Qualitative Business Decision Making

ARTICLE · JULY 1995

Source: CiteSeer

4 AUTHORS, INCLUDING:

Stephen Slade
Yale University
23 PUBLICATIONS **227** CITATIONS

Ravi Aron
Johns Hopkins University, Carey Business Sc…
42 PUBLICATIONS **754** CITATIONS

Qualitative Business Decision Making

Stephen Slade, Ravi Arunkundram, Michael Fish, and Raghav Madhavan

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 model qualitative business decision making, using both goal-based and case-based reasoning. In this paper, we describe current work in three related areas: normative business goals and beliefs, qualitative business calculus, and belief representation.

Overview

We are developing artificial intelligence programs that incorporate goal-based reasoning (GBR) and case-based reasoning (CBR) to simulate decision making in a variety of business domains. In this paper, we provide describe three current projects.

A description of past work is found in [Slade, 1992b; Slade, 1994]. 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].
- Securities analysis [Madhavan, 1994].

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.

We have begun to change domains from politics to business. 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]

In extending VOTE to business domains, we have identified three key research areas.

- *Normative Business Goals and Beliefs.* If we want our programs to reason based on goals, we must first develop a vocabulary or taxonomy of those goals. Business decisions are predicated on a set of implicit goals and beliefs. A computer program must represent this knowledge explicitly. Therefore, we need an inventory of normative business goals and beliefs. We have developed a case-based reasoning tool for cataloging this type of business knowledge. This program may also be used for case analysis and teaching.

- *Qualitative Business Calculus.* Though business decisions rely on goals and beliefs, to a great extent such decisions are driven by numbers. Most existing business decision models are quantitative, applying rigorous analytical methods to numeric data. We suggest that qualitative methods can complement the traditional quantitative methods, by providing both an initial justification for applying a quantitative analysis, and a meaningful interpretation of the quantitative results. We have developed a qualitative business calculus that begins to bridge the gap between numbers and decisions.

- *Belief Representation: An Alternative to Truth.* The VOTE program used goals as the chief mediating element in decisions. However, business decisions require not only goals, but also knowledge of beliefs, which might be in conflict with goals. Just because an investor wants IBM shares to go up does not mean that it will happen. Artificial intelligence techniques for representing belief include binary logical values, i.e., true or false, and fractional probability or certainty factors. We propose another technique which is complementary to our existing goal representation, and in keeping with the qualitative nature of our decision model.

Below we present extended abstracts for each of these research efforts.

Normative Business Goals and Beliefs

A fundamental precept of case-based reasoning is the use of a rich vocabulary for indexing and retrieving relevant cases. Goals are often the most useful indices.

A typical business school case requires a student to read between the lines. For example, a case which focusses on a new order processing system will probably not state that the company is interested in cutting overhead, improving productivity, reducing errors, or decreasing the time required to process an order. These goals are axiomatic in business and implicit in the case.

In this paper we present a preliminary inventory of normative business goals, beliefs, and relationships. A normative goal would be for a company to increase market share. A normative belief would be that low tax rates are good. A normative relationship would be for a company to adopt the goals of its customers. By explicitly representing this knowledge, we can develop a descriptive vocabulary for indexing business cases which allows programs to reason about outcomes of business decisions.

Rather than derive the goals and beliefs top-down, starting at a root goal, such as MAXIMIZE PROFITS, we have adopted a bottom-up process of looking at specific cases. In deriving our inventory, we analyzed one week's worth of page 1 stories from *The Wall Street Journal*, using the Case Explorer indexing and retrieval program.

Case-Based Reasoning and Business Cases

For most of this century, leading business schools have been using the case method of teaching. The case method developed as a more realistic and practical alternative to lectures and textbooks. Case-based reasoning developed as a psychologically more realistic alternative to rule-based systems.

The obvious point of comparison between the case method and case-based reasoning is that they each focus on a real episode, rather than abstract principles or rules. The case method is based on the idea that students learn better from concrete cases than from abstract principles. Case-based reasoning asserts that learning cases is more natural and compelling than learning rules, for both computers and people.

The other major features of case method cases are consistent with CBR systems: agent perspective, specific problem, implicit goals, and the role of explanation. The real cases are more memorable and have a richer set of consequences and inferences than abstract principles or rules.

Given the considerable overlap in the fundamental nature of CBR and the case method, it seems natural to apply case-based reasoning techniques to the case method.

We can make business cases richer and more memorable by providing the student with tools for exploring a library of previous cases, making explicit the paradigm of case-based reasoning. We are developing a case-based reasoning tool, the Case Explorer, to be used by MBA students preparing business cases. This tool also provides a convenient method for obtaining an initial inventory of normative business goals and beliefs.

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 making connections with 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. This is a fundamental premise of case-based reasoning [Simpson, 1985; Hammond, 1986; Slade, 1991a].

