Software Project Selection Methods

Dr. Efrim Mallach

Note: This material was first posted here via a direct cut-and-paste from the book manuscript. Certain items didn't come across properly. Hopefully this isn't enough to impact the usefulness of the material.

11.3. Choosing and prioritizing: the project portfolio

As you read at the beginning of this chapter, every organization is faced with a wealth of possible information systems projects. Some are large, others small. Some are certain of success, others border on basic research. Some will make the firm more efficient, others are intended to enhance its competitive position. Some can be completed in a matter of days, others may take years. Some used tried-and-true technologies, others introduce new ones. Some support the manufacturing function; others are for finance, marketing, engineering, or a group of departments. Some run on personal computers, others require a large-scale mainframe. Some ... the list of potential differences goes on and on.

In the face of all these differences, management must decide where to invest finite resources. This is the *project selection* dilemma. The group of projects that an organization decides to fund—its *project portfolio*—must meet the needs and satisfy, at least to a degree, the desires of everyone in that organization. To make sure that this happens, information system users and their managers must get into the act. In this section you will learn about some of the ways projects are selected and how users are involved in the selection process.

Cost-benefit analysis

(Refer to separate article in Quantitative Project Selection section.)

The balanced portfolio approach

A financial approach is appropriate when all the important information about a project can be reduced to financial terms. However, non-financial factors must also often be taken into account. One approach to considering them is to strive for a *balanced portfolio* of information systems projects.

The concept behind the balanced portfolio approach is that selecting projects based on any single criterion cannot yield the best possible results. Its logic is similar to that of an investor who wishes to balance a portfolio of financial investments. Such an investor usually holds some low-risk investments, such as bonds and public utility stocks, and some higher-risk investments, such as the stock of small firms or options. Low-risk

investments will provide some benefits while protecting against major losses. High-risk investments, while they may turn out to be worthless, may also increase in value dramatically. The correct mix varies from investor to investor but the concept is nearly universal.

We can apply the same concept to selecting information systems projects. The same factors, risk and expected financial return, are used. We can plot all the proposed projects on a grid such as this:

An organization might have the proposed projects indicated by the dots. Most fall near a line from the bottom left to the top right of the figure. This is not a coincidence. Risk and return are related. Safe bets tend to offer modest returns. Projects that aim at higher goals have a higher chance of missing their mark altogether.

Projects that fall into the top left quadrant of the grid offer high return with low risk. They should be given the highest priority. Unfortunately, there usually aren't many of these.

Projects that fall into the bottom right quadrant of the grid offer high risk with little potential return. They should be examined in an effort to reduce their risk or to increase their potential benefits. If neither is feasible, they should be given low priority.

The balancing act comes into play along the heavily populated diagonal axis of the grid, from the bottom left to the top right. An organization should carry out some high-risk projects, because not doing so will eliminate the possibility of obtaining their potential rewards. Yet it can't afford to carry out all high-risk projects, because it needs confidence that many of its development projects will succeed. By choosing a mix of projects spread along this line, the firm can undertake a balanced portfolio of information systems projects.

While the traditional balanced portfolio concept balances risk and potential return, the idea of balance can be extended to other areas as well. Additional areas in which organizations should strive to balance their information systems portfolio include:

. • Technology. A certain fraction of the organization's resources should be devoted to trial projects using new technologies. (This is related to, but not identical with, risk. Most information system project risks have nothing to do with technology.)

. • Functional areas supported. While it is not necessary to allocate resources to different parts of the organization—marketing, manufacturing, etc.—in precise proportions, no part of the organization should feel totally neglected.

. • Staff capabilities required. The projects selected must be balanced with respect to the programming languages, database systems, etc., with which the staff is familiar, taking both hiring plans and feasible training schedules into account.

. • Time frame. It's a good idea to have some projects that can be completed quickly and others that take longer.

While selecting a balanced project portfolio is not the entire answer to the project selection question, the concept of balance is good. It is a goal to strive for and a test that should be applied to any tentative selection of projects.

The Powell-Barkhi Method

Another approach to balancing multiple factors is due to N. Powell and R. Barkhi [POWE97]. This method considers two factors in selecting the projects to pursue. One is estimated project profitability, generally taken as its NPV or its ROI. The other is an estimate of the potential strategic and other intangible benefits of the project, collectively referred to as "fuzzy factors" in their work. The two figures allow information systems planners to locate each candidate project on a graph, as in the figure.

