Liesiö, Juuso; Kee, Taeyoung; Malo, Pekka

Modeling project interactions in multiattribute portfolio decision analysis : Axiomatic foundations and practical implications

Published in:
European Journal of Operational Research

DOI:
10.1016/j.ejor.2024.02.015

Published: 01/08/2024

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

European Journal of Operational Research 316 (2024) 988–1000

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Modeling project interactions in multiattribute portfolio decision analysis: Axiomatic foundations and practical implications

Juuso Liesiö *, Taeyoung Kee, Pekka Malo

Department of Information and Service Management, Aalto University School of Business, Finland

ARTICLE INFO

Keywords:
Decision analysis
Multiattribute utility theory
Project portfolio selection
Project interactions

ABSTRACT

A common approach to modeling project interactions in multiattribute project portfolio selection is to augment the additive portfolio utility function, in which portfolio utility is the sum of the included projects' utilities, with additional terms representing the synergy/cannibalization effects triggered by selecting specific subsets of projects. In this paper we develop a set of sufficient and necessary assumptions for representing preferences among multiattribute project portfolios with a quasi-symmetric multilinear utility function and show how this function gives rise to interpreting interaction effects as additional terms in the additive portfolio utility function. To foster practical applicability of these theoretical contributions, we also develop techniques to elicit such portfolio utility functions as well as optimization models to identify the feasible portfolio that satisfies relevant resource and other constraints with the maximal expected utility. In recognition that incorporating project interactions necessitates increased involvement of decision makers in assessing the interaction effects and results in computationally more challenging portfolio optimization problems, we analyze the importance of modeling interactions through series of simulation studies based on randomly generated and real-world data sets. Specifically, we examine the impact that omitting project interactions has on the project-level decision recommendations and on the expected utility of the recommended portfolio.

1. Introduction

Organizations typically seek to achieve their objectives through selecting and implementing a portfolio of projects such as investments, actions, or initiatives. These portfolio decisions are often made in the presence of multiple objectives and limited availability of monetary, human, and other resources to carry-out the projects (see, e.g., Barbati et al. 2018, Grushka-Cockayne et al. 2008, Kleinmuntz 2007, Lopes and de Almeida 2015, Mavrotas and Makryvelios 2021, Montibeller et al. 2009, Phillips and Bana e Costa 2007). This has motivated the development theory, methods, and practices commonly referred to as portfolio decision analysis (PDA; Liesiö et al. 2021, Morton et al. 2016, Salo et al. 2010) to (i) capture projects' multiattribute outcomes, (ii) model decision makers' preferences among uncertain multiattribute outcomes, and (iii) produce decision recommendations as to the optimal project portfolio composition through the use of mathematical optimization. Perhaps the most common approach is to use the additive portfolio value/utility function (see, e.g., Golabi et al. 1981, Mild et al. 2015, Parnell et al. 2002), in which portfolio utility/value is modeled as the sum of the multi-attribute utilities/values of those projects that are included in the portfolio. Such an additive model does not require any additional preference elicitation on top of specifying the multiattribute project value/utility function and makes it possible to identify the optimal project portfolio through standard integer linear programming (ILP). Moreover, such additive portfolio preference models are well-motivated from the decision-theoretic perspective because they have rigorous axiomatic foundations both in the case of deterministic and uncertain outcomes (Clemen and Smith 2009, Golabi et al. 1981, Liesiö 2014, Liesiö and Punkka 2014, Liesiö and Vilkkumaa 2021, Morton 2015).

A significant shortcoming with the additive portfolio preference models is that they do not allow for interactions among the projects. Specifically, the preference model does not provide mechanisms through which the utility/value of a specific combination of projects could be higher or lower than the sum of the individual projects' utilities/values if these projects have synergy or cannibalization effects. The lack of such mechanisms has motivated the development of ad hoc approaches to handle interactions. Usually this involves addition of dummy projects whose utility/value captures the magnitude of the interaction effect. These dummy projects do not require any resources to be implemented, but additional portfolio feasibility constraints are used to ensure that they are included in the portfolio if and only if the combination of actual projects that triggers the interactions is also included (see,

* Corresponding author.
E-mail addresses: juuso.liesio@aalto.fi (J. Liesiö), taeyoung.kee@aalto.fi (T. Kee), pekka.malo@aalto.fi (P. Malo).

https://doi.org/10.1016/j.ejor.2024.02.015
Received 22 June 2023; Accepted 11 February 2024
Available online 12 February 2024

e.g., Liesiö et al. 2008, Stummer and Heidenberger 2003). By utilizing this approach, the optimization model remains linear, although one additional binary decision variable and generally two additional constraints are required for each project interaction. Moreover, if the interaction affects only attributes with natural measurement scales (e.g., monetary value) and its effects can be quantified on these scales (e.g., implementing a pair of projects increases the monetary value by an additional \$100k), then this approach is also theoretically sound since the interactions are related to the (objective) outcomes rather than the (subjective) preferences.

This notwithstanding, PDA applications often utilize attributes that capture less tangible qualities of projects such as their strategic fit or the competencies of the project team (see, e.g., Abbassi et al. 2014, Eilat et al. 2006). In such cases, it can be difficult to estimate the interaction effects on the constructed scales, let alone assign the total synergy effect to some specific (subset of) attributes. Thus, it is more common to rely on holistic estimates of what the effects would be on the level of overall project utilities/values. From a computational perspective, such holistic estimates can be handled in exactly the same manner as those related to specific attribute outcomes. However, from a decision-theoretic perspective such holistic estimates of the magnitude of the interaction effects are problematic as there is no axiomatic foundation that would justify introduction of interaction effects as additional terms to an additive portfolio utility/value function. Thus, it is unknown what are the sufficient and necessary assumptions that the decision maker's preferences must satisfy for them to be representable with such an augmented additive portfolio value/utility function. Moreover, without a solid decision-theoretic foundations, it is difficult to develop preference assessment techniques and processes to quantify the interaction parameters since it is not clear how these parameters are linked to preferences between portfolios.

This paper lays out an axiomatic foundation for the commonly used approach of adding project interaction terms to additive portfolio preferences models. In particular, we first develop a set of preference assumption such that preferences among uncertain portfolio outcomes satisfy these assumptions if and only if they are represented by a quasi-symmetric multilinear portfolio utility function. We then show that this quasi-symmetric multilinear portfolio utility function can be interpreted as an additive portfolio utility function with extra terms capturing the effects of project interactions.

In addition to these theoretical contributions, we also study the practical implications of the quasi-symmetric multilinear portfolio utility function. In particular, we develop preference assessment approaches that enable the theoretically sound quantification of synergy and cannibalization effects based on the decision maker's preference between (uncertain) portfolio outcomes. We also develop stochastic ILP models to maximize the expected quasi-symmetric multilinear portfolio utility subject to linear portfolio feasibility constraints, which can also be utilized if the projects' outcomes are not stochastically independent.

As the third contribution this paper examines the project interactions' influence on the composition and expected utility of the most preferred portfolio. In particular, as the identification and quantification of the interactions in most applications requires intensive involvement from the decision makers, it seems unrealistic to assume all interactions can always be included in the portfolio model, especially since the number of possible interaction grows exponentially with the number of projects. Therefore, we present computational experiments based on randomly generated and real-world problem instances to analyze how the true optimal portfolio obtained when one considers all interaction effects differs from those portfolios obtained when only the largest interaction effects are considered. The results from these experiments suggest that the prevailing practice of using an additive portfolio utility/value function augmented with extra terms capturing the ex ante most significant interactions can in some cases offer a reasonable approximate approach for the real-world applications in which the time and effort that can be devoted to estimating the project interactions is limited.

Our research links to several strands of literature that are important to recognize here. First, our theoretical development utilizes the work of Fishburn (1974) and Fishburn and Keeney (1975) on multiattribute utility theory, or MAUT. This work examines the implications of relaxing commonly deployed utility independence assumptions and replacing them with the less restrictive assumption of generalized utility independence. Second, Liesiö (2014) and Liesiö and Vilkkumaa (2021) develop axiomatic preference models for multiattribute project portfolios under deterministic and uncertain outcomes, respectively. However, these models assume that preferences are fully symmetric among the projects, which prohibits the existence of the types of project interactions studied in this paper. Third, our computational experiments are motivated partly by the research of Durbach et al. (2020) into the performance of behavioral decision heuristics in project portfolio selection under project interactions. Finally, this paper adds to the growing literature on the applications of multilinear preference models (Abbas 2009, Bordley and Kirkwood 2004, Keller and Simon 2019, Montiel and Bickel 2014).

The rest of the paper is structured as follows. Section 2 introduces the use of the additive portfolio utility function in decision support and the standard approach of augmenting this function to incorporate project interactions. Section 3 develops the axiomatic basis for the quasi-symmetric multilinear portfolio utility function. Section 4 describes approaches for assessing the parameters of this portfolio utility function to quantify project interactions. Section 5 formulates optimization models to identify the most preferred portfolio when preferences are captured by the quasi-symmetric portfolio utility function. Section 6 reports on computational experiments to analyze the effect of interactions on the utility and composition of the optimal portfolio. Section 7 concludes.

2. The additive portfolio utility function and its extension to capture project interactions

We consider a decision setting in which a portfolio is selected from m project candidates, with indices $j \in J = \{1, \dots, m\}$, and each project candidate is evaluated with regard to n attributes. The measurement scales for these attributes, denoted by $Y_1, \dots, Y_n$, can correspond to quantitative measures (e.g., present value or sales; see, e.g., Solak et al. 2010) or qualitative evaluations (e.g., expertise of the project team or quality of the proposal; see, e.g., Chowdhury and Quaddus 2015, Clemen and Smith 2009, Kleinmuntz 2007). The uncertain outcome of the jth project is denoted by $\tilde{x}_j$, which is a random variable whose vector-valued realizations $y = (y_1, \dots, y_n)$ belong to the set

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

This set is assumed to include the baseline outcome $y^B \in Y$ that is obtained if a project is not selected (Clemen and Smith 2009, Liesiö and Punkka 2014).

Each portfolio thereby corresponds to a random variable $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_m)$ with realizations $x = (x_1, \dots, x_m)$ in the set

$$X = \underbrace{Y \times \cdots \times Y}_{m \text{ sets}}, \tag{2}$$

where Y is given by (1) and we use $\mathcal{X}$ to denote the set of all such random variables $\tilde{x}$, i.e., portfolios. The notation $x \in \mathcal{X}$ is used for portfolios that have a single deterministic outcome $x \in X$ (cf. a degenerate random variable) and notation x_j to highlight that the outcome of the jth project is deterministic. As an example, the portfolio $\tilde{x} = (x_1, y^B, \tilde{x}_3, y^B, \dots, y^B) \in \mathcal{X}$ consists of the first and third project candidates, where the outcome of the first project is equal to $x_1 \in Y$ with probability of one.

If preferences among the portfolios $\mathcal{X}$ satisfy specific assumption (for details, see Liesiö and Vilkkumaa 2021), they can be represented by the additive portfolio utility function

$$U(x) = U(x_1, \dots, x_m) = \sum_{j=1}^{m} u(x_j), \tag{3}$$

where $u : Y \to \mathbb{R}$ is the project-level utility function that maps multiattribute project outcomes in Y to a single-dimensional utility. For instance, the additive multiattribute utility function $u(y) = \sum_{i=1}^{n} w_i u_i(y_i)$ is often deployed in practice as it requires assessing only the marginal utility function $u_i : Y_i \to \mathbb{R}$ and the importance weight w_i for each attribute $i \in \{1, \dots, n\}$. If the projects' outcomes are deterministic, an additive portfolio value function (Golabi et al. 1981, Liesiö 2014) can be used, which has a functional form equivalent to (3).

