Decision Analysis:
Concepts, Tools and Promise

A Fair Isaac White Paper

Zvi Covaliu

May 2001

1 800 999 2955 *from the US* 1 415 472 2211 *from anywhere* info@fairisaac.com *email*

www.fairisaac.com

Table of Contents

Introduction

In the last two decades, the field of decision analysis has had a fast-increasing impact in the way organizations, private and public, are making strategic and tactical decisions. Major advances in theory, modeling tools and computational techniques, made possible by leaps in computer science, have made decision analysis increasingly essential in business and government decision making. These techniques are being used routinely by some of the largest firms in business and in a myriad of applications, from yield management to product development to investment to medicine, to name a few.

At Fair Isaac, some powerful and unique decision analytic tools have been developed in recent years. These tools add significant value to clients who use our products and services to optimize their strategies. While traditionally the emphasis at Fair Isaac had been on sophisticated prediction technologies, it has shifted in recent years to developing and using a powerful, synergetic coupling of these time-proven tools with state-of-the-art decision modeling and techniques. The combination of new decision analytic expertise, models, and tools with existing prediction and exploratory analysis techniques, promises substantial added value to clients. This is achieved through explicit modeling of multiple business objectives and their tradeoffs, of the sequential nature of many decisions, through a proactive, continuous learning from observations, and implementing constraints. Recently, for example, tests conducted in the area of credit line strategies with data from several top-tier credit card issuers, have shown potential profit increases of 5–35% over 18 months.

The advent and spread of the Internet and its related technologies not only allow Fair Isaac to provide specialized decision analytic services in ASP mode, but also enable an unprecedented system of collecting and analyzing consumer data, to better model and update consumer preferences and to offer customized products in real-time.

What is Decision Analysis?

Decision analysis refers to the broad quantitative field, overlapping *operations research* and *statistics*, that deals with modeling, optimizing and analyzing decisions made by individuals, groups and organizations.

The purpose of decision analysis is to assist decision makers in making better decisions in complex situations, usually under uncertainty. The quality of the decisions is measured by their *expected consequences* and the stated *preferences* of the decision maker(s). The decision analytic framework helps the decision maker think systematically about his or her objectives, preferences, and the structure and uncertainty in the problem, then model quantitatively these and other important aspects of the problem and their interrelationship.

The Typical Decision Analysis Problem

In the generic but typical *decision problem*, the decision maker needs to choose an *alternative* from a *domain* of possible alternatives. The *consequences* to the decision maker, in terms of objective achievement, depend not only on that choice but also on the "*state of nature*," one or more events about which the decision maker is uncertain, has no control on, but holds some beliefs in the form of assessed probabilities. To obtain more information about the state of nature, the decision maker may choose, prior to making the main decision, to conduct one or more *experiments* (tests, surveys, samplings, etc.) at known costs. The *optimal solution* to the problem prescribes for each decision, and each combination of alternatives chosen and outcomes realized prior to that decision, which is the "best" alternative to pursue. The resulting *optimal strategy* optimizes, subject to the decision maker's *risk attitude*, some previously-established measure of objective attainment, sometimes referred to as a *utility* function.

About this paper

This paper provides a brief overview of the concepts, methodology and tools of modern decision analysis. We start with a broad perspective of the scope of decision analysis, its many applications, and its relations with other disciplines. We then define the main elements of a decision problem and describe influence diagrams and decision trees as the main graphical paradigms. The key concepts we discuss and illustrate include decision rules, risk and risk attitude, multiple objectives and tradeoffs, sequential decisions, experimentation and Bayes inference. The importance of tools such as sensitivity analysis, value of information, and constrained optimization is also addressed.

Throughout, we attempt to illustrate the concepts, methodology and tools, mainly through a two-stage simplified example of credit initiation and acquisition.

Applications and Scope of Decision Analysis

Common applications

Decision analysis is widely used in business and government decision making. Following is a non-exhaustive list of most common applications:

Business

- Airline and hotel yield management
- Oil exploration
- Quality assurance and control
- Reliability and maintenance
- Crop protection
- Credit and loan portfolio management
- Project selection
- New product development
- New venture launching

Government

- Emergency management
- Environmental risk management
- Choice of new energy sources
- Research and development programs

Common

- Medical diagnosis and treatment
- Bidding and negotiation
- Litigation

The scope of decision analysis

Decision analysis is a *prescriptive* discipline that is designed to assist people in making better decisions. It is prescriptive in the sense that is uses a normative framework, and a set of tools and procedures to help the decision maker model, optimize and analyze complex, hard decisions. In contrast, there exists a *descriptive* view of decision making, which focuses on how people actually make decisions. This view, which relies heavily on psychology, provides ample experimental evidence that people generally process information, assess probabilities, and make decisions in ways inconsistent with the rational prescription of decision analysis. These findings only emphasize the importance of using the tools of decision analysis in making good decisions, particularly when they are difficult and important.

