

Qualitative Reasoning for Financial Assessments: A Prospectus

Peter E. Hart, Amos Barzilay, & Richard O. Duda

Syntelligence, 1000 Hamlin Court, Sunnyvale, California 94088

Most high-performance expert systems rely primarily on an ability to represent surface knowledge about associations between observable evidence or data, on the one hand, and hypotheses or classifications of interest, on the other. Although the present generation of practical systems shows that this architectural style can be pushed quite far, the limitations of current systems motivate a search for representations that would allow expert systems to move beyond the prevalent "symptom-disease" style. One approach that appears promising is to couple a rule-based or associational system module with some other computational model of the phenomenon or domain of interest. According to this approach, the domain knowledge captured in the second model would be selected to complement the associational knowledge represented in the first module.

Simulation models have been especially attractive choices for the complementary representation because of the causal relations embedded in them (Brown & Burton, 1975; Cuena, 1983). Most investigators interested in this subject appear to have selected, as research vehicles, domains involving the diagnosis or assessment of biological or physical systems. Weiss *et al.*, (1978); Patil *et al.*, (1981); and Pople, (1982) provide examples of incorporating causal models within expert systems applied to medical diagnosis. A recent example of this approach applied to a mechanical diagnosis problem is given by Fink *et al.*, (1984). In this article we ask how these ideas can be extended to allow expert systems to improve their ability to assess human organizations—in particular, their ability to assess the future prospects of a business corporation. We introduce the problem of forming assessments about cor-

porations, describe the reasoning styles currently used by people, and show how some of these assessments can be addressed by extending existing AI techniques. A sound approach to the most significant challenge—incorporating qualitative causal models in an expert system—remains a speculative subject.

The Corporate Assessment Domain

There are more than 5000 publicly listed business corporations in the United States and perhaps 100 times that number of private corporations. The fates of each of these firms is a matter of great interest to the business and investment community. The larger firms are subject to intense scrutiny by armies of financial analysts, and even the smaller corporations have creditors of various sorts who hope that the firms are financially sound. The details of the procedures used to make assessments vary according to the specific objective of the analyst. It might be that an equity investment is under consideration, that a loan request has been made, that a merger is being contemplated, and so forth.

Abstract

Historically, the evolution of expert systems has been driven by scientifically based fields such as medicine, geology, and computer engineering. More recently, expert system developers have turned their attention to the highly judgmental decision tasks found in business and finance. We introduce the *corporate assessment* problem, point out the limitations of current expert system approaches to the solution to this problem, and suggest that a more fundamental approach based on recent work in qualitative physics might be fruitful.

We thank John Seely Brown, Johan de Kleer, and Magnus Ljungberg for their insightful comments.

Regardless of the details, however, the common elements are an analysis of the firm's past performance; an analysis of the industry in which the firm operates; and, based on these and other considerations, a projection of the firm's likely performance over some future period of time.

A person analyzing a corporation relies on several qualitatively different classes of evidence to represent the history of the firm and its industry. Each of these classes has associated with it certain characteristic styles of reasoning that would appear desirable in an expert system. Unfortunately, it is not clear that all of these styles can be represented using current methodologies. Moreover, there is a close coupling among them, which raises difficult issues of how to exploit complementary perspectives and how to maintain consistency among them.

Conventional Financial Statements

The time-honored means for representing the condition of a corporation is with a set of financial statements. Because these statements are usually the starting point for any analysis of a firm's prospects, we digress briefly to describe them.

Because "financials" have been used so long, and because financial reporting is subject to strict regulation, financial statements have a highly constrained canonical form. They consist of two main parts: the *balance sheet* and the *income statement*. (A third part of financial statements, the statement of changes in financial position, is derivative and will not be discussed here.)

The balance sheet is a financial snapshot of the corporation at a given point in time. One side of the balance sheet, *assets*, represents everything the corporation has of value, tabulated under a number of mutually exclusive and exhaustive categories. Assets can range from cash in the bank to buildings and equipment owned to goods in the warehouse ready for sale. The other side of the balance sheet represents claims on the corporation's assets. These claims fall into two main categories: The first, *liabilities*, is money that the firm has a legal obligation to repay to outsiders. Liabilities include such items as unpaid bills from suppliers and long-term mortgages on property. The second claim on the firm's assets, *equity*, represents the ownership interest of the shareholders. The shareholders' equity is simply whatever remains after all outstanding obligations are subtracted from the assets.