The portfolio that maximizes expected additive utility can be identified through integer linear programming (ILP) as long as the relevant portfolio constraints, such as budget, are linear. This requires introducing a vector of binary decision variables $z = (z_1, \dots, z_m)$, where $z_j = 1$ if the jth project is included in the portfolio and $z_j = 0$ otherwise. Then the most preferred portfolio corresponds to the optimal solution to the ILP problem

$$\max_{z\in\{0,1\}^m} \left\{ \sum_{j=1}^{m} \big(z_j \mathbb{E}[u(\tilde{x}_j)] + (1 - z_j)u(y^B)\big) \middle| Az \le B \right\}, \tag{4}$$

where the elements of matrix $A \in \mathbb{R}^{q\times m}$ and vector $B \in \mathbb{R}^q$ code the parameters of the q portfolio feasibility constraints.

To illustrate the notation consider an environmental agency that is selecting projects to restore a recreational area closer to its natural state.[1] Suppose there are $m = 10$ project candidates with implementation costs $c_1, \dots, c_{10}$ and a limited budget b to fund the projects. The agency uses $n = 3$ attributes that capture the projects' impacts on greenhouse gas emissions ($i = 1$), biodiversity ($i = 2$) and recreational use ($i = 3$). Specifically, the attributes measure the reduced annual greenhouse gas emissions in CO_2 equivalent tons ($Y_1 = [0, 100]$), the number of new species for which the habitat would become favorable ($Y_2 = \{0, 1, \dots, 50\}$) and the increase in recreational use potential ($Y_3 = \{\text{'none'}, \text{'modest'}, \text{'high'}\}$). An additive multiattribute utility function $u(y) = \sum_{i=1}^{3} w_i u_i(y_i)$ is specified by utilizing standard decision analysis techniques to assesses attribute-specific utility functions $u_1 : [0, 100] \to [0, 1]$, $u_2 : \{0, 1, \dots, 50\} \to [0, 1]$ and $u_3 : \{\text{'none'}, \text{'modest'}, \text{'high'}\} \to [0, 1]$ as well as attribute weights $(w_1, w_2, w_3) \in [0, 1]^3$ such that $\sum_{i=1}^{3} w_i = 1$. Moreover, the utility function is scaled so that $u(0, 0, \text{'none'}) = 0$ and $u(100, 50, \text{'high'}) = 1$. Estimating the outcomes of the jth project candidate thus yields a vector $x_j \in Y_1 \times Y_2 \times Y_3$ in the case of deterministic outcomes, or, in the general case, a vector-valued random variable $\tilde{x}_j$ characterized by a probability distribution over $Y_1 \times Y_2 \times Y_3$. The baseline outcome, $y^B = (0, 0, \text{'none'})$, is obtained if a project is not implemented. The project portfolio that maximizes the expected additive portfolio utility is identified by solving the ILP problem $\max_{z\in\{0,1\}^{10}} \{\sum_{j=1}^{10} \mathbb{E}[u(\tilde{x}_j)]z_j \mid \sum_{j=1}^{10} c_j z_j \le b\}$, where $\mathbb{E}[u(\tilde{x}_j)] = \mathbb{E}[u(\tilde{x}_{j1}, \dots, \tilde{x}_{j3})] = \sum_{i=1}^{3} w_i \mathbb{E}[u_i(\tilde{x}_{ji})]$. The simpler form of the objective function is a result of scaling the utility function such that the utility of the baseline outcome is zero, i.e., $u(y^B) = 0$.

Clearly, the additive portfolio utility function (3) implicitly assumes that there are no project interactions. This is evident also in optimization problem (4), where the added utility of including the jth project into the portfolio (i.e., changing the value of z_j from zero to one) is always equal to $\mathbb{E}[u(\tilde{x}_j)] - u(y^B)$ irrespective of which other projects are included in the portfolio (i.e., the values of z_k, $k \ne j$). Commonly, project interactions are handled by directly extending optimization problem (4) rather than replacing the additive portfolio utility function (3) with a suitable non-additive utility function (see, e.g., Carazo et al. 2010, Li et al. 2020, Perez and Gomez 2016, Stummer and Heidenberger 2003). To illustrate this approach, let us suppose there are K interactions in total and that the kth interaction is triggered if all the projects with indexes in the set $J_k \subseteq \{1, \dots, m\}$ are included in the portfolio. Furthermore, suppose that the kth interaction, if triggered, changes the additive portfolio utility by s_k utility units. Extending the optimization problem (4) to incorporate such interactions yields the ILP problem

$$\max_{\substack{z\in\{0,1\}^m \\ \zeta\in\{0,1\}^K}} \sum_{j=1}^{m} \big(z_j \mathbb{E}[u(\tilde{x}_j)] + (1 - z_j)u(y^B)\big) + \sum_{k=1}^{K} \zeta_k s_k \tag{5}$$

$$Az \le B \tag{6}$$

$$m\zeta_k \le \sum_{j\in J_k} z_j - |J_k| + m \;\forall\; k \in \{1, \dots, K\} \tag{7}$$

$$m\zeta_k \ge \sum_{j\in J_k} z_j - |J_k| + 1 \;\forall\; k \in \{1, \dots, K\}. \tag{8}$$

This optimization problem introduces one additional binary decision variable ζ_k for each of the K interactions and $2K$ additional constraints which ensure that $\zeta_k = 1$ if and only if all the projects with indexes in J_k are selected. Specifically, if $\zeta_k = 1$, then for constraint (7) to hold all projects with indexes in J_K must be included in the portfolio (i.e., $\sum_{j\in J_K} z_j = |J_k|$). In turn, if $z_j = 1$ for all $j \in J_k$, then constraint (8) is equivalent to $m\zeta_k \ge 1$, which holds only if the kth interaction is included in the objective function (i.e., $\zeta_k = 1$).

To demonstrate the application of this approach in handling interactions we revisit the environmental management example. Suppose that projects $j = 5$, $j = 8$, and $j = 9$ affect areas that are geographically close to each other and the agency predicts that implementing all three projects would thus increase the recreational value even more than suggested by the outcomes of the individual projects' outcomes with regard to attribute $i = 3$. To quantify this effect, the agency estimates that implementing these three projects would increase the portfolio utility by an amount equal to the utility of a hypothetical project with the outcome $y = (0, 0, \text{'high'})$. In this case, the project portfolio that maximizes the expected portfolio utility corresponds to the optimal solution of the ILP problem

$$\max_{\substack{z\in\{0,1\}^{10} \\ \zeta_1\in\{0,1\}}} \sum_{j=1}^{10} \mathbb{E}[u(\tilde{x}_j)]z_j + \zeta_1 s_1$$

$$\sum_{j=1}^{10} c_j z_j \le b$$

$$z_5 + z_8 + z_9 - 2 \le 10\zeta_1 \le z_5 + z_8 + z_9 + 7,$$

where $s_1 = u(0, 0, \text{'high'})$.

Although this approach is relatively straightforward in terms of implementing the model as an ILP problem, from a decision analytic viewpoint it has both theoretical and practical shortcomings. First, it is not clear what is the functional form of the portfolio utility function whose expectation optimization model (5)–(8) seeks to maximize. Second, without a precisely determined functional form it is not possible to specify the preference assumptions under which the most preferred portfolio actually corresponds to the optimal solution of optimization model (5)–(8). Third, the lack of theoretical foundations, which link a particular set of preference assumptions to a family of portfolio utility functions representing preferences that satisfy these assumptions, makes it difficult to interpret the parameters s_k that quantify the interaction effects. Such an interpretation is required for developing rigorous techniques to assess these parameters based on the decision makers' preference statements. To address these shortcomings the following sections offer an axiomatic foundation for portfolio utility functions that incorporate project interactions and utilize this foundation to develop approaches for preference assessment and the identification of the most preferred portfolio.

[1] This illustrative example is loosely based on the real-world application of Marttunen et al. (2023).

3. An axiomatic basis for portfolio utility functions capturing project interactions

Decision maker'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 that relation $\succeq$ satisfies a set of standard axioms, such as those developed by von Neumann et al. (1944) or (Savage, 1972), there exists a utility function $U : \mathcal{X} \to \mathbb{R}$ that represents the decision maker's preferences $\succeq$ in the sense that

$$\mathbb{E}[U(\tilde{x})] \geq \mathbb{E}[U(\tilde{x}')] \Leftrightarrow \tilde{x} \succeq \tilde{x}' \ \forall \tilde{x}, \tilde{x}' \in \mathcal{X}. \tag{9}$$

This utility function is unique up to positive affine transformations, and hence it can be scaled by fixing the utility of any two portfolios that are not equally preferred. Here we will use a scaling in which an empty portfolio (i.e., all projects have the baseline outcomes y^B) receives the zero utility and a portfolio consisting of a single project with a specific outcome $y^+ \in Y$ receives a unit utility. However, the standard utility axioms cannot rule out the possibility that these two portfolios are equally preferred or that preference for the latter portfolio depends on the index of the project that receives the outcome y^+. Hence, before formally fixing the scaling we need to introduce three preference assumptions.

The first preference assumption is motivated by the fact that in portfolio decision analysis the decision alternatives are combinations of projects and hence the preference model should provide equal treatment for all project candidates. Specifically, preference for a particular project should not be contingent on which index the project candidate receives in the preference model. To illustrate this assumption suppose that the environmental management example includes a project candidate with the outcome $y = (30, 40, \text{‘high’})$. In this case, a portfolio x consisting of this project alone, i.e., $x = (y^B, \ldots, y^B, y, y^B, \ldots, y^B)$, should be equally preferred regardless of the index $j \in \{1, \ldots, 10\}$ assigned to the project.

Assumption 1. Single project preferences are independent from project indexing in the sense that portfolios $x = (y, y^B, \ldots, y^B)$ and $x' = (y^B, \ldots, y^B, y, y^B, \ldots, y^B)$ are equally preferred for any $y \in Y$.

Note that the index of the project that receives outcome y in portfolio x' is arbitrary and the assumption explicitly restricts only preferences among deterministic project outcomes. Assumption 1 is a weaker version of Assumption 1 in Liesiö and Vilkkumaa (2021), which requires that any portfolio remains equally preferred if the outcomes of its projects are permuted. This stronger assumption is also implied by the additive portfolio utility function (3). Here, we assume only that such permutations do not affect the preferability of portfolios consisting of a single project, which does not rule out the possibility of interactions between projects. In the environmental management example, for instance, Assumption 1 allows the portfolio $(y, y, y^B, y^B \ldots, y^B)$ to be preferred to the portfolio $(y, y^B, y, y^B, \ldots, y^B)$, thereby allowing synergistic interactions between projects $j = 1$ and $j = 2$.

The second preference assumption can be seen as a requirement that projects yield utility as independent entities. More specifically, it rules out the possibility that a single project would be entirely worthless itself and produce utility only if combined with some other projects. Obviously, this assumption does not rule out interactions (i.e., that more utility is produced if the project is a part of such combinations), but it ensures that evaluating preferences among outcomes of an individual project makes sense.

Assumption 2. There exists a project outcome that is preferred to the baseline, i.e.,

$$(y^+, y^B, \ldots, y^B) \succ (y^B, y^B, \ldots, y^B).$$

Note that it is sufficient to require that an outcome preferred to the baseline exists only for the first project, as Assumption 1 then implies that the same outcome is preferred to the baseline for any project. In the environmental management example, for Assumption 2 to hold, there must be a project outcome $y \in [0, 100] \times \{0, 1, \ldots, 50\} \times \{\text{‘none’}, \text{‘modest’}, \text{‘high’}\}$ such that a portfolio consisting of a single project with this outcome (i.e., $x = (y, y^B, \ldots, y^B)$) is preferred to an empty portfolio (i.e., $x = (y^B, y^B, \ldots, y^B)$).

Instead of Assumption 2 one could require that the portfolios $(y^+, y^B, \ldots, y^B)$ and $(y^B, y^B, \ldots, y^B)$ are not equally preferred for some y^+ and obtain a functional form for the portfolio utility function equivalent to the one being developed here. However, this would require working with two alternative ways of scaling the portfolio utility function (i.e., $U(y^+, y^B, \ldots, y^B) = \pm 1$), which we want to avoid here for the sake of clarity.

