

Identifying, prioritizing, and using multiple objectives

Ralph L. Keeney

Received: 28 January 2013 / Accepted: 6 March 2013 / Published online: 3 May 2013

© Springer-Verlag Berlin Heidelberg and EURO - The Association of European Operational Research Societies 2013

Abstract Analyzing decisions with multiple objectives requires thought about issues not relevant to single-objective decisions. The two key issues are identifying an appropriate set of objectives and prioritizing them. Both tasks may appear to be almost trivial: just write down your objectives and state your priorities. Experiments and experience indicate that this approach is inadequate for decisions worthy of thought. Identifying a complete set of independent objectives and prioritizing those objectives requires concepts, logical foundations, and thought-provoking procedures to create quality results. Those concepts, foundations, and procedures are presented and illustrated. Their use is described on two significant problems, allocating funds to alleviate different homeland security hazards and evaluating potential customers for American Express charge and credit cards.

Keywords Decision analysis · Problem structuring · Preference elicitation · Multi-criteria decision making · Applications

Introduction

The basic challenge to an individual or group facing a complex decision is to solve it. Solving it usually means finding a good solution, a very good solution, or the best solution. The only thoughtful way that one can solve any problem is to make a decision to do so. Indeed, decisions are the only way to have a purposeful influence on anything.

Decisions can be made with or without the benefit of analysis. Whether to do analysis or not is itself a decision. The intent of and potential benefits of analysis are

R. L. Keeney (✉)
Fuqua School of Business, Duke University, Durham, USA
e-mail: KeeneyR@aol.com

the insights provided about the problem and its potential alternative solutions. These benefits include the creation of innovative alternatives, better understanding of key problem issues and interrelationships of problem components and variables, better communication among those concerned with the problem, and an evaluation of the proposed alternative solutions. The negative aspects of proceeding with the analysis are its time and cost, the risk that its quality is mediocre, and hence that some perceived insights are misleading.

When we say that one problem is more complex than another, we mean that there is more that we do not understand that is relevant to make a good decision about the former problem. Good analysis can provide such understanding. Thus, the more complex a problem, the more potential value of insights from analysis. However, the ease and ability to convert that potential value into real value is a separate issue. Typically, the substance relevant to complex problems comes from various fields and disciplines and from various individuals and organizations. There is no way to logically integrate all of this relevant information without a model that can be used for analysis.

Complex problems frequently involve multiple objectives, and indeed, it is these multiple objectives that are a significant contributor to the complexity of the problem. The focus of this paper is on the unique aspects of analysis that arise because of multiple objectives. Specifically, it concerns both the theoretical foundations and the practical procedures to implement multiple-objective analysis. It categorizes the types of issues that must be addressed in this process, provides guidelines for how to address them logically, and illustrates their use. The specific issues addressed are identifying objectives, structuring objectives for analysis, measuring the degree to which objectives are achieved, and prioritizing objectives. The illustrations address a public/governmental problem concerning homeland security and a business decision concerning the acquisition of more customers.

Identifying objectives

As there are different definitions and concepts in the literature for the term objective, it is important to define what the term objective means in this paper. I define an objective as a statement of something that the decision-maker, an individual or a group, hopes to achieve in solving the problem being faced. Furthermore, for consistency and clarity, objectives should be stated in the form of a verb followed by an object. Examples of objectives are to maximize profits, increase product sales, limit environmental degradation, reduce the dependency on foreign energy sources, enhance the quality of education, and decrease vehicle fatalities.

Identifying an appropriate set of objectives for a complex decision is not a simple task. It requires hard thinking. As the essence of a complex decision is to find a solution that well achieves a chosen set of multiple objectives, it is worthwhile to focus attention on identifying those objectives.

If one has no experience in working with decision-makers, policymakers, and stakeholders to elicit their objectives for a complex problem that they may face, it may seem like one could just ask them to write down their relevant objectives for

the problem and allocate about 10 min for this task. This is plenty of time to produce a list of objectives. However, my practical experience and recent research indicates that such lists are often woefully inadequate.

Research on identifying objectives

Bond et al. (2008, 2010) have done a series of experiments that asked individual decision makers facing significant personal decisions to list their objectives, in any form they wished, for those decisions. The decisions involved individuals selecting an MBA program, current MBA students selecting a summer internship in between the 2 years of study, and doctoral students selecting a dissertation topic. In each case, the decision makers were given plenty of time to do this. The details vary in the experiments, but one group was given a week over which to list objectives and routinely encouraged by e-mail to add any newly recognized objectives to their list. Other groups had the thinking process stimulated, after providing an original list, by suggestions of possible categories of objectives (e.g., personal, professional), examples of objectives for the specific decision, or challenges to create a more complete list.

In most cases, individuals originally listed less than half of their personally relevant objectives that were subsequently identified from a master list of almost all possible objectives that the researchers previously developed and provided. Some stimulation of thought provided impressive improvements, but still in the best cases usually a third of the objectives subsequently identified on the master list were not previously recognized by the individuals. The best thought stimulant was to challenge individuals by informing them that their original list of objectives was likely seriously incomplete and that we knew they could do better with harder thinking, and then giving them a reasonable but demanding target for the number of new objectives to create.

Different groups of individuals in a large MBA internship experiment got a different target: specifically, each participant was asked to identify at least 3, 6, or 9 additional personally relevant objectives. For each target, approximately 90 % of the individuals added at least one additional objective. In the original lists, the average number of objectives listed was 6.39 and the average number eventually chosen from the master list of 34 objectives was 19.6. Hence the 10 % who added no new objectives likely just chose not to accept the challenge. Of those who did add at least one new objective, an average of 3.36 objectives was added for the “add 3” case, 4.15 for the “add 6” case, and 5.62 for the “add 9” case. The latter is an 88 % improvement over the original 6.39, which led to 12.01 objectives. Such an improvement is impressive. Yet, on average only 12 of the relevant 19.6 objectives were listed, which means that over 35 % of the individual’s personally relevant objectives were unrecognized in the best case. This is important because there rarely is a master list of objectives for decisions that are not part of experiments.

In all of the experiments, the importance of the listed objectives and those that were only recognized later from the master list were assessed. The results indicated that the importance of the objectives that were only recognized after seeing the master list was essentially the same as the importance of the objectives originally

listed by the individuals. In other words, the research definitively shows that it is not the case that decision-makers know their important objectives and can list them and are only unaware of their relatively unimportant objectives. Comparing alternatives in terms of only some of the relevant objectives, especially when it is not recognized that other relevant objectives are missing, can at best provide relatively little insight and at worst guide the decision maker to inappropriate choices and their consequences.

The fact is that a logically sound reasonably complete set of objectives cannot be developed for an important decision without a focused effort on the task of developing those objectives and using concepts and procedures that will facilitate doing that task well. Since one cannot make a good decision without knowing the objectives appropriate for evaluating alternatives, it is essential to do the best job you can in identifying objectives for decisions. The following illustrates many tested procedures that help significantly in identifying better lists of objectives for your decisions.