(Figure not yet drawn. It is a two-dimensional graph with one axis indicating financial return and the other indicating other benefits. A project in the top right corner scores high on both.)

Projects are then selected based on the relative desirability of different points on the grid. First, one draws the "envelope" of good projects, as indicated by the dotted line in the figure2. Any project not on this line is inferior to at least one of the projects on the line in terms of both factors: that is, its NPV or ROI is lower *and* it has lesser fuzzy benefits. Such a project should not be chosen, at least not at this stage.

The organization can then choose one or more of the projects on the envelope. It may have a minimum required ROI. This translates into a horizontal threshold on the graph. Any project below this threshold will not be chosen, no matter how great its potential fuzzy benefits. Similarly, there may be a vertical threshold representing a minimum level of fuzzy benefits. Any project that does not provide this level of strategic advantage (or other fuzzy benefits) will not be chosen, no matter how great its ROI or NPV. Other organizations may have different selection strategies.
If the selected project does not exhaust the organization's potential budget for IS projects, it is removed from the grid and the process repeats without it. Removing a project which was on the envelope changes the envelope. The new envelope will, in all likelihood, include one or more projects that were previously inside it. This allows projects to be considered for selection even if they were eliminated in the first pass through the set of candidates.

The Powell-Barkhi method can take risk into account as well. Risk is expressed as the standard deviation, or some other measure of statistical variability, of the position of a project's point on the chart. In the case of ROI or NPV, standard statistical methods can be used to obtain this figure. In the case of fuzzy factors, the way one determines the variability range depends on how projects are measured along this axis in the first place. If a project's fuzzy-factor "figure of merit" is, for example, the average of several people's ratings, the standard deviation of the mean can be used. However this is done, the result is a variability figure in each direction. These can be combined to create an oval

within which the project is expected to lie, rather than the single point of the figure. How one uses this added information depends on one's attitude toward risk. A risk-prone decision maker would tend to use the upper right side of the oval, while a risk-averse one would tend to use its lower left. Representing the projects in this form helps decision makers understand how they use multiple factors in the decision and how their approaches differ.

(This type of decision making is often called **multiple-criteria decision analysis**. If it interests you, you can learn more about it in any book on decision making methods or decision support systems.)

Project scoring

Project scoring is a decision method which recognizes that many factors affect the desirability of choosing one project over another, that candidate projects differ with respect to these factors, and that the firm must combine all these factors into an overall "desirability rating."

Project scoring begins by listing the attributes of a project that affect its desirability. The list will depend on whether we are choosing alternative ways to meet one stated need or alternative projects competing for their share of a finite budget. In selecting one of several ways to meet one need, the list will include such items as cost, schedule, degree to which the project meets user needs, fit to existing technologies in the organization, compatibility with existing data bases, quality of documentation and support, and so on. In prioritizing from a list of projects, it will include financial factors, consistency with corporate strategies, potential to create strategic advantage, fit with available resources, degree of urgency, and more. The leftmost column of the figure below shows one possible list.

SAINT ALEXIS HOSPITAL MEDICAL CENTER

PROJECT PRIORITIZATION WORKSHEET

PROJECT: Laboratory Instrument Interface

Indicator	Wgt.	Indicator Scores A	B	C	D	Ext.
Return on investment	4	>$10,000 0.50 Ã	>$3,000 0.33 ___	>$0 0.17 ___	No return 0.00 ___	2.00
Medical staff impact	4	Attracts new physicians 0.57 ___	Retains current physicians 0.50 Ã	Retains specific physicians 0.10 ___	No medical staff impact 0.00 ___	1.32
Matches hospital goals	3	Very helpful short-term 0.60 Ã	Very helpful long-term 0.30 ___	Somewhat helpful, long or short term 0.10 ___	No impact, long or short term 0.00 ___	1.80
Promotes CQI (Continuous Quality Improvement)	3	Very helpful, hospital-wide 0.60 ___	Somewhat helpful, hospital-wide 0.30 ___	Very helpful, individual team 0.10 Ã	Somewhat helpful, individual team 0.00 ___	0.30
Meets external reporting requirements	3	Meets legal requirements 0.50 ___	Meets corp. or external requirements 0.33 ___	Helpful in serving ext. requirements 0.17 ___	No external requirements 0.00 Ã	0.00
Meets departmental plans	2	Very helpful short-term 0.60 Ã	Very helpful long-term 0.30 ___	Somewhat helpful, long or short term 0.10 ___	No impact, long or short term 0.00 ___	1.20
Ease of completion	1	<8 hours 0.70 ___	<24 hours 0.20 ___	<40 hours 0.10 ___	>40 hours 0.00 Ã	0.00
					Total Score:	6.62

Once the factors are determined they are assigned importance weights. One method gives the most important factor(s) a weight of 10 and scales the others down from there. Another, which has the advantage of forcing decision makers to make hard trade-offs, is to have the weights sum to 100 or to some other predetermined figure. With this method, a point added to one factor must come off another. The second column of the example scoring sheet shows the factor weights assigned to the factors listed to their left.