The income statement, rather than being a snapshot, expresses the flow of funds resulting from the company's operations over a given period of time, such as a year. It is a record, under standard categories, of the payments received by the corporation in return for its goods or services and of the costs incurred by the corporation in supplying these goods or services. The difference between the payments and the costs represents the profit or loss resulting from the firm's operation during the year.

A variety of standard techniques are used by analysts to interpret financial statements. For example, most of us would be concerned if we were personally burdened by a level of debt that is high in comparison to our income. Analysts are similarly concerned if a corporation has a high debt load, where "high" might be relative to the amount of equity interest owned by the shareholders. Accordingly, the debt-to-equity ratio, often called the *leverage*, is computed from the balance sheet as one measure of a firm's financial health at a given point in time.

Analysts typically compute one or two dozen other numerical ratios in an effort to understand the current condition of a corporation and, therefore, to be in a better position to assess the corporation's prospects. Some ratios, like leverage, are balance-sheet ratios. Others, like *net margin* (profits as a percentage of sales), are computed from the income statement. These ratios might be examined over a period of recent history, might be compared with the ratios for other companies in the same industry, and so forth. An expert system for assessing corporate prospects will undoubtedly have to make use of similar numerical procedures. However, as the following sections illustrate, standard analyses of financial ratios by no means tell the whole story.

Associational Knowledge

The number and diversity of factors that can affect the prospects of a corporation are very large. Almost any event reported in the daily newspaper can have an economic effect on some enterprise in the short run. In the longer term, underlying trends, which can develop over a period of time before they are clearly recognized, can affect whole industries. Because an exhaustive analysis of all theoretically possible events is out of the question, analysts have developed a usable body of associational knowledge to apply to individual cases.

For specificity, imagine a banker considering the request of an established corporation for a loan. A dominant thought running through the banker's mind is, "Can anything go wrong with this company?" This global question can, conceptually at least, be partitioned into questions about various phases of the company's operations and environment. If, for example, the principal product of the company is widgets, the banker would want to consider factors that could affect the demand for widgets, the company's ability to manufacture widgets, the competitive threat posed by other widget makers, and so forth.

Knowledge required to perform an analysis at this level can often be effectively captured in an associational representation, *i.e.*, a representation linking case data more or less directly with the confirmation or reputation of a hypothesis. For example, suppose that at some point in the analysis we were examining the ability of the firm to continue manufacturing widgets. Because any manufacturing process brings together materials, people, and processing

equipment, one would want to consider the likelihood that any of these might in the future prove to be a source of weakness, *e.g.*, through shortages of raw materials, labor unrest, or equipment obsolescence. It is not difficult to imagine adapting or extending existing expert system approaches to represent knowledge of this sort.

Of course, such an associational representation is, in principle, not independent of numerical representations used in the analysis of financial ratios. Quite the contrary. The connections between associational representations and numerical representations are deep, because each provides a perspective on a single underlying entity—the corporation being analyzed. The following section explores the nature of these connections, argues the need to represent them in future expert systems, and speculates on what some of the characteristics of a solution might be.

Needed: An Integrating Representation

The central problem is to find a representation that integrates the various perspectives taken by corporate analysts. To develop intuition about the requirements such a representation must satisfy, we proceed by means of examples.

The Problem

A person assessing the prospects of a corporation will, of course, be alert for any signs of imminent catastrophe. If one notices the corporate equivalent of a dramatically overdrawn checking account or sees that the corporate strategic plan contemplates a major move into buggy whips, it does not require deep expertise to recognize that the future is not bright. However, although the absence of signs of incipient bankruptcy is a necessary condition for a favorable assessment, it is far from sufficient.

By far the most challenging task in a corporate assessment problem is to estimate what might be called the robustness of the firm, *i.e.*, the ability of the firm to withstand economic shocks and to manage growth. It is far more difficult to assess the robustness of a firm than it is to identify signs of immediate catastrophe, because doing so requires forming an integrated view of all aspects of the firm. The distinction between identifying signs of disaster and assessing robustness is subtle, but we believe it is crucial to extending the power of future generations of expert systems. An analogy with medical diagnosis problems makes this clearer.

