# Applying Goals and Cases to Business Decision Making

**Stephen Slade**
Information Systems Department
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012-1126
sslade@stern.nyu.edu

**Abstract**

We are developing artificial intelligence programs that incorporate goal-based and case-based reasoning to simulate decision making in a variety of business domains. In this paper, we first provide a summary introduction to our past work. We contrast our general approach with traditional decision analysis and utility theory. We provide brief descriptions of three current projects.

## Introduction

Our research has focussed on the task of decision making by modelling human cognitive behavior. Our programs explicitly represent goals and relationships. This requirement was proposed initially in our description of an *advisory system.*

> It is critical that the program examine a problem in the particular context of the present client. The program has to know about the client's goals and needs; what specific ramifications each alternative may have for the client; how the decision might affect the client's employees, owners, competitors, customers, and suppliers; and what priorities the client places on these possible effects. [Schank and Slade, 1984, page 251]

A description of past work is found in [Slade, 1992b; Slade, 1993]. Specific aspects of our research include the following.

- Qualitative reasoning [Slade, 1991e].
- Interpersonal relationships [Slade, 1990].
- Case-based reasoning [Slade, 1991a; Slade, 1991b].
- Decision strategies [Slade, 1991d].
- Natural language generation [Slade, 1991c].
- Explanation of decisions [Slade, 1992a].

We incorporated these features into the VOTE program, which simulated the roll call voting decision making of members of the United States House of Representatives. Given a member of Congress and a specific bill, VOTE would try to determine how that member would vote and then produce a natural language explanation of the resulting decision in English or French.

Below is an example of the VOTE program simulating Congressman Morris Udall voting on a bill banning the desecration of the American flag.

```
> (vote 'udall 'hr-2978)

* Member:        Morris K. Udall
* Bill:          Flag Desecration
* Bill banning the desecration of the flag.
```

— OMITTING INTERMEDIATE OUTPUT —

```
* English rationale:
```

> *Morris K. Udall votes against bill HR-2978, the flag desecration bill. After weighing the implications, he believes that provisions of this bill are not constitutional. He completely supports the United States Constitution and the Bill of Rights. Udall readily endorses the right of freedom of speech. Even so, Udall realizes that members of the Democratic party oppose the right of burning the American flag in protest.*

```
* French rationale:
```

> *Morris K. Udall s'oppose au projet de loi HR-2978, la loi de la profanation du drapeau. Après une considération approfondie, il croit que les dispositions de ce projet de loi ne sont pas constitutionelles. Il est un champion de la Constitution américaine et de la déclaration des Droits. Udall désire vivement appuyer le droit de libre expression. Cependant, Udall comprend que les membres du parti Démocratique s'opposent au fait de*

*brûler le drapeau américain lors d'une manifestation.*

The natural language explanation above is not canned text, but is generated automatically by VOTE. Similarly, the French text is not simply a translation of the English text, but is generated from the underlying knowledge representations.

The VOTE model of decision making was tested in the domain of Congressional roll call voting. However, there are many parallels between political decision making and business decision making. Both domains share the following attributes.

- **Abundance of goals:** numerous objectives which cannot all be satisfied.
- **Limited resources:** scarce resources lead to conflicts among goals.
- **Relationships:** a multitude of other agents through which additional goals are adopted.
- **Cases and experience:** a history of past decisions that might bear on the current choice.
- **Explanations:** a need to explain or justify the decision, particularly to other agents for whom the outcome may have adverse consequences.

We shall next contrast our descriptive, qualitative decision model with traditional prescriptive decision analysis. Then we shall briefly describe three ongoing projects at New York University which apply the VOTE decision making method to business problems.

## Prescriptive Decision Models

As Abelson and Levi (1985) have observed, there is a basic dichotomy in the decision making literature between prescriptive and descriptive models. Prescriptive or normative models focus on how people *should* make decisions, while descriptive theories explore how people *do* make decisions. The greater the observed difference between the two approaches, the greater the questions raised about human rationality. That is, if people's observed behavior is markedly different from optimal behavior, then we must question the assumption of rational thought.