The Case Explorer

The Case Explorer was first developed in HyperCard for the Apple Macintosh, and subsequently ported to

ToolBook for Windows. Our current work uses the ToolBook version. The Case Explorer is organized into five interrelated databases:

- *Case.* A case is a business episode. It may be a Harvard Business School case, or simply an article or capsule summary from a newspaper or magazine. Here is a sample case from *The Wall Street Journal.*

  Sega captured 63% of the most crucial segment of the U.S. market for video-game players during the holiday season, as Nintendo's slide accelerated, falling to a 37% market share.

- *Issue.* Issues are the normative business goals and beliefs. All other things being equal, what does an executive want to achieve? A typical issue is INCREASE MARKET SHARE. That issue would serve as a primary index for the Sega story above. In addition, the issue INCREASE MARKET SHARE would itself have an index to other justifying issues, such as GENERATE ADDITIONAL INCOME.

  Finally, other issues would justify themselves by pointing to INCREASE MARKET SHARE. These would include:

  - LOW COST PRODUCER - PRICE LEADER
  - PRODUCT DIFFERENTIATION
  - PRODUCT QUALITY
  - REDUCE CYCLE TIME
  - GLOBAL MARKETS
  - BUSINESS GROWTH
  - IMPROVE MARKETING EFFORT
  - OFFER A NEW PRODUCT OR SERVICE
  - RESPOND TO COMPETITOR
  - ADVERTISE PRODUCTS

  Each of these goals might be explained in terms of increased market share.

- *Relation.* In VOTE, relations served as a mechanism for goal adoption. That is, an agent who had a positive relationship with another agent would adopt the goals of the other agent. A pro-labor Congressman would adopt labor's issues. In business, there are generic relations, such as customer, employer, and competitor, as well as specific relations, such as between Sega and Nintendo. We may reason that while Sega and Nintendo will compete against each other, they might very well cooperate in establishing voluntary standards for video games to limit government regulation of the industry.

  The Case Explorer contains generic relations, such as customer, which can contain indices to specific issues, such as SIMPLIFY CUSTOMER ORDER PLACING and PROVIDE STATUS INFORMATION TO CUSTOMERS.

- *Industry.* The Case Explorer Industry database reflects a hierarchy or network of indices. Specific companies would have links to their related industries. For example, Sega would be linked to the VIDEO GAME industry which in turn is linked to CONSUMER ELECTRONICS, and so forth. Each of these entries would have associated issues or cases or relations or technologies.

- *Technology.* One initial application of Case Explorer was for information technology cases. For these cases, particular technologies, such as local area networks or client-server applications, were salient features. Moreover, many technologies could be justified in terms of specific business issues, such as REDUCING COSTS, REDUCING CYCLE TIME, REDUCE PROCESS RESPONSE TIME, or INCREASE PRODUCTIVITY.

Two exhibits are attached. Exhibit 1 depicts a ToolBook screen from the Case Explorer Issue database, showing the page for the issue INCREASE MARKET SHARE with its associated features. The **Issue** and **Features** menus provide the user access to the main system commands. The tabs at the bottom provide links to the other databases. Exhibit 2 is a screen image from Case Explorer's on-line help system, which uses the standard Windows Help program WINHELP.

Initial Results

We have begun testing the Case Explorer with MBA students, and will make it available for other researchers as well. We have analyzed one week's worth of page one business stories from *The Wall Street Journal* (WSJ) and developed an initial set of indices.

Issues

Our preliminary analysis indicates that the WSJ stories seem to present a higher level of issue from those found in traditional business school cases.

There appears to be a spectrum of concerns from the macroeconomic (IMPOSE TRADE SANCTIONS, DEFLATE DOLLAR, REDUCE DISCOUNT RATE) to the microeconomic (FORM STRATEGIC ALLIANCE, RENEGOTIATE DEBT, CONTROL BENEFIT EXPENSES).

We also observe opposite goals (INCREASE DIVIDEND and DECREASE DIVIDEND), which is consistent with our experience in VOTE. There we observed that members of Congress often had inconsistent sets of goals. On average, a member of Congress had ten conflicting stances. We expect to discover similar inconsistencies in the business world.

A fundamental assumption underlying our work is that business decision-making is a goal-driven process. In other words, managers evaluate and justify various courses of action on the basis of how well these alternatives align with the manager's goals. Thus, goals (i.e., issues) provide a critical index for business cases. Our assumption appears to be at least partially validated by the fact that we found numerous explicit references to goals in the WSJ articles.

Proper analysis and justification of business decisions requires not only that business issues be identified, but also that the instrumental relationships between goals be defined. In [Slade *et al.*, 1993], we developed a goal chain for use in evaluating decisions to invest in information technology (IT). This goal chain is patterned after the chain of causality which Porter [Porter, 1985] developed to explain the sequence of conditions and events that leads to a firm's financial success. Each link in the goal chain depicts how one goal is believed to be instrumental to achieving another. For example, managers may believe that an investment in office automation software will reduce paperwork, which provides a means of increasing productivity, which leads to reduced costs, and so on. By navigating through the goal chain, the specific consequences of an IT investment can be assessed within the context of a firm's overall set of business goals.