Process for an individual to identify objectives

The basis for identifying objectives is the statement of the problem to be solved or resolved. It is useful to state that problem by phrasing what needs to be decided to solve it. An appropriate problem statement would complete the following: we want to decide ... For businesses, completion of this statement might be who is the best candidate to be our new CEO, where to locate a new manufacturing facility, how to become the leader in our product class in the EU, or how to reorganize to avoid bankruptcy. For public agencies, completion of this statement may be how could immigrants be better integrated into a national society, how should expensive medical procedures be allocated in a constrained system, where should a new local school be located, how to ensure reliable electricity supply in a country that is phasing out nuclear power, and what route is best for a lengthy new high-voltage line.

All objectives must originally be thought of in a mind. As only an individual has a mind, and a group does not, an individual's mind is the original source of all objectives. A group may subsequently agree with and accept such an objective for their group decision. Thus, we first outline the procedure to generate objectives from individuals and then consider how to combine the lists to construct a group or organizational perspective.

To begin, ask an individual to list anything that he or she hopes is achieved by solving the problem being addressed. To facilitate retrieving any relevant thoughts from an individual's mind, no restrictions on the form of response should be used. Each item on the list is what I refer to as a value, which can readily be converted into the verb-object format of an objective later. Once an individual appears to be done, which rarely takes more than 10 min, challenge them as discussed above to double their list of values. You can further motivate them by mentioning that since the problem is important to solve, and typically implementing a solution may be very expensive and take months or years, is it not worth more than 10 min of your thoughts to develop a good list of why you are interested in solving this problem?

After an individual has added as many additional values as she can, suggest ways to stimulate further thinking. You could begin by asking her to think of a very desirable (undesirable) alternative that may be a practical possibility or hypothetical. Then ask why this alternative is good (bad). The response will likely indicate additional values. Ask her to consider a scenario that could occur that would render good (bad) consequences, and then ask what consequences these are to generate additional potential values. Without ever defining goals and constraints, you can if ask any goals or constraints are relevant to solving the problem. Again, responses indicate potential values. Keeney (2007) provides details of this probing process and several illustrations of its use.

Now you want the individual to review each value on her list and articulate why she cares about it. The response may indicate another value that is relevant to the problem solution. For example, she may be concerned about runoff water from an agricultural area into lakes. One reason for caring about this may be to minimize the amount of chemicals from fertilizer going into the lakes, which is a different value. And why does she care about that? The response may be to reduce lake pollution, another value. Following this process, we soon arrive at values of providing quality drinking water, avoiding any health effects to individuals, and do not harm the fish in the lakes.

At this stage, either the individual or an analyst can translate each value into the format of an objective. Some values are initially stated as an objective and the translation of others is obvious. The value 'profit' can become the objective 'maximize profits'. In other cases, it is important for the individual who stated the value to ensure that the corresponding objective captures her concern. For the value stated above of do not harm the fish, the objective may be one or more of minimize degradation of the fish habitat, avoid negative effects on fish reproduction, and avoid killing any fish.

After all values are translated into objectives, one more effort to recognize any missing objectives is useful. There are two cases to consider. The first case is when only one individual's objectives are relevant to solve a problem. In this case, it is often useful for the individual to have a few friends and/or colleagues review this list of objectives and try to generate any potentially missing objectives that the individual may wish to include. The second case is when the objectives of more than one individual are relevant to a complex problem. This is the more common case for complex problems, such as when individuals in different departments of a company or organization are concerned with the decision and/or there are stakeholders who are influenced by the alternative chosen to solve the problem. In such cases, the values generated by each of these individuals are essentially a check on potentially missing objectives on the other's lists. In addition, after each individual has provided a list of objectives, it can be useful to convene the group to collectively brainstorm in an attempt to create any additional relevant objectives.

Process to organize objectives

Once the list of objectives is complete, they should be organized into different types of objectives, as different types of objectives have different roles for clear thinking

and for analysis. The substantial aspects of organizing an individual's or a group's objectives are the same and it begins with a complete list of the objectives to be organized. When more than one individual provides objectives, all of the lists of objectives must be combined into a group list.

The discussion below describes the various intellectual tasks necessary to construct a useful organization from a list of relevant objectives. In describing these tasks, it may appear that there is a step-by-step process that provides the resulting objectives network. In practice, there will likely be much iteration between tasks and additional creative thinking that leads to the identification of previously missing objectives and a revision of the current evolving structure. A current structure may suggest that the connection between two categories of objectives is accomplished by achieving newly recognized objectives. This may result in modifications to the initial categorization of objectives and a revised structure. Organizing objectives requires hard thought-provoking work, which stimulates thoughts that improve the eventual result.

As different individuals will likely have stated some of the same objectives, the first step is to eliminate such redundancies in the group list. For example, in a decision involving setting policy to address homeland security, you would expect that almost everyone would have included an objective to minimize loss of lives, although it may have been stated as save lives, avoid fatalities, or protect citizens.

Given a single list of possible objectives void of redundancies, the next step is to categorize them into groups that address the same concern. Suppose the problem statement concerned the national preparation and response to minimize impacts of anticipated flooding. There would likely be categories of objectives that concerned loss of life, health effects, environmental effects, property damage, citizens' response, integrity of flood control systems, emergency resources, communication, and so forth. Once categories are identified, each can also be stated as an objective such as to minimize loss of life, maintain emergency resources, enhance timely citizen response, and provide quality communication.

The way these categories are initially generated is by aggregating the objectives that address a common concern and then naming a category for that group of objectives. The first pass through the list of objectives will likely get 50–80 % of the objectives into readily identified categories. Those remaining that are not yet in categories are a much smaller list, so it is now easier to identify related objectives among them. The process continues until all objectives are in a category, and it is possible that a category may include only one objective.

It is usually the case that the process of categorizing objectives allows one to recognize some stated objectives that are not relevant for evaluating alternatives to solve the specific problem being addressed. To illustrate, for the flooding illustration above, the list of objectives may include items such as 'collect the relevant data', 'involve people affected by flooding', and 'build heliports near areas where severe flooding may occur'. Each of these is a useful idea. However, the objective of collecting the relevant data pertains to the decision of what analysis to do, not what alternatives to prepare for and respond to flooding. The objective to include people who will be influenced by flooding is relevant to the decision of who to involve in the decision process and whose values should be elicited, and the proposed

objective about building heliports is not really an objective, but an alternative that possibly should be part of the preparation strategy for floods.

Once categories are identified, the third step in organizing objectives is to develop relationships among categories of objectives and among the objectives within categories. Two types of relationships are useful to distinguish: means–ends relationships and inclusion relationships. These can be illustrated using the flooding example above. The objective of providing quality communication will contribute to a timely citizen response (e.g. evacuation), which may influence the subsequent loss of life. Thus, providing quality communication is a means objective to the ends objective of a timely citizen response, which is itself a means objective to the ends objective concerning loss of life. You can proceed through a series of means–ends reasoning by asking for each potential means objective, why do you care about achieving that? For example, you care about quality communication to enhance a timely citizen response. When asking why do you care about minimizing loss of life, the response may be something like “that is one of the fundamental reasons for being concerned about this problem”. Hence, I refer to such an objective as a fundamental objective, as it indicates a fundamental reason for being interested in the problem.

It is useful to note that a means objective can influence more than one ends objective, as maintaining emergency resources is a means to avoid property damage and to minimize loss of life. Also, more than one means objective can influence the same ends objective, as a timely citizen response and maintaining emergency resources can influence loss of life.