Since most decisions are made under uncertainty, it is essential to distinguish between a good (high-quality) decision and a lucky outcome. The former is the result of following the rational approach prescribed by decision analysis, based on thorough understanding and proper modeling of all relevant aspects of the decision problem, while the latter may be just a fluke of luck—an unlikely favorable consequence that is realized regardless of decision quality.

Related disciplines and methodologies

Decision analysis borrows from, relies on or overlaps a number of related disciplines and methodologies. The following is a non-exhaustive list:

- Probability theory
- Statistical decision theory
- Bayes learning and inference
- Graph theory and paradigms (e.g., Bayes networks)
- Stochastic processes (e.g., Markov processes)
- Stochastic dynamic programming
- Information theory
- Utility theory, including multi-attribute utility theory
- Game theory
- Negotiations
- Mathematical programming
- Goal programming

Main Elements of a Decision Model

All complex decision problems include the following main elements:

- *Decisions*, to be made by the decision maker;
- *Uncertain events*, whose outcomes determine the state of nature;
- *Consequences*, which result from the decisions and the outcomes of the uncertain events; and
- *Objectives* and *preferences*, which determine how the decision maker feels about the consequences.

Decisions and alternatives

A *decision*, D_i, $i = 1, 2, ...$ refers to a point in time when the decision maker has to choose one *alternative*, d_i, out of a domain of available alternatives, that could be discrete or continuous, simple or combined. For example, in the context of a credit granting decision, it can simply be the set {accept, reject}; in the context of a credit line assignment decision, it can be any value between, say, $500 and $50,000. An alternative can be a combination of values to be chosen at the decision point for a number of *decision variables*. Table 1 illustrates this point in the context of a pre-screened credit card offer decision, where each alternative consists of a combination of four values, one for each decision variable.

TABLE 1: OFFERS AS ALTERNATIVES

Alternative d_i	Promotional purchase rate (APR)	Promotional balance transfer rate (APR)	Go-to rate (APR)	Promotional period (months)
1	2%	6%	14%	4
2	4%	3%	16%	6
3	0	9%	18%	9

What separates one decision, say D_1, from another, D_2, is the difference in the *information set* available to the decision maker before each decision is made. The information set corresponding to a decision is the set of all observations available to the decision maker prior to making that decision. For example, prior to making an offer decision, say D_1, the information set may contain a credit bureau risk score and a revenue score for each target individual; prior to making a subsequent credit granting and credit line decision, say D_2, the information set would include, in addition to that available at D_1, some application information, notably the income of the applicant.

Uncertain events

Uncertain events, X_j, $j = 1, 2, ...$, typically occur interspaced between subsequent decisions. If an uncertainty realizes before a decision is made, its outcome, x_j will typically be observed by the decision maker before that decision is made. For example, an applicant's income and other credit application information, while uncertain at the time an offer decision, say D_1, is made, will have been realized and observed before the subsequent credit line decision, say D_2, is made. Still, at the time D_2 is made, there would be a number of unresolved uncertainties, like the true credit worthiness of the applicant and his or her future use of the credit and repayment patterns. Together with the decisions already made, such remaining uncertainties determine the consequences to the decision maker, and therefore need to be estimated upfront in the form of probability distributions.

Consequences

The *consequences*, V_k, $k = 1, 2, \ldots$, to the decision maker are the results of, and determined by, the alternatives chosen at all decision points D_i, and the outcomes of the uncertain events X_j [1]. They are themselves uncertain at the time all decisions need to be made and are closely related to the objectives of the decision maker. For example, the cost of a marketing campaign, and the size, revenue and loss of a resulting portfolio, are all consequences.

Objectives and preferences

The decision maker's *objectives* in solving a decision problem are the quantities he or she cares about, including their preferred direction. Maximizing a portfolio's size is an example of an objective; minimizing a portfolio's loss is an example of another.

Rarely in realistic decision problems is there a single objective. When there are two or more objectives, they typically *conflict*, in the sense that some strategy is *optimal* (performs best) with respect to one objective, while a different strategy is optimal with respect to another. In these situations the decision maker needs to articulate his or her preferences in terms of *tradeoffs* among the objectives. For example, in the context of credit portfolio decisions, the decision maker may have to articulate how much expected loss he or she is willing to accept to increase the volume by 10%.[2]

Even with a single objective, say maximizing profit, the decision maker should state his preferences with respect to risk, or *risk attitude*, whenever uncertainty is involved. Figure 1 shows two hypothetical profit distributions, or *risk profiles*, for two different strategies, (a) and (b). Clearly, even though the expected profit with strategy (a) is larger than that with strategy (b), most decision makers would prefer (b) to (a) because of the higher uncertainty and much higher probability of losses in (a).