We found this taxonomic structure to be useful in indexing the WSJ articles as well. Thus, we integrated the 25 new issues discovered from an analysis of the WSJ articles into the chain of 130 issues that we had previously identified. For illustrative purposes, a portion of the goal chain is shown in Figure 1. A WSJ article on 3/10/94 indicated that "some of the biggest U.S. companies are joining forces in a well-financed push for new legislation to curb jury awards and discourage frivolous lawsuits." This is hardly a surprising move considering that seven lawsuits made the front page of the WSJ in the one week we investigated, with two of those lawsuits involving multi-billion dollar settlements. The decision to form a lawsuit lobby is shown at the bottom of Figure 1. The explicit goal of the corporations involved in this lobby is to limit their legal liabilities. Limiting legal liabilities is one means that firm's have to reduce their costs (a few others are also shown in Figure 1). Reducing costs is one of two ways in which firms can improve their financial position vis-a-vis rivals (successfully differentiating their products or services is the other way). An attractive relative position within an industry manifests itself in terms of higher profitability. Increased profitability, in turn, leads to increased shareholder value (i.e., higher stock price or increased dividends).

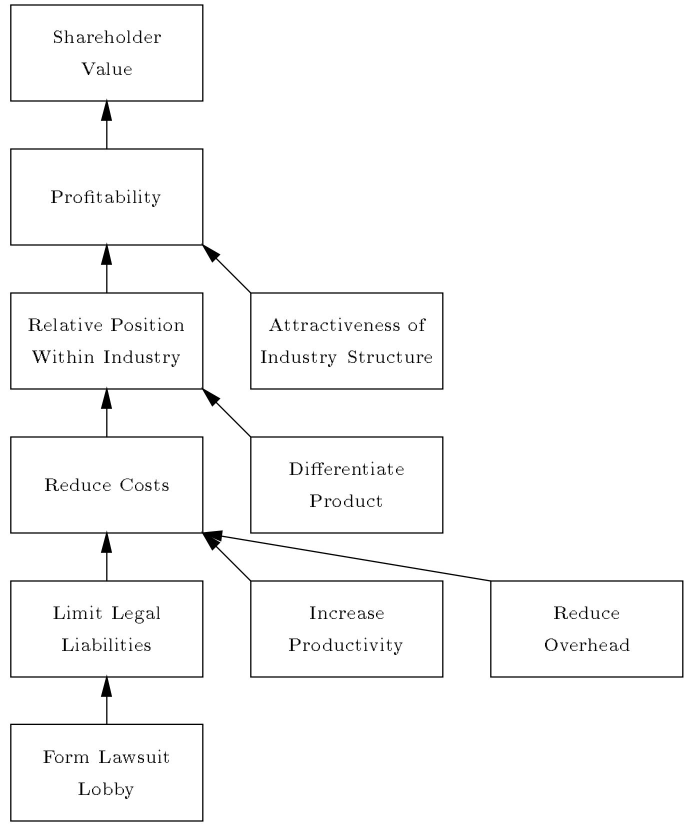

Figure 1: Goal Chain for Business Decisions

Industry

Both the analysis of the WSJ articles and the class exercise with MBA students revealed the importance of indexing cases by the companies involved, as well as mapping companies to industries. We found that companies in the same or similar industries are subject to similar competitive, technological, and regulatory conditions. Thus, cases involving these companies are likely to provide pertinent remindings for each other. For example, one of the problems in the class exercise called for the MBA students to consider the case of the planned joint cable venture between Southwestern Bell and Cox Enterprises. When the students indexed this case by the cable industry, the remindings mechanism found the earlier case of Bell Atlantic's planned acquisition of TCI. The cases were similar in that both deals were jeopardized by the FCC's recent decision to reduce cable rates, which reduced the attractiveness of investing in the cable industry.

We chose to use the Standard Industrial Classification scheme as the primary mapping of companies to industries. This scheme essentially maps compa-

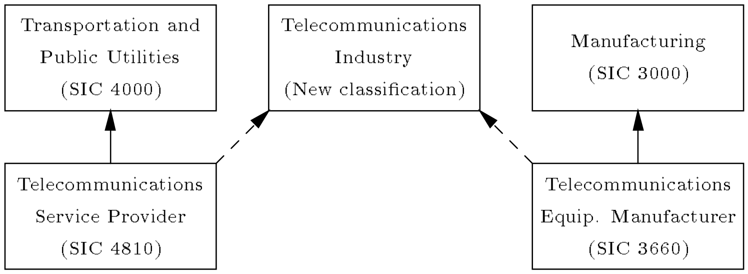

Figure 2: Industry Classifications

nies to divisions of the economy (e.g., manufacturing, construction, retail trade, etc.). We are investigating the possibility of including additional mapping schemes that are orthogonal to the SIC scheme.