The inclusion relationship between objectives is when one objective is a part of another objective that specifies in more detail the concern of the broader objective. There could have been a category objective of minimize health and safety impacts, which would then include minimize loss of life, minimize injuries, and minimize illness. Also one could treat the objective of minimize loss of life as a category and define minimize life loss of children, seniors, handicapped, and others as four separate objectives. There may also be a distinction made to minimize the loss of life to the public and to flood control and flood emergency workers. The reason for such specification is to contribute to building an appropriate decision model (Phillips 1984), collecting data to use in a model, creating alternatives, and communicating clearly.

The set of fundamental objectives for a decision should be complete, meaning that they should not only adequately describe all of the consequences of a decision that concern the decision-maker, but also avoid any overlap of value that would induce double counting of the desirability of consequences. To illustrate the distinctions here, consider a decision of an environmental agency that was intended to promote the future of a large mammal (e.g. bear) in a specified remote area. Three objectives might be to minimize cost of the alternative, maintain desirable bear habitat, and maximize the bear population. They might be measured, respectively, by metrics of cost in euros, hectares of desirable bear habitat in 5 years, and the increase in the number of bears in the area in 5 years. The cost objective clearly captures different concerns than the habitat and population objectives. If the impact of all of the alternatives was to maintain habitat, which is a means to population

increase, then one might select the habitat objective as a second fundamental objective and not include the population objective as a fundamental objective. In assessing preferences for different areas of habitat, the perceived relationship between habitat area and population increase should be considered. If some alternatives directly affected habitat and other directly affected the bear population, such as limiting human interaction with bears, then a better fundamental objective would be the population objective without the habitat objective because of the potential for double counting.

There are cases where all three objectives could be fundamental objectives, but clarity about exactly why each objective is to be valued is essential to avoid double counting and yet count all that matters. Suppose that bear habitat is also habitat for some smaller animals and birds and for some plant species. Then in addition to cost, the population objective would capture the direct impact on bears, and the habitat objective could be used to capture the habitat value for animal and plant species other than bears. In any subsequent value judgments, care must be taken to avoid enhancing the stated value of habitat because of the implications for increased bear population, as this is accounted for with the population objective.

Several examples that illustrate the process of creating structures of objectives are found in Keeney (1992). Both fundamental objectives hierarchies and means–ends objectives networks are illustrated in detail for individual, company, and organizational decisions.

Managing homeland security

Terrorist attacks, transnational crime (meaning illegal immigration and drug smuggling), natural disasters, pandemics, and major industrial accidents are hazards that threaten security and quality of life. In the United States, the Department of Homeland Security (DHS) has an annual budget of over \$50 billion to address these issues. An important problem is how to allocate these funds to the various hazards to best serve the nation and its citizens. This problem was initially addressed by a DHS working group in 2009 led by S. P. Bennett, R. P. Kolasky, and E. S. Levine of the DHS Office of Risk Management and Analysis.

An overall framework to compare the relative seriousness of the hazards and subsequent allocations of resources to address these hazards that was developed by the DHS working group is shown in Table 1. When completed, this table would indicate the annual possible consequences, with the uncertainties as appropriate, for each of several different scenarios within each of the hazard categories. For example, scenarios of terrorist attacks could range from a suicide bomber in a mall to the use of the nuclear bomb. Transnational crime scenarios may include different annual levels of illegal immigration. Natural hazard scenarios may include a major hurricane directly hitting a large city and major flooding throughout a region.

The task described here was to develop a thorough characterization of the set of possible consequences of concern. These consequences would certainly include lives lost and economic costs, which were used in Table 1 to illustrate that aspect of the framework. The basis for each of consequences of concern is the fundamental

Table 1 Framework to compare homeland security hazards

Fundamental Consequence Categories	Hazards (with several scenarios in each hazard category)										
	Terrorist attacks		Transnational crime			Natural disasters		Pandemics		Industrial accidents	
	Nuclear attack	Suicide bomber	2 million illegal immigrants			Major hurricane hits city	Major regional flood				
Lives lost											
Economic cost											
Other ...											
Other ...											
Annual probability or frequency											

objectives for the allocation decisions to be made. Hence, our effort described below, which draws heavily on Keeney et al. (2010), was to identify and structure a set of objectives for the DHS problem.

Thirteen members of the DHS working group individually completed a questionnaire designed to elicit values for the problem. There were five questions, each with essentially the same form. The first question was the following: if a major terrorist event occurred in the U.S., list the possible consequences and/or impacts that you think would matter to our nation and its citizens. Questions 3 through 5 were identical with terrorist event replaced by natural disaster, pandemic, and a major industrial accident (e.g. nuclear, chemical). As transnational crime is not a set of single events, but an ongoing process, question 2 was stated as follows: with regard to major illegal trafficking and/or transnational crime (e.g. illegal immigration, drug smuggling) in the U.S., list the possible consequences and or impacts that you think should matter to our nation and its citizens. Directions requested that participating individuals include anything that they thought might matter to themselves or others, to be as complete as possible, and to include both broad and narrow consequences. To expand their thinking, we suggested ideas using scenarios and alternatives as outlined in “Introduction”. Participants were told that if in doubt about whether a possible impact was appropriate, it should be added and that there were no wrong responses. We also mentioned that they need not repeat any responses given to a previous question.

Collectively, there were approximately 350 responses to the five questions. Table 2 lists some responses to indicate the typical range of responses and range of formats for the responses. I then converted each response to an objective and proceeded to categorize the objectives, which required two passes through the objectives. Table 3 provides the final set of categories; those from the first pass are in the first column and categories in the second column are from the second pass.

Table 2 Some responses of concerns about homeland security hazards

Immediate deaths
Delayed deaths and illnesses due to toxic exposures
Increased levels of fear
Violation of national sovereignty
Affect on GDP
Environmental damage
Military actions implemented in response
Displacement of people from homes and communities
Loss of confidence in the federal government
Property damage
Failure of infrastructure system
Overwhelming medical facilities
Increase in profiling based on the terrorist group ethnic background
Decrease in travel to the United States
Increased deficit spending to counter the actual/perceived new threat(s)
Threat to present vision of personal liberty
Coarsening of society
Contamination of food or water supply
Public health impacts—High morbidity and mortality, transmissibility, pathogenicity, medical countermeasure effectiveness—if extreme, potential follow on of disease and problem with dealing with corpses, including contaminated agriculture and food products or infected plants and animals
Psychological impact (short term, long term)
Encourages copycat attacks
Imposed martial law
Results in politicized or irrational decision-making
Increase in gang activity or organized crime
Increase in human smuggling and kidnapping
Potential for weapons/WMD movement across borders
Illegal immigration effect on the community (schools, hospitals impacted)
Spread of communicable diseases into the community (e.g. TB, cholera)
Not having enough medication to handle the pandemic
General panic within the population

The objectives were then structured by identifying means–ends relationships and inclusion relationships. The resulting structure is shown in Fig. 1. The fundamental consequences are those that indicate the impact in terms of the fundamental objectives identified for the problem. Those designated with a capital letter refer to the objective categories and those designated by lower-case letters indicate specific objectives that are included in the respective category.