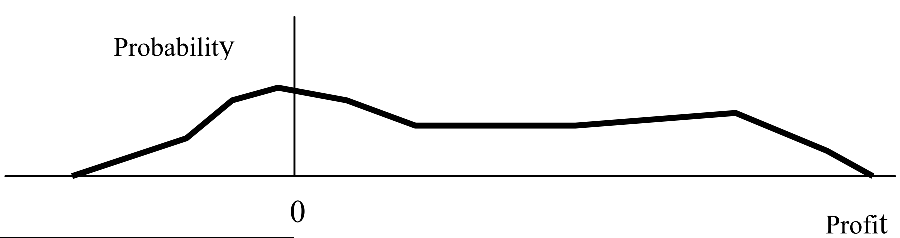

[1] That is, each V_k is a function of all decisions made $\mathbf{D} \equiv \{D_i,\ i = 1, 2, \ldots\}$ and all realized events $\mathbf{X} \equiv \{X_j,\ j = 1, 2, \ldots\}$: $V_k = V_k(\mathbf{D}, \mathbf{X}),\ k = 1, 2, \ldots$.

[2] This tradeoff factor will typical depend on the current size of the portfolio, among other things.

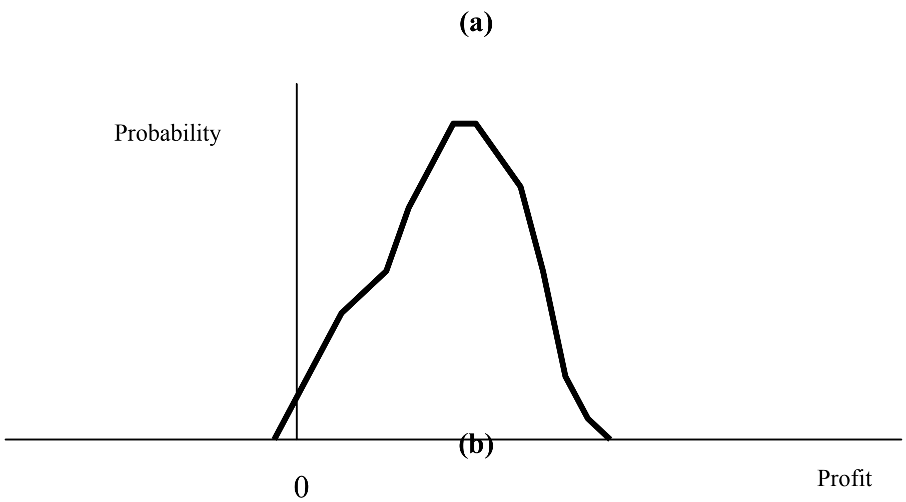

FIGURE 1: PROFIT DISTRIBUTIONS

Graphical Paradigms and Modeling

Graphical paradigms play an important role in modeling and structuring decision problems. The two most commonly used graphs are *decision trees* and *influence diagrams*. In both types of graphs, *decision nodes*, represented by rectangles, label decisions; *chance nodes*, represented by ovals, label uncertain events; and *consequence* or *value nodes*, represented by rounded rectangles or diamonds[3], represent consequences.

Influence Diagrams

An *influence diagram* is an acyclic directed graph in which each node labels a single variable of a decision problem and the arcs represent two main types of relationships among the variables. Arcs into decision nodes represent that all variables labeled by the nodes from which the arcs emanate (called *direct predecessor* or *parent* nodes) are observed by the decision maker before the decision is made. These are sometimes called *information arcs*. Arcs into chance or consequence nodes represent possible probabilistic dependence on their direct predecessor and are usually referred to as *dependence arcs*.

Figure 2 shows an influence diagram that illustrates these notions and the modeling power of influence diagrams. It is a simplified model for a 2-stage credit card campaign decision problem. Before making the Offer decision, the only information available to the decision maker consists of the Risk Score and the Revenue Score of the candidate, a fact captured by the two (information) arcs into the Offer node. Alternatives at the Offer node may be of the type listed in Table 1. At the Credit Limit decision, the decision maker has observed, in addition to the two scores, whether the customer responded to the offer, and, if so, the Income on his credit application. The arc from Offer to Credit Limit conveys that the decision maker remembers and is aware of the alternative chosen (which offer has been made) at the previous decision.[4]

[3] There is less convention for value nodes representation, and sometimes they are represented by triangles or octagons.

[4] Such an information arc, between two decision nodes, is called a *no-forgetting arc*.