Case

When analyzing the WSJ articles, we found that a single story line frequently spanned multiple cases (i.e., in the case of the WSJ articles, a case represents a single article appearing in a daily issue). For example, the U.S. - Japan trade summit was occurring during the week we analyzed WSJ articles. A report of trade summit proceedings appeared on the front page of the WSJ during each day of this week. To capture the temporal progression of a single story line over multiple cases, we indexed the cases by each other.

Future Work

We plan to expand the Case Explorer in several ways. First, we plan to take advantage of the structure of issues captured in the causal chain in order to provide more sophisticated remindings. Currently, the remindings mechanism in the Case Explorer performs simple issue matching to locate relevant cases. A more sophisticated remindings mechanism could navigate along the goal chain to locate potentially relevant cases with similar, but not necessarily identical issues.

Second, we plan to explore the possibility of including several orthogonal mappings of company to industry. The SIC scheme is useful for finding related cases in the same division of the economy. Industry groupings that complement the SIC scheme might also be useful. For example, telecommunications service providers and telecommunications equipment manufacturers are both affected by Judge Harold Greene's rulings, even though they are classified in different divisions of the economy according to the SIC scheme (public utilities and manufacturing, respectively). In this case, a *telecommunications industry* grouping would be useful (see Figure 2) for indexing purposes.

Qualitative Business Calculus

Goal-based reasoning in general and VOTE in particular provide a paradigm for reasoning about decisions based on goals and relationships. We have contrasted VOTE with traditional quantitative models such as decision analysis, pointing out that decision analysis often relies on knowledge of probability or payoff numbers that may not in fact be known. A qualitative, goal-based analysis may often be more realistic than the comparable quantitative analysis.

However, there are still many situations, particularly in business, in which it is not merely propitious, but advisable to take the numbers into account. Business decisions are full of quantities such as prices, rates, margins, shares, and volume. A robust business decision making system needs to be able to handle the numbers.

Rather than create a system which has hundreds of special rules for specific situations, we propose the development of a general qualitative business calculus to reason about business data. This effort reflects previous AI work in qualitative physics, which resulted in symbolic models of physical phenomena.

In this paper, we present our basic motivation and framework for the qualitative business calculus, together with examples. We discuss a computer implementation of the calculus which can perform a rudimentary financial analysis.

Quantitative versus Qualitative Models

Standard decision theory provides a quantitative approach to decision making [Raiffa, 1968]. Specific quantities in the form of payoffs and probabilities are used to arrive at a quantitative expected value. The decision maker simply selects the alternative that has the highest expected value. Some researchers have applied decision theory to AI problems [Hanks, 1990; Holtzman, 1989]. Sycara [Sycara, 1987] has combined case-based reasoning with a quantitative utility theory.

However, we may observe certain drawbacks to the quantitative approach to decision making through a comparison with AI approaches to the study of physics. The epitome of a quantitative science is physics, which is replete with precise equations for describing a wide range of physical phenomena such as motion, energy, and electricity.

One would expect that such a precise quantitative science would lend itself well to computational modeling, that is, to produce programs that reason about physical phenomena. However, it has turned out that it is not computationally feasible to create AI programs that *do* physics. Instead, AI researchers have developed qualitative theories for reasoning about physics [Forbus, 1985; de Kleer and Brown, 1985]. There are several motivations for pursuing a qualitative approach.

- It is often difficult to obtain the data required for modeling the necessary states of the world. For example, we may not know an object's precise mass or velocity or coefficient of friction.
- It is often computationally infeasible to calculate the answer. For example, even if we know the exact state of the world at time $T = 0$, we may not be able to compute the state for $T = 1$ within a reasonable amount of time due to the complexity of the calculations.
- A qualitative analysis of a problem is usually logically prior to a quantitative analysis. For example, if we let go of an object, we can be fairly sure that it will fall to the ground, even if we do not know how long it will take or what velocity it will achieve.
- A qualitative model can serve to prune the computation space of the quantitative approach. The qualitative analysis can eliminate certain computations and focus attention on others. In some cases, the qualitative approach may be sufficient.
- It is psychologically inappropriate to suggest that people reason about physics in a purely quantitative fashion. By proposing a mixture of qualitative and quantitative reasoning, we can arrive at a model that is both computationally feasible, and psychologically satisfying.

Thus, in the field of physics for which quantitative reasoning would seem well-suited, AI researchers have discovered compelling reasons for developing qualitative theories. We suggest that a similar argument holds for decision making.