The means impacts relate to the achievement of means objectives. Figure 1 mainly illustrates categories of means objectives, each of which would contain specific means objectives that are not shown. The possible responses, indicated at the left of Fig. 1, are reactive decisions that could follow any specific scenario. Not shown are the possible proactive decisions that are a major reason to address this

Table 3 Categories of homeland security objectives

Health and safety	Possible response actions
Economic impacts	Impact on ability of government to function
Socioeconomic impacts	Disruption of business
Emotional/psychological impacts	Required goods and services
Freedoms/rights	International relations
Environmental impacts	Criminal activity
Loss of confidence in government	Disease
Loss of national symbol	Contamination
Disruption of services	Disruption of workforce
Infrastructure damage	Degradation of communities
Impact on emergency response systems	Ethnic/racial strife
Importation of undesired material	

problem. These concern implementing alternatives, prior to the occurrence of a specific scenario, that may lessen the likelihood of its occurrence and/or the severity of its consequences if it occurs.

A structure of objectives such as Fig. 1 has many uses. It helps transform a large amorphous problem, sometimes referred to as a mess, into an organized structure that helps define, communicate about, and understand the problem and to suggest alternatives that may be useful to pursue. This can occur whether there are plans to proceed with any analysis or not. For any subsequent analysis, the structure indicates what information is needed for that analysis. The fundamental objectives in Fig. 1 should be placed in the framework of Table 1. Then to compare the seriousness of each scenario, we need information of its annual probability of

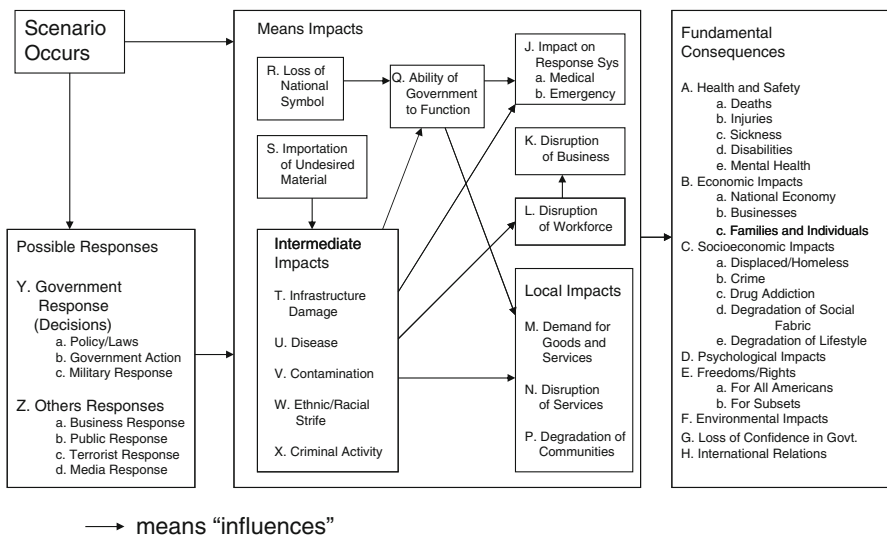


Fig. 1 Structure of homeland security objectives

occurrence (or the number of occurrences expected annually) and assuming this scenario occurs, descriptions of its possible consequences in terms of each of the fundamental objectives. The overall consequences of a hazard are those of the aggregated expected consequences of the scenarios that specify that hazard. Alternatives to eliminate or alleviate scenarios can be compared by describing how the possible consequences may change as a result of those alternatives.

Prioritizing objectives

When one refers to a decision problem as having multiple objectives, it means that there are multiple fundamental objectives. A decision with multiple means objectives that influence a single fundamental objective, such as to maximize profit, is not a multiple-objective decision. In multiple-objective decisions, the logical prioritization of the fundamental objectives is necessary to evaluate alternative courses of action.

Prioritizing objectives refers to specifying the relative importance of achieving each of the fundamental objectives relevant to a decision. It is obviously a unique feature of multiple objective analyses, as one cannot prioritize objectives in a decision with only one objective. Prioritizing objectives appears to many people as a straightforward intuitive task, namely to ask the decision-maker to prioritize the objectives for a problem. However, such a lack of attention in prioritizing objectives results in a number of important logical and practical errors summarized in Keeney (2002). Most of these errors result from and contribute to an ambiguous meaning for the concept of ‘importance of an objective’ and the lack of a logical theoretical basis for the prioritization task (Salo and Hämäläinen 1997).

Shortcomings of intuitive prioritization

To illustrate these issues, consider a simplified decision problem where a company is considering several alternatives to ship large numbers of a tangible product from Dresden to Amsterdam. They have three fundamental objectives for the shipping process: to be fast, inexpensive, and good. To be more explicit using the verb-object characterization of an objective, the objectives are to minimize time for delivery, minimize cost of delivery, and maximize quality of the delivery process. They will measure these objectives, respectively, in terms of days, thousands of euros, and number of products broken en route.

Suppose the CEO for the company was asked to prioritize these objectives and he intuitively stated that time was the most important objective, then quality, and last low-cost. Furthermore, suppose he said time was twice as important as quality and six times as important as cost. A logical check with the CEO may indicate that he feels that quality is three times as important as cost. An implication of these judgments is that the relative importance of time is 1.5 times as important as quality and cost together. The basic problem with such judgments is that they are not grounded on a clear meaning for ‘the importance of an objective’ and, therefore, can provide no meaningful guidance for making decisions (Weber 1997).

A careful analysis of shipping time may indicate that shipping company A can deliver the product faster than any of the other shipping companies. Suppose another company was better than company A on both quality and cost. Would it matter? Since company A is better on the time objective that is 1.5 times as important as the other two objectives together, would not this imply that company A is better in all cases? If yes, the CEO should consider information about the cost and quality of other companies to be irrelevant? If no, what does it mean that time is 1.5 times as important as the other objectives together?

To resolve this apparent dilemma, suppose the CEO had data indicating that the quality is equal for shipping companies A and B, and that the shipping time of A is 3 days and the cost is €720,000, whereas the shipping time of company B is 4 days for a cost of €250,000. Does it seem reasonable that the CEO should choose company A because it is better on his most important objective? Given this data, the CEO may not hesitate in stating that he prefers company B to company A because it saves €470,000 and results in only one extra day of delivery time. In other words, the logical issue is whether €470,000 is more important than one additional day of delivery time or not.

The basis for this conundrum is that there is no clear meaning for the concept that ‘one objective more important than another objective’. There is a clear meaning for the concept that ‘a specific change in the level of achievement on one objective is more important than a specific change in the level of achievement on another objective’. There is a clear unambiguous meaning to the statement that a reduction in cost of €720,000 to €250,000 is more important than an increase in delivery time in 3 to 4 days. The CEO may also feel that a reduction in cost from €720,000 to €705,000 is less important than an increase in delivery time from 3 to 4 days. In this situation, in one case the CEO prefers money over time and in the other case, he prefers time over money. So what is more important to that CEO: time or money? The fundamental point is that talking about the relative importance of changes in the levels of achievement with different objectives makes sense, whereas talking about the relative importance of different objectives does not.