A typical expert system for diagnosing diseases begins with the presumption that there are “presenting symptoms” and that one or more diseases will be diagnosed. It might turn out that there is insufficient evidence to support a diagnosis of any particular disease in which case the patient might be said to be “healthy.” However, the state of one’s health is not accurately represented by a binary predicate. In fact, even in the absence of presenting symptoms we undergo physical examinations from time to time

in order to obtain an assessment of our state of health. If we are fortunate, the exam will reveal an absence of any specific disease, but we might nonetheless receive a rather extensive report assessing our general state of health. The suggestion made here is that an expert system aimed at assessing an overall degree of health is likely to differ in important ways from a system aimed at diagnosing the presence of specific diseases. Similarly, an ideal system for assessing the robustness of a firm will differ from one aimed at noticing the presence of serious, but specific, corporate maladies.

To illustrate this distinction in the corporate assessment domain, let us return to the notion of leverage. A corporation’s leverage is important because it is an indicator of the firm’s ability to borrow additional funds, should that become necessary, and because it reveals something about the drain on the firm’s funds to pay interest and principal on its borrowings. However, although leverage is routinely reviewed in an analysis, it does not by itself tell the complete story. It might be that leverage for a particular company is high compared with other firms in the same industry (often taken as a warning sign) but that the company has a large stream of income which enables it to handle the debt. An analyst would then want to examine the income stream to assess whether it is likely to continue into the future. Doing so requires an examination of revenues (all funds taken in from the firm’s operations) and expenses (all funds paid out). Analyzing sources of revenue entails a review of the markets the firm serves, the competition in these markets, and so forth. Similarly, a review of the firm’s expenses requires an analysis of the operating efficiencies experienced by the firm, which in turn might require an analysis of the degree of obsolescence of equipment, labor relations, and so forth. Thus, examination of a simple ratio indicator leads rapidly into a global analysis of virtually all the operations of the company.

The ability to form a global picture of a firm and its prospects is one of the distinguishing abilities of highly experienced analysts. For example, in the case of one small computer manufacturer that we studied, the leverage was so high (compared with other firms in its industry) that a loan request was given a negative recommendation by a junior officer. A more fundamental analysis by a senior officer, roughly along the lines outlined here, revealed that the company was leasing rather than selling its equipment, and, therefore, its financial statements combined the characteristics of a leasing company with the characteristics of a manufacturing company. As it happens, leasing companies are (for quite fundamental reasons) highly leveraged as a class, and in this particular case there was no unusual risk.

Assessing the robustness of a corporation is complicated by several additional factors. First, it is rare for an analyst to have access to complete data about the corporation and its industry, because for all but the smallest firms

the amount of possibly relevant numerical and symbolic data is effectively unlimited. Second, there is always the likelihood that data about the company will be presented in the most favorable light by those providing it; in rare extreme cases, this tendency can become deliberate fraud. Finally, most corporations are not homogeneous but can be subdivided into distinctive components. In the extreme case of a conglomerate, these components represent qualitatively different businesses. Even in more typical situations, one cannot view a corporation as a monolithic, stereotypical entity.

What Is Needed?

We are far from defining an expert system architecture that addresses, in a direct and natural way, the problem of forming an integrated assessment of the overall health or robustness of a modern corporation. We can, however, use the powerful stimulus of real-world problems and examples to build our intuitions about what the properties of such a system would be. In the following sections, we elucidate some of these properties and contrast them with existing associational expert systems. Our intent, of course, is not to denigrate the performance that such systems have demonstrated but to use associational systems as a widely known point of departure that facilitates discussion.

Structured Constraint Checking Analysts routinely perform cross-checks of case data to test for validity and to highlight areas for further investigation when validity is questionable. For example, if symbolic data is received asserting that a firm's productivity is high, an analyst can examine the numerical information in the financial statements to see whether investment in processing equipment has been high enough to support this assertion. If not, *i.e.*, if the source of claimed high productivity cannot be found in up-to-date process equipment, then the analyst might investigate further. For example, the income statement might suggest that unusually low labor costs are responsible for the firm's high productivity.

Other useful constraints might involve purely symbolic information. For example, suppose a firm is planning to grow by expanding its market share. This might be quite reasonable in "mature markets" such as soft drinks but would raise questions in new markets (say, AI hardware and software) where the overall growth of the market is the dominant consideration.

