Case-based Reasoning for Financial Decision Making

Stephen Slade
Information Systems Department
Stern School of Business
New York University
40 West 4th Street
New York, NY 10003
sslade@stern.nyu.edu

Abstract

There is a qualitative component to decision making that is complementary to the more quantitative approaches usually adopted in the financial world. Case-based reasoning provides a computational paradigm for simulating the role of experience and intentions in decision making. We describe a program that provides a framework for modeling financial decision making.

1 Introduction

Many financial problems may benefit from reasoning based on past experience. Case-based reasoning (CBR) provides a principled paradigm for the creation of computer programs that can solve new problems based on prior episodes. In contrast with traditional rule-based expert systems, CBR programs provide a psychologically coherent approach to simulating expert problem solving or decision making behavior. CBR systems provide constructive approaches for knowledge acquisition, learning from experience, and robust performance.

In the financial world, there are enormous amounts of data and clear measures of the success or failure of investment decisions. The challenge is to develop systems that can meaningfully assimilate vast quantities of information automatically and provide timely, accurate financial advice, produce pertinent explanations, and modify future behavior based on feedback from current performance. We have previously termed such programs *advisory systems* [8].

In the present paper, we shall discuss how advisory or CBR systems might be applied to financial decision making. First, we suggest the use of qualitative decision models as a computationally feasible complement to the traditional quantitative approaches to financial decision making. Second, we contrast the CBR approach to standard rule-based expert systems. Third, we propose an analysis based on the goals and relationships among the various economic agents in the financial community.

2 Qualitative Decision Models

Much of modern financial theory, such as the capital asset pricing model and arbitrage pricing theory, is based on quantitative, statistical assumptions. In [13], we have argued that computational models of decision making need not be purely quantitative, but can benefit from the addition of qualitative reasoning ability, such as afforded by CBR.

For example, standard decision theory provides a quantitative approach to decision making [6]. 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.

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 physics [1, 2]. There are several motivations for pursuing a qualitative approach.

TH0399-6/91/0000/0232$01.00 © 1991 IEEE

- 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 expect a stock to go up in value, we are likely to invest in the stock. We may not know exactly how much the stock will rise or how soon.

- 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.

In the next section, we discuss some general principles underlying one computational approach for qualitative decision making: case-based reasoning.

3 Case-based Reasoning

A freshly minted MBA knows the rules of finance. She knows about alphas and betas, CAPM and APT, P/E's and Black-Scholes, purchasing power parity and exchange rates. At the same time, a managing director of an investment house knows all this, and more. That "more" is experience. The managing director has witnessed hundreds or thousands of financial episodes. Those episodes have been assimilated, and provide a rich context for recognizing and analyzing new problems and opportunities.

By way of comparison, we may consider the new MBA's textbook knowledge to be analogous to a traditional rule-based system, and the experienced managing director to be a CBR system.

In [10], we have presented a set of three basic problems with traditional rule-based expert systems. The first problem was *knowledge acquisition*. In order to build an expert system, a computer programmer (or knowledge engineer) had to sit down with the human expert informant to determine what rules were appropriate for the given domain. This knowledge was difficult to uncover. The human expert could not simply make a list of the hundreds of rules he used to solve problems. Often the informant would articulate a set of rules that in fact would not accurately reflect his own problem solving behavior. For these reasons, this difficult knowledge acquisition process became known as a bottleneck in constructing rule-based expert systems [3].

Second, the rule-based systems did not have a *memory*. For example, if a medical diagnosis program is presented with a patient with a certain set of symptoms, the program may fire dozens or hundreds or thousands of rules and come up with a diagnosis or treatment. Subsequently, if the program is presented with another patient displaying the same set of symptoms, the program will fire the same set of rules as before. The program will not *remember* having previously seen a similar patient. One might argue that this observation is of little consequence beyond some argument for computational efficiency. However, efficiency can be a significant concern in many situations. Moreover, a program without a memory will not remember its mistakes, and thus, will be destined to repeat them. Thus, both accuracy and efficiency are related problems for a system without a memory.

Third, rule-based systems were not *robust*. If a problem were presented to the system that did not match any of the rules, the program could not respond. The system's knowledge base was limited to its rules, so if none of the rules could apply, the system had no alternatives. It was brittle.

