LECTIO PRAECURSORIA **5.6.2014**

EEVA VILKKUMAA

Doctor Custos, Doctor Opponent, ladies and gentlemen.

In our everyday lives, we often make decisions about selecting a subset of available alternatives subject to limited resources. For instance, when shopping for groceries, we select food items to purchase such that the combined price of these items does not exceed some predetermined budget. We typically select these items based on multiple criteria such as taste and nutritional content. When considering the relative importance of these criteria, and how well the different food items perform on these criteria, we may need to consider the views and preferences of other people living in the same household. Depending on our past experiences, we may feel the need to be prepared for uncertain events such as surprise visits by friends or relatives. We may also wish to consider the ease at which the selected food items can be combined into meals, and take into account some intriguing offers.

Despite these complicating factors, most of us manage to do our grocery shopping without having to apply formal decision support methods. However, grocery shopping belongs to a more general class of so-called *project portfolio selection problems*, which can be formulated as follows: which alternatives or *projects* should be selected such that this combination or *portfolio* of projects would offer the highest possible value in return for the resources spent? Other examples of such problems include the development of a strategic research agenda at a research institution, the selection of R&D projects in a company, or the selection of bridge maintenance projects by a governmental transportation agency.

These kinds of decision problems often deal with large capital investments and may involve dozens or even hundreds of project candidates. For instance, Intel and Samsung each spent over 10 billion dollars on R&D last year. The use of formal methods for supporting such large-scale decision problems typically leads to significant improvements in cost-efficiency, but also improves the transparency and defensibility of the decision process.

So, how exactly do formal methods contribute to decision-making? First of all, methods of mathematical programming are needed to tackle the issue of combinatorial explosion, which stems from the exponential growth in the number of possible project portfolios with the number of project candidates. Every time a project candidate is added, the number of possible portfolios is doubled. With 20 projects, the number of possible portfolios is over one million. With 33 projects, the number of portfolios exceeds the number of people on earth. With 60 projects, there are over one billion billon portfolios, which is roughly the number of insects on earth.

With so many possibilities to choose from, it is not surprising that relying purely on intuition usually results in selecting a suboptimal portfolio, even if the information about the projects' values was perfectly accurate. The increasing computational power of modern-day computers and the development of powerful optimization algorithms have made it possible to optimally solve relevant-sized portfolio selection problems within a reasonable time. Moreover, formulating the portfolio selection decision as a mathematical optimization problem makes it easy to accommodate different kinds of constraints and interdependencies between the projects – for instance, that project A can only be selected if project B is selected, or that the combined cost of selecting both projects C and D is lower than the sum of their individual costs.

Besides mathematical programming, project portfolio selection processes can also benefit from methods of decision analysis. Such methods help decision makers to systematically consider their objectives, values**,** assessments about uncertainty, and attitudes toward risk, and to quantify these as parameters for a mathematical decision model which generates decision recommendations. These parameters may correspond to, for instance, importance weights for multiple criteria with regard to which the projects are evaluated, or to the projects' criterion-specific values. This Thesis, in particular, develops decision analytic models which help synthesize the views and preferences of multiple stakeholders, and support the optimal design of project evaluation and selection processes when the values or costs of these projects are uncertain.

Consider first a situation in which two stakeholders need to decide which four projects to select out of 13 candidates. The stakeholders have different preferences so that their individually preferred portfolios consist of different projects. One way to obtain a decision recommendation for this group of two people would be to assign weights for the stakeholders' views, such as 40% and 60%. This, however, would imply making trade-offs between the stakeholders, which can prove to be difficult – starting with the question of *who* in fact would be in the position to make such trade-offs.

This Thesis develops a method to generate portfolio recommendations under incomplete information about the stakeholders' importance weights. That is, instead of precise estimates for these importance weights, one can simply state that the weight of both stakeholders' views needs to be at least 25%, for instance. With such incomplete information, no single best portfolio can usually be determined. However, this information can be used to identify so-called 'non-dominated' portfolios which are not outperformed by any other portfolio for all importance weights compatible with the information.

These non-dominated portfolios can then be used to generate project-specific recommendations. In particular, the set of non-dominated portfolios divides the projects into three categories: (i) core projects that are included in all non-dominated portfolios and should thus be selected, (ii) exterior projects that are not included in any of the non-dominated portfolios and should therefore be rejected, and (iii) borderline projects that are included in some non-dominated portfolios, but not all. Such a three-way classification helps the group of stakeholders to identify points of agreement (that is, core and exterior projects), and focus further negotiations on the points of disagreement (that is, borderline projects).

The future values of project candidates are typically uncertain. In some cases, these uncertainties are best captured by a set of scenarios describing alternative futures such that the projects' future values are contingent on which scenario is realized. In this setting, the problem is to select a project portfolio which would perform relatively well across the different scenarios. One way to do this is to assign probabilities for the scenarios, evaluate the projects' performance in each scenario, and select the portfolio with the highest expected performance.

The recommended project portfolio resulting from such an approach can be very sensitive to even small changes in scenario probabilities. Also, this approach does not take into account the impact that the selected projects may have on the probabilities of reaching the different scenarios. For instance, projects undertaken to prepare for a world of conflicts (such as increased armament and protectionism) may well increase the odds of this scenario being realized. If such interdependencies between the selected projects and scenario probabilities are not accounted for, then the selected portfolio is likely to be suboptimal.

This thesis develops a portfolio selection method that admits incomplete and project-dependent scenario probability information through statements such as 'Scenario 2 is more probable than scenario 1' and 'The selection of project B makes the probability of scenario 4 higher than 15%'.