We view prescriptive theories as top-down models, in which a formal reasoning paradigm is applied directly to decision making problems. The goal is to arrive at a correct or optimal outcome. The resulting model usually requires significant assumptions to achieve compatibility with the formal theory.

By contrast, descriptive theories can be viewed as bottom-up models, in which the data define the significant features and dimensions of the model. The resulting theory is derived to match the data.

We consider the different approaches, prescriptive versus descriptive, as not simply a dichotomy, but as a dialectic. That is, decision making theories develop in relationship and in response to other theories.

Economics has been the main field concerned with decision making. One significant dimension for decision making is the certainty of outcomes. Decision theory distinguishes among certainty, risk, and uncertainty as follows [Luce and Raiffa, 1957].

- *Certainty.* If the consequences of a decision are known with certainty, then the decision problem becomes one of direct optimization. The methods of linear programming and operations research are appropriate for such problems. Airline scheduling or production planning are examples of decision making under certainty.
- *Risk.* Some decisions have indeterminate outcomes that nonetheless have known probabilities. Most lotteries and casino gambling games have known risks – specific probabilities and payoffs. The expected value method from decision analysis is effective for these domains.
- *Uncertainty.* In many real world domains, the decision maker lacks accurate assessments of both the probabilities of success and the payoffs. That is, the agent is uncertain not merely of the outcome, but of the actual likelihood of a possible outcome and the actual value of the outcome.

The methods of decision making under risk may be applied to decision making under uncertainty, but rely on estimates of probabilities and payoffs. The field of decision analysis [Raiffa, 1968] provides a prescriptive mathematical foundation for selecting among alternatives. The theory rests on assumptions of quantitative measures for both preference and uncertainty. Preferences for different outcomes are assigned numerical utility values, and judgments about uncertainties are given numerical probabilities.

Decision theory has been applied to a variety of AI problems, including projecting future events in planning [Hanks, 1990], and advisory systems in the domain of genetic counseling [Holtzman, 1989].

The applicability of decision theory depends on the accuracy of the probability and payoff estimates. In many cases, payoffs cannot be measured on a uniform scale, such as money. In such cases, the models must reflect the trade-offs which the agent is willing to make among different types of outcomes. Multi-attribute decision models [Rennels *et al.*, 1987; Keeney and Raiffa, 1976] provide a framework for analyzing such decisions.

The mathematical basis of decision theory requires quantitative measures for payoffs. Economics, by its very nature, tends to focus on money. However, economists have recognized that the outcomes of many decisions cannot be reduced to monetary values. They have then postulated the notion of *utility* to account for such decisions. Utility theory makes a number of assumptions [Luce and Raiffa, 1957], including the following.

- *Transitivity.* If A is preferred to B, and B is preferred to C, then A is preferred to C. Indifference is likewise transitive.
- *Monotonicity.* If an agent prefers A, and likewise prefers a choice having probability $P(A)$ over a choice with $P'(A)$, then $P(A) > P'(A)$.

These properties are nicely met by numbers, and many economists readily and willingly assume that numbers provide the best means for measuring the utility of outcomes. However, utility theory does not, in fact, require the use of numbers, as pointed out by Luce and Raiffa.

> One may contend that introducing numbers does no harm, that they summarize the ordinal data in a compact way, and that they are mathematically convenient to manipulate. But, in part, their very manipulative convenience is a source of trouble, for one must develop an almost inhuman self-control not to read into these numbers those properties which numbers usually enjoy. For example, one must keep in mind that it is meaningless to add two together or to compare magnitudes of difference between them. If they are used as indices in the way we have described, then the only meaningful numerical property is order. We may compare two indices and ask which is larger, but we may not add or multiply them. [Luce and Raiffa, 1957, page 16]

The VOTE program complies with this ordinality requirement. VOTE ranks goals and relationships using simple ordinal values of A, B, and C. We could have chosen to use numbers instead, but we felt that the semantics of numbers was too strong. We did not want to find ourselves adding or multiplying or averaging levels of importance. Our usage is therefore consistent with the precepts of utility theory.

We further note that the standard model of decision making relies on weights and evaluation functions for calculating an optimal decision. For example, assume we use a 10 point scale, in which 10 is good and 0 is bad. If one choice involves two equally likely outcomes, we can see that an expected value of 5 could result from a number of different sets of values for the respective branches, such as the following.