We may compare the behavior of the rule-based expert system with the behavior of the human expert. The central feature of *expertise* is *experience*. An expert is someone who has vast, specialized experience, who has witnessed numerous cases in the domain, and who has generalized this experience to apply it to new situations. When confronted with a problem, the expert is reminded of previous, similar problems and their respective resolutions. It may be that the expert has so many exemplars for a given problem that the experiences have been distilled into a general rule to be applied. Still, that general rule has its roots in actual experience.

Thus, the human expert derives knowledge from experience. The basic unit of knowledge is not the *rule*, but the *case*. Human experts acquire knowledge by assimilating new cases, either first-hand or through the reports of others. Furthermore, it is easier for people to articulate knowledge in the form of experience than as rules. This observation suggests the psychological hypothesis that expert knowledge may in fact be encoded primarily as episodes, rather than as rules. We contrast this acquisition of knowledge from experience with the knowledge acquisition bottleneck given above as the first problem of rule-based systems.

Second, human experts remember their own experience. The doctor who fails to treat a case effectively should remember that case when another patient presents the same symptoms. The doctor can learn from his mistakes.

Third, human experts can reason by analogy. If our doctor sees a patient who presents symptoms that are unlike anything in his experience, the doctor need not simply give up. The doctor might be reminded of various previous cases that were similar in one way or another, and devise a treatment accordingly.

These arguments suggest an alternative to the rule-based system: a case-based system. An expert system that can extract information from its experience will be able to grow and acquire knowledge on its own.

The technology of case-based systems directly addresses problems found in rule-based systems.

- *Knowledge acquisition.* The unit of knowledge is the case, not the rule. It is easier to articulate, examine, and evaluate cases than rules.
- *Performance experience.* A case-based system can remember its own performance and modify its future behavior to avoid repeating prior mistakes.
- *Adaptive solutions.* By reasoning from analogy with past cases, a case-based system should be able to construct solutions to novel problems.

In the following section we present a specific framework for developing CBR programs in financial domains.

4 Goals, Relationships, and Explanations

Two key questions to be addressed in a qualitative decision making system are:

- *Choice:* What does the agent want to do?
- *Explanation:* How can the agent justify the decision?

Financial decisions are based on the intentions and expectations of the economic agents. In a given situation, an agent will have numerous goals and limited resources. Some goals are more important than others. Goals are often in conflict, and thus decisions require trade-offs among conflicting goals. Answering the question of choice requires an understanding of these goals.

For example, in order to maximize profits, a businessman may have a multitude of goals for his company, such as increasing market share and revenues and stock price, while decreasing overhead and taxes and cost of capital. In general, he will not be able to achieve all these goals at once. Furthermore, there are other goals or constraints on the businessman, such as decreasing the market share of the competition and complying with appropriate regulations and statutes. An advisory system must recognize the relevant set of economic goals, their relative priorities, and the associations among the different goals.

In addition, economic agents do not act unilaterally. Agents have a multitude of relationships with other agents. Relationships have priorities, and goals are adopted from relationships. These adopted goals impinge on decisions and may result in conflicts and trade-offs [9].

For example, our businessman may have relationships with numerous interested parties, including stockholders, employees, unions, suppliers, distributors, customers, competitors, bankers, the neighboring community, and regulators. Some of these relationships are friendly, while others are adversarial. Each of these groups has its own set of goals that may impinge upon the decisions taken by the businessman. An advisory system must understand these relationships.

In making a decision, an economic agent must be cognizant of the consequences associated with his actions. The prediction and analysis of outcomes and expectations is consistent with the CBR paradigm. For example, in contemplating an action, a businessman will be reminded of similar previous actions. Our advisory system should contain an experiential knowledge base to permit the program to infer the consequences of its intended actions.

In making decisions, the businessman must recognize conflicts among his own goals and the goals of the groups with whom he has positive relationships, and address those conflicts in justifying or explaining the decision. The businessman is accountable for his actions. An automatic advisory system must also be accountable. It must be able not only to offer advice, but also to justify its position by providing an explanation of its reasoning.

We note that this type of explanation is different from the usual sense of explanation found in the case-based reasoning literature [7, 4, 5]. Previous researchers have focused on explanation of anomalous observed events as part of the process of learning. Our present use of explanation is complementary to that process: decision makers offer explanations for the benefit of observers who may find the decision to be anomalous.

