

# Model-based reasoning in financial domains

WALTER HAMSCHER

Price Waterhouse Technology Centre, 68 Willow Road, Menlo Park, CA 94025, USA.

## Abstract

Finance is a challenging yet appropriate domain for model-based reasoning, an area of research otherwise grounded in classical physics. Among the many features that suggest a model-based approach are that firms have formal internal structures, business entities have idealizable behaviours, and there is a history of formal analysis of business problems. This article discusses the motivations and foundations of the model-based approach, and surveys several existing artificial intelligence programs that exploit its advantages. The survey shows that there are ample opportunities for useful systems and significant research in this area. However, accomplishing either of these goals depends crucially upon moving beyond qualitative models based only on accounting information, which tend not to capture the actual complexities of the domain.

## 1 Overview of model-based reasoning

Model-based reasoning first became an identifiable subfield of artificial intelligence research around 1980 as an effort to address limitations of the contemporary generation of expert systems. Then, as today, most expert systems encoded empirical knowledge extracted from interviewing and analysing the problem solving behaviour of experts. Expert knowledge was encoded in a declarative manner, while problem-solving was performed by a separate inference engine relatively independent of the specific domain knowledge. The architecture has great appeal and commercial ability (Feigenbaum et al., 1988), but by itself provides little guidance as to the appropriate content and organization of the knowledge. The model-based approach retains this basic architecture, but focuses on the content of the knowledge base by emphasizing the formulation and use of *models of structure and behaviour*. Two features tend to characterize the model-based approach:

**Categorical knowledge** The content of the knowledge base in a model-based system tends to concern categorical causes and effects rather than probabilistic associations among problem features. A *model* of an electronic circuit, for example, consists of the network of its components and their connections (structure) along with knowledge about how components interact (behaviour). There is no uncertainty about the facts about the circuit design that are represented in such a model. This stands in contrast, to say, accumulated experience associating faults and their external symptoms, which are inherently uncertain.

**Modularity** Model-based systems aspire to use general-purpose models, and thus achieve an even stricter separation between *what* is known from *how* the knowledge will be used than the comparable separation found in a typical expert system. For example, the electronic circuit model might be used in several ways; to simulate the effects of hypothesized faults in a diagnosis task, or to propose structural changes in a redesign task. It is important to realize that this is a methodological aspiration, not dogma; for computational efficiency, both the content (level of detail and abstractions used) and form of knowledge (rules, constraints, frames, and so forth) present in all model-based reasoning systems are invariably influenced by the intended use to which the knowledge will be put.

These features of a model-based approach enhance *robustness*, the ability to solve a broad range of problems; *compositionality*, the ability to combine knowledge about different problems within a

single system; and *reusability*, the ability to use a knowledge base acquired to solve one class of problems in order to solve a somewhat different but related class of problems. A disadvantage is that when tackling average or typical problems, such systems are likely to be less efficient, since as a practical matter the system will involve computational overhead to support the solving of atypical cases. Note also that there is nothing to prevent building a model-based system using traditional expert system techniques such as production rules (Karp & Wilkins, 1989). What is important is the nature and content of the knowledge in the system, not its form or computational realization.

The primary features of model-based reasoning—emphasis on categorical knowledge and separation of domain knowledge from problem solving knowledge—mean that most work in this area is grounded in physics, and thereby inherits common attributes. First, the domain of discourse consists of decomposable structures and constrained interactions between the elements of those structures. Second, the task to be performed reduces in principle to making and evaluating predictions about the evolution of the aggregate states of such decomposable systems. Third, predictions are made for a virtually closed system whose initial state and all relevant subsequent exogenous influences are knowable. Although the literature contains exceptions, these features characterize the mainstream (Weld & de Kleer, 1990; de Kleer & Williams, 1991).

The ability to model is a prerequisite for model-based reasoning, and model-based reasoning is well exercised in physics-based domains because modelling of physical systems is—at least in principle—well understood. However, this traditional emphasis on physics does not preclude a model-based approach to programs in non-physical domains. On the contrary, this article will show how the principles of model-based reasoning have been applied in financial domains as well. Not every task domain lends itself to a model-based treatment, but transfer should be possible given a domain in which an appropriate theory exists. The key requirement for this transfer of techniques lies in finding and exploiting correspondence between tasks of interest—financial tasks in this case—and the characteristics of physical domains.