A logical foundation for prioritization

The logical foundation for prioritization can be illustrated using the notion of an objective function that is either a utility function or a measurable value function (see Keeney and Raiffa 1976; Dyer and Sarin 1979; Salo and Hämäläinen 1997). As the prioritization concepts are the same for both types of objective functions, I will use the utility function to indicate those concepts since the decision in the next section uses a utility function.

Let u be a utility function for evaluating achievement in terms of the three fundamental objectives discussed above, namely minimize time, minimize cost, and maximize quality, where (t, c, q) represents a particular achievement in terms of these objectives with t measured in days, c measured in thousands of euros, and q measured in broken units of product. The utility $u(t, c, q)$ is an indicator of the desirability of consequence (t, c, q) . If (t_1, c_1, q_1) is preferred to (t_2, c_2, q_2) , then $u(t_1, c_1, q_1) > u(t_2, c_2, q_2)$ and vice versa. If one begins at a consequence (t_0, c_0, q_0) , it is

logical to say that an improvement to (t_1, c_1, q_1) is more important than an improvement to (t_2, c_2, q_2) if and only if $u(t_1, c_1, q_1) > u(t_2, c_2, q_2)$.

A utility function also allows one to characterize all the value tradeoffs among fundamental objectives that are necessary to consider in a particular decision. Value tradeoffs specify how much a specific achievement in terms of one objective is worth in terms of achievement on another objective. Suppose $u(t_1, c_1, q_0) = u(t_2, c_2, q_0)$, so the consequences (t_1, c_1, q_0) and (t_2, c_2, q_0) are indifferent to each other. Then, with q_0 fixed, a change from t_1 to t_2 is compensated for by a change from c_1 to c_2 , which is referred to as a value tradeoff.

As discussed in “[Introduction](#)”, the set of fundamental objectives are chosen to value mutually exclusive components of the overall value of potential consequences. This provides the logical basis for a multiple-objective utility function to be represented by an additive form (Keeney 1981), which for our illustrative problem is

$$u(t, c, q) = w_T u_T(t) + w_C u_C(c) + w_Q u_Q(q), \quad (1)$$

where u_T , u_C , and u_Q are component utility functions and w_T , w_C , and w_Q are weighting factors specified by a set of assessed value tradeoffs. Fishburn (1965) provided the formal assumptions required for a utility function to be additive.

The value tradeoffs between time and cost mentioned above may depend on the common level q_0 . If it does, then the utility function u cannot be represented as an additive utility function. There are several functional forms for multiple-objective utility functions other than the additive form (e.g., Keeney and Raiffa 1976), although they are more complicated to use than the additive utility function. One of the main factors that leads to the inappropriateness of the additive utility function is the inclusion of means objectives either entirely or in conjunction with some fundamental objectives in the utility function. Partially for that reason, the process to identify and structure objectives articulated in “[Introduction](#)” distinguishes between means and fundamental objectives to produce a complete list of fundamental objectives, in which case the additive utility function is appropriate.

Customer acquisition at American Express

American Express (AX) is one of the largest and most respected companies providing charge cards and credit cards to customers worldwide. Charge cards are the original AX product and are to be paid off in full each month. Credit cards allow for borrowing money to purchase products and subsequently pay off the principal and interest over time. By 1995, mainly because of the rapid growth of rivals Visa and MasterCard in the 1980s and early 1990s, the AX share of the total market for such cards had been decreasing each year for over 10 years in a row. Over that time period, AX had built sophisticated analytical models that evaluated potential customers who could apply for a charge card or credit card. These models used the expected contribution to the net present value of profits as a basis for evaluating prospective customers.

Because of the 10-year decline in market share, AX strategically decided to include market share as a consideration in evaluating prospective customers. So

what had been a model with a single objective needed to be converted to a model that had multiple objectives. Hence, a multiple-objective evaluation function was needed. This section describes the expansion of the AX customer acquisition model and its use. The presentation is based on information presented at a professional conference and a concise published summary (Keeney and Lin 1999, 2000).

We sequentially consider identifying and structuring the objectives, selecting fundamental objectives and metrics to measure them, constructing a multiple-objective utility function to prioritize ranges of objectives, use of the multiple-objective model to evaluate prospective customers, and subsequent results. All of the information needed to modify the customer acquisition model was developed based on judgments of executives and managers at AX and included both the president and CFO of Travel Related Services, which was the largest group in AX and responsible for charge and credit cards among other services.

Identifying and structuring objectives

The process of identifying objectives began with separate discussions with several individuals at AX about their objectives for customer acquisition for charge and credit cards. First I ask each individual to state their values and I wrote them down. Then I used ideas discussed in “[Introduction](#)” to expand their thinking to provide additional values. The individual’s lists of values were the basis for constructing a single list of AX objectives.

The several individual lists were combined and redundant values eliminated. Sometimes this was trivial as separate individuals had stated the same value. Usually it was obvious even when different words are used. Values stated as ‘make money’ and ‘profit’ clearly referred to the same thing. After categorizing sets of common values, I translated their meanings into objectives. The example above became maximize profits.

Once I had a list of the AX objectives, they were organized using means-ends relationships as discussed in “[Introduction](#)”. It is easiest to describe this process by viewing the result in [Fig. 2](#). Each box in the figure represents an objective, although I left off the verb to reduce clutter as one can readily imagine what the complete objective is. For example, the stated ‘acquisition expenses’ means the objective ‘minimize acquisition expenses’ and the stated ‘market share of charge volume’ means ‘maximize the market share of spending using the American Express charge card’. It is useful to describe the meaning of network effects and market effects. The network is all of the companies that accept AX cards as payment for goods and services. Arrangements between AX and these companies specify what percentage of each transaction is paid to AX. When there are more prospective customers who use AX cards, the negotiated transaction percentage for AX typically will be a bit higher, which increases overall AX profit. Market effects essentially refer to the effect on the stock price of AX. It is also important to distinguish between short-term profits from the current business and long-term profits. The former is defined as profits over the next 3 years directly from any charge or credit card use. Long-term profits are defined as those accruing after 3 years from any sources including investments, network effects, and card use.

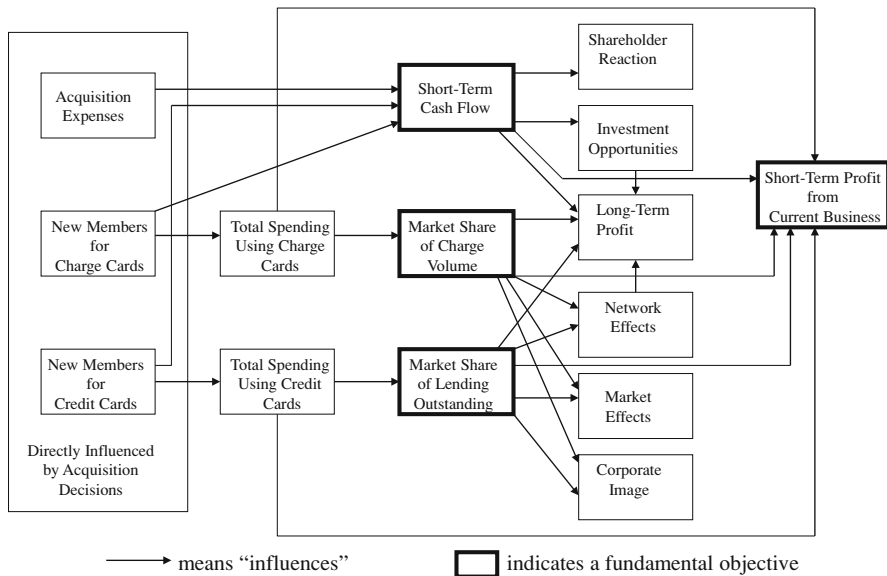


Fig. 2 Relationships among American Express customer acquisition objectives

The arrows in Fig. 2 indicate the means–ends relationships among objectives, most of which are rather straightforward. More new members for charge cards will lead to more total spending on charge cards, which will result in a higher market share of charge card volume, which will contribute to higher short-term profits.