The final assumption restricts the way in which preferences among portfolios that differ with regard to the uncertain outcome of a single project are affected by changes in the project outcomes that are common to all portfolios. In particular, suppose that, given two pairs of alternative uncertain outcomes for the jth project (e.g., $\tilde{x}_j^a$, $\tilde{x}_j^b$ and $\tilde{x}_j^c$, $\tilde{x}_j^d$) there is a strict preference for one of the outcomes in both pairs (e.g., $\tilde{x}_j^a$ and $\tilde{x}_j^c$) when all the other projects in the portfolio have some fixed deterministic outcomes ($x_k, k \neq j$). Now, if these fixed outcomes are changed, we require that (i) the preferred outcome in both pairs remain the same, (ii) the preferred outcome is changed in both pairs, or (iii) all outcomes become equally preferred. Essentially, this assumption forbids a situation where a change in the fixed outcomes reverses the preference order in one pair but not in the other.

Assumption 3. The preference order of uncertain project outcomes remains the same, it is fully reversed, or all of these outcomes become equally preferred when the outcomes of other projects are changed, i.e., if

$(\ldots, x_{j-1}, \tilde{x}_j^a, x_{j+1}, \ldots) \succ (\ldots, x_{j-1}, \tilde{x}_j^b, x_{j+1}, \ldots)$ and
$(\ldots, x_{j-1}, \tilde{x}_j^c, x_{j+1}, \ldots) \succ (\ldots, x_{j-1}, \tilde{x}_j^d, x_{j+1}, \ldots)$,

then for any $x_k' \in Y, k \neq j$, one of the following holds:

(*i*) $(\ldots, x_{j-1}', \tilde{x}_j^a, x_{j+1}', \ldots) \succ (\ldots, x_{j-1}', \tilde{x}_j^b, x_{j+1}', \ldots)$ and $(\ldots, x_{j-1}', \tilde{x}_j^c, x_{j+1}', \ldots) \succ (\ldots, x_{j-1}', \tilde{x}_j^d, x_{j+1}', \ldots)$,

(*ii*) $(\ldots, x_{j-1}', \tilde{x}_j^a, x_{j+1}', \ldots) \prec (\ldots, x_{j-1}', \tilde{x}_j^b, x_{j+1}', \ldots)$ and $(\ldots, x_{j-1}', \tilde{x}_j^c, x_{j+1}', \ldots) \prec (\ldots, x_{j-1}', \tilde{x}_j^d, x_{j+1}', \ldots)$,

(*iii*) $(\ldots, x_{j-1}', \tilde{x}_j^a, x_{j+1}', \ldots) \sim (\ldots, x_{j-1}', \tilde{x}_j'', x_{j+1}', \ldots)$ for any $\tilde{x}_j''$.

Assumption 3 is less restrictive than the standard assumption of each project being utility independent of the others, which is required by the additive portfolio utility function (3) and the non-additive portfolio utility functions developed by Liesiö and Vilkkumaa (2021). These earlier utility models represent preferences that always satisfy condition (i). Thus, Assumption 3 should be acceptable in most real-world applications of multiattribute portfolio decision analysis.

To demonstrate the types of preferences allowed by Assumption 3, we revisit the environmental management example. Suppose the environmental agency prefers adding a project with the outcome $x_1 = (80, 30, \text{‘high’})$ into an empty portfolio rather than a project with the outcome $x_1' = (50, 40, \text{‘high’})$ (i.e., $(x_1, y^B, \ldots, y^B) \succ (x_1', y^B, \ldots, y^B) \succ (y^B, y^B, \ldots, y^B)$). If the agency also prefers adding x_1 into a non-empty portfolio with project outcomes $x_2, \ldots, x_{10}$ (i.e., $(x_1, x_2, \ldots, x_{10}) \succ (y^B, x_2, \ldots, x_{10})$), then condition (i) implies that adding x_1' would also be preferred, although not as strongly as adding x_1 (i.e., $(x_1, x_2, \ldots, x_{10}) \succ (x_1', x_2, \ldots, x_{10}) \succ (y^B, x_2, \ldots, x_{10})$). In turn, if there are interaction effects such that the agency would not prefer to include x_1 into this non-empty portfolio, then condition (ii) implies also preference for not including x_1' (i.e., $(y^B, x_2, \ldots, x_{10}) \succ (x_1', x_2, \ldots, x_{10}) \succ (x_1, x_2, \ldots, x_{10})$). Thus, Assumption 3 allows for interactions among projects but prohibits these interactions from making arbitrary changes to the preference order of single project outcomes.

Preferences among portfolios that satisfy Assumptions 1–3 are represented by an multilinear quasi-symmetric portfolio utility function as formally established by the following theorem.

Theorem 1. *Assumption 1, 2, and 3 hold if and only if the portfolio utility function $U : X \to \mathbb{R}$ is multilinear and quasi-symmetric, i.e.,*

$$U(x_1,\dots,x_m) = \sum_{J\subseteq\{1,\dots,m\}} \kappa(J)\prod_{j\in J} u(x_j), \tag{10}$$

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

$$u(y) = U(y^B,\dots,y^B,y,y^B,\dots,y^B) \tag{11}$$

and $\kappa : 2^{\{1,\dots,m\}} \to \mathbb{R}$ satisfies

$$\sum_{J'\subseteq J} \kappa(J') = U(\chi(J))\ \forall\ J \subseteq \{1,\dots,m\}, \tag{12}$$

where $\chi(J) = x \in X$ is a portfolio outcome such that $x_j = y^+$ if $j \in J$ and $x_j = y^B$ otherwise. Thus, $\kappa(\emptyset) = U(y^B,\dots,y^B) = u(y^B) = 0$ and $\kappa(\{j\}) = U(y^B,\dots,y^B,y^+,y^B,\dots,y^B) = u(y^+) = 1$.

The detailed proof for the theorem is provided in the appendix, but the overall proof strategy can be summarized as follows. To establish the sufficiency of the preference assumptions, the first step is to interpret each project outcome as an attribute. Then Assumption 3 implies that each of these attributes is generalized utility independent of the others, which in turn implies that preferences across the attributes are represented by a quasi-additive utility function (Fishburn and Keeney 1975, Theorem 4). Assumption 2 forbids the case where all outcomes of a single-project portfolio are equally preferred and thus the utility function must be multilinear. Finally, Assumption 1 implies that the multilinear utility function must be symmetric with regard to project-specific outcomes, i.e., quasi-symmetric. In turn, the necessity of Assumptions 1–3 requires lengthy but relatively straightforward algebraic manipulations of the quasi-symmetric multilinear utility function, which are detailed in the appendix.

The quasi-symmetric multilinear portfolio utility function (10) captures preferences through two parameters: the project utility function u that maps the space of multiattribute project outcomes Y to real values and the interaction function κ that associates a real number with every project subset $J \subseteq \{1,\dots,m\}$. Intuitively, the project utility function captures the utilities of project outcomes when there are no other projects in the portfolio (i.e., all other projects have the baseline outcome y^B). The chosen scaling thus fixes the utilities of two project outcomes y^B and y^+, since $u(y^B) = U(y^B,\dots,y^B) = 0$ and $u(y^+) = U(y^+,y^B,\dots,y^B) = 1$.

The interpretation for the project interaction function κ can be established through Eq. (12). Evaluating this equation for $J=\emptyset$ yields $\sum_{J'\subseteq\emptyset}\kappa(J') = \kappa(\emptyset) = U(y^B,\dots,y^B) = 0$, where the last equality follows from the chosen scaling of the utility function. In turn, evaluating Eq. (12) for set $J = \{j\}$, where $j \in \{1,\dots,m\}$, gives $\sum_{J'\subseteq\{j\}}\kappa(J') = \kappa(\emptyset)+\kappa(\{j\}) = U(y^B,\dots,y^B,y^+,y^B,\dots,y^B)$, which, together with the fact that $\kappa(\emptyset) = 0$ and the chosen scaling, implies $\kappa(\{j\}) = 1$. Thus, $\kappa(\emptyset)$ corresponds to the utility of an empty portfolio and $\kappa(J)$, where $|J| = 1$, to the utility of a portfolio with a single project having the deterministic outcome y^+. When $x_j = x_{j'} = y^+$, evaluating (12) for set $J = \{j, j'\}$ results in

$$\begin{aligned}\sum_{J'\subseteq\{j,j'\}} \kappa(J') &= \kappa(\{j,j'\})+\kappa(\{j\})+\kappa(\{j'\})+\underbrace{\kappa(\emptyset)}_{=0}\\ &= U(y^B,\dots,y^B,x_j,y^B,\dots,y^B,x_{j'},y^B,\dots,y^B)\\ \Leftrightarrow \kappa(\{j,j'\}) &= U(y^B,\dots,y^B,x_j,y^B,\dots,y^B,x_{j'},y^B,\dots,y^B) - \underbrace{(\kappa(\{j\})+\kappa(\{j'\}))}_{=1+1=2}.\end{aligned}$$

This intuitively shows that $\kappa(\{j,j'\})$ captures the synergy/cannibalization effect that needs to be added to the sum of the project utilities $\kappa(\{j\})+\kappa(\{j'\}) = 2$ to obtain the portfolio utility $U(y^B,\dots,y^B,x_j,y^B,\dots,y^B,x_{j'},y^B,\dots,y^B)$. In general, consider the portfolio $\chi(J)$ consisting of projects with indexes $J \subseteq \{1,\dots,m\}$, each having the outcome y^+. Then $\kappa(J)$ captures the synergy/cannibalization effect that is obtained only if all projects in the set J are included in the portfolio, i.e., it would not be realized if any projects were removed from the portfolio. Formally, this can be seen from rewriting Eq. (12) in the form

$$\kappa(J) = U(\chi(J)) - \sum_{J'\subset J}\kappa(J').$$

The quasi-symmetric portfolio utility function provides theoretical justification for the commonly suggested approach of extending the sum of project utilities with additional terms to capture interaction effects (cf. ILP problem (5)–(8)). In particular, utilizing the fact that $\kappa(\{j\}) = 1$ for all $j \in \{1,\dots,m\}$ makes it possible to write the quasi-symmetric multilinear portfolio utility function as

$$U(x_1,\dots,x_m) = \sum_{j=1}^{m} u(x_j) + \sum_{\substack{J\subseteq\{1,\dots,m\}\\|J|\ge 2}} \kappa(J)\prod_{j\in J} u(x_j). \tag{13}$$

For those non-singleton project subsets $J \subseteq \{1,\dots,m\}$ that do not yield any interaction effects it holds that $\kappa(J) = 0$ and, hence, the corresponding term in (13) would be equal to zero. Moreover, if the jth project is not included in the portfolio, then $x_j = y^B$, and thus each term that corresponds to a project subset J that includes j would evaluate to zero because one of the multipliers is $u(y^B) = 0$.

4. Assessment of quasi-symmetric multilinear portfolio utility functions

Specification of the quasi-symmetric multilinear portfolio utility function (Theorem 1) requires assessment of the multiattribute project utility function u and the interaction function κ. The project utility function u can be assessed with standard approaches (see, e.g., Keeney and Raiffa 1976) by assuming that all other projects in the portfolio are fixed to some specified levels. Perhaps the most intuitive choice is to assume that the portfolio does not contain any other projects, i.e., the outcomes of all other projects are equal to y^B.

Specification of the interaction parameters $\kappa(J)$, $J \subseteq \{1,\dots,m\}$, can be carried out by asking the decision maker to state her preferences between specific pairs of portfolios. We develop here preference elicitation processes in which the values of $\kappa(\cdot)$ are specified iteratively by starting from parameters capturing interactions between pairs of projects, then moving on to interactions involving three projects, four projects, etc. At each stage, the elicitation questions involve portfolios consisting of projects with indexes in specific subsets of $\{1,\dots,m\}$, but these questions do not require that the deterministic outcomes of the projects take specific values (e.g., $x_j = y^+$). Instead, the project outcomes of the portfolios being compared can be freely chosen from the outcomes of the actual project candidates in the application. Arguably, it is more convenient for the decision maker to compare portfolios with realistic project outcomes rather than be forced to consider hypothetical portfolio outcomes where each project included obtains the same fixed outcome y^+, for instance.