## 2 Physics and finance

Reasoning tasks involving businesses are grounded not in physics, but in disciplines such as economics, finance, accounting, marketing, human resource management, and so forth. Such tasks include:

- **Financial assessment** Given financial data such as historical balance sheets, statements of income, and market surveys, predict the future prospects for a firm. *Going-concern evaluation*, in which the key question is whether the firm will be able to meet its obligations, is a special case of financial assessment.
- **Auditing** Provide reasonable assurance that the financial statements of the firm present a fair picture of its operating results and financial position.
- **Tax planning** Given complex tax laws and the planned activities of a firm, recommend actions to minimize the overall tax burden.
- **Cost control** Given records of production costs, find areas where process changes are likely to improve profit margins.

As noted, earlier, the knowledge used in these tasks, and more generally in the business domain, suggest that it may be as appropriate for model-based techniques as any physical domain. Here we consider these similarities in more detail.

### 2.1 Decomposable structures

Decomposable structures that constrain relationships are ubiquitous in finance: financial statements, accounts, reporting and ownership structures, flows of goods and materials, market segments, all are widely used abstractions. The relationship between starting and ending inventory, production cost, and cost of goods sold, the relationship between credits and debits, and other

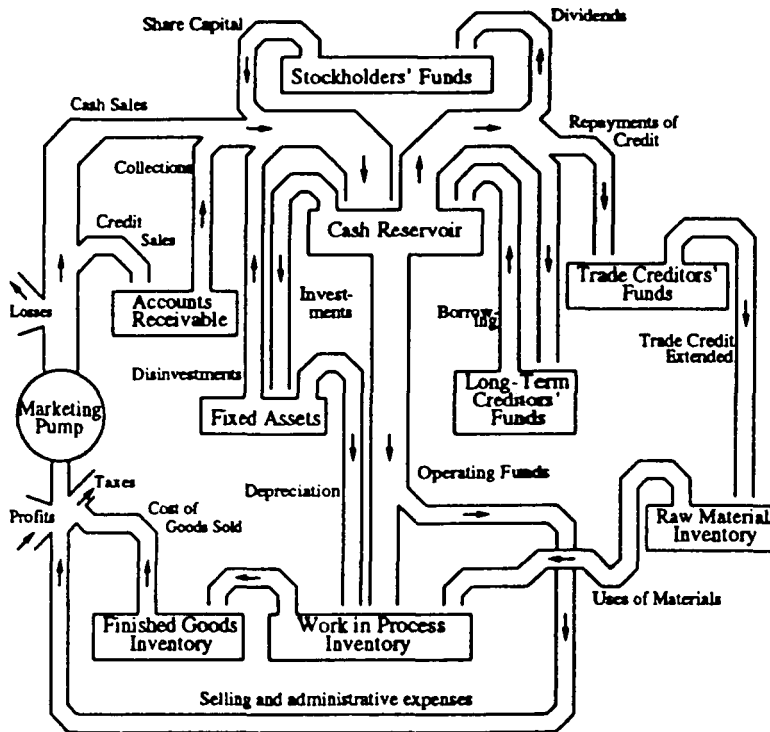


Figure 1 Flow of funds in a firm (after Helfert 1986)

relationships in the financial domain play the same role as the equations governing the interaction of components in an electronic circuit: to formalize the limited interactions that are possible between separate entities. Figure 1 shows an exceptionally direct illustration of this idea: assets and liabilities are shown as tanks, income and cash flow statement items are shown as pipes, so that the arithmetic relationships between these items correspond to their relationships on a financial statement. Teleology provides further constraints on the behaviour of entities in structures that are designed to achieve a certain purpose, such as accounting procedures and manufacturing processes. The role of internal controls in accounting procedures is similar at least metaphorically to the use of error detection and feedback mechanisms in engineering domains, and the same style of analysis might be used in both.

### 2.2 Behaviour prediction

Predicting the behaviour of business entities over time is key to many tasks essential to managing a business. Tax planning requires estimates of future cash flows; auditing requires understanding of how errors are manifested in accounts; evaluating a cost control scheme requires detailed predictions of its eventual effects on operations and profitability. Simulation is a common form of behaviour prediction not only in engineering domains, but also in management domains, as reflected by its importance in the field of operations research. The form of knowledge used in simulation—localized knowledge about the output response of component entities given varying inputs—is similar, although in business domains there is usually more uncertainty associated with the responses of individual components.

### 2.3 Idealizations and assumptions

Although open systems are the rule rather than the exception in the real world, in business just as in physics, it is possible to construct powerful and useful models of nearly-closed subsystems. In particular, there seems to be little difference in kind between the idealizations and assumptions

used for reasoning in physics-based domains and those in economics-based domains. There are several motivations for such assumptions.

