On the Choice of Baselines in Multiattribute Portfolio Analysis: A Cautionary Note

Robert T. Clemen, James E. Smith
Fuqua School of Business, Duke University, Durham, North Carolina 27708 {clemen@duke.edu, jes9@duke.edu}

In multiattribute portfolio optimization, a decision maker must evaluate a number of projects on multiple dimensions and then select the set of projects that optimizes the portfolio's overall value. In this note, we discuss the importance of establishing an appropriate baseline score for *not doing a project* in multiattribute portfolio analysis. We believe that practitioners often implicitly assume that not doing a project results in the worst possible score on all attributes. We argue that this assumption is often inappropriate and may lead to incorrect recommendations.

1. Introduction

When an organization has more ideas for projects than resources for doing them, they must decide which projects to pursue. For example, a pharmaceutical company must decide which compounds in its research pipeline to fund, a hospital must choose which capital budgeting projects will move forward, and a government agency must determine which of a number of activities deserve part of the agency's budget. Our interest is in the multiattribute portfolio problem where the decision maker must evaluate options on several attributes, with the goal of choosing the set of projects that provides the best overall value given the resource constraints. In this setting, projects typically have both financial and nonfinancial attributes, and the decision maker must make tradeoffs among the attributes when selecting projects. Kleinmuntz (2007) provides a review of recent work on portfolio optimization, including the multiattribute portfolio optimization problems of the kind we discuss here; Kirkwood (1997) provides a textbook treatment. We will discuss some specific applications in §7 below.

In this note, we discuss the importance of choosing an appropriate baseline score for *not doing a project* in these multiattribute portfolio optimization models.

We believe that practitioners using standard multiattribute scoring techniques often implicitly adopt a baseline that assumes not doing a project results in the worst possible score on all attributes. We argue that this implicit baseline is often inappropriate and may lead to incorrect recommendations. Although our discussion and examples focus on project selection problems, the issues we discuss are somewhat more general and may arise whenever multiattribute models are used to select several options from a larger set of possible options. For example, one could use a multiattribute portfolio model to choose staff to form a team.

In §2, we introduce the "CBA Associates" case (Kleinmuntz 2000), which we will use as an illustrative example throughout this note. In §§3 and 4, we present two different approaches to analyzing this example, the first based on a standard multiattribute value analysis and the second based on "pricing out" the nonfinancial attributes. The contrasting recommendations from these approaches will highlight the importance of the assumed baseline, which we discuss in §§5 and 6. In §7, we identify some specific examples of what we believe to be erroneous choices of baselines in published applications. In §8, we offer some final thoughts.

2. The CBAA Example

Within the fictional consulting firm CBA Associates (CBAA), the information technology group is in the process of choosing a set of software development projects to pursue. We will assume that CBAA has a fixed staff of programmers with a total of 2,500 person days that can be allocated to the chosen projects.[1]

The data for the problem are shown in Table 1. There are eight projects (A–H) to consider, and each project is scored on three attributes. The first attribute is the project's financial contribution, measured in dollars. The other two attributes are risk and fit. Risk reflects the likelihood that the project will lead to a marketable product. Projects are placed in one of three risk categories: the least risky projects are categorized as safe, the intermediate ones as probable, and the riskiest projects as uncertain. Fit, scored on a scale from 1 (worst) to 5 (best), is a measure of a project's compatibility with CBAA's primary consulting business and can be thought of as a rough measure of the incremental consulting revenue that a project might generate. The final column in Table 1 shows the number of days of programmer effort required by each project.

Inspecting Table 1, we note that two projects, B and E, have negative financial contributions. However, each of these projects scores well on one of the other dimensions, which could offset the negative financial contribution enough so that these projects would be included in the optimal portfolio. Also, although E is dominated by both C and D in terms of the three specified attributes, E requires fewer programmer days. Finally, F is dominated by G and requires more programmer days than G. Thus, we would not expect to see F in the optimal portfolio unless G is also included.