- It is often difficult to obtain the data required for modeling the necessary states of the world. For example, we may not know an outcome's precise payoff or probability.
- It is often computationally infeasible to calculate the answer. For example, an accurate estimation of the behavior of a complex system, such as the stock market, requires thousands of probability and payoff estimates for each security at different points in time.
- A qualitative analysis of a decision problem is usually logically prior to a quantitative analysis. For example, if we learn of good news for a company, we expect its stock to go up in value. We may not know exactly how much the stock will rise or how soon.

Feature	Bank One	Bank Two	Bank Three
Assets ($Bil)	48.6	31.3	136.2
Ratio 1	13.4	10.8	13.3
Ratio 2	2.5	6.9	9.8
Ratio 3	1.6	1.2	1.8
Ratio 4	1.5	0.4	0.7

Table 1: Selected Bank Financial Data

- A qualitative model can serve to prune the computation space of the quantitative approach. As with physics, the qualitative analysis can be used to eliminate certain computations and focus attention on others. A qualitative analysis may even obviate a quantitative analysis.
- It is psychologically inappropriate to suggest that people reason about decisions in a purely quantitative fashion. Most decision theory avoids this problem by stating that the quantitative approach is prescriptive, rather than descriptive.

There are both theoretical and practical reasons for pursuing a qualitative model of business decision making. Our work with VOTE has demonstrated the computational feasibility of a qualitative decision model. However, we recognize that business decisions are often predicated on numerical information, such as prices, revenues, profits, tax rates, interest rates, exchange rates, and market share. It is necessary for a business decision making program to handle the numbers.

However, it is possible to have a *qualitative* analysis of quantitative data, as the work in qualitative physics has demonstrated. We are developing a qualitative business calculus which permits us to bridge the gap between business statistics and qualitative goals and beliefs.

Consider a simple decision comparing the quality of three bank holding companies. For each institution, we may have numerous statistics, such as depicted in table 1. The ratios are as follows.

- *Ratio 1:* (Equity + Reserves) / (Loans + Standby Letters of Credit).
- *Ratio 2:* (Non-performing Assets) / (Loans + Other Real Estate Owned)
- *Ratio 3:* (Net Charge Offs) / Loans
- *Ratio 4:* Return on Assets (using Income Before Security Transactions).

The question is: which bank is best? The answer is going to depend on several factors. First, we need

Operand 1	Operand 2	+	−	*	/
high	high	high	?	high	?
high	low	?	high	?	high
low	high	?	low	?	low
low	low	low	?	low	?

Table 2: Qualitative Arithmetic

to make sense of these ratios. Is it good to be high or low? Second, we need to consider the perspective of the decision maker. An investor, a depositor, a takeover specialist, and a regulator may each have a different view.

To analyze the ratios, we appeal to basic qualitative arithmetic, as depicted in table 2. The terms *high* and *low* refer to whether we desire the given quantity to be high or low. For example, in business, we generally want revenues to be high and overhead to be low. Given two high quantities, we prefer the larger. Given two low quantities, we choose the smaller.

We may use the information in the table to determine preferences for derived quantities. For example, profit, which is revenue (high) minus overhead (low), results in a high. This procedure may be applied recursively.

Other highs would include sales, interest earned, assets, equity, volume, market share, and market size. Lows would include costs, interest paid, taxes, tax rates, and bad debts.

In our bank data, ratio 3 is a low (Net Charge Offs) divided by a high (Loans), resulting in a low. By this measure, we prefer Bank Two's performance.

We note that half of the cells in table 2 are unspecified, as indicated by the question mark. For example, we do not know *a priori*, if a high divided by a high should be high or low. In the bank data, ratio 4, return on assets, is such a quantity: earnings (high) divided by assets (high). We can resolve the ambiguous ratio by ascribing a relative importance to the underlying quantities. That is, we can assert that earnings are more important than assets, resulting in a preference for a high return on assets.

Importance is the primary means of ranking goals in the VOTE decision making program [Slade, 1994]. By incorporating importance into our qualitative calculus, we can provide a principled connection with the decision making model. Using the relationships in table 2 together with the ordinal importance of the underlying quantities, we can perform qualitative assessments of quantitative data.

Furthermore, this approach provides a means to account for the subjective interpretation of the data by different decision making agents. For example, in the bank data, should *reserves* be a high or a low? An investor concerned about the amount of capital available for investment would prefer low reserves, whereas a regulator concerned about guaranteeing the safety and liquidity of the deposits would prefer reserves to be high.

These different perspectives can each be accomodated by our qualitative business calculus.