The idea of using multiple constraints involving numerical and/or symbolic data is of course not new. (As a matter of fact, the Renaissance invention of double-entry bookkeeping, which is universally used today, is at root a highly structured system of numerical constraint checks.) It is equally true that many AI systems have been based on the idea of constraint checking and constraint propagation (Waltz, 1975; Tenenbaum & Barrow, 1975; Mackworth & Freuder, 1985). What, then, is special about applying

constraint-based techniques to the corporate assessment task?

An important consideration is that the number of constraints which might be usefully exploited is great enough to demand they be carefully structured. Although it is possible to use associational representations to express constraints, these representations do not especially invite such expression, do not provide natural mechanisms for imposing a structure on constraints, and do not integrate constraint analysis into the overall assessment process, *e.g.*, by allowing constraint-based reasoning to affect the flow of control. What is needed is a representation that provides an appropriate global framework for organizing constraint knowledge.

Local Versus Global Reasoning An important characteristic of associational systems is that they facilitate attacking an assessment or classification task by subdividing it into successively smaller, or more local, subtasks. For example, the tasks of troubleshooting a complex man-made system, of estimating the likelihood of a mineral deposit, and of diagnosing the presence of a disease can all be approached in this fashion. This widely used approach assumes that a global assessment can be adequately constructed by combining a hierarchy of local assessments through the use of certainty factors, Bayesian updating, or other analogous techniques. As noted earlier, however, there are tasks such as corporate assessment for which this key assumption is problematic. If the assumption is not warranted, then the question of how global assessments are to be formed is very much open.

It is interesting in this regard to reflect on work done more than 20 years ago on perceptrons (Rosenblatt, 1962). Perceptrons are networks of summation-and-threshold elements that were applied to many pattern-recognition tasks, including two-dimensional pattern classification. They operate by first making a large number of local decisions about a pattern and then propagating the local decisions through the network to arrive at a global decision. Although not Turing machines, perceptrons were initially thought to be powerful mechanisms for classifying patterns. Then, in a landmark work (Minsky & Papert, 1969), it was mathematically shown that perceptrons cannot in principle make distinctions that rely on global geometric properties such as connectivity. It would be interesting indeed if one could formalize the distinction between local and global assessments in expert systems and use that distinction to suggest architectural styles.

Causality The inability of associational systems to represent and exploit causal relations has been widely noted, and this lack has been felt in corporate assessment applications. As is the case with all artifacts, corporations and their constituent activities have goals and participate in cause-and-effect relations. However, the coupling of cause

and effect is generally looser and less defined in human organizations than it is in engineered systems. Thus, dealing with causality is, if anything, more challenging in corporate assessment problems than in electrical or mechanical domains. However, we believe that causal models—even relatively weak ones—are equally important for developing powerful expert systems for corporate assessment.

Qualitative Reasoning: A Starting Point

It might be that recent work on qualitative reasoning provides an attractive approach to these and related problems.¹ To motivate the discussion, consider Figure 1, taken from a standard text on financial analysis (Helfert, 1982). This figure, which illustrates the flow of funds within a manufacturing firm, is worth careful examination. It is provocative in its clear depiction of causal relations, yet it contains only a subset of the information that must be considered in analyzing a firm. We begin by describing some of the relations implicit in the figure and then suggest some key research issues that need to be addressed if an approach based on qualitative reasoning is to be successful.

The funds flow model pictures assets and liabilities as reservoirs, among which move either funds or the financial equivalent of physical assets (such as inventory). The main flow of funds through the system is in a clockwise direction. Beginning at the top left corner of the diagram, we see cash and credit sales. These sales are an important (but not the only) contributor to the cash reservoir in the center of the figure. Outflow from the cash reservoir supports the purchase of raw materials, illustrated on the right, and also supplies operating funds needed by the firm. The operating funds are used for two main purposes: They support the manufacturing payroll, and they support R&D, administrative, and marketing expenses. Goods produced by manufacturing are sold by marketing, and the “marketing pump” produces revenue in the form of cash and credit sales that allow the cycle to continue. The valves in the figure show points at which management decisions affect the flow of funds through the system.