[5, 5], [4, 6], [1, 9], [0, 10]

Each of these alternatives has the same expected value, but they differ considerably in their variance. The expected value masks the degree of consensus in the underlying choice components. The derived weights do not distinguish between conflict and indifference. By relying on symbolic ordinal values, we reduce the opportunity for averaging away significant information, while adhering to the ordinality requirement of utility theory.

However, VOTE does not enforce the transitivity requirement of utility theory. There is no requirement that a member's stances be consistent. Psychologists using factor analysis and conjoint analysis to measure utilities often observe inconsistent utility functions. For example, a subject's responses may suggest both

$a > b > c$

and

$c > a$

Often, such results are thrown out as experimental error, with the assumption that the subject had not properly understood the directions. Luce and Raiffa discuss this issue.

> Reported preferences [by experimental subjects] almost never satisfy the axioms, e.g., there are usually intransitivities. Furthermore, if the same pair is offered several times, then in some cases the subject will not be consistent in his reports. One cannot expect the data to fit the model perfectly, but how does one determine which model they fit most closely and how does one measure how good the agreement is? ... Every indication now is that the utility model, and possibly therefore the game model, will have to be made more complicated if experimental data are to be handled adequately. [Luce and Raiffa, 1957]

Goal-based decision theory does not require consistency. VOTE's database of members of Congress contains examples of inconsistent stances. We suggest that the goal-based model begins to address some of the problems associated with empirical application of utility theory.

Another problem with utility theory is the question of analyzing the consequences of an action. The AI problem of controlling inference arises when predicting a chain of events following a decision. How far should one follow the chain when evaluating alternative outcomes?

Stone (1988) discusses this problem of valuing intangible outcomes in the context of a cost-benefit analysis for a public health decision.

> A child vaccination program, for example, can be made to appear highly cost-effective if one counts as a benefit the number of lives saved, and values each life as a person's expected lifetime earnings. But now, count as part of the program's costs all the future medical expenses of the people "saved," and the program appears to save less money. Add in their children's schooling and medical care costs — children they would not have had if they had died without the vaccine — and these beneficiaries of our policy become burdens on the public treasury. Now, count as benefits of the vaccination program the taxes paid by the people we rescued and the program begins to look better; add their children's taxes and it looks better still. Why not go to the next generation, including in the analysis the costs and benefits to society they generate? And why not include psychic consequences, such as the security of knowing you and your loved ones are protected? Or tilt the other way, and include the insecurity of worrying that you or a loved one might be the rarity who actually catches the fatal disease from the vaccine? For the creative mind, the possibilities are endless. [Stone, 1988, page 203]

We observe that the prescriptive methods of decision analysis are mathematically precise, but for many real world domains lack principled ways of estimating the probabilities and payoffs of outcomes.

## Qualitative Securities Analysis

We now turn to current applications of goal-based and case-based reasoning to business domains.

We are working with Raghav Madhavan to develop a program, SAP, that simulates the behavior of a securities analyst [Madhavan, 1993]. In VOTE, the program was told to simulate a given member of Congress voting on a particular bill, and to provide a natural language explanation justifying the decision.

SAP simulates a specific securities analyst's recommendation on a given stock based on a new piece of information, and provides a natural language explanation for the recommendation.

The recommendation is either BUY, SELL, or HOLD. The new information might be an earnings report or news such as the resolution of some contingent liability, like the settlement of a law suit.

Unlike VOTE, SAP models not only the analyst's goals, but also her beliefs. For example, suppose that the news is that company X's earnings will be $1.50 per share. Is this considered good news or bad?

For analyst A, who had predicted earnings of $1 per share, this is good news, and she might upgrade her recommendation for company X.

By contrast, if analyst B had predicted earnings of $2 per share, then the news is bad, and the recommendation might need to be downgraded.

We note that a regular truth maintenance system [Doyle, 1979] may not be appropriate in this domain since the program will likely contain many inconsistent beliefs. By comparison, every member in the VOTE program possessed conflicting goals. There were always issues for which a given member could find reasons for either supporting or opposing.