To formalize such a preference elicitation process, we use X_J to denote the set of outcomes of portfolios that consist of projects with indexes in the set J, or

$$X_J = \{x \in X \mid x_j = y^B\ \forall\ j \notin J\}. \tag{14}$$

Moreover, we use $(x', x''; p, 1-p) \in \mathcal{X}$ to denote an uncertain binary portfolio outcome that yields outcome $x' \in X$ with probability p and outcome $x'' \in X$ with probability $(1-p)$, or, more formally,

$$(x',x'';p,1-p) = \begin{cases} x'', & \text{with probability } 1-p\\ x', & \text{with probability } p.\end{cases} \tag{15}$$

Assessment of $\kappa(J)$ for some $J \subseteq \{1,\dots,m\}$ such that $|J| \ge 2$ (recall that $\kappa(\{j\}) = 1$) can be operationalized by choosing two subsets of projects $J', J'' \subset J$, such that $|J'| = |J''| = |J| - 1$, and deterministic outcomes of three portfolios that consist of projects in these subsets, denoted by $x \in X_J$, $x' \in X_{J'}$, and $x'' \in X_{J''}$ such that $x' \nsim x''$. The decision maker is then asked to provide a preference ranking of these three portfolio outcomes and to consider an uncertain binary portfolio outcome that yields either the least or most preferred of these ranked outcomes (see Eq. (15)). Finally, the decision maker adjusts the probability p in this binary outcome until it is equally preferred to the portfolio outcome ranked second. This preference statement can be used to solve for the value of parameter $\kappa(J)$ for any ranking of original three portfolio outcomes as established by the following theorem.

Theorem 2. *Take any $J \subseteq \{1, \dots, m\}$, $|J| \geq 2$, and $J', J'' \subset J$, $J' \neq J''$, such that $J = J' \cup \{j'\} = J'' \cup \{j''\}$, and consider portfolios $x \in X_J$, $x' \in X_{J'}$, and $x'' \in X_{J''}$. Then the following equivalences hold:*

(i) $x \succeq x' \succ x''$ *and* $x' \sim (x, x''; p, 1-p)$ *for some* $p \in (0, 1]$ *if and only if*

$$\kappa(J) = \frac{-\sum_{\substack{I \subseteq J'' \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j) + (1-p) \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subset J \\ j', j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{p \prod_{j \in J} u(x_j)}, \tag{16}$$

(ii) $x' \succ x \succ x''$ *and* $x \sim (x', x''; p, 1-p)$ *for some* $p \in (0, 1)$ *if and only if*

$$\kappa(J) = \frac{(p-1) \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subseteq J'' \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - \sum_{\substack{I \subset J \\ j', j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{\prod_{j \in J} u(x_j)}, \tag{17}$$

(iii) $x' \succ x'' \succeq x$ *and* $x'' \sim (x', x; p, 1-p)$ *for some* $p \in [0, 1)$ *if and only if*

$$\kappa(J) = \frac{-\sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) + p \sum_{\substack{I \subseteq J'' \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - (1-p) \sum_{\substack{I \subset J \\ j', j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{(1-p) \prod_{j \in J} u(x_j)}. \tag{18}$$

Eqs. (16)–(18) are linear in parameters $\kappa(\cdot)$. Thus, a collection of these types of preference statements would result in a system of linear constraints that could be solved with appropriate techniques to specify the values of these parameters. However, note that the right-hand sides of Eqs. (16)–(18) contain only parameters $\kappa(I)$ in which I is a proper subset of J. Thus, if $\kappa(I)$ has been assessed for each $I \subset J$, then these equations provide a formula for directly computing the value of $\kappa(J)$. Therefore, an elicitation process that starts with the assessment of interactions between all pairs of projects and then at each stage increases the size of the considered interaction subsets by one project, avoids the need to actually solve a system of linear equations to estimate the values of parameters $\kappa(\cdot)$.

To illustrate this preference assessment process, let us consider a portfolio selection problem with $m = 4$ candidate projects with utilities $u(x_1) = 0.7$, $u(x_2) = 0.5$, $u(x_3) = 0.3$, and $u(x_4) = 0.6$. The decision maker is first asked to rank the three portfolio outcomes (x_1, x_2, y^B, y^B), (x_1, y^B, y^B, y^B), and (y^B, x_2, y^B, y^B). Suppose the decision maker states that $(x_1, x_2, y^B, y^B) \succ (x_1, y^B, y^B, y^B) \succ (y^B, x_2, y^B, y^B)$. The decision maker is then asked to adjust probability p until the portfolios

$$(x_1, y^B, y^B, y^B) \text{ and } \tilde{x} = \begin{cases} (y^B, x_2, y^B, y^B), & \text{with probability } p \\ (x_1, x_2, y^B, y^B), & \text{with probability } 1-p \end{cases}$$

are equally preferred. If the decision maker states that $p = 0.8$, then

$$\kappa(\{1,2\}) = \frac{0.2u(x_1) - u(x_2)}{0.8u(x_1)u(x_2)} \approx -1.29$$

according to Eq. (16) with $J = \{1, 2\}$, $J' = \{1\}$, $J'' = \{2\}$, $j' = 2$, and $j'' = 1$.

Suppose the values $\kappa(\{1,3\}) = 0.9$ and $\kappa(\{2,3\}) = 1.1$ for the remaining project pairs are obtained through similar questions, after which the process proceeds by determining $\kappa(J)$ for subsets of three projects ($|J| = 3$). Specifically, the decision maker is asked to consider a preference ranking of three deterministic portfolio outcomes (x_1, x_2, y^B, y^B), (x_1, y^B, x_3, y^B), and (x_1, x_2, x_3, y^B) and the decision maker states that $(x_1, x_2, y^B, y^B) \succ (x_1, x_2, x_3, y^B) \succ (x_1, y^B, x_3, y^B)$. The decision maker is then asked to choose a probability p such that she would be indifferent between the portfolios

$$(x_1, x_2, x_3, y^B) \text{ and } \tilde{x} = \begin{cases} (x_1, x_2, y^B, y^B), & \text{with probability } p \\ (x_1, y^B, x_3, y^B), & \text{with probability } 1-p. \end{cases}$$

If the decision maker states that $p = 0.7$, then Eq. (17) with $J = \{1, 2, 3\}$, $J' = \{1, 2\}$, $J'' = \{1, 3\}$, $j' = 3$, and $j'' = 2$ yields

$$\kappa(\{1,2,3\}) = \frac{-0.3(u(x_2) + \kappa(\{1,2\})u(x_1)u(x_2)) - 0.7(u(x_3) + \kappa(\{1,3\})u(x_1)u(x_3)) - \kappa(\{2,3\})u(x_2)u(x_3)}{u(x_1)u(x_2)u(x_3)} \approx -4.97,$$

where the values of $\kappa(J')$ and $\kappa(J'')$ are known because they were assessed on the previous stage in the assessment process. Continuing in this vein can determine all values of $\kappa(J)$, $J \subset \{1, \dots, 4\}$, without the need to solve systems of linear equations.

The preference assessment process outlined above requires the decision maker to consider three portfolios, two of which are obtained from the third via removal of a single project. There might be a smaller cognitive burden, however, if the comparison involved three portfolios such that the first is a subset of the second, which in turn is a subset of the third. We thus develop an alternative process in which the decision maker is asked to consider three portfolios that consist of projects with indexes in the subsets $J'' \subset J' \subset J$ and each subset differs from its superset with regard to a single project, i.e., $J'' \cup \{j''\} = J'$, and $J' \cup \{j'\} = J$. Again the decision maker first determines a preference ranking of the three portfolios outcomes denoted by $x \in X_J$, $x' \in X_{J'}$, and $x'' \in X_{J''}$ (cf. Eq. (14)) and then adjusts the probabilities of a binary outcome (cf. Eq. (15)) that yields either the least or most preferred outcome until it is equally preferred to the second-ranked portfolio outcome. The value of $\kappa(J)$ can then be determined for any preference ranking of the three portfolio outcomes via the following theorem.

Theorem 3. *Take any $J \subseteq \{1, \dots, m\}$ and $J'' \subset J' \subset J$, such that $J'' \cup \{j''\} = J'$, and $J' \cup \{j'\} = J$. Consider portfolios $x'' \in X_{J''}$, $x' \in X_{J'}$, and $x \in X_J$. Then the following equivalences hold:*

(i) $x \succeq x' \succ x''$ *and* $x' \sim (x, x''; p, 1-p)$ *for some* $p \in (0, 1]$ *if and only if*

$$\kappa(J) = \frac{(1-p) \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subset J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{p \prod_{j \in J} u(x_j)}, \tag{19}$$

(ii) $x' \succ x \succ x''$ *and* $x \sim (x', x''; p, 1-p)$ *for some* $p \in (0, 1)$ *if and only if*

$$\kappa(J) = \frac{(p-1) \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - \sum_{\substack{I \subset J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{\prod_{j \in J} u(x_j)}, \tag{20}$$

(iii) $x' \succ x'' \succeq x$ *and* $x'' \sim (x', x; p, 1-p)$ *for some* $p \in [0, 1)$ *if and only if*

$$\kappa(J) = \frac{\sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) + (1-p) \sum_{\substack{I \subset J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{(p-1) \prod_{j \in J} u(x_j)}, \tag{21}$$

(iv) $x \geq x'' \succ x'$ *and* $x'' \sim (x, x'; p, 1-p)$ *for some* $p \in (0, 1]$ *if and only if*

$$\kappa(J) = \frac{\sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subseteq J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{p \prod_{j \in J} u(x_j)}, \tag{22}$$

(v) $x'' \succ x \succ x'$ *and* $x \sim (x'', x'; p, 1-p)$ *for some* $p \in (0, 1)$ *if and only if*

$$\kappa(J) = \frac{-\sum_{\substack{I \subseteq J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{\prod_{j \in J} u(x_j)}, \tag{23}$$

(vi) $x'' \succ x' \geq x$ *and* $x' \sim (x'', x; p, 1-p)$ *for some* $p \in [0, 1)$ *if and only if*

$$\kappa(J) = \frac{(1-p) \sum_{\substack{I \subseteq J \\ j' \in I}} \kappa(I) \prod_{j \in I} u(x_j) - p \sum_{\substack{I \subseteq J' \\ j'' \in I}} \kappa(I) \prod_{j \in I} u(x_j)}{(p-1) \prod_{j \in J} u(x_j)}. \tag{24}$$

To illustrate the second assessment process, we revisit the portfolio selection problem with $m = 4$ projects, with utilities $u(x_1) = 0.7$, $u(x_2) = 0.5$, $u(x_3) = 0.3$, and $u(x_4) = 0.6$. The decision maker is first asked to consider a preference ranking of three portfolio outcomes, (x_1, x_2, y^B, y^B), (x_1, y^B, y^B, y^B), and (y^B, y^B, y^B, y^B). Suppose the decision maker states $(x_1, y^B, y^B, y^B) \succ (x_1, x_2, y^B, y^B) \succ (y^B, y^B, y^B, y^B)$. Next, the decision maker is asked to set probability p such that portfolios

$$(x_1, x_2, y^B, y^B) \text{ and } \tilde{x} = \begin{cases} (x_1, y^B, y^B, y^B), & \text{with probability } p \\ (y^B, y^B, y^B, y^B), & \text{with probability } 1-p \end{cases}$$

are equally preferred. If the decision maker states that $p = 0.6$, then Eq. (20) with $J = \{1, 2\}$, $J' = \{1\}$, $J'' = \{2\}$, $j' = 2$, and $j'' = 1$ yields

$$\kappa(\{1, 2\}) = \frac{-0.4u(x_1) - u(x_2)}{u(x_1)u(x_2)} = \frac{(-0.4)(0.7) - 0.5}{(0.7)(0.5)} \approx -2.23.$$