Candidate projects are then scored on each of the factors. Doing this earlier allows subconscious bias to creep into the process. There are several ways of doing this:

. • The candidate that rates highest on each factor is given a maximum score such as 10. The other scores are scaled down from there. Each other project's score reflects the degree to which that project is lacking on that attribute.
. • Predetermined definitions, as in columns 3–6 of the example, can be used to assign scores. In this example the scoring system uses a high score of 1.

Next, calculate the weighted scores of each project on each factor. A project that scored 10 on a factor with a weight of 5 receives a weighted score of 50 for that factor. The rightmost column of the example shows the weighted scores of a sample project. Finally, the weighted scores for each candidate are added. That gives each one's overall score—6.62 for this example, as shown in the bottom right corner of the example.

The scoring method recommends selecting the alternative with the highest score. If we are using this method to prioritize projects, we would select projects from the highest-scoring one down until we run out of resources. The laboratory instrument interface project used as an example might be the top priority, the bottom one or anywhere in between. We can't tell until we see the others.

Pure project scoring is not the entire decision making process. **Filter criteria** or “musts” should be checked before a project is eligible to be scored. For example, a project might have to meet a financial hurdle rate of 15% to be considered at all. If we proceed directly to scoring, a project with lower IRR might receive 0 on the financial factor but rank so high on all the other factors that it will be selected. Subjecting all candidates to the filter criteria prevents this from happening. If a filter of 15% IRR is applied a project with less than this will not be scored at all.

An attribute such as IRR is often both a filter criterion and a scoring factor. To continue the example of the previous paragraph, an IRR of 15% is a bare minimum but higher is better. A project with a 15.0% IRR would proceed to the scoring, but would score 0 for IRR as it only reaches the minimum acceptable level. Projects with higher IRRs would receive higher scores, up to 10 or whatever other maximum is used.

The lists of criteria, weights, and scores used in this method should be a consensus of all decision makers. If one person is making the decision, that person can pick any weights he or she chooses - using the method, perhaps, only to explain how a decision was

reached. If a committee is making the decision, it will have to agree on weights. Different committee members can then be assigned to score the candidates on different factors, with one (for example) rating financial aspects while another member rates potential contribution to increased market share. This approach reflects the different capabilities and interests that members of any committee are likely to have.

Project selection approaches compared

Both portfolio balancing and project scoring start with cost-benefit analysis and extend it to make up for its limitations. However, they extend it in different directions. The two extensions are not compatible with each other. Selecting a balanced portfolio will generally not yield the top-scoring projects. Selecting the top-scoring candidate projects will also not generally result in a balanced portfolio, even when risk is considered as one of the scoring factors. What, then, is one to do?

First, recognize that each approach has merit. It is good to select a balanced portfolio of projects. It is also good to select projects that score high on the consensus criteria of the organization. That it is hard to do both at the same time is a fact of life.

Second, try to accommodate both needs. Start with the top-ranked projects in the project scoring. As projects are selected, note where they fall in terms of portfolio balancing. Soon the group of selected projects will become unbalanced in one way or another. At that point, skip down the list to find projects that will improve portfolio balance in that respect. Some compromises will be necessary. Fortunately, it is usually possible to maintain acceptable levels of both portfolio balance and project scores.

The steering committee

Many organizations use a **steering committee** (or, when the context is not obvious, an **IS** or **MIS steering committee**) to make project selection decisions. A steering committee typically consists of representatives from the major functional areas of the firm and from the information systems group. It meets regularly, usually once a month to once a quarter with additional meetings arranged if necessary. At its meetings it reviews project proposals from all parts of the firm and recommends which should be taken forward. It may also set information technology policies for the firm and deal with other aspects of enterprise-level IS coordination and planning.