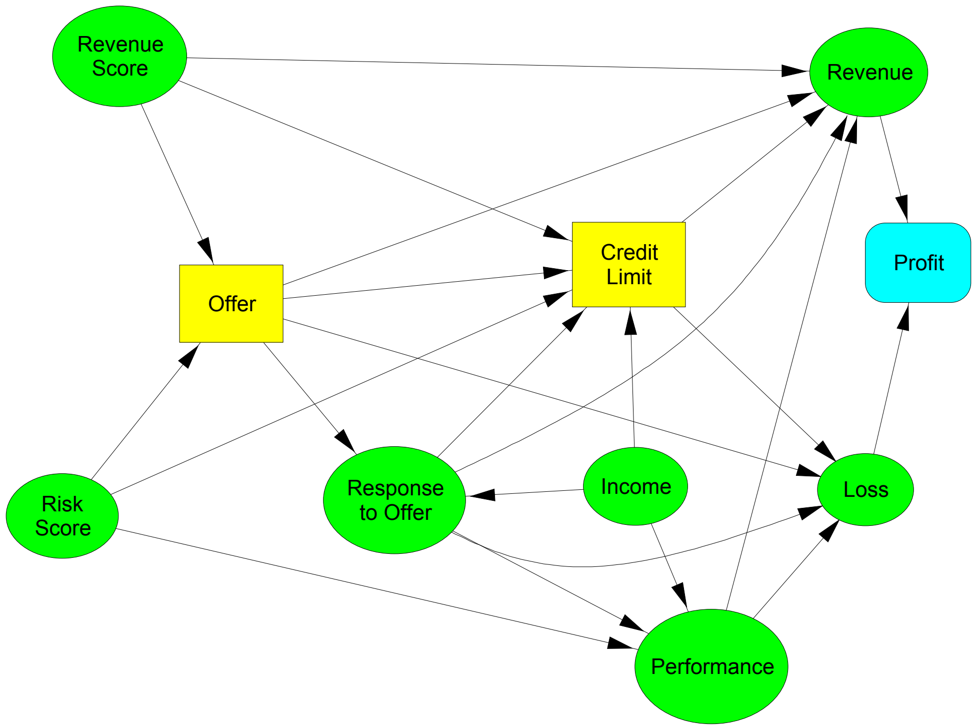

FIGURE 2: INFLUENCE DIAGRAM FOR TWO-STAGE PROBLEM

The rest of the arcs in the influence diagram are dependence arcs. A strong statement in an influence diagram is made by the lack of such arcs, which implies *conditional independence*. For examples, some conditional independence statements made by the model in Figure 2 (whether realistic or not) are:

- Given the Revenue and Loss, the Profit is conditionally independent of everything else, in particular the scores;
- Given the Risk Score, the (positive) Response and the Income, the credit Performance ("Good" or "Bad") of the customer is conditionally independent of the Revenue Score and both decisions;
- Given the Revenue Score, both decisions, the (positive) Response, and the customer's credit Performance, the Revenue is conditionally independent of both the Risk Score and the Income.

Strengths of influence diagram

Influence diagrams are a most powerful tool in modeling decision problems. This is because they allow one to structure and visualize fairly complex problems in a compact graph that conveys explicitly the assumed dependence or independence among variables, the sequence of decisions, and the flow of information to the decision maker. They are most effective in the early stages of modeling an unstructured problem, when data and other details are unavailable, as a communication tool between a decision analyst (consultant) and a decision maker (client). [In conjunction with sensitivity analysis, they allow the determination of what matters in a problem and what does not, and thus the construction of tractable models that provide insight into the problem and its solution.

More advantages of influence diagrams are mentioned below.

Limitations of influence diagrams

The most significant limitation of conventional influence diagrams is their inability to capture the asymmetric structure of a decision problem. *Asymmetry* in a decision problems refers to the very common situation where different *scenarios* do not have the same realization of variables or the same order of variables realized. In the example above, for instance, if the customer does not respond to the Offer, then the Credit Limit decision is never made; and other events, like Income and Performance, are never realized.

The downside of the compactness of influence diagrams is that the level of function and number behind each node are not apparent on the graph.

Decision Trees

In contrast to influence diagrams, *decision trees* explicitly show any asymmetry in the structure of a decision problem. They also show the functional and numerical details for each node on the corresponding *branches*. Each branch emanating from a decision node corresponds to an alternative, and each branch emanating from a chance node corresponds to a possible outcome.

Figure 3 shows a small portion of a decision tree for the credit origination problem described above. It explicitly shows the following asymmetries:

- When there is no Response, the immediate realization of Profit[5] is the end of this scenario; other events, like Income and Performance, are never realized, and the Credit Limit decision never gets to be made;
- When the customer applies but, based on the Income information, the decision maker decides to not grant credit, the immediate realization of Profit[6] is similarly the end of this scenario.