3. A Multiattribute Value Model and Analysis

We first consider a standard multiattribute value formulation of the portfolio optimization problem,

[1] Note that the CBAA case initially assumes that there are 2,000 person days available and later asks students to consider the possibility of hiring additional programmers to increase the days available to 2,250 or 2,500. We focus on the 2,500-person-day scenario, because it highlights the effect of the baseline in a more striking way.

Table 1 **The Basic Data for the CBAA Example**

Project	Financial contribution ($)	Risk	Fit	Days required
A	200,000	Uncertain	5	800
B	−13,750	Probable	5	250
C	125,000	Safe	4	700
D	307,500	Safe	3	650
E	−1,250	Safe	2	350
F	393,000	Uncertain	2	800
G	442,500	Uncertain	2	600
H	265,000	Probable	1	400

following the standard approach described in, for example, Kirkwood (1997, Chap. 8). In this approach, we begin by normalizing the scores for each attribute to range from 0.0 (worst) to 1.0 (best). The normalized scores for fit's intermediate levels are assumed to be proportional to their original scores in Table 1. For risk, the intermediate level, probable, is assumed to fall halfway between 0 (for uncertain) and 1 (for safe). The resulting normalized scores are shown in Table 2.

The next step is to assign weights for each attribute. There are a variety of ways to assess these weights (e.g., by using the "swing weight" approach), but we need only concern ourselves with the final result of the assessments. Following the case, we will assume that fit has a 50% weight, and risk and financial contribution each have 25% weights. These weights imply that going from the worst level to the best level on the fit attribute is worth twice as much as going from the worst level to the best level on either of the other two attributes. The value score for each project is then given by the weighted sum of the normalized scores.

Table 2 **A Standard Multiattribute Value Model for the CBAA Example**

Project	Financial contribution	Risk	Fit	Value score	Days required	Go?
A	0.47	0.00	1.00	0.617	800	0
B	0.00	0.50	1.00	0.625	250	1
C	0.30	1.00	0.75	0.701	700	1
D	0.70	1.00	0.50	0.676	650	1
E	0.03	1.00	0.25	0.382	350	1
F	0.89	0.00	0.25	0.348	800	0
G	1.00	0.00	0.25	0.375	600	0
H	0.61	0.50	0.00	0.278	400	1
Weight	0.25	0.25	0.5			
Total value						2.66
Total days required						2,350

The final step is to select a portfolio of projects to pursue by finding the set of projects that maximizes the total value score, subject to the constraint that the number of programmer days required does not exceed the total programmer days available. To calculate the total value of the portfolio and the total days required, we simply sum the value scores and days used for those projects that CBAA chooses to pursue. To formalize the optimization problem, let v_i and d_i denote the value score and programmer days required, respectively, for project i. We then define a binary variable x_i that is equal to 1 if project i is chosen, and 0 otherwise. The optimization problem can then be written as

$$\max_{x_i \in \{0,1\}} \sum_i x_i v_i$$

$$\text{subject to } \sum_i x_i d_i \leq 2{,}500. \qquad (1)$$

The optimal settings for the binary decision variables for this case are shown in the "Go?" column of Table 2. According to this analysis, the optimal portfolio of projects is B, C, D, E, and H. This solution uses a total of 2,350 programmer days, 150 less than the total available.

Note that this multiattribute value analysis assumes that CBAA's preferences are additive across the attributes for the individual projects (no interactions among the attributes) and across the projects (no interactions among projects). These additivity assumptions are quite standard. For a careful discussion of these assumptions in the context of the portfolio problem, see Golabi et al. (1981).

