European Journal of Operational Research 263 (2017) 225–239

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Decision Support

Modeling project preferences in multiattribute portfolio decision analysis

Tommi Tervonen[a,*], Juuso Liesiö[b], Ahti Salo[c]

[a] Evidera, London, United Kingdom
[b] Department of Information and Service Economy, Aalto University School of Business, Finland
[c] Systems Analysis Laboratory, Department of Mathematics and Systems Analysis, Aalto University School of Science, Finland

ARTICLE INFO

Article history:
Received 17 February 2016
Accepted 25 April 2017
Available online 2 May 2017

Keywords:
Multiple criteria analysis
Project portfolio selection
Multi-attribute value theory
Preference modeling

ABSTRACT

When choosing a portfolio of projects with a multi-attribute weighting model, it is necessary to elicit trade-off statements about how important these attributes are relative to each other. Such statements correspond to weight constraints, and thus impact on which project portfolios are potentially optimal or non-dominated in view of the resulting set of feasible attribute weights. In this paper, we extend earlier preference elicitation approaches by allowing the decision maker to make direct statements about the selection and rejection of individual projects. We convert such project preference statements to weight information by determining the weights for which (i) the selected project is included in all potentially optimal or non-dominated portfolios, or (ii) the rejected project is not included in any potentially optimal or non-dominated portfolio. We prove that the two complementary selection rules will exclude exactly the same set of weights. However, analyses that apply the dominance structure often lead to multiple, mutually exclusive feasible weight sets, and therefore the approach based on potential optimality is more relevant for practical decision analysis. We also propose ex ante value of information measures to guide the elicitation of project preference statements, and illustrate our results by analyzing a real case on the selection of infrastructure maintenance projects.

1. Introduction

The decision of allocating resources to a subset of many proposals is important in public administration and private firms which launch new products, invest in infrastructure projects and make commitments to policy actions. For this problem class, Portfolio Decision Analysis (PDA) (Salo, Keisler, & Morton, 2011) offers a collection of theory and methods. The use of PDA methods for project portfolio selection is based on (i) the development of a decision model which captures the salient properties of the available project proposals and the preferences of the Decision Maker (DM) for risk and multiple objectives, and (ii) the solution of a mathematical (integer) optimization problem which helps to determine the most preferred portfolios subject to the relevant constraints. PDA methods are widely employed in practice, and numerous high-impact applications have been reported in contexts such as R&D project selection (Grushka-Cockayne, de Reyck, & Degraeve, 2008; Phillips & Bana e Costa, 2007; Toppila, Liesiö, & Salo, 2011), healthcare capital budgeting (Kleinmuntz, 2007), military resource allocation (Ewing, Tarantino, & Parnell, 2006), and infrastructure asset management (Mild, Liesiö, & Salo, 2015).

Project portfolio selection usually involves multiple attributes for evaluating the proposals. In order to lower the DM's cognitive load in providing information about the exact attribute trade-offs (weights), much research has been carried out to develop methods in which the DM can provide incomplete preference information (Argyris, Figueira, & Morton, 2011; Fliedner & Liesio, 2016; Liesiö, Mild, & Salo, 2007; 2008; Lourenço, Morton, & Bana e Costa, 2012). Many of these methods resemble those for choosing the best alternative out of many proposals (Argyris, Morton, & Figueira, 2015; Hazen, 1986; Punkka & Salo, 2013; Salo & Hämäläinen, 1992; Weber, 1987). For instance, instead of requiring the DM to provide exact attribute weights, she can make a holistic assessment of two (real or hypothetical) projects and state that she prefers the first project to the second. Such a statement corresponds to a linear weight constraint that bounds the set of feasible weights.

* Corresponding author.
E-mail addresses: tommi.tervonen@evidera.com, tommi@smaa.fi (T. Tervonen), juuso.liesio@aalto.fi (J. Liesiö), ahti.salo@aalto.fi (A. Salo).

http://dx.doi.org/10.1016/j.ejor.2017.04.051

With incomplete preference information, no portfolio is typically optimal for all feasible weights. Hence, plenty of research has focused identifying portfolios that are defensible alternatives in view of incomplete information. Probably the two most widely used concepts are non-dominated and potentially optimal portfolios (see, e.g. Liesiö et al., 2007; Lourenço et al., 2012; Liesiö & Punkka, 2014; Argyris et al., 2015; Fliedner & Liesio, 2016). The concepts are not identical: A feasible portfolio is non-dominated (or efficient) if no other feasible portfolio provides greater or equal value for all feasible weights, whereas a potentially optimal (also convex efficient) maximizes the overall value for some feasible weights. Both of these solution concepts can be used to provide well-founded decision recommendations on the project-level. In particular, there usually exists projects that are included in all of the potentially optimal or non-dominated portfolios. Such projects should be selected, because if the available incomplete information were to be refined so that the feasible weight space would contain a single weight vector, the resulting optimal portfolio for this weight vector would contain all such projects. Conversely, based on the same rationale, all projects which do not belong to *any* potentially optimal or non-dominated portfolio should be rejected.