We can summarize the contents of our decision making framework as follows.

- Explicit representation of goals of economic agents.
- Explicit representation of relationships among agents.
- Case-based analysis of consequences of actions.
- Automatic processing of decisions by matching preferences with consequences and detecting conflicts.
- English explanation of resulting advice or decision.

5 Computer Programs

We have implemented this basic framework in a program which processes thousands of decisions in the domain of Congressional roll call voting, and provides English justifications for its decisions [12, 11]. We are currently adapting this program to financial decision making.

In the VOTE program, goals are political issues, such as gun control, affirmative action, or the death penalty. Members of Congress have stances on various issues, as well as relationships with constituency groups, such as organized labor or the National Rifle Association. In arriving at a decision, VOTE must recognize conflicts among disparate stances and generate an explanation. Here is the output from VOTE simulating the decision of the republican Representative Newt Gingrich voting on a clean water bill.

> *Newt Gingrich votes for bill HR-8, the Water Quality Renewal Act. He believes this bill to be in the best interests of the people. He feels strongly in favor of limiting federal regulation of industry and society. Gingrich is committed to free enterprise and capitalism.*

Our first step has been to apply this same framework directly to financial decisions. We have started with a few dozen financial issues including cash flow, debt, dividends, earnings, growth, interest rates, liquidity, market share, price/earnings ratio, regulations, risk, and taxes. For example, the issue of debt is is represented as follows.

```
Debt
  Norm:
  ((CON C ISSUE:ISSUE.847 DEBT))
   PRO Stances:
  (((PRO B ISSUE:DEBT LEVERAGE)))
   CON Stances:
  (((CON B ISSUE:DEBT LEVERAGE)
    (PRO B ISSUE:DEBT EQUITY)))
```

Popular opinion is opposed to increased debt. Support of increased debt is important for increased leverage. Opposition to increased debt opposes leverage. It upholds increased equity.

Norms are meant to reflect the conventional wisdom on a given issue. The letter "C" indicates the strength or importance accorded that view on a scale where "A" is very important and "D" is unimportant. Here, PRO stances are reasons in support of increasing debt, and CON stances are reasons to decrease debt. The program expresses these symbolic representations in English through an automatic natural language generation facility.

Economic groups include individual investors (small shareholders), institutional investors (such as pension funds), employees, rating agencies (such as S&P), specific industries (such as utilities or banks), and regulatory agencies (such as the SEC). Each group has its own set of stances on various economic issues. For example, here is the program's representation of an individual investor.

```
Individual Investor
  Norm:
  ((PRO B GROUP:GROUP.235 INDIVIDUAL))
   Stances:
  (((PRO B GROUP:INDIVIDUAL ROI)
    (CON B GROUP:INDIVIDUAL RISK)
    (PRO B GROUP:INDIVIDUAL STOCK-PRICE)
    (CON B GROUP:INDIVIDUAL TAXES)
    (PRO B GROUP:INDIVIDUAL DIVIDENDS)
    (CON B GROUP:INDIVIDUAL INSIDER-TRADING)
    (PRO B GROUP:INDIVIDUAL LIQUIDITY)
    (PRO B GROUP:INDIVIDUAL PROFITS)))
```

Most stockholders strongly support the individual investor. The individual investor is an opponent of increased risk, taxes, and moreover increased insider trading. He believes in increased liquidity, profits, dividends, increased stock price, in addition to increased shareholder return on investment.

We note that the language generation program was designed for the political domain, and will benefit from future refinements to adapt to financial concepts.