4. A "Pricing-Out" Model and Analysis

Rather than normalizing the scores and assigning weights for the attributes, we can instead "price out" the raw scores for risk and fit by converting them to their financial equivalents. For the pricing-out analysis to be comparable to the earlier multiattribute value analysis, it is important to choose "prices" for the nonfinancial attributes so changes in attribute levels have the same relative impact on the overall value measure as in the multiattribute value analysis. As mentioned earlier, the weights in the multiattribute value analysis imply that going from the worst level

Table 3 A "Pricing-Out" Analysis for the CBAA Example

Project	Financial contribution ($)	Risk ($)	Fit ($)	Net contribution ($)	Days required	Go?
A	200,000	−456,250	912,500	656,250	800	1
B	−13,750	−228,125	912,500	670,625	250	1
C	125,000	0	684,375	809,375	700	1
D	307,500	0	456,250	763,750	650	1
E	−1,250	0	228,125	226,875	350	0
F	393,000	−456,250	228,125	164,875	800	0
G	442,500	−456,250	228,125	214,375	600	0
H	265,000	−228,125	0	36,875	400	0
Total value					$2,900,000	
Total days required					2,400	

to the best level on fit is worth twice as much as going from the worst level to the best level on either of the other two attributes. Going from the worst to best levels of financial contribution corresponds to an increase of $456,250. Thus, we set the contribution of poor fit (1) to be $0 and best fit (5) to be $912,500. Similarly, we take the cost of doing a risky project to be equal to $456,250, and the cost of a safe project to be $0. As in the multiattribute value analysis, the intermediate levels for all attributes are assumed to be proportional between the extremes.

The net contribution for a project is then calculated by totaling the contributions from the financial and nonfinancial attributes. These values are shown in Table 3. With a little algebra, we see that the net contribution values shown in Table 3 are linearly related to the value scores shown in Table 2:

$$\text{Net Contribution} = \$1{,}825{,}000 \times \text{Value Score}$$
$$- \$470{,}000. \qquad (2)$$

Thus, as required, the rankings of the individual projects are unaffected by this change in scale in the overall value measure.

To determine the optimal portfolio of projects, we then solve an optimization problem like (1), but with v_i representing the net contribution shown in Table 3 rather than the value score shown in Table 2. The optimal solution is shown in the last column in Table 3. Note the difference between this solution and the one shown in Table 2: E and H have been dropped and A has been added.

5. Why the Baseline Matters

Although the multiattribute value and pricing-out analyses are consistent in their rankings of individual projects, this does not ensure that the same portfolio of projects will be included in the optimal portfolio. Note that optimization problem (1) implicitly assumes that *not doing* a project leads to a value score of zero (in the multiattribute value analysis) or net contribution of zero (in the pricing-out analysis). The key issue is that these two analyses assume different baseline values for not doing a project. In the multiattribute value analysis, the assumed value score of zero for not doing a project reflects an implicit assumption that not doing each project gives the worst possible result on each attribute. Specifically, in the CBAA example, the multiattribute value analysis assumes that not doing a project

- results in a financial contribution of −$13,750,
- is in the worst possible risk category (uncertain), and
- is a very poor fit with CBAA's consulting business.

In contrast, the zero levels for the pricing-out analysis assume that not doing a project

- results in a financial contribution of $0,
- is in the best possible risk category (safe), and
- is a very poor fit with CBAA's consulting business.

The first two assumptions are inconsistent, and this is what leads to different optimal solutions for the two formulations.

Mathematically, it is easy to reformulate the optimization problem (1) to allow nonzero values for not doing a project. If we let v_i° represent the value of not doing project i, the optimization problem can be rewritten as

$$\max_{x_i \in \{0,1\}} \quad \sum_i x_i v_i + (1 - x_i) v_i^\circ$$
$$\text{subject to} \quad \sum_i x_i d_i \le 2{,}500. \tag{3}$$

In the objective function, the binary variable x_i is equal to one when the project is selected, and $(1 - x_i)$ is equal to one when the project is not selected. Thus, the ith term of the summation is equal to v_i if project i is selected and v_i° otherwise. Alternatively (and equivalently), rearranging terms in the objective function

and dropping a constant, we can take the objective function to be $\sum_i x_i (v_i - v_i^\circ)$; this formulation emphasizes the fact that the portfolio optimization problem seeks to maximize the total *incremental* value associated with doing projects versus not doing them.[2]