After the values of $\kappa(\{1, 3\}) = 0.9$, $\kappa(\{2, 3\}) = 1.1$ have been obtained through similar elicitation questions, the decision maker can move on to assessing $\kappa(J)$ for subsets that contain three projects. Suppose, for instance, that when considering the portfolio outcomes (x_1, y^B, y^B, y^B), (x_1, x_2, y^B, y^B), and (x_1, x_2, x_3, y^B) the decision maker states that $(x_1, x_2, x_3, y^B) \succ (x_1, y^B, y^B, y^B) \succ (x_1, x_2, y^B, y^B)$. The decision maker is then asked to adjust probability p until she is indifferent between portfolios

$$(x_1, y^B, y^B, y^B) \text{ and } \tilde{x} = \begin{cases} (x_1, x_2, x_3, y^B), & \text{with probability } p \\ (x_1, x_2, y^B, y^B), & \text{with probability } 1-p. \end{cases}$$

The probability $p = 0.3$ would yield

$$\begin{aligned} &\kappa(\{1, 2, 3\}) \\ &= \frac{u(x_2) + \kappa(\{1, 2\})u(x_1)u(x_2) - 0.3(u(x_3) + \kappa(\{1, 3\})u(x_1)u(x_3) + \kappa(\{2, 3\})u(x_2)u(x_3))}{0.3u(x_1)u(x_2)u(x_3)} \\ &\approx -15.13, \end{aligned}$$

by Eq. (22). The process can be continued in a similar manner until the rest of the κ values have been elicited.

5. Optimization models for maximizing expected quasi-symmetric multilinear portfolio utility

In general, the most preferred project portfolio that maximizes expected utility can be identified by solving a stochastic optimization problem. Specifically, if the uncertain outcomes of the m project candidates are captured by the random variables $\tilde{x}_1, \ldots, \tilde{x}_m$, then this problem can be formulated as

$$\max_{z \in \{0,1\}^m} \; \mathbb{E}[U(\hat{x}_1, \ldots, \hat{x}_m)] \tag{25}$$

$$\hat{x}^j = \begin{cases} \tilde{x}_j, & \text{if } z_j = 1 \\ y^B, & \text{if } z_j = 0 \end{cases}$$

$$Az \leq B,$$

where binary decision variables $z = (z_1, \ldots, z_m) \in \{0, 1\}^m$ indicate which projects are included in the portfolio, while matrix A and vector B encode the parameters of the relevant constraints (e.g., budget).

In principle, optimization problem (25) can be solved by enumerating all possible portfolios (i.e., binary vectors $z \in \{0, 1\}^m$), checking which ones are feasible (i.e., satisfy $Az \leq B$), and then evaluating the expected utility of each feasible portfolio. However, in applications the number of projects m is often is high enough to make this approach impossible in practice. Hence, the practical applicability of portfolio decision analysis models is enhanced if the general problem (25) can be formulated as a type of optimization problem for which standard solvers are readily available. With the quasi-symmetric multilinear portfolio utility function (Theorem 1), such a formulation exists as demonstrated by the following lemma for the special case of stochastically independent project outcomes.

Lemma 1. *Assume the project outcomes $\tilde{x}_1, \ldots, \tilde{x}_m$ are stochastically independent. Then z^* is an optimal solution to optimization problem (25) if and only if there exists ζ^* such that (z^*, ζ^*) is an optimal solution to the ILP problem*

$$\max_{\substack{z \in \{0,1\}^m \\ \zeta \in \{0,1\}^K}} \sum_{j=1}^{m} z_j \mathbb{E}[u(\tilde{x}_j)] + \sum_{k=1}^{K} \zeta_k \kappa(J_k) \prod_{j \in J_k} \mathbb{E}[u(\tilde{x}_j)] \tag{26}$$

$$m\zeta_k \leq \sum_{j \in J_k} z_j - |J_k| + m \;\; \forall \; k \in \{1, \ldots, K\}$$

$$m\zeta_k \geq \sum_{j \in J_k} z_j - |J_k| + 1 \;\; \forall \; k \in \{1, \ldots, K\}$$

$$Az \leq B,$$

where $\{J_1, \ldots, J_K\} = \{J \subseteq \{1, \ldots, m\} \mid |J| \geq 2, \kappa(J) \neq 0\}$.

Note that, since $u(y^B) = 0$, denoting $s_k = \kappa(J_k) \prod_{j \in J_k} \mathbb{E}[u(\tilde{x}_j)]$ renders the optimization problem of Lemma 1 equivalent to the ILP problem (5)–(8), which represents the standard approach for handling project interactions. This observation suggests that in the standard approach uncertain project outcomes need to be stochastically independent and that the interaction effect coefficients s_k should be contingent on the utilities of the projects involved in each interaction. Specifically, if the projects' outcomes and, thereby, their utilities change, the values of the interaction coefficients s_k should be updated accordingly.

Stochastic dependencies among project outcomes can be incorporated into the optimization model (26) by constructing a finite state-space such that the project outcomes are conditionally independent given these states. Formally, if $\tilde{\omega}$ denotes the integer-valued random variable indicating which of the d states is realized, then the probability distributions of the project outcomes must satisfy $\mathbb{P}(\tilde{x}_j = x_j | \tilde{x}_{j'} = x_{j'}, \tilde{\omega} = \omega) = \mathbb{P}(\tilde{x}_j = x_j | \tilde{\omega} = \omega)$ to be conditionally independent. The resulting optimization problem, which is formalized by the following theorem, includes as input parameters the conditional expected project utilities in each of the states ($\mathbb{E}[u(\tilde{x}_j)|\tilde{\omega} = \omega], \omega \in \{1, \ldots, d\}, j \in \{1, \ldots, m\}$) as well as the state probabilities ($\mathbb{P}(\tilde{\omega} = \omega), \omega \in \{1, \ldots, d\}$).

Theorem 4. *Assume the project outcomes $\tilde{x}_1, \ldots, \tilde{x}_m$ are conditionally independent given the state $\tilde{\omega}$. Then z^* is an optimal solution to the optimization problem (25) if and only if there exists ζ^* such that (z^*, ζ^*) is an optimal solution to the ILP problem*

$$\max_{\substack{z \in \{0,1\}^m \\ \zeta \in \{0,1\}^K}} \sum_{j=1}^{m} z_j \mathbb{E}[u(\tilde{x}_j)] + \sum_{k=1}^{K} \zeta_k \kappa(J_k) \left(\sum_{\omega=1}^{d} \mathbb{P}(\tilde{\omega} = \omega) \prod_{j \in J_k} \mathbb{E}[u(\tilde{x}_j)|\tilde{\omega} = \omega] \right) \tag{27}$$

$$m\zeta_k \leq \sum_{j \in J_k} z_j - |J_k| + m \;\; \forall \; k \in \{1, \ldots, K\}$$

$$m\zeta_k \geq \sum_{j \in J_k} z_j - |J_k| + 1 \;\; \forall \; k \in \{1, \ldots, K\}$$

$$Az \leq B,$$

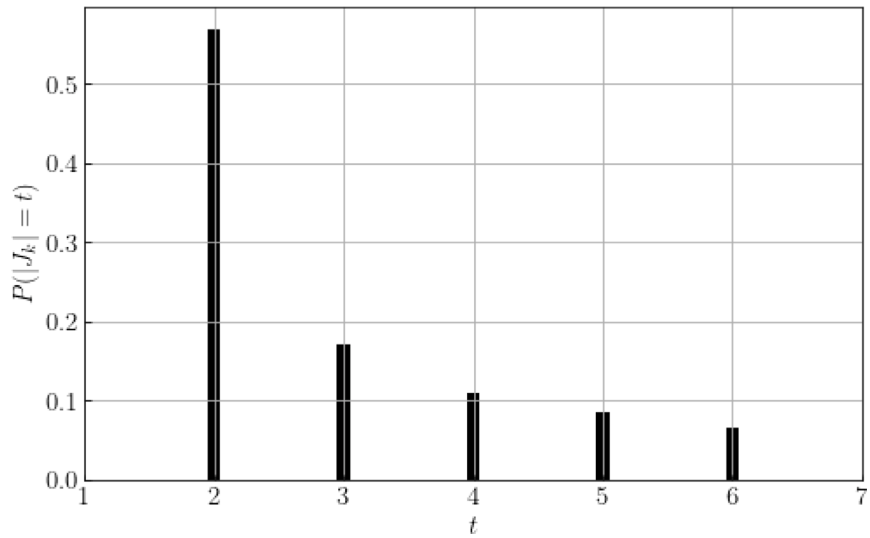

Fig. 1. Probability distribution of the number of projects $|J_k|$ involved in the kth interaction.

where $\{J_1, \dots, J_K\} = \{J \subseteq \{1, \dots, m\} \mid |J| \geq 2, \kappa(J) \neq 0\}$.

From a computational perspective the ILP problems (26) and (27) are equivalent since they have an equal number of decision variables and an equal number of constraints, and they differ only with regard to the values of the objective function coefficients. The computational effort required to solve these ILP problems is contingent mainly on the number of binary decision variables, which itself depends not only on the number of project candidates m but also on the number of project interactions, i.e., project subsets $J \subseteq \{1, \dots, m\}$ for which $\kappa(J) \neq 0$.

6. Impacts of modeling project interactions to decision recommendations

Although the quasi-symmetric multilinear portfolio utility function can be used to identify optimal portfolio decisions that account for interactions among all project subsets, in real-world applications it might not often be possible to quantify all existing interactions, since the assessment process is likely to require heavy involvement of the decision makers or other experts (see Section 4). Moreover, each additional project interaction increases the computational effort required to identify the portfolio that maximizes expected utility (see Section 5). Thus, it is important to examine how the true optimal portfolio that accounts for all project interactions differs from portfolios recommended by approximate models in which only a subset of the interactions is considered. In this section we compare the project compositions and expected utilities of the true optimal portfolio to those recommended by the approximate models by using problem instance based on randomly generated and real-world data.

6.1. Simulation setup

The analysis is based on randomly generated samples each consisting of 100 project portfolio selection problem instances. Each problem instance has $m = 30$ project candidates with project interactions and a single budget constraint that limits the total cost for the portfolio to be no more than 50% of the sum of all candidate project costs c_j, i.e., $\sum_{j=1}^{30} z_j c_j \leq 0.5 \sum_{j=1}^{30} c_j$. The joint distribution of each project's cost c_j and utility $\mathbb{E}[u(\tilde{x}_j)]$ follows a Gaussian copula with correlation coefficient ρ and uniform marginal distributions. The samples generated will be varied in terms of four dimensions: (i) the correlation between project utilities and costs (ρ), (ii) the number of project interactions (K), (iii) the sizes of the interaction subsets ($|J_k|, k \in \{1, \dots, K\}$), and (iv) the magnitude of the interaction effects ($\kappa(J_k), k \in \{1, \dots, K\}$).

For the sizes of the interaction subsets $|J_k|$ we will consider two cases. The first includes only pairwise interactions, i.e., $|J_k| = 2$ for all $k \in \{1, \dots, K\}$. In the second case, the size of each interaction subset $|J_k|$ follows a shifted beta-binomial distribution such that $|J_k| = \tilde{\beta} + 2$, where $\tilde{\beta} \sim \text{BetaBin}(n^*, \alpha, \beta)$ with parameter values $n^* = 4$, $\alpha = 0.3$, and $\beta = 1$ (see Fig. 1). The parameter values are chosen to represent a realistic setting in which small-sized interaction subsets are more frequent than larger ones. The projects included in each subset J_k are randomly sampled from the set $\{1, \dots, m\}$.