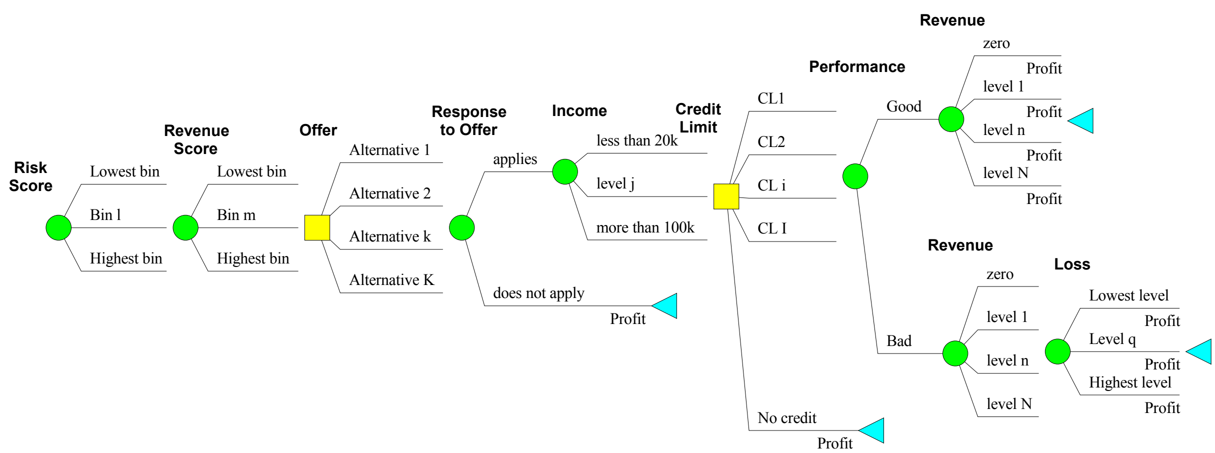

FIGURE 3: PARTIAL DECISION TREE FOR TWO-STAGE PROBLEM

[5] The fixed cost of sending the offer, say, which is implicitly modeled.

[6] An implicitly-modeled fixed cost that also would include in this case the evaluation of the application.

The explicit depiction by decision trees of the asymmetric and detailed structure of a problem comes at a huge cost: a decision tree's *size*, usually measured by the number of *leaves* (end nodes) grows exponentially in the number of variables in the problem. In Figure 3, for example, for each possible path from the root of the tree, the Risk Score node to any of the Response nodes, there would be a *sub-tree* similar to the two shown in the figure, starting at the Response nodes. For example, for 25 bins of the Risk Score, 20 bins of the Revenue score and 10 alternatives (offers) to choose from, there would be 5,000 such sub-trees!

This major limitation of decision trees necessitates the use of *schematic decision trees*, where branches and sub-trees are denoted schematically rather than explicitly. It also has led to a number of frameworks, generalizing the influence diagram paradigm, to include arcs, additional nodes or notation to capture asymmetry. One such framework is implemented in Fair Isaac's proprietary decision analytic software.

Key Concepts and Tools of Decision Analysis

Decision Rules, Risk Attitude and Utility—Single Objective

It is rare for the decision maker to have a single objective in mind when confronted with a decision problem, but even then, she needs to choose among strategies based on their corresponding risk profiles, as illustrated in Figure 1 on page 7. To make such a choice rationally, there is need for an agreed *decision rule*, a rule that specifies how different strategies should be evaluated in achieving the stated objective.

Maximize Expected Value

The most commonly used decision rule when faced with a single objective is the *expected value*. If profit alone is the objective, this is known as the *EMV* (Expected Monetary Value) decision rule. Taking into account only the first moment of the objective's distribution, it simply ignores any risk in selecting a strategy, and as such is a poor decision rule. A number of decision rules that consider risk are listed below.

Maximize Expected Value and Minimize Variance

This decision rule measures the risk through the variance of the objective's distribution, and is commonly used in selecting among security portfolios and many other investment decisions.

Figure 4 illustrates the two criteria for a portfolio composed of two securities, 1 and 2, where 2 has a higher expected return, $E_2 > E_1$, but also a higher risk, $\sigma_2 > \sigma_1$, and the decision maker needs to select her optimal mix (i.e., the fractions of each of the securities, x_1, x_2, in the portfolio, where $x_1 + x_2 = 1$). Points along the solid line in the figure form an *efficient frontier*, but the rule does not prescribe which out of these should be chosen, since it does not incorporate the decision maker's subjective attitude towards risk.