To illustrate the difference in the formulation of the optimization problem in the CBAA example, let us assume that for each project, not doing that project has no financial contribution ($0), is safe, and has a fit score of 0, as assumed in the pricing-out analysis. Thus, on the 0–1 scales in the multiattribute value analysis, doing nothing has a financial contribution value of 0.03, a risk value of 1 and a fit value of 0, resulting in $v_i^\circ = 0.258$. A corrected multiattribute value analysis that focuses on the incremental contributions of projects is shown in Table 4. The earlier multiattribute value analysis in Table 2 used the project value v_i instead of the incremental value $(v_i - v_i^\circ)$, and consequently, all of the values are overstated by 0.258. Adjusting the values downward by this amount, it becomes clear that E and H together are worth considerably less than A and require similar numbers of programmer days (750 days for E and H, and 800 days for A). Without this adjustment, E and H together appear to be worth more than A,

Table 4 Corrected Scores for the CBAA Example

Project	Financial contribution	Risk	Fit	Value score	Inc. value score	Days required	Go?
A	0.47	0.00	1.00	0.617	0.360	800	1
B	0.00	0.50	1.00	0.625	0.367	250	1
C	0.30	1.00	0.75	0.701	0.443	700	1
D	0.70	1.00	0.50	0.676	0.418	650	1
E	0.03	1.00	0.25	0.382	0.124	350	0
F	0.89	0.00	0.25	0.348	0.090	800	0
G	1.00	0.00	0.25	0.375	0.117	600	0
H	0.61	0.50	0.00	0.278	0.020	400	0
Weight	0.25	0.25	0.5				
Total value							1.59
Total days required							2,400

[2] Note that in some applications the cost of not doing a project may also not be zero, as is assumed in the optimization problems of (1) and (3). This can be easily incorporated by rewriting the cost function in the same way we have rewritten the benefit function, e.g., by replacing the left side of the constraint in (3) with $\sum_i x_i d_i + (1 - x_i) d_i^\circ$, where d_i° denotes the costs (here programmer days required) if you do not do project i.

which is why E and H, rather than A, appear in the multiattribute value solution.

Rather than solve an optimization problem of the form of (1), some authors prefer to prioritize projects by their benefit-to-cost ratios (here, v_i/d_i) and recommend pursuing those projects with higher ratios before pursuing those with lower ratios (see, e.g., Phillips and Bana e Costa 2007). This benefit–cost approach can be seen as an intuitive, approximate approach to solving the optimization problem (1). The benefit–cost approach to portfolio optimization also implicitly assumes that not doing a project has zero value, and therefore has the same issues as the multiattribute value approach. To counter this, in the benefit–cost approach one should focus on incremental benefit-to-cost ratios (here $(v_i - v_i^\circ)/d_i$) rather than the benefit-to-cost ratio (v_i/d_i). This leads to an approximate solution of the optimization problem (3). Again v_i°, the value of not pursuing a project, should play a key role in calculating project priority.

6. How to Evaluate Not Doing a Project

This reformulation of the portfolio optimization problem begs the question of how to score not doing a project? Indeed, the main message of this note is that this question must be considered carefully in portfolio applications. In general, we believe it is best to consider each project individually by explicitly scoring not doing the project on each attribute and using these scores to calculate an overall project-specific score, v_i°, for not doing project i. By explicitly scoring both doing and not doing a project on each attribute, the decision maker must explicitly consider the incremental benefit of doing versus not doing a project.

To illustrate some of the subtleties and complexities of this issue, consider how not doing a project might be scored in the CBAA example:

• *Financial contribution.* Here it seems most natural to assume that not doing a project yields a financial contribution of zero, as assumed in the pricing-out analysis. However, we can imagine cases where other assumptions would seem natural. For example, if a project allows the firm to reduce or avoid a loss, we might assign a negative contribution for not doing the project. Though there is some flexibility in how

one measures financial contribution, in our view, it is hard to justify the implicit assumption in the multiattribute value analysis of the CBAA case that not doing a particular project results in a financial contribution of −$13,750.