The funds flow model can be used to construct qualitative causal relations. As a first simple example, consider the “raw materials inventory” reservoir. The value of this inventory is increased by purchasing raw materials (illustrated by the supplier extending trade credit to the buyer) and is decreased by using the material in the manufacturing process. Following the approach of De Kleer & Brown (1984), we can express this relation as a *confluence*. Using a notation similar to theirs we would write

$$\begin{aligned} d(\text{materials-used}) - d(\text{trade-credit}) \\ + d(\text{raw-materials-inventory}) = 0 \end{aligned}$$

¹See the special volume of *Artificial Intelligence*, 24:1–3 (December 1984) for a collection of papers on qualitative physics.

where the differentials are understood to represent qualitative changes. This confluence states that trade credit positively influences materials used, that materials used negatively influences the raw materials inventory, and so forth.

As a second example, consider the “fixed assets” reservoir, and assume for the moment that there are no disinvestments or losses. If management increases its investment in automation equipment, the value of the fixed assets (which include such items) will rise. This will cause depreciation expenses to rise, which in turn will increase the value, *i.e.*, the cost, of the “work in process inventory.” If no other changes occur, the firm will have a higher production cost. However, if the investment in automation allows the firm to produce goods more efficiently, we would expect operating expenses to decrease, hopefully by an amount greater than the depreciation increase. If we neglect inventory losses, we can express the “work in process inventory” reservoir by the following:

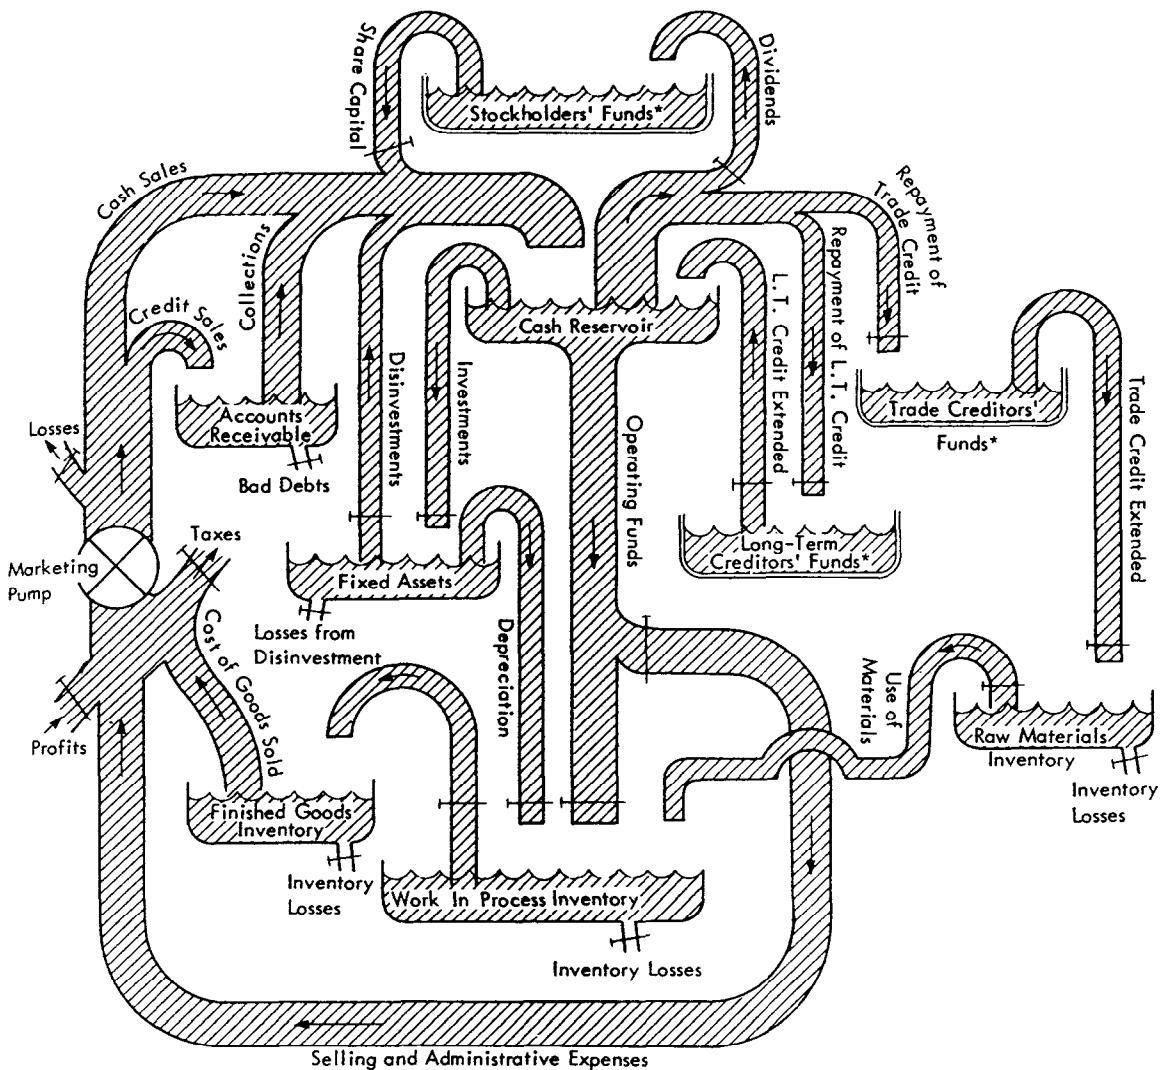
$$\begin{aligned} d(\text{finished-goods-produced}) - d(\text{depreciation}) \\ - d(\text{operating-expenses}) - d(\text{raw-materials}) \\ + d(\text{work-in-process-inventory}) = 0, \end{aligned}$$