Many decision support tools applied in practice allow the DM to iteratively select or reject projects included in some but not all of the non-dominated/potentially optimal portfolios, to construct the final portfolio (Kleinmuntz, 2007; Mild et al., 2015). These tools are heuristic in the sense that they do not model what implications these project preference statements have on the set of feasible weights and, in fact, the current literature offers no formal models to capture such implications. This can be seen as a major shortcoming, because such models could be used to inform DM about implicit judgments on the attributes' importance that are implied by the project preference statements, assuming that the model is consistent. Furthermore, such information would be needed to examine whether the project preference statements are consistent with preferences elicited through standard trade-off questions involving project or portfolio consequences.

The importance of analyzing the implications of project preference statements on the set of feasible attribute weights is further motivated by the apparent cognitive complexity of making such statements. Even in a simple setting with a linear portfolio value function, a DM selecting a project into the portfolio has to, in theory, simultaneously take into account the project's score profile across all attributes, how this profile is in line with attributes' importance, and consider the project resource consumption. In more complicated problems with non-linear portfolio value function and project interactions, the DM may also have to consider how well the project consequences supplement those of other projects included in the portfolio, and whether including the project enables utilizing some synergy effects. Despite these general challenges, it is possible that some decision support processes could benefit from the use of project preference statements if proper methodological support was available. In fact, behavioral research on standard multiattribute single alternative choice problems suggests that holistic preference elicitation can lead to more consistent weights than direct methods (Korhonen, Silvennoinen, Wallenius, & Öörni, 2013).

In this paper, we take the first step to bridge this apparent gap in the PDA toolset by developing approaches for modeling project preference statements as sets of feasible weights. Specifically, we consider two alternative approaches based on analyzing sets of (i) potentially optimal portfolios and (ii) non-dominated portfolios. We identify challenges with the approach based on analyzing sets of non-dominated portfolios, and show that it is as informative as the approach that analyzes potentially optimal portfolios. We also show how commonly used project performance indexes can be extended to build useful *ex ante* measures that support the elicitation of additional project preference statements, and illustrate how these indexes can be used to guide the preference elicitation process. Finally, we demonstrate our approaches by analyzing a high-impact application on infrastructure asset management.

Our contributions advance the theory and practice of PDA in several ways. First, to our best knowledge, we provide the first theoretical basis for modeling project preference statements by using the concepts of dominance and potential optimality to derive weight information in portfolio problems. Given that modeling preference statements concerning selection and rejection of multi-attribute alternatives in choose-one-out-of-many decision problems has attracted much methodological and applied research (Corrente, Greco, Kadziński, & Słowiński, 2013; Greco, Mousseau, & Słowiński, 2008; Kadziński & Słowiński, 2015; Kadziński & Tervonen, 2013a; 2013b; Kadziński, Tervonen, & Figueira, 2015; Spliet & Tervonen, 2014; Tervonen, Sepehr, & Kadziński, 2015), this contribution has the potential to open a new stream of methodological PDA research. Second, modeling project preference statements as constraints on the feasible attribute weights makes it possible to use these statements in combination with other approaches for eliciting incomplete weight information (see e.g. Liesiö et al., 2007). Finally, our methods can be readily implemented to enhance existing processes and decision support tools for multi-attribute project portfolio selection.

The approach developed here is based on the assumption that project preference statements reveal meaningful information about the attribute weights. By meaningful we mean, that the set of feasible weights implied by the inclusion of a project in, or exclusion of a project from the portfolio, is consistent with the DM's preferences on trade-offs among the attributes. Whether this assumption holds in practice can be debated, but the theory developed here provides techniques for testing this assumption empirically. In particular, the models developed in this paper can be used to translate project preference statements into a set of feasible attribute weights. This set can be then compared to that obtained from the DM's preference statements on attribute trade-offs.

The rest of the paper is structured as follows. Section 2 introduces the additive value model for multi-attribute project portfolio selection and defines the concepts of potential optimality and dominance. Section 3 models project preference statements in terms of constraints on the set of feasible weights, and examines the structure of the resulting feasible weight set. Section 4 develops measures for assessing the 'strength' of these statements in providing additional preference information, and shows how these measures can be used to support portfolio decision processes. Section 5 addresses computational aspects. Section 6 presents an example analysis based on real-life data, and Section 7 concludes by discussing the main results.

2. Multiattribute project portfolio selection with incomplete preference information

Let there be m project proposals $X = \{x^1, \ldots, x^m\}$ which are evaluated on multiple attributes $i = 1, \ldots, n$, and denote the performance (score) of project x^j on attribute i by v_i^j. A project portfolio $p \subseteq X$ is a subset of the m project proposals, and the set of all possible portfolios is the power set $P = 2^X$. In what follows, we assume that the overall value of a portfolio can be expressed as

$$V(p, w) = \sum_{i=1}^{n} w_i V_i \left(\sum_{x^j \in p} v_i^j \right), \tag{1}$$

where the attribute-specific portfolio value functions $V_1, \ldots, V_n$ are assumed to be strictly increasing. The functional (1) form can model non-constant marginal attribute-specific portfolio values, and is thus more general than the widely applied additive-