Here is a brief transcript from a program which implements the qualitative business calculus, and applies it to the analysis of bank data. We first input the data. Banc One, Bank of Boston, and Chemical Bank are banks 1, 2, and 3, respectively, from table 1.

```
(defvar banc-one-92
 (make-instance 'bank
  :bname "Banc One Corp." :bdate "1992"))
(set-bank-val
  banc-one-92 48.6 13.4 2.5 1.6 1.5)

(defvar banc-one-91
 (make-instance 'bank
  :bname "Banc One Corp." :bdate "1991"))
(set-bank-val
  banc-one-91 46.2 13.0 2.6 1.4 1.2)

(defvar bank-boston-92
 (make-instance 'bank
  :bname "Bank of Boston Corp." :bdate "1992"))
(set-bank-val
  bank-boston-92 31.3 10.8 6.9 1.2 0.4)

(defvar bank-boston-91
 (make-instance 'bank
  :bname "Bank of Boston Corp." :bdate "1991"))
(set-bank-val
  bank-boston-91 32.7 10.7 8.2 1.5 -0.2)

(defvar chemical-bank-92
 (make-instance 'bank
  :bname "Chemical Banking Corp." :bdate "1992"))
(set-bank-val
  chemical-bank-92 136.2 13.3 9.8 1.8 0.7)

(defvar chemical-bank-91
 (make-instance 'bank
  :bname "Chemical Banking Corp." :bdate "1991"))
(set-bank-val
  chemical-bank-91 138.9 11.2 9.6 4.2 0.1)
```

We now can compare bank performance from year to year from a particular perspective. In this case, we have a `regulator` and a `depositor`.

```
> (compare-banks
    regulator banc-one-92 banc-one-91)

Raw goals for Regulator are:
assets: High(A)  earnings: High(C)  NCO: Low(A)
   NPS: Low(B)      Loans: High(A) OREO: Low(C)
Equity: High(C)  Reserves: High(B) SLOC: High(C)
```

```
Inferring bank-data goals from raw goals
Inferred goals for Regulator are:
 assets: High(A) roa: Low(B) chargeoff: Low(A)
       nonperforming: Low(A)    eqloan: Low(C)

Comparing Banc One Corp.(1992)
      and Banc One Corp.(1991),
Applying Strategy: Simple Majority...done.

The winner is Banc One Corp.(1991)

> (compare-banks
    depositor banc-one-92 banc-one-91)

Raw goals for Depositor are:
assets: High(C)  earnings: High(A) NCO: Low(C)
   NPS: Low(C)      Loans: Low(C) OREO: Low(B)
Equity: High(A)  Reserves: Low(C) SLOC: High(B)

Inferring bank-data goals from raw goals
Inferred goals for Depositor are:
 assets: High(C) roa: High(B) chargeoff: High(D)
       nonperforming: Low(C)     eqloan: High(C)

Comparing Banc One Corp.(1992)
      and Banc One Corp.(1991),
Applying Strategy: Simple Majority...done.

The winner is Banc One Corp.(1992)
```

We can also compare different banks.

```
> (compare-banks
    depositor chemical-bank-91 banc-one-92)
Raw goals for Depositor are:
assets: High(C)  earnings: High(A) NCO: Low(C)
   NPS: Low(C)      Loans: Low(C) OREO: Low(B)
Equity: High(A)  Reserves: Low(C) SLOC: High(B)

Inferring bank-data goals from raw goals
Inferred goals for Depositor are:
 assets: High(C) roa: High(B)  chargeoff: High(D)
nonperforming: Low(C)   eqloan: High(C)

Comparing Chemical Banking Corp.(1991)
      and Banc One Corp.(1992),
please wait...
Applying Strategy: Simple Majority...done.

The winner is Banc One Corp.(1992)

> (compare-banks
    regulator chemical-bank-91 banc-one-91)
Raw goals for Regulator are:
assets: High(A) earnings: High(C)  NCO: Low(A)
   NPS: Low(B)     Loans: High(A) OREO: Low(C)
Equity: High(C) Reserves: High(B) SLOC: High(C)

Inferring bank-data goals from raw goals
Inferred goals for Regulator are:
 assets: High(A)  roa: Low(B) chargeoff: Low(A)
        nonperforming: Low(A)    eqloan: Low(C)

Comparing Chemical Banking Corp.(1991)
      and Banc One Corp.(1991),
please wait...
Applying Strategy: Simple Majority...done.

The winner is Chemical Banking Corp.(1991)
```

Belief Representation: An Alternative to Truth

The VOTE decision making model is driven primarily by goals. Modeling economic decisions seems to require an additional capability of reasoning about beliefs. For example, an investor may want interest rates to go down, but believes that rates will actually go up. A program that simulates economic decisions must be able to cope with the differences between goals and beliefs.

Philosophers and logicians have wrestled over the years with ways to represent knowledge and truth. AI researchers have developed mathematical techniques for handling certain and uncertain propositions. Abelson (1979) describes seven features that served to contrast belief and knowledge systems.

In this paper, we propose a unified representation of belief and knowledge that addresses some of the shortcomings of traditional logical or probabilistic representation systems, as well as the points raised by Abelson. The proposed belief representation is complementary to the goal representation used in VOTE.