Selecting fundamental objectives and metrics

There are two important and related tasks that must be completed before any prioritization of objectives. First, the fundamental objectives, which are those objectives that will subsequently be prioritized, must be identified. Second, metrics for those fundamental objectives must be specified to indicate the possible levels of achievement of those objectives. The AX customer acquisition decision involves subtle issues in addressing these tasks. The process can best be explained by beginning with the objectives in the boldface boxes of Fig. 2, which were chosen as the fundamental objectives for this decision.

The implementation of the multiple-objective consumer acquisition model at AX was planned for, and subsequently occurred in, the latter part of 1996. Thus, to evaluate different alternatives for selecting prospective customers, the potential consequences focused on the next 3 years 1997–1999. The selected metrics listed in Table 4 recognized this time period. Short-term profit of the card business was measured with the metric of the net present value (NPV) in 1997 dollars due to card use during the 3 years of 1997–1999. The short-term cash flow was also measured in terms of the NPV in 1997 dollars over the same period of time. The charge card market share was measured with the AX percent of charge volume on all charge cards at the end of the 3-year period in December 1999 and the credit card market

share was the AX percent of the total amount of lending outstanding on credit cards in December 1999.

To understand how these four fundamental objectives cover consequences of concern in the acquisition decision and avoid double counting, consider what each of the four fundamental objectives account for of value in the model. The short-term profit objective measures profits from the volume of card use given current operations (e.g. investments, network arrangements) remain as they are. The cash flow objective is valued for its impacts on shareholder reactions, investment opportunities, and long-term profits. Each of the market share objectives directly affects the network, the market, corporate image, and long-term profit. The relative values of different market share levels are to account for those effects. Note that the short-term profitability of AX can increase directly from the arrangements with network members who accept AX cards. Those short-term profit effects are valued in the market share assessments as they are part of the network implications, and they are not again valued in the fundamental objective of short-term profit from card use. Double counting is avoided with this careful categorization.

Prioritizing ranges of objectives

A multiple-objective utility function was assessed over the four fundamental objectives using value judgments agreed upon by a group of executives in meetings set up to construct an appropriate utility function. Because of the way that the fundamental objectives were chosen to be complete and avoid double counting of consequences, the logic discussed in “[Managing homeland security](#)” suggested that an appropriate utility function is the additive form. Hence,

$$u(x_1, x_2, x_3, x_4) = w_1u_1(x_1) + w_2u_2(x_2) + w_3u_3(x_3) + w_4u_4(x_4), \quad (2)$$

where $x_i, i = 1, 2, 3, 4$, is a specific consequence in terms of the metric X_i , which measures the achievement of fundamental objective O_i , u_i is a single-attribute utility function over X_i scaled from 0 to 1, and w_i is a weighting factor indicating the relative priority of the potential range of possible consequences on metric X_i . Table 4 defines the scales for the metrics X_i , presents the possible ranges of

Table 4 Metrics and constructed priorities for the American Express utility function

Metric (scale)	Range of consequences		Priority of range	
	Worst	Best	Ranking	Rating
$X_1 =$ Short-term cash flow (\$NPV for 1997–1999 in 1997 billions)	\$1.2	\$3	4	5.2
$X_2 =$ Short-term profit (\$NPV for 1997–1999 in 1997 billions)	\$6	\$12	1	62.5
$X_3 =$ Charge volume market share (percent in December 1999)	12 %	20 %	2	20.4
$X_4 =$ Lending outstanding market share (percent in December 1999)	3 %	6 %	3	11.9

consequences for each, and provides the prioritization results discussed below. The ranges were chosen by the executives to represent the full range of what the possible consequences may be by the end of 1999.

As discussed in “[Managing homeland security](#)”, prioritization of objectives must account for the range of achievement of the respective objectives. We initially ranked the importance of changing each objective from its worst level in Table 4 to its best level, as ranking requires less precise value judgments than rating. Specifically, I ask each member of the executive group separately to assume that all consequences for the four objectives in Table 4 were at their worst levels. Then, if they could move one objective only to its best level, which change would most improve the situation for AX? I then asked, of the remaining objectives, which one would now be most important to move from its worst level to its best. The same process followed for the remaining two objectives. This provides a ranking for each individual. Then, the rankings of each individual were displayed to the group and it happened there was almost unanimous agreement on the ranking collectively accepted by the group presented in Table 4.

In order to rate the ranges of the four objectives, which prioritizes the ranges of the objectives by specifying the weighting factors in (2), value tradeoffs among the objectives and the single-objective utility functions needed to be assessed. In both cases, I again began by asking each individual for his or her own assessments and then proceeded to develop an agreement on a group assessment. The type of questions to ask to assess a value tradeoff can be illustrated by considering the top two ranked objectives concerning short-term profit and the market share of charge volume.

Figure 3 presents the consequence space for all possible combinations of consequences that differ only in terms of short-term profit and charge card market share. We assume the other two metrics for credit card market share and cash flow are fixed. It is pointed out to all group members that it has been established by the group ranking of the ranges of objectives that consequence A, namely ($x_2 = \$12\text{B}$, $x_3 = 12\%$), is preferred to consequence B, ($x_2 = \$6\text{B}$, $x_3 = 20\%$). Also, clearly, consequence B is preferred to consequence C, ($x_2 = \$6\text{B}$, $x_3 = 12\%$). Hence, for each group member, there must be a consequence between A and C, with $x_3 = 12\%$, that must be indifferent to consequence B. All group members were indifferent with the level of x_2 between $\$8\text{B}$ and $\$10\text{B}$, except for one member who was indifferent with x_2 at $\$7\text{B}$. After consideration of each person’s reasoning for the indifference level, the group chose $\$9\text{B}$. Thus, the consequence B was indifferent to the consequence D, ($x_2 = \$9\text{B}$, $x_3 = 12\%$), in Fig. 3. This value tradeoff is that AX would trade an increase in profit from $\$6\text{B}$ to $\$9\text{B}$ in exchange for a drop in charge card market share from 20 to 12% or vice versa.