The values for the parameters $\kappa(J_k), k \in \{1, \dots, K\}$, capturing the magnitudes of the interaction effects are generated by using the following approach. First, note that in case there is no interaction associated with the project subset J_k (i.e., $\kappa(J_k) = 0$), the utility of a portfolio $x \in X_{J_k}$ consisting of only projects in subset J_k (Eq. (14)) is

$$U(x) = \sum_{J \subseteq \{1,\dots,m\}} \kappa(J) \prod_{j \in J} u(x_j) = \sum_{J \subset J_k} \kappa(J) \prod_{j \in J} u(x_j),$$

since $u(y^B) = 0$. Now suppose that the relative change in the utility of portfolio x due to interaction J_k is captured by the coefficient $\phi > 0$ in the sense that

$$U(x) = \sum_{J \subseteq \{1,\dots,m\}} \kappa(J) \prod_{j \in J} u(x_j) = \phi \sum_{J \subset J_k} \kappa(J) \prod_{j \in J} u(x_j), \tag{28}$$

where a value $\phi \in [0, 1)$ yields a cannibalization effect and a value $\phi > 1$ yields a synergy effect. Solving the parameter $\kappa(J_k)$ from this equation yields

$$\kappa(J_k) = \frac{(\phi - 1) \sum_{J \subset J_k} \kappa(J) \prod_{j \in J} u(x_j)}{\prod_{j \in J_k} u(x_j)}, \tag{29}$$

which specifies the magnitude of each interaction effect as a function of ϕ. By varying the distribution from which the value of ϕ is drawn for each subset $\kappa(J_k)$, we can control which types of project interactions are present in the problem instances. Specifically, instances with both synergies and cannibalization effects are generated by drawing values $\phi \in [0, \infty)$ from a log-normal distribution with median value equal to one. In turn, in problem instances having only synergies values $\phi \in [1, \infty)$ follow unit-shifted exponential distribution (i.e., $\phi - 1$ follows an exponential distribution). Finally, problem instances having only cannibalization effects are generated by drawing values $\phi \in [0, 1]$ from a beta distribution.

For each problem instance the following computations are carried out. First, for each $k \in \{0, \dots, K\}$, we use Lemma 1 to solve for a portfolio z^k that maximizes the quasi-symmetric multilinear portfolio utility function when the k interactions with the largest magnitude

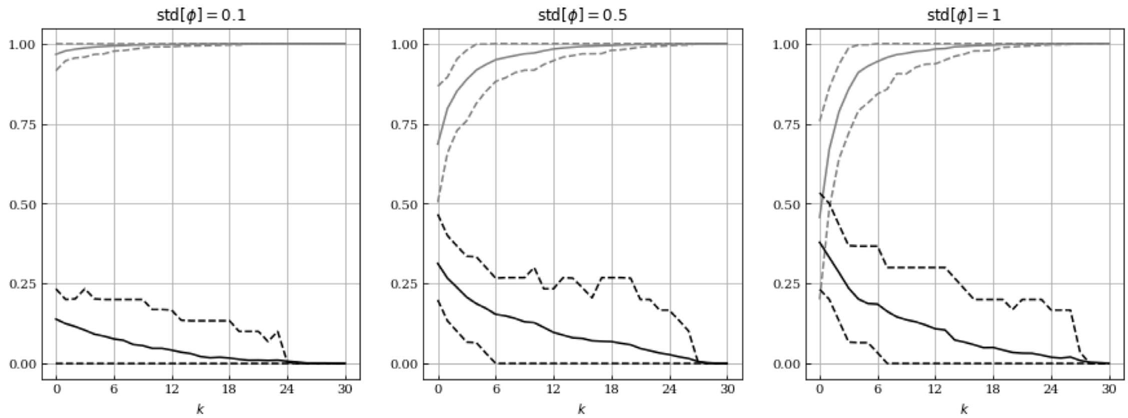

Fig. 2. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when the k largest interactions are considered. Interaction magnitude $\phi \in [0, \infty)$ (Eq. (28)) follows a log-normal distribution with unit median.

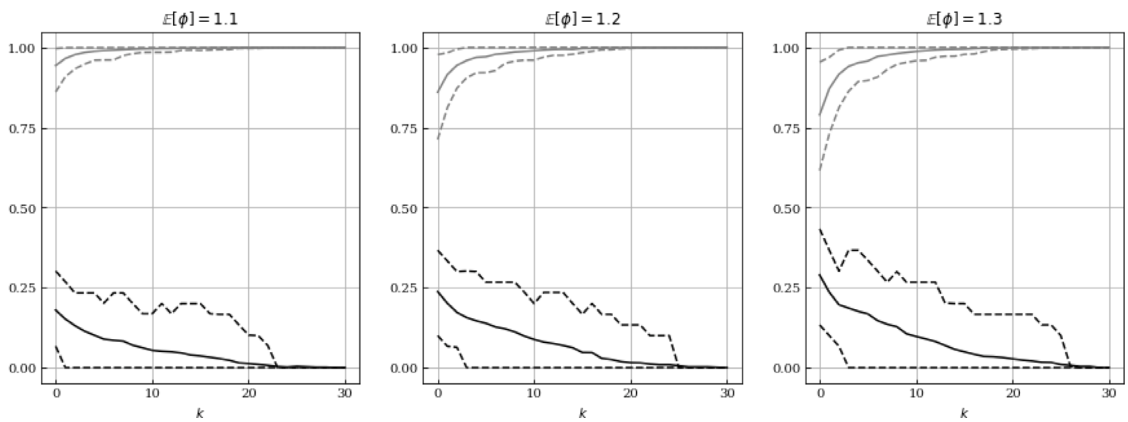

Fig. 3. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when the k largest interactions are considered. Interaction magnitude $\phi \in [1, \infty)$ (Eq. (28)) follows a unit-shifted exponential distribution.

$|\kappa(J_k)\prod_{j\in J_k} u(x_j)|$ are considered. Specifically, z^0 thus denotes the portfolio maximizing the additive portfolio utility function that does not include any interactions (Eq. (3)), while z^K denotes the true optimal portfolio that maximizes the quasi-symmetric multilinear portfolio utility function with all interactions taken into account. Next, we calculate two measures to analyze the differences between the true optimal portfolio z^K and the portfolio z^k for $k \in \{0,\dots,K\}$: (i) the share of different project decisions between the two portfolios, i.e., $\Delta(k) = \sum_{j=1}^m |z_j^k - z_j^K|/m$, and (ii) the utility achieved by the portfolios z^k relative to that of the true optimal portfolio z^K, i.e., $\Delta'(k) = EU(z^k)/EU(z^K)$. Here $EU(z)$ denotes the expected utility of portfolio z when evaluated with the utility function including all the interaction terms, which implies that $EU(z^K) \geq EU(z^k)$ for all $k \in \{0,\dots,K\}$. Thus, it holds that $\Delta(K) = 0$ and $\Delta'(K) = 1$.

6.2. Simulation results

Figs. 2–6 illustrate the differences between the optimal portfolio accounting for all project interactions and those portfolios that are optimal when only k interactions with the largest magnitudes are considered. In particular, the figures' black lines show the share of different project decisions ($\Delta(k)$), with the solid line corresponding to the mean and the dashed lines corresponding to the 5th and 95th percentiles. The gray lines show the share of true optimal expected utility achieved by the portfolios that consider only the k largest interactions ($\Delta'(k)$) with the solid line representing the mean and the dashed lines the 5th and 95th percentiles.

Fig. 2 shows the results for $K = 30$ project interactions including both synergy and cannibalization effects. The correlation coefficient between project utilities and costs is $\rho = 0.3$, and the parameter ϕ capturing interaction magnitude follows a log-normal distribution with unit median and varying standard deviation (std[ϕ]). Note that the realization $\phi = 1$ implies a zero interaction effect $\kappa(J_k) = 0$ (see Eq. (29)), while realizations $\phi \neq 1$ imply non-zero effects. Consequently, increasing the standard deviation of ϕ increases the average absolute magnitude of the interaction effects and thus amplifies the impacts of considering only the k largest interactions. For instance, when std[ϕ] = 1, the portfolio obtained when no interactions are considered ($k = 0$) yields on average only half of the true optimal expected utility ($k = K = 30$). Moreover, the average share of project decisions changed is nearly 40%, which corresponds to some 12 out of 30 projects. In turn, when std[ϕ] = 0.1, the average loss in expected utility and the average share of changed project decisions are significantly lower.

Figs. 3 and 4 show the results when all project interactions have either synergy or cannibalization effects, respectively, while other parameters of the problem instances remain unchanged. Specifically, in Fig. 3 parameter $\phi \in [1, \infty)$ specifying the interaction magnitudes (see (29)) follows a unit-shifted exponential distribution with varying mean value. Because greater values of ϕ correspond to larger synergy effects, increasing this mean value results in larger differences between the true optimal portfolio z^{30} and portfolios z^k accounting only for the k largest interactions. In Fig. 4 the values of $\phi \in [0, 1]$ follow a beta distribution with varying mean value. Since small values of ϕ result in large cannibalization effects, decreasing this mean amplifies

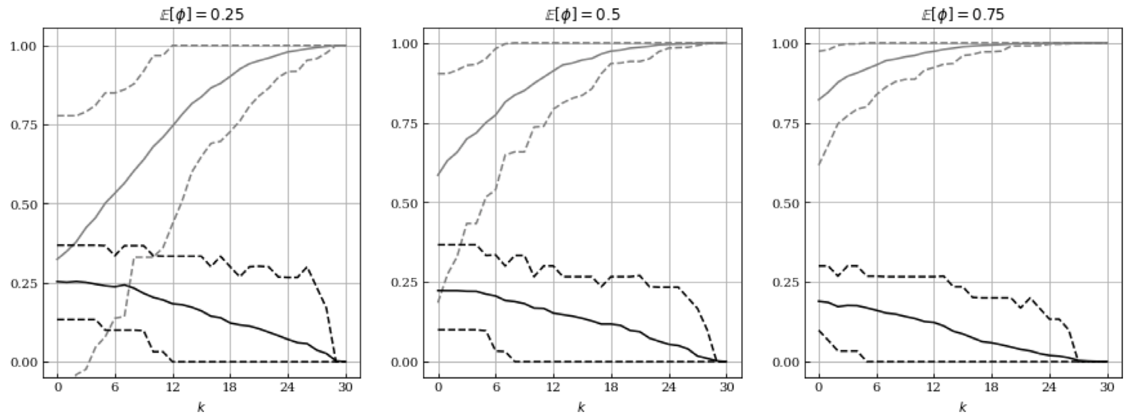

Fig. 4. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when the k largest interactions are considered. Interaction magnitude $\phi \in [0, 1]$ (Eq. (28)) follows a beta distribution.

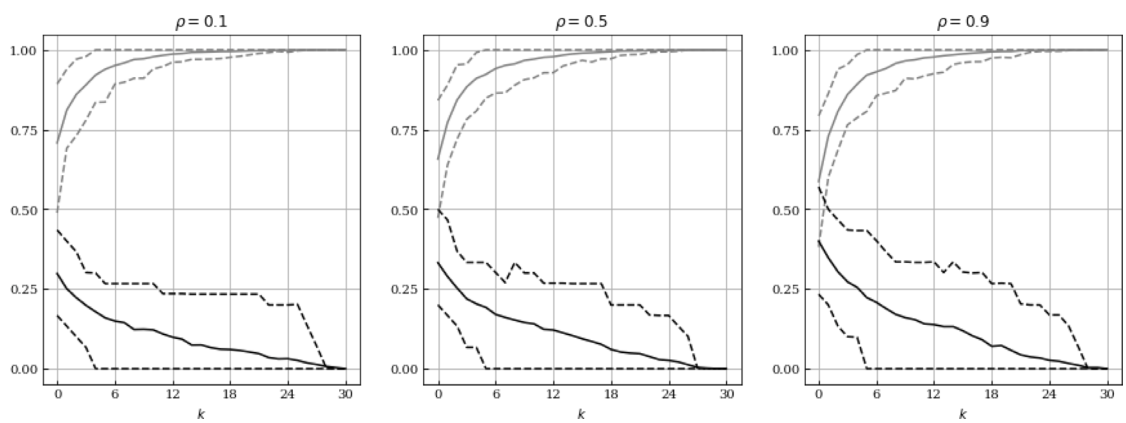

Fig. 5. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when the k largest interactions are considered for different correlations between projects' utilities and costs (ρ).

the differences between the optimal portfolio and those in which only some of the cannibalization effects are considered. In fact, with the mean equal to 0.25, the portfolio obtained when no interactions are considered (z^0) is actually less preferred than the empty portfolio in some of the problem instances. This leads to the 5th percentile of the relative utility dropping below zero.