which shows that depreciation, operating expenses, and raw materials all contribute positively to the cost of goods completed but that a decrease in the work-in-process inventory will cause an increase in finished goods produced.

In general, it appears that one confluence could be written for each node in the funds flow model. Collectively, this set of conflences constitutes a qualitative funds flow model. (We are here glossing over a myriad of important details, such as the distinction between the total cost of goods produced and the unit cost of production. We assume that these details can be worked out without formal difficulty.) For the most part, this model emphasizes the movement of funds within a single firm. If we restrict attention exclusively to this, we have the choice of appealing to an alternative, highly quantitative model, *i.e.*, the financial accounting model described earlier. In fact, a major part of financial accounting is concerned precisely with problems of representing the funds flow model in great detail. Thus, a qualitative funds flow model stands in the same relation to the financial accounting model as qualitative physics stands to physics; both qualitative models offer a unique perspective on their subjects and complement existing well-developed quantitative models.

Extensions

There are, however, more compelling reasons to search for ways in which qualitative models can be applied to the corporate assessment problem. These reasons have to do with the nature of the interface between the firm and its environment and are most easily described by reexamining Figure 1.



* These represent funds held by others, as contrasted to the resources of the business

Dynamic Funds Flow Model (From Helfert, 1982).

Figure 1.

There are several indications in the figure of the interactions between the firm and its environment. At the top, for example, we see a "stockholders' funds" reservoir that does not in fact lie within the direct control of the firm's management. The firm can attempt to access these funds, if needed, by offering new stock for sale, but the attractiveness of the stock is only partially affected by the firm's actions. Many other factors, such as current and expected interest rates, have an affect on the ability of the firm to acquire new funds through the sale of its stock.

A more complicated example is afforded by the so-called marketing pump. In the funds flow model, the pump is in effect a surrogate for a complex set of relations between the firm and its market. The ability of the firm to

sell its product depends on a great many factors: the perceived need of potential buyers for products of this type, the awareness of buyers of the firm's products, the availability of the product, its price, the nature of the competition, and so on. Many of these factors can be influenced, at least in part, by actions taken by the firm. For example, advertising can increase the perceived need for the product, R&D programs can improve the product with respect to competitor's offerings, automation programs can lower the cost of manufacturing the product, and so forth. Of course, there are many feedback loops involving these factors. A primary one involves demand and price. The demand for a product is usually a strong function of its price, which is a function of its manufacturing cost, which

is a function of production volume, which is a function of demand.

The study of such functional relations is the central concern of microeconomics or, as it is sometimes called, the theory of the firm. For example, a well-explored topic in microeconomics is the study of the *elasticity function* relating demand to price. (See, for example, Thompson, 1981.) Microeconomics, as a well-developed discipline, is likely to be a fertile source of concepts and theories that could inform a qualitative model.

It would seem, then, that there is an opportunity to apply qualitative modeling methods to represent, in an integrated way, the internal operations of a firm together with the interactions of the firm and its environment. The funds flow model described here provides an excellent starting point for modeling internal operations. Microeconomics theory provides the basis for modeling the interactions with the environment, at the very least through its modeling of individual relations and conceivably through quantitative modeling of complex economic systems.