As consequences B and D must be indifferent, the utilities of those two consequences calculated using (2) must be equal, so.

$$w_2 u_2(\$6\text{B}) + w_3 u_3(20\%) = w_2 u_2(\$9\text{B}) + w_3 u_3(12\%). \quad (3)$$

We subsequently proceeded with the same process to assess value tradeoffs between cash flow and short-term profit and between charge card market share and short-term profit. The assessed value tradeoffs were ($x_1 = \$3\text{B}$, $x_2 = \$6\text{B}$) is

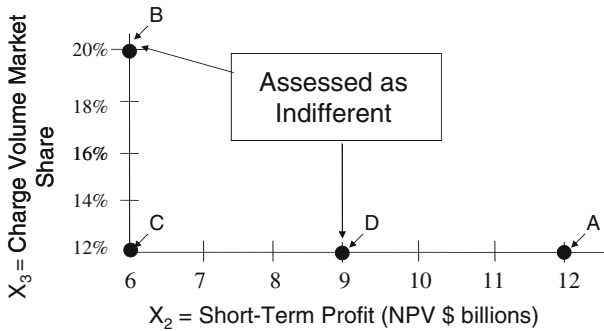


Fig. 3 Illustrated American Express value tradeoff

indifferent to $(x_1 = \$1.2B, x_2 = \$7B)$ and $(x_2 = \$6B, x_4 = 6 \%)$ is indifferent to $(x_2 = \$8B, x_4 = 3 \%)$. Equating utilities of the indifferent consequences yields the corresponding equations.

$$w_1u_1(\$3B) + w_2u_2(\$6B) = w_1u_1(\$1.2B) + w_2u_2(\$7B), \tag{4}$$

and

$$w_2u_2(\$6B) + w_4u_4(6\%) = w_2u_2(\$8B) + w_4u_4(3\%). \tag{5}$$

Because the scale of a utility function can be arbitrarily set, we chose the weighting factors to sum to 100, so

$$w_1 + w_2 + w_3 + w_4 = 100. \tag{6}$$

Once the single-objective utility functions $u_i, i = 1, 2, 3, 4$, are assessed as discussed below, the three equations (3) to (5) constructed from value tradeoffs and the scaling equation (6) provide four independent equations with four unknowns. Solving these equations yields the weighting factors (i.e., rating) in Table 4 that prioritize the relative importance of the ranges of possible consequences of the objectives.

The procedures to assess single-objective utility functions are well known (Keeney and Raiffa 1976; Kirkwood 1997; Clemen and Reilly 2001; Eisenführ et al. 2010). Using these procedures, that included consistency checks, the utility functions in Fig. 4 were constructed from values expressed by the AX management team. For each objective, at least four of the five individuals participating in constructing these corporate utility functions initially expressed values consistent with these utility functions. It is useful to discuss the logic used.

The s-shaped utility function for the market share charge volume is most interesting. At the time of assessing this utility function, this market share for AX was about 17 %. It is evaluated because of the effects of market share on networks, the market, corporate image, and long-term profit. Changes from the status quo of 17 %, as well as the resulting market share, are both important. In other words, a 0.5 % change from 16.75 to 17.25 % is much more desirable than either a change from 15 to 15.5 % or from 18 to 18.5 %. Each of these three changes has a 0.5 %

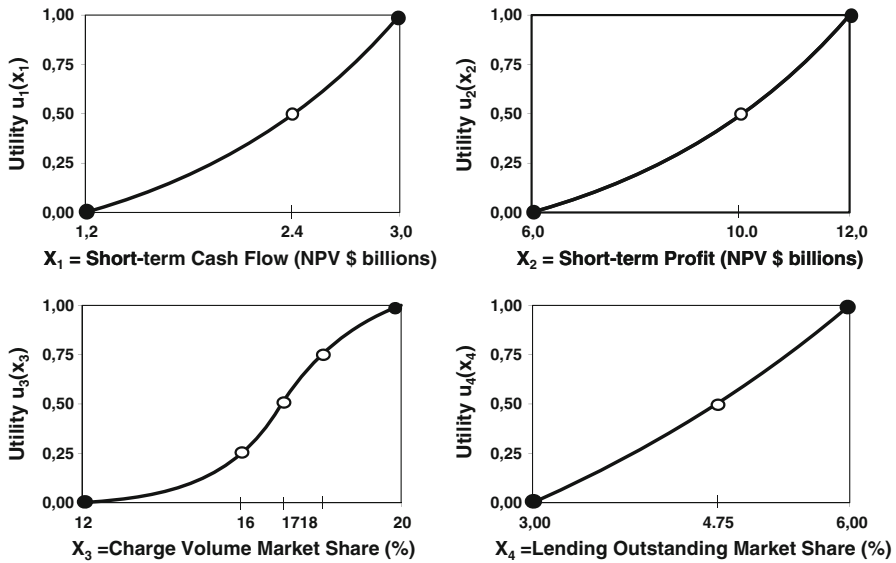


Fig. 4 American Express single-objective utility functions

increase in market share, but the first also changes from continuing a historic decrease in market share to a significant turnaround to an increase in market share.

The other three single-objective utility functions are slightly risk prone. Cash flow is valued as a fundamental objective for its usefulness on investments, shareholder reaction, and long-term profit. It was thought that each additional hundred million dollars would offer slightly better investment opportunities. As the charge card was a relatively new AX product, the main purpose was growth. Growth was expected to occur over time, but there was uncertainty about the rate of growth. The utility function indicates a willingness to take risks to possibly speed up the growth rate. Profits would be distributed to shareholders and used for corporate investments. Considering both of these uses, and that more profitable investments cost more, the utility function was constructed as shown in the Fig. 4.

For each of the single-objective utility functions, the general shape of risk-prone or risk-averse segments were checked and confirmed with the management team. In Fig. 4, each of utility functions was scaled from 0 to 1 indicated by the solid circles, and the open circles indicate the assessed values consistent with the scaling. The equations representing the individual utility functions are

$$u_1(x_1) = -0.3.51 \left(1 - e^{0.8074(x_1-1.2)} \right), 1.2 \leq x_1 \leq 3.0, \tag{6}$$

$$u_2(x_2) = -0.3051 \left(1 - e^{0.2422(x_2-6)} \right), 6 \leq x_2 \leq 12, \tag{7}$$

$$u_3(x_3) = \begin{cases} -0.01955 \left(1 - e^{0.656(x_3-12)} \right), & 12 \leq x_3 \leq 17, \\ +0.5 + 0.6526 \left(1 - e^{0.4844(17-x_3)} \right), & 17 \leq x_3 \leq 20, \end{cases} \tag{8}$$

$$u_4(x_4) = -1.0268 \left(1 - e^{0.2267(x_4-3)} \right), 3 \leq x_4 \leq 6. \quad (9)$$

The prioritization of the ranges of the fundamental objectives are calculated by substituting the appropriate u_i , $i = 1,2,3,4$, levels from (6) to (9) into the four Eqs. (3) to (6) and solving for the weighting factors w_1 . These weighting factors are the prioritization of the ranges of the objectives shown in Table 4.

Evaluating customer acquisitions

The cost of acquiring customers and the implications of such acquisitions in any particular marketing program is small relative to the ranges of the objectives in Table 4 used to construct the utility function for AX. All other decisions at AX and exogenous factors related to the economy, the market, and political affairs can affect overall AX performance in terms of these objectives. Hence, to evaluate alternatives to attract different potential customers to various AX products (i.e. charge cards or credit cards), it is necessary to compare the desirability of those different customers relative to the current performance at AX.