Fig. 5 illustrates the effect that the correlation between project costs and utilities (ρ) has on the importance of capturing project interactions. Again, each problem instance has $K = 30$ interactions and the parameter ϕ follows a log-normal distribution ($\mathrm{std}[\phi] = 0.5$) generating both synergy and cannibalization effects. The results suggest that the increased correlation between projects utilities and costs increases the importance of considering project interactions. This can be explained by the observation that in cases of low correlation there likely exists a subset of low-cost, high-utility projects that form at least the core of the optimal portfolio. In this case taking into account cannibalization effects among these projects, or synergies between them and other projects, would likely not result in major changes in the composition of the optimal portfolio. In turn, if the correlation between project utilities and costs is high, then the existence of several low-cost and high-utility projects is less likely and, therefore, interactions have a greater impact on the composition of the optimal portfolio.

The last results (Fig. 6) focus on problem instance involving only pairwise interactions the number of which takes values $K \in \{1, 50, 100\}$. The problem instances were generated with the correlation coefficient $\rho = 0.3$ and by using the same log-normal distribution for ϕ as in Fig. 5. Notably, regardless of the total number of interactions K, it seems that increasing the number of interactions considered from $k = 0$ to some $k = K/5$, results in a sharp decrease in the differences between the true optimal portfolio z^K and the optimal portfolios z^k in terms of expected utility. This indicates that by considering only the largest 20% interactions it is possible to obtain the majority of the total expected utility achieved by considering all the interactions.

6.3. Application to R&D project portfolio selection

In this section, we analyze the impact of not modeling all the project interactions by using a problem instance based on a real-world application in R&D (de Almeida & Duarte, 2011). In this application, a decision maker selects a portfolio of R&D projects to maximize the multiattribute portfolio utility subject to multiple constraints. In particular, $m = 10$ project candidates are evaluated with regard to $n = 4$ attributes: expected return, success probability, strategic impact, and operational impact. The deterministic attribute-specific evaluations are aggregated with an additive multiattribute utility function to obtain project utilities $u(x_1), \ldots, u(x_{10})$ and the utility of not selecting a project is $u(y^B) = 0$. Moreover, there are four resource constraints, limiting the labor ($l = 1$), equipment ($l = 2$) and energy ($l = 3$) usage of the selected projects and their total cost ($l = 4$). Let a_j^l denote the jth project's consumption of the lth resource and B_l the availability of the lth resource. Table 1 shows the projects' utilities, resource consumption, and the resource availability. Finally, out of the $9 \cdot 10/2 = 45$ project pairs, 32 pairs have synergies (see Table 2). Note that de Almeida and Duarte (2011) have quantified the synergy effect obtained if both the jth and the j'th project are included in the portfolio in terms of the resulting increase

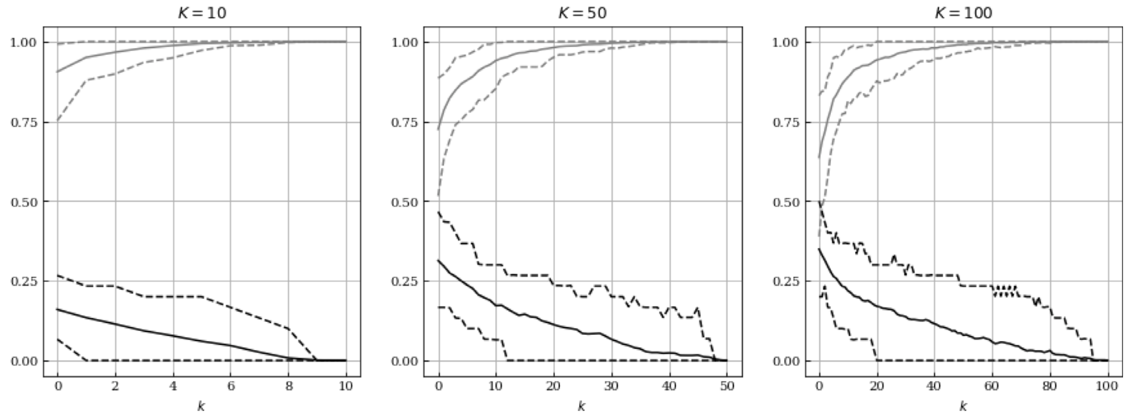

Fig. 6. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when there are K pairwise interactions of which k are considered.

Table 1
Projects' utilities, resource consumption, and resource availability.

Project j	Utility $u(x_j)$	Labor (# of people) a_j^1	Equipment usage (h) a_j^2	Energy (MW) a_j^3	Cost (k\$) a_j^4
1	0.485	10	39	65	190
2	0.5425	15	30	70	166
3	0.49	18	38	63	205
4	0.645	35	45	80	250
5	0.4025	8	20	53	107
6	0.435	8	18	58	112
7	0.43	5	20	58	97
8	0.455	5	12	60	83
9	0.43	3	16	54	85
10	0.4025	3	12	55	40
	B_l	60	160	380	1000

in portfolio utility, which in our framework corresponds to the term $\kappa(\{j,j'\})u(x_j)u(x_{j'})$.

The resulting project portfolio selection problem corresponds to the non-linear binary optimization problem

$$\max_{z\in\{0,1\}^{10}} \sum_{j=1}^{m} z_j u(x_j) + \sum_{J\subseteq\{1,\ldots,10\}} \kappa(J)\prod_{j\in J} z_j u(x_j) \tag{30}$$
$$\sum_{j=1}^{10} a_j^l z_j \leq B_l \ \forall\, l \in \{1,\ldots,4\},$$

where $u(x_j)$, a_j^l, and B_l are specified in Table 1 and each non-zero $\kappa(J)\prod_{j\in J} z_j u(x_j)$ in Table 2. This problem can be formulated as an ILP problem per Lemma 1.

Similarly to the analysis of the randomly generated problem instances (Sections 6.1 and 6.2), we solve problem (30) by considering the k largest synergy effects for $k \in \{1,\ldots,32\}$. The resulting portfolios z^k, $k \in \{1,\ldots,32\}$, are then compared to the optimal portfolio z^{32} obtained when all $K = 32$ synergies are taken into account in terms of both the relative expected utility ($\Delta'(k)$) and changed project decisions ($\Delta(k)$). However, since the magnitudes of the synergy effects are now based on a real application, it is also of interest to analyze how robust the decision recommendations are to variations in these estimates. Thus, we carry out the aforementioned comparisons also for cases in which all the interaction magnitudes are tenfold and hundredfold. Technically, this corresponds to replacing each parameter $\kappa(J)$ in (30) with $\theta\kappa(J)$, where $\theta = 1$ (the base-case), $\theta = 10$, or $\theta = 100$.

Fig. 7 shows the share of changed project decisions ($\Delta(k)$, black) and the achieved relative expected utility ($\Delta'(k)$, gray) as a function of the number of interactions considered (k) when different multipliers (θ) are applied for the magnitudes of the synergy effects. For the case

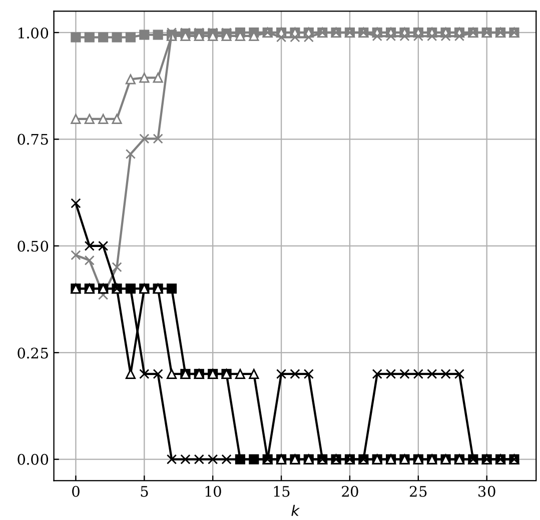

Fig. 7. The share of changed project decisions ($\Delta(k)$, black) and the achieved relative utility ($\Delta'(k)$, gray) when the k largest interactions are considered with base case ($\theta = 1$, squares), tenfold ($\theta = 10$, triangles), and hundredfold ($\theta = 100$, crosses) synergy effects.

in which $\theta = 100$, decision recommendations differ from the optimal ones even when $k = 28$ out of the $K = 32$ synergy effects are considered. Moreover, the optimal portfolio from the additive portfolio utility function (i.e., z^0) yields less than half of the optimal expected utility. However, as k increases from 0 to 7, the expected utility from portfolio z^k increases rapidly and becomes very close to that of the optimal portfolio. For the cases in which $\theta = 1$ and $\theta = 10$, correct project decision recommendations are obtained when at least half of the interactions are considered (i.e., $k \geq 14$).

Fig. 8 shows which projects are included in portfolios z^k for different values of θ. For instance, project 2 is included in z^k for all θ when $k > 2$. Hence, the model would recommend selecting project 2 even if the magnitudes of the synergies were significantly overestimated as long as the two largest synergies are included in the model. In turn, project 7 jumps in and out of the portfolios z^k for all values of θ, which means its selection is sensitive to identification of the synergies as well as quantification of their magnitude.

The correct specification of synergy effects' magnitudes affects the decision recommendations in general. The composition of the optimal portfolio z^{32} is not exactly the same for all values of θ. In particular, the optimal portfolio always contains projects 1, 2, 5, 8, and 10, but the choice of including project 3 or project 7 depends on the value of θ. In fact, the composition of the portfolios z^k is the same across the values

Table 2
Non-zero synergy effects $\kappa(\{j,j'\})u(x_j)u(x_{j'})$.

Project j/ project j'	2	3	4	5	6	7	8	9	10
1	0.0133	0.0046	0.0206	0.0178		0.0037		0.0075	0.0152
2			0.019	0.0162	0.0028	0.0014	0.0182		0.0079
3			0.0141	0.0128		0.0085			0.0102
4				0.0147	0.0066	0.0005	0.0041		0.0044
5					0.0049	0.0077	0.005		0.0138
6								0.0057	0.0063
7							0.0174		
8								0.0059	0.016
9									0.0152
10									

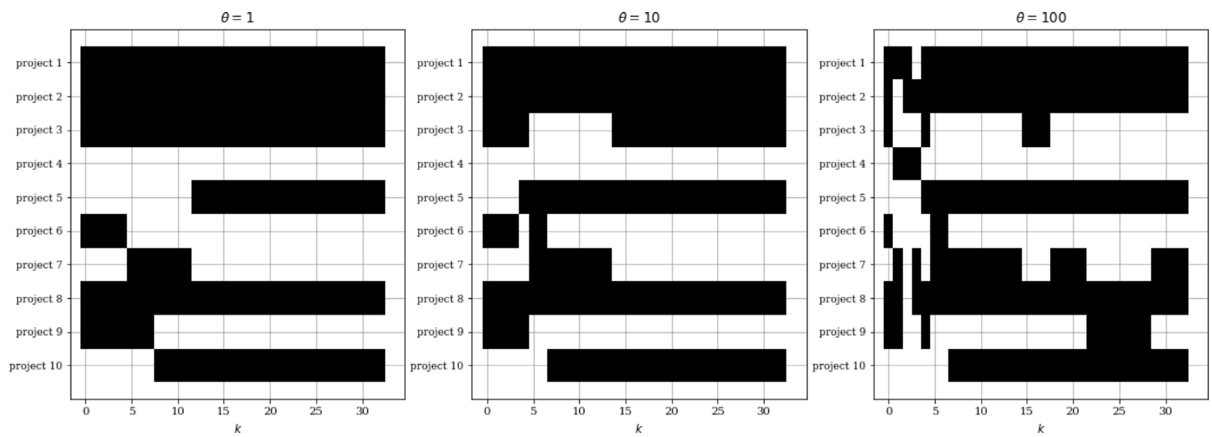

Fig. 8. Portfolio composition when k synergy effects are considered with the base case ($\theta = 1$), tenfold ($\theta = 10$), and hundredfold ($\theta = 100$) synergy effects.

of θ only when $k = 15, 16, 17$ or when no interactions are considered ($k = 0$).