• *Risk.* Here again there is some flexibility in how we might measure the risk associated with doing versus not doing a project. Nevertheless, it seems hard to justify the implicit assumption in the multiattribute value analysis that not doing a project should be put in the worst possible risk category (uncertain). In utility theory terms, we might think of this risk attribute as being a proxy for a "risk premium" that should be subtracted from the (expected) financial contribution to arrive at a certainty equivalent. In our view, given that it is easy to complete a "null project" and there is little uncertainty about its "null value," it would seem more appropriate to assume that not doing a project is not risky—or, more precisely, adds no additional risk to the organization—and should be assigned the "safe" rating.

• *Fit.* If we think about fit as a rough measure of the uncertain follow-on consulting revenues, then it seems natural to assume that not doing a project results in no such revenue, as is assumed in both the multiattribute value and pricing-out analyses. However, some projects could have negative "fits" (e.g., by cannibalizing an existing revenue stream), and in such cases it may be inappropriate to assume that not doing a project would earn the worst possible fit score.

Note that whenever the assumptions have been in conflict (i.e., for financial contribution and risk), we have sided with the pricing-out analysis as making the "correct" assumption. In general, we suspect that the framing of the pricing-out analysis forces users to think more carefully about the implications of not doing a project and to measure differences from the zero reference point consistently. In the multiattribute value approach, the standard practice is to assign a utility score of 0 to the worst possible outcome, and the framing of the optimization problem (1) does not encourage users to think explicitly about how one should score not doing a project. Consequently, in this approach, the implicit assumption that not doing a project leads to the worst possible score on every attribute seems very natural.

7. Examples Noted in the Literature

In hindsight, it is clear that the choice of a baseline score for not doing a project should play an important role in project prioritization. To convince readers that this issue is perhaps not so obvious, in this section we highlight a few places in the literature where we believe leading researchers and practitioners have implicitly made inappropriate assumptions about the value of not doing a project. By highlighting these examples, we hope to help others avoid similar problems in their own work.

7.1. Capital Budgeting in Health-Care Organizations

Kleinmuntz and Kleinmuntz (1999) describe the use of multiattribute portfolio methods to prioritize capital budget expenditures for health-care organizations. Their proposed approach exactly follows the multiattribute value analysis as described above: proposals are evaluated on a number of attributes including financial contribution, market share, physician relations, and operating efficiency. Because these attributes are not described in much detail in their paper, it is hard to think carefully about what level of performance is appropriate for not doing a project. However, we note that, as in the CBAA example, the worst case financial contribution is −$150,000 for one project, whereas the rest of the proposals have positive contributions. All of the projects are assigned normalized scores ranging from 0 (for −$150,000) to 1 (for $1,467,000). Then, just as in the CBAA example, the multiattribute value analysis implicitly assumes—we think incorrectly—that not doing each project leads to a financial contribution of −$150,000.

7.2. Project Selection in a Telecommunications Company

Lindstedt et al. (2008) describe the use of multiattribute portfolio methods to prioritize capital expenditures for a telecommunications company. Their work is novel in that it adopts a robust approach that works with ranges on the attribute weights rather than assuming precise weights, but the basic model underlying their approach is essentially the same as the multiattribute value approach. Attribute scores are normalized to range from 0 (worst) to 1 (best). The objective function totals the scores of the funded projects and implicitly assumes an overall score of zero for projects that are not funded. Although the specific attributes are not described in great detail in the paper, the attributes include a "net profit" attribute that ranges from "large losses" to "large profits," and not doing a project is implicitly assumed to result in "large losses." The authors also consider a risk attribute and implicitly assume that not doing a project results in the worst possible risk level, e.g., not doing a project results in a "severe risk of miscalculation."