The marketing group, in coordination with risk management and financial groups, implemented the model beginning in late 1995 and initially used it for over 50 million prospective customers. As a result, targeting strategies were altered to focus on customers who would be both a higher usage and a higher risk group. This significantly increased customer acquisition volume, exceeding business goals of acquisition in 1996 by a large number. In the first 6 months of 1997, the AX market share of the domestic card market increased from 18.31 to 18.90 %, whereas it lost market share in each of the 6 month periods over the previous 10 years. Calculating the increase in utility due to this increase of market share using the business utility function and then calculating the increase in short-term profit that would provide AX with that same increase in utility indicates that the 0.59 % increase in market share has an equivalent value to AX of an additional \$447 million of profit.

The Wall Street Journal published an article on this impressive turnaround and improvement (Frank 1997). In the article, H. Spencer Nilson, the publisher of the Nilson Report, the trade journal covering the payment systems industry, said “You’ve got to give credit to American Express for overcoming huge obstacles to growth through creative marketing and excellent customer service.” Also in the article, Kenneth I. Chenault, the then president of American Express who is currently the chairman and chief executive officer, was quoted as saying “Market share has been a very key focus of our card business over the last 3 years and I think what you’re seeing is that we have had a very strong effect in increasing spending on our card products and our cards in force”.

Summary and conclusions

Compared with single-objective decisions, the unique features of multiple-objective decisions involve identifying a complete set of those objectives and prioritizing them for any intended use. Simply writing a list of the objectives, without help and

direction for thinking, has been shown to typically overlook more than half of the relevant objectives. Prioritizing objectives, without considering the ranges of possible consequences on those objectives, is logically flawed and misleading. Concepts, logical foundations, and procedures required to do each task well are discussed and their use is illustrated in this paper.

To enhance the identification of objectives from an individual, concepts to stimulate individuals to think more deeply and challenging them to set a higher goal for the number of objectives improves the results. Gathering ideas from friends and colleagues about appropriate objectives is also useful. The concepts of means–ends and inclusion relationships help structure a list of resulting objectives. This structuring facilitates identifying potentially missing objectives and categorizes objectives into two types: fundamental objectives and means objectives. The means objectives are useful for suggesting alternatives to consider and ensuring that all relevant objectives are considered. The fundamental objectives define why one is ultimately interested in the decision and provide a complete and independent set of objectives for evaluating alternatives.

The fundamental objectives are those that should be prioritized for a decision. The logical foundation for prioritization is an objective function over the metrics that measure the degrees of achievement on each of the fundamental objectives. Priorities logically must depend on the changes in the levels of achievement on the various objectives and not on the names of those objectives independent of ranges. To illustrate this critical, and frequently misunderstood and misrepresented point, consider safety improvements that individuals could make to their own private vehicles. Suppose the only two objectives of such improvements are to maximize safety, measured by the annual probability of dying in a vehicle crash in that car, and to minimize cost, measured in euros. If asked for this decision, “What is more important, the safety objective or the cost objective?” Most individuals would respond that safety is most important. Then suppose there are two alternatives, where alternative A costs €5000 and reduces the annual probability of dying by ten chances in a million, and alternative B costs €40 and reduces the annual probability of dying by three chances in a million. The prioritization of safety over cost, if it means anything, suggests that alternative A should be chosen. However, most individuals who stated the priority of safety over cost would conclude based on the consequences that alternative B is easily preferred and justify that choice with “It isn’t worth €5000 to reduce the annual probability of death by ten chances in a million, and it is worth €40 to reduce that probability by three chances in a million”. This reasoning indicates that the ranges of the consequences definitely matter in prioritization.

The use of the concepts, logical foundations, and procedures presented to identify and prioritize objectives are illustrated in two applications. One involves working with the Department of Homeland Security in the U.S. on an all-hazards risk analysis that could be useful to consider allocations of funds across dispersed hazards such as terrorism, natural disasters, illegal immigration, and pandemics. This effort involved developing a reasonable set of fundamental objectives for an analysis and structuring the numerous means objectives that would influence the consequences in terms of these fundamental objectives.

The second application concerns evaluation of prospective customers for charge cards and credit cards at American Express. As American Express had been continuously losing market share to rivals for over a 10-year period prior to the described study, the project was intended halt or reverse this serious detrimental trend. The effort developed a set of fundamental objectives that included profits, market share, and cash flow constructed from judgments of executives at American Express. After carefully defining the ranges of possible consequences for each of those fundamental objectives, a utility function was constructed from value judgments of these executives. Implementation of an analysis using this utility function followed and contributed to a significant improvement in the market share of American Express cards.

References

- Bond SD, Carlson KA, Keeney RL (2008) Generating objectives: can decision makers articulate what they want? *Mgt Sci* 54:56–70
- Bond SD, Carlson KA, Keeney RL (2010) Improving the generation of decision objectives. *Decis Anal* 7:238–255
- Clemen RT, Reilly T (2001) Making hard decisions with decision tools. Duxbury, Pacific Grove
- Dyer JS, Sarin R (1979) Measurable multiattribute value functions. *Ops Res* 22:810–822
- Eisenführ F, Weber M, Langer T (2010) Rational decision making. Springer, Heidelberg
- Fishburn PC (1965) Independence in utility theory with whole product sets. *Ops Res* 13:28–45
- Frank SE (1997) American Express credit-card share increased to 18.9% in first half of year. *Wall Str J*, A3
- Keeney RL (1981) Analysis of preference dependencies among objectives. *Ops Res* 29:1105–1120
- Keeney RL (1992) Value-focused thinking: a path to creative decision making. Harvard University Press, Cambridge
- Keeney RL (2002) Common mistakes in making value tradeoffs. *Ops Res* 50:935–945
- Keeney RL (2007) Developing objectives and attributes. In: Edwards W, Miles RF, von Winterfeldt D (eds) *Advances in decision analysis*. Cambridge University Press, Cambridge, pp 104–128
- Keeney RL, Lin Q (1999) Evaluating customer acquisitions at American Express using multiple objectives. *INFORMS National Meeting*, Philadelphia, November 7–10
- Keeney RL, Lin Q (2000) Evaluating customer acquisition at American Express using multiple objectives. *Interfaces* 30(5):31–33
- Keeney RL, Raiffa H (1976) *Decisions with multiple objectives: preferences and value tradeoffs*. Wiley, New York (Reprinted by Cambridge University Press, 1993)
- Keeney RL, Bennett SB, Levine ES, Kolasky RP (2010) Comparing different types of homeland security hazards. *Society for Risk Analysis Annual Meeting*, Salt Lake City, Utah, December 5–8
- Kirkwood CW (1997) *Strategic decision making: multiobjective decision analysis with spreadsheets*. Duxbury Press, Belmont
- Phillips LD (1984) A theory of the requisite decision models. *Acta Psych* 56:29–48
- Salo AA, Hämäläinen RP (1997) On the measurement of preferences in the analytic hierarchy process. *J Multi-Criteria Decis Anal* 6:309–319
- Weber M (1997) Remarks on the paper ‘On the measurement of preferences in the analytic hierarchy process’ by A. A. Salo and R. PHämäläinen. *J Multi-Criteria Decis Anal* 6:320–321