Greater magnitude of the synergy effects makes the project decision recommendations more sensitive to the number of interactions considered. When $\theta = 1$ or $\theta = 10$, the composition of z^k changes as k increases from 0 to 14, it but remains unchanged for larger values of k. However, when $\theta = 100$, the composition of portfolio z^k changes for almost all unit increases of k until $k = 7$ and the optimal portfolio is obtained only when at least $k = 29$ of the $K = 32$ interactions are considered.

7. Discussion and conclusions

We have shown that if preferences among multiattribute project portfolios satisfy three specific assumptions, then these preferences are represented by a quasi-symmetric multilinear portfolio utility function. In particular, such a representation assumes that preferences among portfolios consisting of a single project are independent of the indexing of that project and that for some project outcome such a portfolio would be preferred to an empty portfolio. The third assumption restricts preferences among portfolios that differ with regard to one project when the outcomes of the identical projects are changed: under any such changes, the preference ranking of the portfolios stays the same, it is fully reversed, or all portfolios become equally preferred.

This quasi-symmetric multilinear portfolio utility function is particularly interesting from the standpoint of applications as it can be written as the sum of a standard additive portfolio utility function and a set of additional terms each capturing the interaction effect of a specific subset of the project candidates. Hence, this portfolio utility function offers a decision-theoretic foundation for the ad hoc approach commonly used in application of PDA, in which the portfolio utility is modeled as the sum of the included projects' utilities augmented with additional terms in case the portfolio's project composition triggers a synergy or a cannibalization effect.

The computational experiments on problem instances based on randomly generated and real-world data highlight the importance of modeling project interactions. Indeed, neglecting to account for interactions can lead to erroneous decision recommendations for project selection and to losses in expected portfolio utility. However, incorporating project interactions into the portfolio model adds challenges to decision support. First, each interaction requires an additional binary variable in the portfolio optimization model, thus making it computationally more demanding. This becomes an issue especially if the number of interactions is in the hundreds. Second, and perhaps more importantly, a large number of interactions necessitates the assessment of a large number of interaction effects, which requires involving the decision makers or other experts. Furthermore, the types of questions that can be utilized to assess these effects are cognitively more demanding than, for instance, the preference elicitation questions used to specify the project-level multiattribute utility function. As the number of interactions can, in the worst case, increase exponentially as a function of the number of project candidates, these assessment processes can quickly become very demanding. However, the computational experiments also show that, generally, incorporating at least some of the most significant interaction into the portfolio model reduces the average number of erroneous project decisions and decreases the resulting loss in portfolios' expected utility.

The results point to several avenues for future research. From a theoretical perspective it would be interesting to examine whether an axiomatic foundation similar to the one developed here could be built based on measurable value functions instead of utility functions (Dyer and Sarin 1979, Liesiö 2014). Such foundations could offer a platform to develop alternative techniques for assessing interaction effects that would not require the decision makers to state preferences between uncertain outcomes but would allow comparing changes in deterministic outcomes. The practice of portfolio decision analysis would benefit from case studies based on real applications to evaluate the effectiveness of various techniques and processes for assessing project interactions. Such studies would also serve building a larger data set of real-world project portfolio selection instances in which synergy and cannibalization effects have been quantified.

Acknowledgments

This research was supported by the Academy of Finland (grant number 323800).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.ejor.2024.02.015.

References

Abbas, A. (2009). Multiattribute utility copulas. *Operations Research*, *57*(6), 1367–1383.
Abbassi, M., Ashrafi, M., & Tashnizi, E. S. (2014). Selecting balanced portfolios of R&D projects with interdependencies: A cross-entropy based methodology. *Technovation*, *34*, 54–63.
Barbati, M., Greco, S., Kadzinski, M., & Slowinski, R. (2018). Optimization of multiple satisfaction levels in portfolio decision analysis. *Omega*, *78*, 192–204.
Bordley, R., & Kirkwood, C. (2004). Multiattribute preference analysis with performance targets. *Operations Research*, *52*(6), 823–835.
Carazo, A. F., Gómez, T., Molina, J., Hernández-Díaz, A. G., Guerrero, F. M., & Caballero, R. (2010). Solving a comprehensive model for multiobjective project portfolio selection. *Computers & Operations Research*, *37*(4), 630–639.
Chowdhury, M. M. H., & Quaddus, M. A. (2015). A multiple objective optimization based QFD approach for efficient resilient strategies to mitigate supply chain vulnerabilities: The case of garment industry of Bangladesh. *Omega*, *57*, 5–21.
Clemen, R. T., & Smith, J. E. (2009). On the choice of baselines in multiattribute portfolio analysis: A cautionary note. *Decision Analysis*, *6*(4), 256–262.
de Almeida, A. T., & Duarte, M. (2011). A multi-criteria decision model for selecting project portfolio with consideration being given to a new concept for synergies. *Pesquisa Operacional*, *31*, 301–318.
Durbach, I. N., Algorta, S., Kantu, D. K., Katsikopoulos, K. V., & Şimşek, O. (2020). Fast and frugal heuristics for portfolio decisions with positive project interactions. *Decision Support Systems*, *138*, Article 113399.
Dyer, J., & Sarin, R. (1979). Measurable multiattribute value functions. *Operations Research*, *27*(4), 810–822.
Eilat, H., Golany, B., & Shtub, A. (2006). Constructing and evaluating balanced portfolios of R&D projects with interactions: A DEA based methodology. *European Journal of Operational Research*, *172*(3), 1018–1039.
Fishburn, P. C. (1974). von Neumann-Morgenstera utility functions on two attributes. *Operations Research*, *22*(1), 35–45.
Fishburn, P. C., & Keeney, R. L. (1975). Generalized utility independence and some implications. *Operations Research*, *23*(5), 928–940.
Golabi, K., Kirkwood, C. W., & Sicherman, A. (1981). Selecting a portfolio of solar energy projects using multiattribute preference theory. *Management Science*, *27*(2), 174–189.
Grushka-Cockayne, Y., Reyck, B. D., & Degraeve, Z. (2008). An integrated decision-making approach for improving European air traffic management. *Management Science*, *54*(8), 1395–1409.
Keeney, R., & Raiffa, H. (1976). *Decisions with multiple objectives: preferences and value trade-offs*. John Wiley & Sons, New York.
Keller, L. R., & Simon, J. (2019). Preference functions for spatial risk analysis. *Risk Analysis*, *39*(1), 244–256.
Kleinmuntz, D. N. (2007). Resource allocation decisions. In W. Edwards, R. F. Miles, Jr., & D. von Winterfeldt (Eds.), *Advances in decision analysis: from foundations to applications* (pp. 400–418). Cambridge University Press.
Li, X., Huang, Y.-H., Fang, S.-C., & Zhang, Y. (2020). An alternative efficient representation for the project portfolio selection problem. *European Journal of Operational Research*, *281*(1), 100–113.
Liesiö, J. (2014). Measurable multiattribute value functions for portfolio decision analysis. *Decision Analysis*, *11*, 1–20.
Liesiö, J., Mild, P., & Salo, A. (2008). Robust portfolio modeling with incomplete cost information and project interdependencies. *European Journal of Operational Research*, *190*, 679–695.
Liesiö, J., & Punkka, A. (2014). Baseline value specification and sensitivity analysis in multiattribute project portfolio selection. *European Journal of Operational Research*, *237*(3), 946–956.
Liesiö, J., Salo, A., Keisler, J. M., & Morton, A. (2021). Portfolio decision analysis: Recent developments and future prospects. *European Journal of Operational Research*, *293*(3), 811–825.
Liesiö, J., & Vilkkumaa, E. (2021). Nonadditive multiattribute utility functions for portfolio decision analysis. *Operations Research*, *69*(6), 1886–1908.
Lopes, Y. G., & de Almeida, A. T. (2015). Assessment of synergies for selecting a project portfolio in the petroleum industry based on a multi-attribute utility function. *Journal of Petroleum Science and Engineering*, *126*, 131–140.
Marttunen, M., Haara, A., Hjerppe, T., Kurttila, M., Liesiö, J., Mustajoki, J., Saarikoski, H., & Tolvanen, A. (2023). Parallel and comparative use of three multicriteria decision support methods in an environmental portfolio problem. *European Journal of Operational Research*, *307*(2), 842–859.
Mavrotas, G., & Makryvelios, E. (2021). Combining multiple criteria analysis, mathematical programming and Monte Carlo simulation to tackle uncertainty in research and development project portfolio selection: A case study from Greece. *European Journal of Operational Research*, *291*(2), 794–806.
Mild, P., Liesi, J., & Salo, A. (2015). Selecting infrastructure maintenance projects with robust portfolio modeling. *Decision Support Systems*, *77*, 21–30.
Montibeller, G., Franco, L. A., Lord, E., & Iglesias, A. (2009). Structuring resource allocation decisions: A framework for building multi-criteria portfolio models with area-grouped options. *European Journal of Operational Research*, *199*(3), 846–856.
Montiel, L. V., & Bickel, J. E. (2014). A generalized sampling approach for multilinear utility functions given partial preference information. *Decision Analysis*, *11*(3), 147–170.
Morton, A. (2015). Measurement issues in the evaluation of projects in a project portfolio. *European Journal of Operational Research*, *245*(3), 789–796.
Morton, A., Keisler, J. M., & Salo, A. (2016). Multicriteria portfolio decision analysis for project selection. In S. Greco, M. Ehrgott, & J. R. Figueira (Eds.), *Multiple criteria decision analysis: state of the art surveys* (pp. 1269–1298). Springer New York.
Parnell, G. S., Bennett, G. E., Engelbrecht, J. A., & Szafranski, R. (2002). Improving resource allocation within the National Reconnaissance Office. *Interfaces*, *32*(3), 77–90.
Perez, F., & Gomez, T. (2016). Multiobjective project portfolio selection with fuzzy constraints. *Annals of Operations Research*, *245*(1–2), 7–29.
Phillips, L. D., & Bana e Costa, C. A. (2007). Transparent prioritisation, budgeting and resource allocation with multi-criteria decision analysis and decision conferencing. *Annals of Operations Research*, *154*(1), 51–68.
Salo, A., Keisler, J., & Morton, A. (2010). *Portfolio decision analysis: improved methods for resource allocation*: *vol. 162*.
Savage, L. J. (1972). *The foundations of statistics*. Wiley Publications in Statistics.
Solak, S., Clarke, J.-P. B., Johnson, E. L., & Barnes, E. R. (2010). Optimization of R&D project portfolios under endogenous uncertainty. *European Journal of Operational Research*, *207*(1), 420–433.
Stummer, C., & Heidenberger, K. (2003). Interactive R&D portfolio analysis with project interdependencies and time profiles of multiple objectives. *IEEE Transactions on Engineering Management*, *50*(2), 175–183.
von Neumann, J., Morgenstern, O., & Rubinstein, A. (1944). *Theory of games and economic behavior (60th anniversary commemorative edition)*. Princeton University Press.