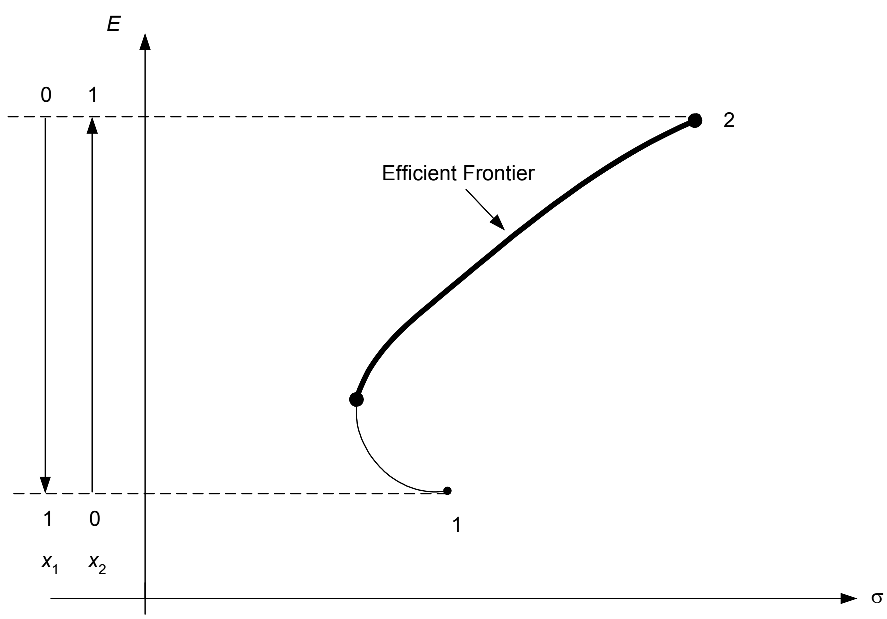

FIGURE 4: EFFICIENT FRONTIER FOR A PORTFOLIO OF TWO SECURITIES

Choose Dominating Strategy

Strategy 1 is said to *stochastically dominate* strategy 2, if, for any value *x*, the probability that the objective, say profit, associated with the former, X_1, exceeds *x* is at least as large as the probability that the latter, X_2, exceeds *x*:

$$\Pr(X_1 > x) \geq \Pr(X_2 > x) \text{ for all } x,$$

that is, for any value of profit, for example, the dominating strategy is always more likely to exceed it than the dominated strategy.

An example of such a situation is shown in Figure 5.

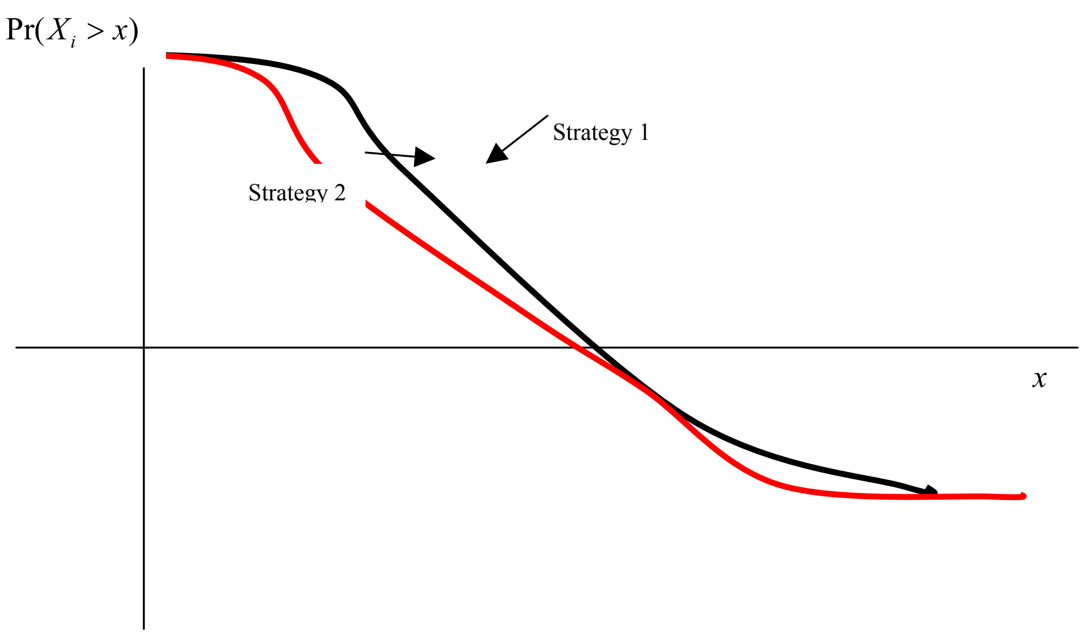

FIGURE 5: STRATEGY 1 STOCHASTICALLY DOMINATES STRATEGY 2

Stochastic dominance[7] is a powerful decision rule, independent of any rational decision maker's attitude towards risk. Unfortunately, it is fairly uncommon for strategies to stochastically dominate each other.

Maximize Expected Utility

The most coherent way to incorporate a decision maker's attitude towards risk is to assess her *utility function* for the relevant consequence (*criterion*), e.g., Profit, and then choose the strategy that maximizes the *expected utility*.

Utility theory provides the procedures for constructing a decision maker's subjective utility function. Most commonly, the decision maker chooses her *certainty equivalent* values for a number of specific *simple lotteries* presented to her, which have outcomes in the range relevant to the decision problem in question. A couple of simple lotteries are illustrated in Figure 6. The

[7] Strictly speaking, this is the definition of *first-order stochastic dominance*.

certainty equivalent for a simple lottery is the consequence (for example dollar amount) for sure that the decision maker is willing to accept in lieu of (be indifferent between it and) playing the lottery. For the lottery in Figure 6(a), the certainty equivalent exceeds the expected value; for the lottery in Figure 6(b), the negative certainty equivalent is far less than the expected value. The decision maker is said to exhibits *risk proneness* in the former case and *risk aversion* in the latter. If given a choice, such a decision maker would prefer:

- Any sure amount higher than 30 to the lottery in (a);
- The lottery in (b) to sure losses only higher than 5; and
- Lottery (a) to lottery (b).