Belief and Knowledge

Abelson [Abelson, 1979] proposes seven features for contrasting belief and knowledge systems. We illustrate each feature in a business context.

- *The elements of a belief system are not consensual.* A basic business transaction has both a buyer and a seller. These two agents should have differing beliefs about the value of the goods.

- *Belief systems are in part concerned with the existence or nonexistence of certain conceptual entities.* Abelson refers to religion, but investors seem to place faith in free enterprise or the free market.

- *Belief systems often include representations of "alternative worlds."* Economists often postulate idealized situations in which agents possess perfect information and there are no transaction costs.

- *Belief systems rely heavily on evaluative and affective components.* Some things are good and some things are bad. Low interest rates are good. High tax rates are bad. At least that is what most people seem to believe.

- *Belief systems are likely to include a substantial amount of episodic material.* The crash of 1987 or the arrest of Michael Milkin are likely to affect investors' beliefs about the stability of the stock market or inside trading.

- *The content set to be included in a belief system is usually highly "open."* Abelson points out that beliefs are subject to unconstrained inference. Thus, a belief that high tax rates are bad would have to weigh the personal benefits gained with the problems associated with lower government revenues. The latter could lead to higher federal deficits, lower spending, increased crime, and so forth.

- *Beliefs can be held with varying degrees of certitude.* The passion or conviction accorded a belief can vary. An investor may feel more certain that tax rates will remain stable, than interest rates are going up, or vice versa.

Beliefs are the building blocks of decisions. The VOTE program used goals as the controlling mechanism for decision making, however, those goals were implicitly the result of beliefs. The role of beliefs can be made explicit by stepping back from decision making and looking at the complementary task of persuasion.

Persuasion

Schank and Abelson (1977) discuss interpersonal planning strategies, such as having someone perform a service for you, or provide you with information, or control of an object, or the authority to perform some act. Generally viewed, these are instances of someone devoting resources to achieve an adopted goal. Schank and Abelson offer a set of specific plans (termed "the persuade package") for getting someone else to act on your behalf.

We may view persuasion as the task of one agent convincing another agent to make a decision. Persuasion may be considered as decision making once removed. In business, *sales* is an example of persuasion. The agent tries to convince the client to buy a product or service.

Goal adoption is one part of the process of persuasion. Below we present several persuasion methods that are consistent with our underlying model of decision making.

- *Establish or increase the importance of the relationship.* As the importance of the relationship increases, so will the relative importance of adopted goals. The salesman may take the client out to dinner.

- *Emphasize the positive consequences of the desired action, or the negative consequences of failure.* The salesman may provide the client with reasons for buying a product.

- *Establish or increase the importance of preferences that match the consequences of the desired action.* The salesman may argue that other important colleagues of the client share his values and concerns, or that the client's previous decisions are consistent with the salesman's point of view.

- *Provide a suitable explanation to justify the desired action or to counter the alternative position.* The explanation must address the concerns of groups adversely affected by the choice. We have identified a number of explanation strategies for the political domain [Slade, 1991d]. Often VOTE can provide justifications for both sides of a given bill.

We note that these methods of persuasion assume a world of incomplete or inexact knowledge. That is, no agent has complete knowledge of all possible outcomes. Persuasion is a process by which one agent selectively alters the decision maker's beliefs to influence a decision or action.

We propose a system in which beliefs are represented similarly to goal stances in VOTE. That is, a belief may be held to be PRO or CON (true or false) at some ordinal level of certainty, e.g., A, B, C, or D. Beliefs may lead to other beliefs or to goals. In this system, knowledge is simply belief held with great conviction, i.e., the A level.

The following rules for ascribing certainty illustrate some of the types of inferences about beliefs.

- **A-level inferences.**

  Perceptions. First-hand experience is preferred to second-hand accounts. Due diligence often requires personal visits.

  Past vs. Future. Events that have transpired are more reliable than events in the future. We know yesterday's stock prices. We can only predict tomorrow's.

- **B-level inferences.**

  Reliable sources. Prefer second-hand accounts from trustworthy sources (*The New York Times*, a priest) over unreliable sources (*The National Enquirer*, a felon).

  Causal Consistency. Prefer beliefs that are causally consistent and do not conflict with other beliefs. We may believe that the Fed controls interest rates, but not that Elvis controls interest rates.

  Volitional Consistency. When a belief reflects an agent's volitional action, we prefer a belief which is consistent with an agent's goals. We question a story about a software company that gives away its product for free, until we learn that their goal is to increase market share for future upgrades and related products.

- **C-level inferences.**

  Skepticism. Question stories whose source stands to gain from other people's acceptance. A stock broker who denies his guilt of insider trader is less credible than one who admits his guilt.

We now consider a specific story from the March 9, 1994 *Wall Street Journal* "Power Play: New Computer Chip Hits Desktop Market With Intel in Its Sights."