These programs permit the representation of conflicting goals and beliefs. The programs are not consistent, in the logical sense. We believe this is a strength of our approach. Real life, especially in business, is inconsistent.

## Explanation-based Project Selection

With Henry Lucas and Michael Fish, we are developing a program to simulate the project selection decision, particularly applied to the acquisition of information technology.

Traditional capital budgeting paradigms, such as net present value, pay-back period, or return on investment, often prove inadequate when used to evaluate investment in new technology. This is in part due to the difficulty in identifying all the costs and benefits associated with new technology.

When a manufacturer builds a new plant, he has a good idea what costs he will incur for construction, financing, new equipment, and training because he or his company has likely gone through this experience before. Similarly, he can identify the benefits in terms of increased production efficiency, lower marginal costs, and increased product quality.

However, investment in high technology differs from traditional capital budgeting problems in that there is often little relevant prior experience. When a company buys a local area network or replaces its mainframe with a client-server architecture, chances are, this is a first-time venture. Thus, it is impossible to evaluate this option based on a previous case, since there is no previous case. Furthermore, given the rapid pace of technological development, investment in new technology may *never* be based explicitly on prior experience.

Typically, investment in information technology is justified based on tangible assessments, such as head-count reduction, and intangible estimates, such as increased productivity or product quality. We have identified dozens of standard justifications, some of which are listed below.

- keep up with the competition.
- lower marginal product cost.
- leverage previous technology investments.
- increase capacity and flexibility.
- reduce future costs of not investing.
- change organizational structure.
- meet customers' requirements.
- comply with regulatory requirements.
- centralize decision making.
- decentralize decision making.
- enhance coordination.
- reduce uncertainty and risk.

These types of justifications may not map directly into net present value calculations. We are developing a case-based reasoning system that will evaluate technology investment decisions based on knowledge of a company's goals and resources, and a library of past cases, indexed by the appropriate types of goals and justifications. The cases will include standard business school cases as well as cases from news stories and other sources.

Through an interactive process, the system will help the user identify and understand the possible costs and benefits associated with a given technology investment decision.

## Case-based reasoning for the case method

With Ken Laudon, we are developing a case-based reasoning tool to be used by MBA students preparing business cases. The case method is a pervasive pedagogical fixture in business schools. In a typical two year program, an MBA student may prepare up to 600 cases.

However, most case preparation focusses on a given business problem in isolation. That is, the case looks at company X and its history with little regard to the choices made by other companies in similar situations.

We are applying case-based reasoning to the case method by providing a case explorer tool that will serve as a repository of business cases with a rich set of indices. The student analyzing the problems of company X could use the case explorer tool to find other companies, perhaps in other industries, which faced similar decisions.

The tool could help the student develop explicit case-based reasoning skills. It would make it easier for students to argue from cases, rather than simply analyzing cases. Most case analysis is focused on issues in a single case, not on bringing out issues from other cases. In the real world, decisions often hinge on the degree to which one can find the best precedent or previous case on which to base a new decision. The initial set of cases is drawn from a textbook on information technology [Laudon and Laudon, 1991].

The case explorer implements the CBR paradigm explicitly as a decision making paradigm. Students would be taught the following.

1. Identify indices for the new case.
2. Retrieve similar cases from case library.
3. Adapt the features of the previous cases for the new problem.
4. Test the new solution.
5. If success, add to case library.
6. If failure, explain failure and cycle again.

By decomposing the process, we can teach students to be more critical and objective in their analysis. We also plan to integrate the VOTE program with the case explorer to permit VOTE to simulate the decision making problem posed in the case.

There are many important types of features that may be used as indices for the cases.

- **Industry type:** manufacturing, transportation, financial services, information technology.
- **Sector:** public, private, non-profit.
- **Level of decision maker:** ceo, division head, manager, consultant.
- **Players:** employees, customers, competitors, vendors, distributors, regulators.
- **Functional unit:** headquarters, manufacturing, marketing, finance, accounting.
- **Type of decision:** investment, technology, marketing, human resources, product design.

- **Industry position:** market share leader, niche player, low price producer, technology leader, start-up.
- **Strategies:** right sizing, customer driven, total quality, reduce cycle time, global expansion.
- **Economic goals:** net present value, return on investment, stock price, earnings.