### *2.3.1 Distinguish exogenous from endogenous variables*

Modelling of any real system must commit to a distinction between those aspects of behaviour that are to be taken as given, and those with which are to be deduced within the model itself. The idealized flow of funds in a manufacturing firm as illustrated in Figure 1, for example, correctly captures the arithmetic relationships between certain key internal accounts even though all interactions with the marketplace are exogenous, as represented by the “marketing pump”. Many financial analyses require an initial commitment to an exogenous “discount rate” that enables the comparison of cash flows occurring in different periods.

### *2.3.2 Distinguish long-term from short-term effects*

In physics-based domains it is common to make assumptions about the relative speeds of interactions. For example, in the lumped element model of electrical circuits one assumes that the circuit is sufficiently small that propagation delay is negligible. Such assumptions come into play at different levels of analysis; for example, in a model of the water and sodium balances in the human body, depending on the level of analysis one may view the levels of water and sodium as either constants or as changing instantaneously (Kuipers, 1987). Economic analysis frequently involves this type of assumption as well. For example, the supply curve of microeconomic market analysis may be viewed as either vertical, upward-sloping, or horizontal depending on whether the analysis is short-, medium-, or long-term, respectively.

### *2.3.3 Treat individuals as a continuous quantity*

Although in physics, matter is composed of discrete and theoretically individuated molecules, a far more common form of analysis makes the simplifying assumption that the aggregate properties of collections of individuals may be measured as continuous and differentiable quantities. The very concept of density, for example, relies on this simplifying assumption. Similarly, in inventory theory, for example, the key idea is to differentiate the total cost of carrying inventory with respect to time and solve for zero (Hillier & Lieberman, 1980). The differentiation operation only makes sense under the assumption that the quantity of items in inventory is a continuous quantity, even though it is composed of discrete items whose count is properly an integer.

### *2.3.4 Permit computational tractability*

Mathematical models routinely trade off accuracy for tractability by making assumptions for which the errors are assumed to be negligible for the purposes of the analysis. For example, a piecewise model of the behaviour of a transistor is usually sufficient for analysis of VLSI circuits. Models of sliding and rolling objects frequently omit friction. This is similar to financial domains in which linearization is a typical approximation step used to permit analysis. Using inventory theory again as an example, the analysis often involves the simplifying assumption that items are sold from inventory at a constant rate, i.e. inventory is assumed to be a linear function of time.

### *2.3.5 Simplify exposition*

Sometimes idealizations and assumptions are made only for pedagogical reasons, since the assumptions have only minor effects on accuracy, on computational tractability, or on the other factors described above. For example, physics textbooks usually explain conservation of momentum using 2D examples, since the 3D version is arrived at merely by a change of coordinate frame with the vector arithmetic remaining identical. Similarly, the capital asset pricing model (Brealey & Myers, 1981) may be explained under the assumption that the portfolio holder may lend and borrow money at the same rate, an unrealistic assumption (compare your own credit card rate to that of your savings account) that is easy to fix without significantly affecting the results of the analysis.

### 2.4 Summary—physics and financial domains

The similarities—behaviour prediction, decomposable structures, assumptions and idealizations—between the way models are used in physics and finance suggest that the model-based techniques of AI are as appropriate for financial domains as they are for physical.

Of course, there are differences. Theories of economic behaviour admit more predictive uncertainty than those of physical behaviour, even when taking into account that physical models are uncertain because of necessary approximations and simplifications. Also, physics is independent of notions such as agency, intention and preference, but these appear prominently in economic theories. While economic theory may ultimately capture these phenomena using close analogues of physical concepts such as feedback, the conceptual gap does exist and it presents a challenge to model-based techniques.

The point is that there *are* similarities, that these similarities are both interesting and not obvious, and that these similarities have been exploited in expert systems taking a model-based approach. The remainder of this article surveys several such expert systems, performing tasks including what-if analysis, explanation, diagnosis, and design using models of financial entities.