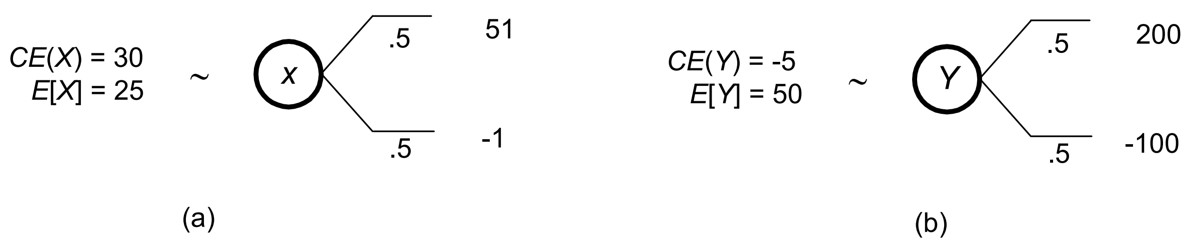

FIGURE 6: TWO 50/50 LOTTERIES

The result of such an assessment procedure is a *utility function* of the type depicted in Figure 7. Concavity of the utility function represents a region of the domain of the decision maker's assets where the decision maker is risk-averse, while convexity represents a region where the decision maker is risk-prone. The shape and locus of the utility function depend, to a large extent, on the current assets of the decision maker. An individual that plays the state lottery, for example, exhibits risk-prone behavior because the expected value of the lottery is lower than the ticket cost. On the other hand, individuals pay insurance premiums, referred to as *risk premium* in the utility theory jargon, because they are typically risk-averse in the range of values associated with houses, cars, etc.

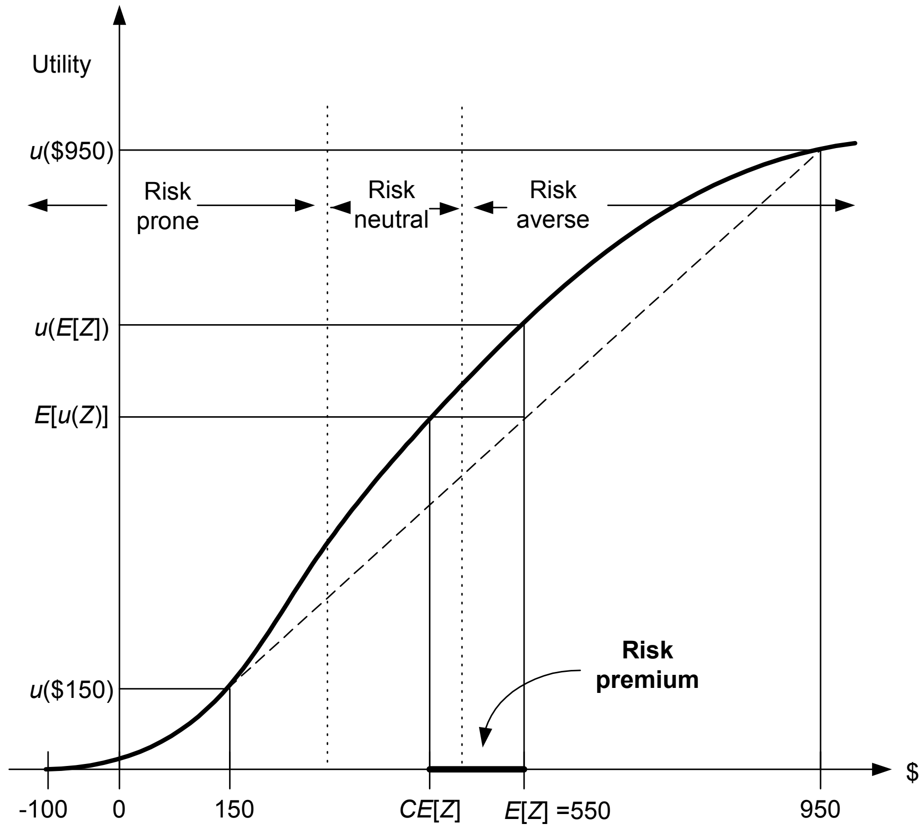

FIGURE 7: A UTILITY WITH CHANGING RISK ATTITUDE

By definition:

Risk premium = Expected value of lottery – Certainty equivalent of lottery.

Figure 7 graphically illustrates these notions for a 50-50 lottery, *Z*, in which the decision maker can win either \$150 or \$950. Clearly, $E[Z] = \$550$. The certainty equivalent of the lottery, $CE[Z]$, is the (certain) dollar value that has the same utility as the expected utility of the lottery, $E[u(Z)] = .5u(\$150) + .5u(\$950)$. The risk premium is then $E[Z] - CE[Z]$, the amount the decision maker is willing to "give up" to avoid the risk.