Part of the analysis is determining which features are salient in a given case. In turn, that will influence what prior cases are retrieved from the case library.

The technology investment advisor described in the previous section can be seen as a focussed subset of the case explorer system. The utility of the case explorer will increase as the number of cases grows. At some point, it may begin to resemble the advisory systems described in [Schank and Slade, 1984].

## Software Availability

VOTE and Case Explorer are not commercial products. They are research tools. In the interest of exploring other decision making domains, the author will make these programs available without charge to other researchers via ftp. VOTE is available in both T [Slade, 1987] and Common LISP versions. A prototype Case Explorer is available as a suite of HyperCard stacks. Interested researchers should contact the author for details and current distribution information.

## Conclusion

We have developed methods for the explicit representation of goals, relationships, experience, explanation, and the process of decision making. Our approach is a qualitative, descriptive alternative to traditional prescriptive decision theory. We are now applying these representational techniques to the specific business decision making tasks of qualitative securities analysis, explanation-based project selection, and case-based reasoning for the case method.

## References


Abelson, R.A. and Levi, A. 1985. Decision making and decision theory. In *The Handbook of Social Psychology*. Random House, New York. chapter 5, 231–309. Third Edition.

Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence* 12:231–272.

Hanks, S.J. 1990. *Projecting Plans for Uncertain Worlds*. Ph.D. Dissertation, Yale University.

Holtzman, S. 1989. *Intelligent Decision Systems*. Addison-Wesley, Reading, MA.

Keeney, R.L. and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. John Wiley and Sons, New York.

Laudon, K. C. and Laudon, J. P. 1991. *Management Information Systems: A Contemporary Perspective*. Macmillan Series in Information Systems. Macmillan, New York, second edition.

Luce, R.D. and Raiffa, H. 1957. *Games and Decisions*. Dover, New York.

Madhavan, R. 1993. Goal based reasoning for securities analysis. In *Proceedings of the Second International Conference on Artificial Intelligence Applications on Wall Street*, New York, NY. IAKE.

Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Addison-Wesley.

Rennels, G.D.; Shortliffe, E.H.; and Miller, P.L. 1987. Choice and explanation in medical management: A multiattribute model of artificial intelligence approaches. *Medical Decision Making* 7(1):22–31.

Schank, R.C. and Slade, S.B. 1984. Advisory systems. In Reitman, W., editor 1984, *Artificial Intelligence Applications for Business*. Ablex Publishing, Norwood, NJ. chapter 14, 249–265.

Slade, S.B. 1987. *The T Programming Language: A Dialect of LISP*. Prentice-Hall, Englewood Cliffs, NJ.

Slade, S.B. 1990. A goal-based model of interpersonal relationships. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Boston, MA. Cognitive Science Society. 836–843.

Slade, S. 1991a. Case-based reasoning: A research paradigm. *AI Magazine* 12(1).

Slade, S.B. 1991b. Case-based reasoning for financial decision making. In *Proceedings of the First International Conference on Artificial Intelligence Applications on Wall Street*, New York, NY. IEEE Computer Society.

Slade, S.B. 1991c. Explaining political decisions. In *Proceedings of the Fourth University of New Brunswick Artifical Intelligence Symposium*, Fredericton, New Brunswick, Canada. University of New Brunswick.

Slade, S.B. 1991d. Goal-based decision strategies. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, Chicago, IL. Cognitive Science Society. 593–598.

Slade, S.B. 1991e. Qualitative decision theory. In *Proceedings of DARPA Workshop on Case-Based Reasoning*, Washington, DC. Defense Advanced Research

Projects Agency, Information Science and Technology Office.

Slade, S. B. 1992a. Generating explanations for goal-based decision making. *Decision Sciences* 23:1440–1461.

Slade, S.B. 1992b. *An Interpersonal Model of Goal-based Decision Making*. Ph.D. Dissertation, Yale University. Technical Report 892.

Slade, S. B. 1993. *Goal-based Decision Making: An Interpersonal Model*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Stone, D.A. 1988. *Policy Paradox and Political Reason*. Scott, Foresman and Company, Glenview, IL.