## 3 Accounting models

An almost liberal interpretation of “model based reasoning in a financial domain” is the idea that the arithmetic relationships between financial accounting quantities form a “model” of a firm that can be used for prediction, diagnosis, and so forth. This will be called an *accounting model*. An accounting model consists of arithmetic relationships between Balance Sheet, Income Statement, and Cash Flow statement items all measured in monetary units (henceforth, dollars) and generally has two parts:

**Definitions** These are standard legal definitions, such as:

$$\text{Inventory}_{\text{start}} + \text{Production} - \text{Cost of Goods} = \text{Inventory}_{\text{end}}$$

$$\text{Receivables}_{\text{start}} + \text{Revenues} - \text{Collections} = \text{Receivables}_{\text{end}}$$

$$\text{Revenues} - \text{Cost of Goods} = \text{Gross Profit}$$

Some of these equations are illustrated in Figure 1.

**Ratios** Analysts use ratios to characterize the relative magnitudes of related financial quantities. These ratios are often normative in the sense that most things being equal, they remain approximately constant over time for a given firm, and also provide a basis for comparison between firms of differing sizes; for example:

$$\text{Return on Equity} = \frac{\text{Income}}{\text{Equity}}$$

$$\text{Gross Profit Margin} = \frac{\text{Gross Profit}}{\text{Revenues}}$$

$$\text{Inventory Turnover} = \frac{\text{Cost of Goods} \times 2}{\text{Inventory}_{\text{start}} + \text{Inventory}_{\text{end}}}$$

The key appeal of accounting models is that the qualities referred to in the equations are often readily available for analysis of legal entities that must make periodic public reports, such as public corporations and government entities. The terminology and meaning of the terms is standardized and their presentation governed by extensive legal requirements.

Accounting models have many of the essential features of the type of models amenable to automated reasoning as discussed above. They are hierarchically decomposable (assets consist of cash, receivables, inventories, and so forth; cash may consist of currency, bank deposits,

marketable securities, and so forth). They rest upon several assumptions, in particular the idealization that diverse aspects of the entity can all be summarized as continuous quantities of dollars and meaningfully compared (for example, both the real estate holdings and the inventory of the goods of a firm are represented by quantities representing the lower of their original cost and their market value). The relationships in the equations are all computationally simple. Finally, they can be used in behaviour prediction, which may involve introducing exogenous variables such as prevailing interest rates, personal income growth, and so forth, and assuming that certain ratios or other quantities will remain constant (Saltzman, 1967; Elliott & Uphoff, 1972; Wild, 1987).

Other notions of modelling in financial domains will be introduced subsequently; for the moment, let us consider the two tasks of *what-if analysis* and *explanation* that have been automated for accounting models.

#### 4 What-if analysis

What-if analysis consists of perturbing one or more variables or parameters of a model to see what their effect on other variables will be. In a model consisting of financial quantities, in principle this can be done through recalculation as occurs in spreadsheet software when cell values in the spreadsheet are changed. For example, if

$$\text{Inventory Turnover} = \frac{\text{Cost of Goods} \times 2}{\text{Inventory}_{\text{start}} + \text{Inventory}_{\text{end}}}$$

were represented in a spreadsheet, then changing the number assigned to  $\text{Inventory}_{\text{start}}$  would change the number assigned to Inventory Turnover. This spreadsheet-based notion of what-if analysis, however, lacks certain desirable features (i) there is no simple and general characterization of how Inventory Turnover changes when  $\text{Inventory}_{\text{start}}$  changes, (ii) if many numbers are affected it is difficult to produce an explanation of those changes that is comprehensible to a human, and (iii) examining a space of possible combinations of perturbations require a commitment to quantitative changes in the variables (for example, increasing  $\text{Inventory}_{\text{start}}$  by 20%), rather than a less precise question such as “what if  $\text{Inventory}_{\text{start}}$  increases?”. The third of these features could be provided to some extent in an interval-based spreadsheet (Hyvönen, 1991) but the other features require a more comprehensive approach.

Technology for solving this class of problem has been in development for some years in the area of *qualitative reasoning about physical systems*. (Weld & de Kleer, 1990; de Kleer & Williams, 1991). Among the key goals of qualitative reasoning is to allow the propagation of imprecisely specified perturbations through a model, ultimately to derive qualitative characterizations of their effect, and to produce explanations of those effects using a simple qualitative vocabulary. Hart et al. (1986) proposed performing exactly this sort of qualitative analysis in an accounting model. For example, one might query the model by asking “what if the liabilities of this firm were to increase”, and the program would respond in terms such as “the assets would increase, and if the shareholder’s equity remained constant, the debt-equity (DE) ratio would increase”, Bailey et al. (1990) report on a program that carries out this idea in the area of analysing cash flow.