The third component of the program is a description of a projected business event – either intentional or circumstantial. Intentional events would include plans for downsizing a company, stock buy-backs, mergers, acquisitions, and public offerings. Circumstantial events would include changes in inflation, interest rates, unemployment, foreign exchange, GNP, or other market conditions. For example, the program represents company downsizing as follows.

```
downsizing
  Stance-FOR:
 (((CON B PLAN:DOWNSIZE EXPENSES)
   (PRO C PLAN:DOWNSIZE PROFITS)
   (PRO B PLAN:DOWNSIZE CASH-FLOW)
   (PRO C PLAN:DOWNSIZE PRODUCTIVITY)
   (PRO B PLAN:DOWNSIZE ROI)
   (PRO B PLAN:DOWNSIZE COMPETITIVENESS)
   (CON B PLAN:DOWNSIZE BUREAUCRACY)
   (PRO B PLAN:DOWNSIZE DECISION-MAKING)
   (PRO C PLAN:DOWNSIZE CUSTOMER-NEEDS)
   (PRO C PLAN:DOWNSIZE SALES)
   (PRO C PLAN:DOWNSIZE MARKET-SHARE)
   (PRO C PLAN:DOWNSIZE PRODUCT-QUALITY)
   (PRO C PLAN:DOWNSIZE TECHNOLOGY)
   (PRO C PLAN:DOWNSIZE INNOVATION)
   (CON C PLAN:DOWNSIZE TAKEOVER)))
  Stance-AGN:
 (((PRO B PLAN:DOWNSIZE JOB-SECURITY)
   (PRO B PLAN:DOWNSIZE EMPLOYEE-NEEDS)
   (PRO B PLAN:DOWNSIZE GROWTH)))
```

Support of the corporate downsizing strategy stands firmly against increased expenses and bureaucracy. It is important for increased shareholder return on investment, competitive advantage, quality of decision making, as well as increased cash flow. It is in opposition to takeover attempts. Support of the corporate downsizing strategy is compatible with increased profits, productivity, customer satisfaction, sales, increased market share, product quality, technological advances, and moreover increased innovation. Opposition to downsizing reinforces increased growth, job security, and moreover increased employee satisfaction.

We can now ask the program to evaluate the downsizing strategy from a particular point of view. We first look at a particular stockholder, Mary Jones.

Mary Jones favors the corporate downsizing strategy. She believes that the majority of people support this measure. She feels strongly in favor of increased shareholder return on investment. Jones is committed to increased profits. Still, she recognizes that Jones is committed to increased growth.

The program recognizes that downsizing has a downside, namely reduced growth. The same downsizing plan gets a different reaction from an employee, Joe Smith.

> *Joe Smith is opposed to the corporate downsizing strategy. He believes this plan not to be in the best interests of the company. He is a defender of increased employee satisfaction. Smith strongly supports job security.*

Each economic agent has a specific interpretation of the plan. This analysis is qualitative. It reveals specific quantitative questions to ask, such as, what is the projected net change in profits from downsizing? and how many jobs will be lost?

6 Future Work

Our program is a work in progress. To apply our decision making framework to financial domains, we require access to real world data. In developing this program, we are interested in establishing cooperative relationships with members of the financial community who can provide case data and expert guidance for this domain. There are two primary types of episodes which we are interested in compiling. The first is a case library of business episodes, such as LBO's, downsizing, and buy-backs. The second library is one of individual financial plans that address common investment goals such as retirement, college tuition, or trusts.

Given these case libraries, the goal of the program is to provide sensible interpretations of the likely consequences of financial events. Case-based reasoning offers a computational paradigm for providing financial advice and justifications tailored to the particular goals of a range of investors.

References

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

[2] K. Forbus. Qualitative process theory. In *Qualitative Reasoning about Physical Systems*, pages 85–168. MIT Press, Cambridge, MA, 1985.

[3] F. Hayes-Roth, D.A. Waterman, and D.B. Lenat (eds.). *Building Expert Systems.* Addison-Wesley, Reading, Mass., 1983.

[4] A. Kass. *Developing Creative Hypotheses by Adapting Explanations.* PhD thesis, Yale University, 1990.

[5] D.B. Leake. *Evaluating Explanations.* PhD thesis, Yale University, May 1990. YALEU/CSD/RR 769.

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

[7] R.C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively.* Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.

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

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

[10] S.B. Slade. Case-based reasoning: A research paradigm. *AI Magazine*, 12(1):42–55, Spring 1991.

[11] S.B. Slade. Goal-based decision strategies. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, Chicago, IL, August 1991. Cognitive Science Society.

[12] S.B. Slade. *An Interpersonal Model of Goal-based Decision Making.* PhD thesis, Yale University, 1991.

[13] S.B. Slade. Qualitative decision theory. In *Proceedings of DARPA Workshop on Case-Based Reasoning*, Washington, DC, May 1991. Defense Advanced Research Projects Agency, Information Science and Technology Office.