The nature of a steering committee varies from one company to another.
Some representative differences are:

. • Is the role of the steering committee limited to project selection, or does it extend beyond that (for example, to setting the overall information technology direction)?
. • Is the role of the steering committee advisory only, or does it make decisions? In other words, does it allocate funds, or does it recommend how they be allocated?

. • Does the steering committee consist of top managers or lower-level representatives from functional areas? If lower-level representatives, are they “thought leaders” in their organizations or just assigned to the committee because they have the time?
. • Does the MIS group or users hold the real power on the committee—or is it shared?
. • Does the committee meet often enough to pay attention to individual projects, or does it just set guidelines and look at the largest ones?
. • Does the committee review only proposed new projects, or does it also look at on-going ones to see if they should be continued or terminated?
. • Does the committee get into technical issues or leave those to professionals?
. • Are committee decisions respected by the user organizations, or do they try to bypass it by funding “bootleg” projects or appealing decisions they don't like?

There is no right or wrong answer to any of these questions. What is important is that you ask them, at least of yourself, before you attend a steering committee meeting or prepare material to go to one. Preparing a highly technical document for a committee that deals with business strategy issues can be a professional disaster. So can preparing a business strategy document for a committee that prides itself on technical competence.

In the project selection context, system users or departments submit project requests to the steering committee. These requests can range from minor enhancements to existing systems to a totally new inter-enterprise project. User-developed systems that fit within the budgeting authority and capabilities of one department may not have to be approved by the committee, though the committee may have the right to review them so that they do not become a vehicle for subverting its decisions.

New project requests are generally submitted to the committee on a form which provides enough information for a preliminary screening of the request. Following its screening, the committee can assign the request to one of three categories:
. • *Approved.* This designation is generally reserved for small projects whose need is obvious and which do not have to be studied further.
. • *Disapproved.* This means a project is clearly not of sufficient priority to justify the allocation of organizational resources.
. • *Feasibility study.* This category is assigned when the organization might want to carry out a project, but the steering committee needs to know more than a one-page form can tell it before deciding. A feasibility study is the first step in developing any major information system. Initiating a feasibility study doesn't delay an urgent project, since the activities carried out in the feasibility study have to be carried out as part of the project in any case.

When the feasibility study is complete the committee has enough information to prioritize the project relative to other candidate projects and corporate needs. This is where the selection methods that we discussed earlier come into play. All are used, often in combination, as the committee considers its choices and renders its decisions.

While steering committees are widely used in larger organizations, they aren't perfect. Perhaps their biggest drawback is that they cannot respond instantly to every request. This is partly because they don't meet every day, partly because their members need time to examine each proposal (while getting their other work done!) and partly because some project selection methods are best applied when there is a group of proposals to choose from. This delay can motivate users to bypass the steering committee by any available means. A common way in which users bypass steering committees is by applying end-user computing outside its realm. One must watch for this tendency and keep it down to an acceptable level.

Ethical considerations in project selection

In the best of all possible worlds all project proposals would be based on ethical principles. It is an unfortunate fact that we do not live in the best of all possible worlds.

Ethical issues come to a head when a system is put into use. In theory, it could be entirely ethical to develop a system with an unethical or illegal purpose as an intellectual challenge, as long as that system was never put to that use. The controlled use of systems whose basic purpose is unethical can be ethical. An example is attempting to break into a computer system on behalf of management to test security measures.

The fact remains that, with very few exceptions, systems whose underlying purpose is unethical are usually put to unethical uses. What's more, few corporations want to waste money developing a system that won't be used. It is therefore important to review the ethical aspects of proposed systems before they get past the feasibility study stage.

The ethical considerations that arise with transaction processing systems generally relate to data collection. It is important to ask if the information being collected is collected with the full awareness of those who provide it as to all the purposes to which it will be put, if people supplying personal data have any control over what they supply, if the system will safeguard the privacy of that information, and what precautions will be taken to prevent erroneous data from entering the system and to detect and correct it if it does enter.

Ethical issues in a decision support system context usually relate to the DSS database and to decisions made on the basis of its content. The DSS database must, like a transaction processing database, only store information that was obtained through ethical channels, share it with those with whom its supplier would agree it may be shared, and combine it only with information that does not create an overall invasion of privacy. It must provide means to detect as many errors as feasible and to correct those that are detected. Data should only be used for the purposes for which they were supplied, humans should review recommended decisions that have the potential to reflect inappropriate use of the data, and the decision making methods supported by the DSS should always be consistent with ethical principles.