This method helps identify non-dominated portfolios that perform relatively well across the scenarios, and help steer the course towards the desirable scenarios. Again, the non-dominated portfolios divide the projects into three categories – core projects that should be selected regardless of which scenario is realized, exterior projects that should definitely not be selected, and borderline projects. This information could be utilized by fully investing in the core projects, but making only small initial investments in the borderline projects, and then abandoning or expanding these projects when more accurate information about the future operational environment becomes available.

Because decisions are always made under uncertainty about the future, it is likely that the selected portfolio will be suboptimal *ex post*, that is, after the uncertainties have been resolved. However, the quality of a decision cannot and should not be evaluated based on the *ex post* portfolio value. Rather, emphasis should be put on evaluating the quality of the decision process; that is, whether all information that is available at the time of the decision is transformed into decision recommendations in a transparent, rational, and logical manner. For instance, you should not feel irrational for having invested time and effort in developing a pandemic contingency plan simply because a pandemic emergency hasn’t occurred yet.

Some issues resulting from uncertainties can, however, be alleviated through careful modeling, such as the systematic overestimation of the value of the selected portfolio, or the systematic underestimation of its cost. Consider, for instance, six project candidates of which the decision-maker wants to select the best two. Each of these projects has equal true value, which is not known by the decision maker at the time of the selection decision. The decision maker observes uncertain but *unbiased* estimates about the projects’ values, meaning that the values of some projects have been overestimated and some underestimated, but, on average the estimation errors cancel each other out.

However, the values of those projects that are *selected* are more likely to have been *overestimated* rather than underestimated, causing the decision maker to feel disappointed after the true values of these projects are eventually realized. Similarly, if the objective is to minimize costs, one is more likely to select those projects whose costs have been *underestimated* rather than overestimated. This so-called post-decision disappointment phenomenon, which applies also when the projects' true values or costs are unequal, has significant practical implications. Indeed, it has been reported that the realized costs of public infrastructure projects are on average 28% higher than estimated, which not only causes mistrust in those responsible for providing the cost estimates, but also makes it difficult to balance the budget.

This thesis shows that by adjusting the uncertain value estimates through Bayesian methods, one can have more realistic expectations about the value of the selected portfolio. In particular, this adjustment is done by associating a prior distribution for the true values of projects that are similar to those considered now, and a distribution for the estimation error. The adjusted value estimate lies somewhere in between the prior mean and the estimate for the candidate project. The larger the estimation error variance compared to the prior variance, the more forcefully the estimate should be adjusted towards the prior mean. By adjusting the uncertain estimates this way, the estimated value of the selected portfolio will, on average, coincide with its true value.

This thesis also shows how the Bayesian modeling of uncertainties helps determine the *expected* value of obtaining additional value estimates prior to actually acquiring these estimates. Consider ten projects of which the decision maker would, based on current information, select projects *A* through *D*.

The re-evaluation of project *E* could reveal that this project is actually more valuable than what was initially thought. If this is the case, the decision-maker would revise her initial decision by including project *E* and excluding project *D,* and could expect to gain 400,000 euros from this exchange. Computing the expected values of obtaining additional estimates for the different projects, and taking into account the costs of acquiring such estimates, the decision-maker can design the evaluation process cost-efficiently. As a general rule, most additional value is obtained from re-evaluating projects that have particularly uncertain initial estimates near the selection threshold. By following this rule, it often suffices to re-evaluate only a small subset of project candidates. Because project evaluation can be time consuming and costly, this result can significantly improve the cost-efficiency of the selection process.

Recall the Bayesian adjustment of uncertain value estimates towards the prior mean. If the value estimates are very uncertain, then the Bayesian adjustment process becomes highly conservative. In other words, by adopting the Bayesian approach, a very uncertain project can be expected to be mediocre at best. Yet, those projects that eventually *do* yield extremely high value are often those whose values are particularly difficult to estimate beforehand. For instance, in early 1980s, hardly anyone believed in the project by Mario Capecchi to develop gene targeting in mammalian cells. This project, however, turned out to be a huge success, earning a Nobel Prize for its instigator in 2007.

The dilemma can be summarized as follows: To be able to correctly predict few exceptionally valuable projects, one must expect to feel disappointed in the realized values of many other projects. On the other hand, if one wishes to minimize this disappointment, it is practically impossible to predict exceptionally valuable projects. The question is, then: what kinds of project funding policies would retain the realism of the Bayesian adjustment process but, at the same time, facilitate the identification and selection of those few projects that yield exceptionally high value?

This thesis develops a multiperiod project funding model in which on-going projects can be abandoned prior to completion based more accurate project evaluations obtained over time. The results of this model suggest that to fund exceptionally valuable projects, one should initially launch a large number of projects, but abandon a large share of them after some time so that funding is continued only for those few projects that, based on observed performance, show the highest potential. Launching more projects decreases the risk of missing out on projects that could have been exceptional; also, decisions about long-term funding can be made based on more accurate information about the projects' potential.

This policy is different from the policy that maximizes average short-term portfolio performance, which is to simply grant full, long-term funding to those projects that appear to be the best based on initial evaluation. This implies that a policy which is optimal in terms of finding exceptional projects can be cost-inefficient in the short-term and, on the other hand, a policy that maximizes the short-term performance may fail to pick those few, exceptional ideas.

This concludes my lectio praecursoria, which, I hope, makes it easier for the audience to follow the examination of my Thesis.

I ask you Professor Jeffrey Keisler, as the opponent appointed by the Aalto University School of Science to make any observations on the thesis which you consider appropriate.