7.3. Resource Allocation with the National Reconnaissance Office

Parnell et al. (2002) describe an application of multiattribute portfolio optimization for selecting projects to support the National Reconnaissance Office's customer support organization. Their multiattribute portfolio analysis also follows the multiattribute value approach: the attribute scores range from 0 for the worst level to 10 for the best, and their objective function implicitly assumes an overall score of zero for projects that are not funded. Though the paper is not clear about exact definitions of the various attributes, there are a few attributes for which it would seem that not doing a project would score better than the worst possible level for that attribute. For example, "management risk" is defined in terms of the likelihood of completing the project on time, and "sustainment costs" represent the ongoing costs of supporting a project that is undertaken. As in the previous examples, not doing a project would seem to be less risky than doing some projects—it is, after all, easy to do a "null project" on schedule! The worst level for the "sustainment costs" attribute is defined as a "costly solution," and the best as being when the initial costs include funding for ongoing support. Here the cost of ongoing support for a project that is not done would seem to be zero, and in terms of future operating costs to the reconnaissance office, this would seem to be equivalent to the case where future operating costs are provided for in the initial evaluation. Thus, we believe that not doing a project should receive the best score for sustainment costs rather than the worst score, as implicitly assumed in Parnell et al. (2002).

In addition to these three specific references, we have surveyed a number of other multiattribute portfolio applications and found similar issues in many

studies. A notable exception is Golabi et al. (1981), the first article to describe the use of multiattribute portfolio techniques. These authors explicitly discuss the value of not doing a project and describe how it was assessed in their application. However, Kirkwood's (1997) influential decision analysis textbook, although it describes the multiattribute value approach for multiattribute portfolio analysis in detail, does not discuss the importance of determining the value of not doing a project, and implicitly assumes that not doing a project leads to the worst possible score on each evaluation measure.[3] Similarly, Phillips and Bana e Costa (2007) describe a multiattribute value approach for multiattribute portfolio analysis. Although they focus on prioritizing projects using benefit–cost ratios rather than explicitly solving an optimization problem of the form of (1), they make the same implicit assumption that not doing a project leads to the worst possible score on each evaluation measure.

8. Conclusion

Though it is clear that the choice of a baseline score for not doing a project should play an important role in project prioritization, we believe that the standard multiattribute value approach (described in texts and many papers) may lead practitioners to inadvertently make inappropriate assumptions about the value of not doing a project. Ultimately, we believe this oversight may stem from incorrectly viewing the multiattribute portfolio problem as a ranking problem rather than a portfolio problem, i.e., as if the problem is to choose one of the projects to pursue rather than to choose whether or not to pursue each project. In the portfolio setting, the challenge is not to identify the "best" project, but to identify the set of projects that maximize the total benefit for the portfolio given the resource constraints. To do this, one

must think carefully about the implications and value of not doing each project.

In conclusion, we emphasize that we believe that the multiattribute optimization approach to selecting a portfolio of projects is very useful as well as sound and defensible. Though we believe that many may have made the mistake we highlight, we would also like to emphasize that such mistakes are easily avoided if one thinks carefully about the implications of not doing a project.

References

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

Kirkwood, C. W. 1997. *Strategic Decision Making: Multiobjective Decision Analysis with Spreadsheets.* Wadsworth, Belmont, CA.

Kleinmuntz, C. E., D. N. Kleinmuntz. 1999. A strategic approach to allocating capital in health-care organizations. *Healthcare Financial Management* **53**(4) 52–58.

Kleinmuntz, D. N. 2000. CBA Associates. Department of Business Administration, University of Illinois at Urbana-Champaign, Champaign, IL.

Kleinmuntz, D. N. 2007. Resource allocation decisions. W. Edwards, R. F. Miles, D. von Winterfeldt, eds. *Advances in Decision Analysis: From Foundations to Applications.* Cambridge University Press, Cambridge, UK, 400–418.

Lindstedt, M., J. Liesiö, A. Salo. 2008. Participatory development of a strategic product portfolio in a telecommunication company. *Internat. J. Tech. Management* **42**(3) 250–266.

Parnell, G. S., G. E. Bennett, J. A. Engelbrecht, R. Szafranski. 2002. Improving resource allocation within the National Reconnaissance Office. *Interfaces* **32**(3) 77–90.

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

[3] After reviewing a draft of this paper, Kirkwood added a discussion of this implicit assumption to his list of additions and corrections for the book (see http://www.public.asu.edu/~kirkwood/SDMBook/sdmadd.htm).