}_1, \ldots, \tilde{x}_m$ are independent. Because logarithm is a monotonically increasing function,

$$\begin{aligned}z^* &\in \arg\max_{z\in Z}\frac{1}{\theta}\log\left(\prod_{j=1}^{m}(1+\theta\mathbb{E}[u(\tilde{x}_j)])\right)\\ &= \arg\max_{z\in Z}\frac{1}{\theta}\sum_{j=1}^{m}\log(1+\theta\mathbb{E}[u(\tilde{x}_j)])\\ &= \arg\max_{z\in Z}\frac{1}{\theta}\sum_{j=1}^{m}\left(z_j\log\left(1+\theta\mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\right.\\ &\quad\left. + (1-z_j)\log\left(1+\theta\mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\right)\\ &= \arg\max_{z\in Z}\frac{1}{\theta}\sum_{j=1}^{m}z_j\left(\log\left(1+\theta\mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]\right)\right.\\ &\quad\left. - \log\left(1+\theta\mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]\right)\right)\\ &= \arg\max_{z\in Z}\frac{1}{\theta}\sum_{j=1}^{m}z_j\log\left(\frac{1+\theta\mathbb{E}\left[u\left(\tilde{x}^F_j\right)\right]}{1+\theta\mathbb{E}\left[u\left(\tilde{x}^B_j\right)\right]}\right).\end{aligned}$$

Acknowledgments

The authors thank James E. Smith, the associate editor, and the referees for their excellent suggestions to help improve the paper.

References

Abbas AE (2009) Multiattribute utility copulas. *Oper. Res.* 57(6): 1367–1383.

Airoldi M, Morton A, Smith J, Bevan G (2011) Healthcare prioritisation at the local level: A socio-technical approach. Working paper, University of Oxford.

Argyris N, Figueira J, Morton A (2011) Interactive multicriteria methods for portfolio decision analysis. Salo A, Keisler J, Morton A, eds. *Portfolio Decision Analysis: Improved Methods for Resource Allocation* (Springer, Berlin), pp. 107-130.

Argyris N, Morton A, Figueira J (2014) CUT: A multicriteria approach for concavifiable preferences. *Oper. Res.* 62(3):633–642.

Barbati M, Greco S, Kadzinski M, Slowinski R (2018) Optimization of multiple satisfaction levels in portfolio decision analysis. *Omega* 78:192–204.

Bordley RF, Kirkwood CW (2004) Multiattribute preference analysis with performance targets. *Oper. Res.* 52(6):823–835.

Clemen RT, Smith J (2009) On the choice of baselines in multiattribute portfolio analysis: A cautionary note. *Decision Anal.* 6(4):256–262.

Dyer JS, Sarin RK (1979) Measurable multiattribute value functions. *Oper. Res.* 27(4):810–822.

Ewing PL Jr, Tarantino W, Parnell GS (2006) Use of decision analysis in the army base realignment and closure (BRAC) 2005 military value analysis. *Decision Anal.* 3(1):33–49.

Farquhar PH (1975) A fractional hybercube decomposition theorem for multiattribute utility functions. *Oper. Res.* 23(5):941–967.

Fishburn PC (1965) Independence in utility theory with whole product sets. *Oper. Res.* 13(1):28–45.

Fishburn PC (1973) Bernoullian utilities for multiple-factor situations. Cochrane JL, Zeleny M, eds. *Multiple Criteria Decision Making* (University of South Carolina Press, Columbia, SC), pp. 47-61.

Fliedner T, Liesiö J (2016) Adjustable robustness for multi-attribute project portfolio selection. *Eur. J. Oper. Res.* 252(3):931–946.

Golabi K, Kirkwood CW, Sicherman A (1981) Selecting a portfolio of solar energy projects using multiattribute preference theory. *Management Sci.* 27(2):174–189.

Grushka-Cockayne Y, De Reyck B, Degraeve Z (2008) An integrated decision-making approach for improving european air traffic management. *Management Sci.* 54(8):1395–1409.

Gurgur C, Morley C (2008) Lockheed Martin Space Systems company optimizes infrastructure project-portfolio selection. *Interfaces* 38(4): 251–262.

Keeney RL, Raiffa H (1976) *Decisions with Multiple Objectives: Preferences and Value Trade-Offs* (John Wiley & Sons, New York).

Keller LR, Simon J (2019) Preference functions for spatial risk analysis. *Risk Anal.* 39(1):244–256.

Kleinmuntz DN (2007) Resource allocation decisions. Edwards W, Miles RF, von Winterfeld D, eds. *Advances in Decision Analysis* (Cambridge University Press, New York), 400–418.

Liesiö J (2014) Measurable multiattribute value functions for portfolio decision analysis. *Decision Anal.* 11(1):1–20.

Liesiö J, Mild P, Salo A (2008) Robust portfolio modeling with incomplete cost information and project interdependencies. *Eur. J. Oper. Res.* 190(3):679–695.

Liesiö J, Punkka A (2014) Baseline value specification and sensitivity analysis in multiattribute project portfolio selection. *Eur. J. Oper. Res.* 237(3):946–956.

Lopes Y, de Almeida A (2015) Assessment of synergies for selecting a project portfolio in the petroleum industry based on a multi-attribute utility function. *J. Petroleum Sci. Engrg.* 126:131–140.

Lourenço JC, Morton A, Bana e Costa CA (2012) PROBE – A multicriteria decision support system for portfolio robustness evaluation. *Decision Support Systems* 54(1):534–550.

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

Montibeller G, Franco LA, Lord E, Inglesias A (2009) Structuring resource allocation decisions: A framework for building multi-criteria portfolio models with area-grouped options. *Eur. J. Oper. Res.* 199(3):846–856.

Montiel LV, Bickel JE (2014) A generalized sampling approach for multilinear utility functions given partial preference information. *Decision Anal.* 11(3):147–170.

Morton A (2015) Measurement issues in the evaluation of projects in a project portfolio. *Eur. J. Oper. Res.* 245(3):789–796.

Morton A, Keisler JM, Salo A (2016) Multicriteria portfolio decision analysis for project selection. Greco S, Ehrgott M, Figueira J, eds. *Multiple Criteria Decision Analysis* (Springer, New York), pp. 1269-1298.

Parnell GS, Bennett GE, Engelbrecht JA, Szafranski R (2002) Improving resource allocation within the national reconnaissance office. *Interfaces* 32(3):77–90.

Phillips L, Bana e Costa CA (2007) Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Ann. Oper. Res.* 154(1):51–68.

Salo A, Keisler J, Morton A (2011) *Portfolio Decision Analysis: Improved Methods for Resource Allocation* (Springer, New York).

Simon J, Kirkwood CW, Keller LR (2014) Decision analysis with geographically varying outcomes: Preference models and illustrative applications. *Oper. Res.* 62(1):182–194.

von Neumann J, Morgenstern O (1947) *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, NJ).

Juuso Liesiö is a tenured associate professor of management science at Aalto University School of Business. His research interests include axiomatic preference models for portfolio selection and spatial decision making as well as computational approaches for handling incomplete preference information and uncertainties in decision support. Recent application areas include environmental decision making, infrastructure asset management, and procurement optimization.

Eeva Vilkkumaa is an assistant professor of management science at Aalto University School of Business. Her research interests include the development and application of mathematical methods to support decision making and resource allocation under uncertainty. Recent application areas include the cost-effective allocation of healthcare resources and scenario-based strategy development. Dr. Vilkkumaa is also the current president of the Finnish Operations Research Society.