> *Next Monday, Apple Computer Inc. will fire the first salvo in a multisided battle for the heart, soul and profits of the burgeoning personal-computer industry.*
>
> *Apple will introduce a powerful line of Macintosh machines built around a new kind of microprocessor (the brains of a computer) called the PowerPC. International Business Machines Corp. has allied with Apple, and will introduce its own PowerPC line this summer. The world's two biggest personal-computer makers and another ally, chipmaker Motorola Inc., have spent billions of dollars developing the PowerPC or enhanced operating systems for it, all with the aim of breaking up the lucrative hegemony of Intel Corp. and Microsoft Corp.*

This story presents facts, opinions, and predictions. The basic question is: will the PowerPC be successful? Related questions include: how will the PowerPC affect Apple, IBM, Motorola, Intel, and Microsoft? Should a company or individual purchase a PowerPC computer? Here are some quotes from the article.

- *Robert Corrigan [president of the IBM Personal Computer Co.] has doubts about the PowerPC. He believes Intel will keep churning out chips that run the vast universe of PC software "better than everyone else." PowerPC "isn't going to easily usurp that space."*

- *Robert W. Stearns [of Compaq Computer] says the PowerPC allies "are smoking dope. There's no way it's going to work."*

- *The most likely near-term effect is to be an outbreak of price wars among the makers of very powerful microprocessors and personal computers.*

- *The biggest and only sure winner is the computer user.*

- *"I welcome the price wars," says Ian W. Diery [of Apple], who vows to keep Macintosh PowerPCs "more than $200" below computers running on Intel's most powerful line of microprocessors.*

- *Officials at Apple say the PowerPC finally gives it room to consistently offer prices for Macs below those for comparable Intel-based machines.*

- *On Monday, Intel introduced new versions of the Pentium that are 50% faster than the first generation, and slightly faster than the PowerPC line.*

- *Andrew Grove, Intel's feisty CEO, says he is even grateful to the PowerPC alliance for making his company more aggressive. "Things like PowerPC don't let you sit on your" rear end, he says.*

Apple, Intel, and the other companies involved are engaged in persuasion. The process of digesting the various statements and assertions should not result in simple truth values or even probabilities, such as "It is 45% certain that the PowerPC will help Apple increase market share."

The result should be a complex knowledge structure linking together the statements, goals, and relationships of the various players. Using our rules for A, B, and C-level inferences, we can rank the credibility of the statements, though we will likely fail to arrive at a simple, consistent conclusion.

We maintain that this knowledge structure has far more information than can be conveyed by a simple probability. Furthermore, this knowledge structure also provides information about the goals of the agents, and such knowledge is useful on its own.

We need to represent beliefs in sufficient complexity to capture the meaningful information about goals and certainty, while maintaining sufficient simplicity to guarantee computational tractability.

Conclusion

We have briefly described three research projects: an inventory of goals and beliefs, a qualitative calculus for reasoning about business statistics, and a system for representing beliefs.

These projects are not unrelated. Reasoning about beliefs can lead to goal development. Understanding the imporance of standard business goals can resolve ambiguities in the interpretation of business data.

Our unifying objective is the simulation of business decisions. We believe that these three projects lie on the critical path.

References

Abelson, R.P. 1979. Differences between belief and knowledge systems. *Cognitive Science* 3(4):355–366.

Kleer, J.de and Brown, J. S. 1985. A qualitative physics based on confluences. In *Qualitative Reasoning about Physical Systems.* MIT Press, Cambridge, MA. 7–83.

Forbus, K. 1985. Qualitative process theory. In *Qualitative Reasoning About Physical Systems.* MIT Press, Cambridge, MA. 85–168.

Hammond, K.J. 1986. *Case-Based Planning: An Integrated Theory of Planning, Learning and Memory.* Ph.D. Dissertation, Yale University. Technical Report 488.

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

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

Madhavan, Raghav K. 1994. Goal based reasoning for securities analysis. *AI Expert* 9(2):22–29.

Porter, M. 1985. *Competitive Advantage.* Free Press, New York.

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

Schank, R.C. and Abelson, R. 1977. *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum Associates, Hillsdale, NJ.

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.

Simpson, R.L. 1985. *A Computer Model of Case-Based Reasoning in Problem-solving: An Investigation in the Domain of Dispute Mediation.* Ph.D. Dissertation, School of Information and Computer Science, Georgia Institute of Technology.

Slade, S.; Lucas, H.; and Fish, M. 1993. Qualitative decision explanation for information technology investment. Working Paper STERN IS-93-18, Information Systems Department, Stern School of Business, New York University, New York.

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. B. 1991a. Case-based reasoning: A research paradigm. *AI Magazine* 12(1):42–55.

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. 1994. *Goal-based Decision Making: An Interpersonal Model.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Sycara, Ekaterini P. 1987. *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods.* Ph.D. Dissertation, College of Computing, Georgia Institute of Technology.