If such models were developed and were computationally tractable, they would provide a powerful means for organizing and using the knowledge required to make corporate assessments. One way in which the model might be used would be to generate expectations, which could then be checked against available data. For example, if a company had relative salary costs that were low, one would expect productivity to be high. This leads to the expectation of low cost of goods sold, which, in turn, suggests that either market share should be high (because goods can be sold inexpensively) or that profit margins should be high (because high prices have been maintained in the face of low production costs). If these expectations are violated, we at the very least have a clue that an anomaly exists which needs further investigation. A causal model that facilitated the systematic generation of all such constraints would represent an advance over a system that could represent constraints but that provided little guidance as to what constraints should be included.

Another use of such models might be to detect and characterize various qualitatively distinct modes of operation of the firm. These modes result from nonlinear effects, such as the emptying of a reservoir, that might be quite difficult to deal with using normal differential equations (even assuming the quantitative data required to drive such equations were available.) This type of analysis of causal models contrasts with the use of simulation models whose execution produces behavior but that cannot readily analyze their behavior automatically.

An extremely important open issue is the question of whether strictly qualitative models provide sufficient power in the corporate assessment domain. One typically has to assess the net effect of conflicting factors, some of which are favorable and some of which are not. It might be that a purely qualitative model leads to too many in-

terpretations to be useful. If this is the case, it might still be possible to augment the qualitative model with enough numerical information to form a semiquantitative representation. This possibility is intriguing but is purely speculative at present.

Conclusion

We have described the corporate assessment problem and have argued that associational representations typical of current expert systems do not adequately capture the knowledge of causal relations and the global perspective which are characteristic of skilled people. Recent work in qualitative reasoning, which to date has focused primarily on physical reasoning, might suggest an approach to this class of problems.

The research issues to be addressed in this area are formidable. It has taken nearly 10 years for qualitative reasoning to reach its present state of development, with all of classical physics to appeal to as an existing guide. The one thing we can be certain of is that many challenges lie ahead.

References

- Brown, J S & Burton, R R (1975) Multiple representations of knowledge for tutorial reasoning In D G Bobrow & A Collins, (Eds), *Representation and Understanding* New York: Academic Press, 311-349
- Cuena, J (1983) *The use of simulation models and human advice to build an expert system for the defense and control of river floods* IJCAI-8, 246-249
- De Kleer, J , & Brown, J S (1984) A qualitative physics based on confluences *Artificial Intelligence*, 24(1-3):7-83
- Fink, P K , Lush, J C., & Duran, J W (1984) A general expert system design for diagnostic problem solving In *IEEE Workshop on Principles of Knowledge-Based Systems* Silver Spring, Md : IEEE Computer Society Press, 45-52
- Helfert, E. A (1982) *Techniques of financial analysis*. Homewood, Ill.: Richard D Irwin
- Mackworth, A K , & E C Freuder (1985) The complexity of some polynomial network consistency algorithms for constraint satisfaction problems *Artificial Intelligence*, 25(1):13-64
- Minsky, M., & Papert, S (1969) *Perceptrons* Cambridge, Mass.: MIT Press
- Patil, R S , Szolovits, P , & Schwartz, W B (1981) *Causal understanding of patient illness in medical diagnosis* IJCAI-7, 893-899
- Pople, H E., Jr (1982) Heuristic methods for imposing structure on ill-structured problems: The structuring of medical diagnostics. In P Szolovits, (Ed), *Artificial Intelligence in Medicine* Boulder, Col : Westview Press, 119-1902
- Rosenblatt, F (1962) *Principles of neurodynamics: Perceptrons and the theory of brain mechanisms* Washington, D C : Spartan Books
- Tenenbaum, J M , & Barrow, H. G (1975) *MSYS: A system for reasoning about scenes*. Menlo Park, Calif : SRI International
- Thompson, A A (1981) *Economics of the firm: Theory and practice*. Englewood Cliffs, N J.: Prentice-Hall
- Waltz, David (1975) Understanding line drawings of scenes with shadows In P H Winston, (Ed.), *The Psychology of Computer Vision* New York: McGraw-Hill, 19-91
- Weiss, S , Kulikowski, C , Amarel, S , & Safir, A (1978) A model-based method for computer-aided medical decision making *Artificial Intelligence*, 11(2):145-172