Decisions with Multiple Objectives

When multiple objectives are at stake, they must somehow be aggregated into a single measure of performance, to which a decision rule can be applied, unless a subjective, judgmental decision rule is left to the decision maker's choice. One way to reconcile conflicting objectives is through tradeoffs.

Tradeoffs

Explicit tradeoff factors allow the decision maker to specify how much she is willing to give up in one objective to gain a unit in another. For example, in a credit accounts portfolio, where both loss and volume are important, a tradeoff could measure how much the portfolio manager is willing to risk in expected loss in order to increase expected volume by, say, 1000 accounts.

One limitation of tradeoff factors is that their value is implicitly constant throughout the applicable range of the objectives, which typically is not true. For example, in the portfolio management example, the manager may be willing to increase expected losses by $10,000 to increase volume from 15,000 to 16,000 accounts, but only by $5,000 to increase volume from 100,000 to 101,000 accounts. Another shortcoming of tradeoff factors is in their failure to capture *interactions* among objectives.

Efficient Frontiers

A simple and very effective way to graphically depict the tradeoffs among objectives is by using *efficient frontiers*. Given a decision model, a *frontier* represents, in the space of two or more objectives or attributes, the set of all achievable points by a specific strategy.

Figure 8 shows, for example, the expected-Volume-expected-Profit frontier associated with an accept-reject strategy in a credit portfolio, using a single risk score. It illustrates that the lower the score cutoff, above which applicants are accepted, the higher the volume. In the high range of score values, where decreasing the cutoff mainly accepts "Good" applicants, the profit also increases. In the low range, however, continued decrease of the score cutoff results in reduced profit because more and more "Bad" applicants are accepted. Only the green portion of the frontier is *efficient* in the sense that, for any given level of expected Profit, all decision makers would prefer a higher, rather than a lower, expected Volume.

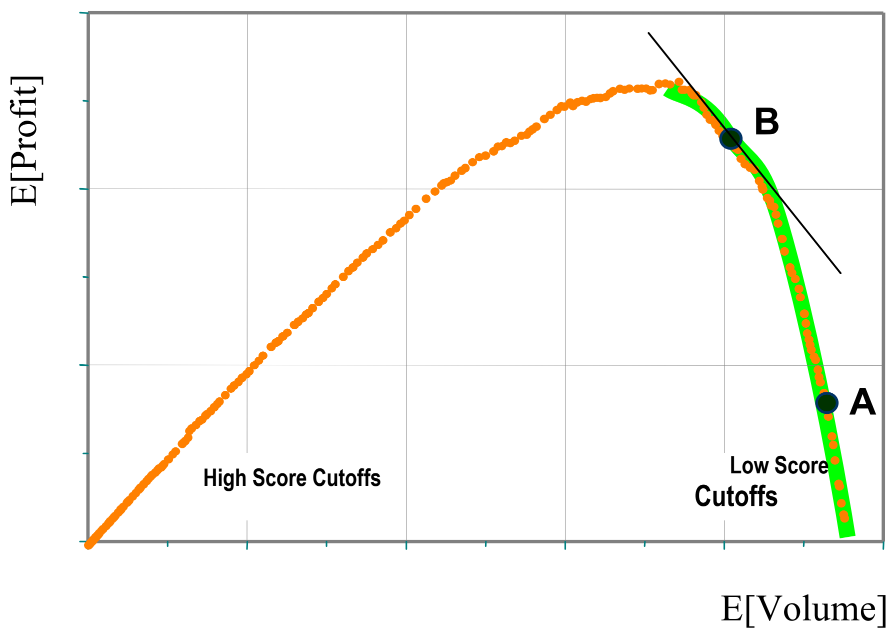

FIGURE 8: EFFICIENT FRONTIER IN PROFIT-VOLUME SPACE

About Fair Isaac

Fair Isaac Corporation (NYSE:FIC) is the preeminent provider of creative analytics that unlock value for people, businesses and industries. The company's predictive modeling, decision analysis, intelligence management, decision management systems and consulting services power more than 25 billion mission-critical customer decisions a year. Founded in 1956, Fair Isaac helps thousands of companies in over 60 countries acquire customers more efficiently increase customer value, reduce fraud and credit losses, lower operating expenses and enter new markets more profitably. Most leading banks and credit card issuers rely on Fair Isaac solutions, as do insurers, retailers, telecommunications providers, healthcare organizations and government agencies. Through the *www.myfico.com* Web site, consumers use the company's FICO® scores, the standard measure of credit risk, to manage their financial health. As of August 2003, HNC Software Inc., a leading provider of high-end analytic and decision management software, is part of Fair Isaac. For more information, visit *www.fairisaac.com*.

Corporate Headquarters:
200 Smith Ranch Road
San Rafael, CA 94903-5551
1 800 999 2955 *from the US*
1 415 472 2211 *from anywhere*
info@fairisaac.com *email*

Offices Worldwide:
Brazil, Canada, France, Germany,
Japan, Singapore, Spain,
United Kingdom, United States

www.fairisaac.com

1076WP 05/01 PDF