DISSERTATION EXAMINATION OF JUUSO LIESIÖ
OPENING REMARKS BY OPPONENT

Research Professor Don N. Kleinmuntz
University of Southern California, USA
November 14, 2008

Resource allocation problems are ubiquitous in both business and government organizations. Stated quite simply, organizations typically have more good ideas for projects and programs than funds, capacity, or time to pursue them. Further, projects and programs require significant initial investments in the present, with the anticipation of future benefits. This necessitates balancing the promise of a return on the investment against the risk that the promised benefits do not materialize. An added complication is that organizations often have complex and poorly articulated objectives and no consistent methodology for determining how well alternative investments measure up against those objectives.

The field of decision analysis (or DA) has recognized the ubiquity of resource allocation problems for the last three decades, at least. The uncertainty of future benefits clearly provides an opportunity for application of standard approaches for modeling uncertainties—such as decision trees, influence diagrams, and Monte Carlo simulation, and for using utility functions to model decision makers' risk preference. The problem of many objectives is an opportunity to apply multiattribute utility or value models—or similar techniques—to decompose a difficult multidimensional problem into a series of relatively simply judgments about one-dimensional preferences and relative trade-off weights. The use of these models in real-world organizations also requires state-of-the-art techniques and processes for eliciting probabilities and preferences, performing sensitivity analyses, and presenting clear and compelling results and recommendations. As such, these problems clearly fall within the main-stream of DA methods and practice.

However, in contrast to traditional DA, it is also important to recognize that resource allocation problems are *portfolio* problems, where the decision makers must choose the best subset of possible investments subject to resource constraints. As such, more traditional techniques of operations research (OR), such as constrained optimization and mathematical programming, are directly relevant. With respect to trading off risk and return, the considerable body of work on applications of OR to financial portfolios, including the seminal work of Markowitz, is obviously relevant. Notwithstanding this connection, finance and DA use different approaches and make different assumptions for assessing and modeling risk preferences. This may explain why portfolio optimization is not typically considered part of routine DA practice. One notable exception is in the area of resource allocation problems with multiple objectives, where the possibility of linking multiattribute value/utility models to portfolio optimization was established in the late 1970s with work done by Keefer, Kirkwood, and their colleagues.

Despite these promising developments, applications that combine the best of decision analytic models with mathematical programming seems to have received only limited attention. It is much more common for DA practitioners to solve portfolio problems with heuristic approaches, such as prioritization based on benefit-cost ratios. Even more surprisingly, there has been limited development of new theory or methods for portfolio DA. One reason may be that DA is an interdisciplinary field, and many decision analysts have backgrounds in statistics, psychology, or other disciplines unfamiliar with the significant recent advances in optimization tools. Another possibility is that these decision analysts are more concerned with data-driven probability models or elicitation of expert judgments.

More fundamentally, however, there is the very real possibility that portfolio methods have seemed inapproachable because of the sheer size of the problems. In large organizations, there can easily be hundreds or thousands of projects and programs being considered for inclusion in a portfolio. Assessing multiple evaluation criteria across hundreds or thousands of projects can be a daunting task. In addition, if probability assessments are required, then good DA practice typically emphasizes careful and rigorous assessment and modeling of the uncertainties. This amounts to building hundreds of models and assessing probability distributions for many thousands of uncertain parameters. In my own experience working on resource allocation in hospitals and large healthcare systems in the U.S., I have found that many of those organizations struggle to develop the most straightforward deterministic financial models for more than a small number of projects. Conducting a full-blown decision analysis for each project is simply infeasible.

I have made similar observations in other resource allocation settings, including prioritizing R&D project portfolios at large manufacturing firms and evaluating the effectiveness of counter-terrorism risk management strategies. Stated quite simply, the most critical scarce resources are the resources required to obtain and evaluate data, assess probabilities, construct models, and evaluate alternative portfolios of projects. Often, decision makers start with only limited information about each project and portfolio, and they need guidance about which projects to analyze, and at what level of detail. Turning to this dissertation, then, there is a clear need for novel methods that combine DA modeling of probabilities and preferences with an appropriate OR-based portfolio optimization framework. Methods that offer a practical approach to guiding decision makers about which projects to analyze further would make these methods even more useful than they already are.