Briefly, a key idea in qualitative reasoning is to reformulate quantitative equations like

$$\text{Liabilities} = \text{DE} \cdot \text{Equity}$$

$$\frac{d}{dt} = \frac{d}{dt} \text{DE} \cdot \text{Equity} + \text{DE} \cdot \frac{d}{dt} \text{Equity}$$

as constraints over the signs  $\{-, 0, +\}$  of the variables. In the following let  $[x]$  denote the sign of variable  $x$ . Setting a variable to a particular sign constrains the signs of other variables. For

example, setting  $[d/dt \text{ Liabilities}] = +$  (“the liabilities increase”) forces  $[d/dt \text{ DE}] = +$ , given  $[\text{Equity}] = +$  and  $[d/dt \text{ Equity}] = 0$ . This allows reasoning to proceed even when quantities are vaguely specified, and also when all of the relationships in the system are not completely known. For example, suppose that overhead expenses tend to increase when investment in property, plant and equipment (PP&E) increase; this can be described as  $[d/dt \text{ Overhead}] = [d/dt \text{ PP\&E}]$ .

Although the qualitative reasoning approach has the useful features described above, ambiguity results from the weakness of interval arithmetic (of which sign arithmetic is a special case); in particular, the sum of a positive number and a negative number could be any number at all. As a result of this, qualitative analysis is frequently unable to derive useful conclusions without augmentation with further quantitative information, such as the relative orders of magnitude of quantities, or by further partitioning the reals beyond  $\{-, 0, +\}$  (Weld & de Kleer, 1990; de Kleer & Williams, 1991).

### 5 Explanation

The work by Hart et al. and Bailey et al. above focused on the problem of determining the consequences of a proposal perturbation to an accounting model. When humans analyse accounting numbers they are often more interested in finding out why the numbers are as they are. The *explanation* problem—given some observations, to find the underlying perturbations that could be responsible—has been treated in a number of ways.

A direct approach to explaining accounting numbers is to simply reproduce the computation that produced the number. To take an example from the ROME system (Kosy & Wise, 1984; Kosy, 1989), suppose one observes that revenues have increased 20% but profits have declined 1%. Table 1 shows some example figures. A naive explanation would simply reproduce equations and variable values, for example:

$$\begin{array}{l} \text{Profit declined 1\% because} \quad \text{and} \quad \text{Revenues increased 20\% because} \\ \\ \% \text{ Change of any } V = \frac{V_t}{V_{t-1}} - 1 \quad \% \text{ Change of any } V = \frac{V_t}{V_{t-1}} - 1 \\ \\ \text{Profit}_{t_1} = 99 \quad \quad \quad \text{Revenues}_{t_1} = 1200 \\ \text{Profit}_{t_0} = 100 \quad \quad \quad \text{Revenues}_{t_0} = 1000 \end{array}$$

These are poor explanations for two reasons. First they are shallow: the form of the question itself suggested that there are some common parameters influencing both variables. In this case, since

$$\text{Profit} = \text{Revenues} - \text{Cost of goods} - \text{Other Expenses}$$

while Cost of goods generally varies proportionally with the Revenues, one would like an

**Table 1** Example accounting numbers to be explained. Following financial statement conventions, time periods are displayed right to left.

	Period	
	$t_1$	$t_0$
Revenues	1200	1000
Cost of goods	981	800
Other expenses	120	100
Profit	99	100

explanation that mentions Cost of goods. This problem could be partially fixed by expanding the explanation of computed figures such as Profit.

Second they fail to discriminate significant effects. In this example, Other Expenses increased proportionally with Revenues (20%), but Cost of Goods (CoG) increased by 22%—proportionally more than the increase in Revenues (R), and a bigger contributor to Profit (P) besides. A focused explanation would probably not mention Other Expenses (OE) at all.

The ROME system produces explanations improved in both these and other ways by

“measuring the effect of a variable on another variable by what the result would have been without the influence of that variable, leaving all other influences intact” (Kosy, 1989).

where *effect* is given a precise quantitative definition and only those sets of variables with the greatest effect on a given variable appear in its explanation. In this case the effect  $\epsilon$  of variables  $X$  on  $y$  is  $\epsilon(X, y) = |y_2 - f(\mathbf{Z})|$ , where  $\mathbf{Z}$ , contains the values of variables in  $X$  at  $t_0$  and values for the other arguments to  $f$  at  $t_1$ . Using the example in Table 1.

$$\begin{aligned} \epsilon(\{\text{Revenues}\}, \text{Profit}) &= |P_{t_1} - (R_{t_0} - \text{CoG}_{t_0} - \text{OE}_{t_0})| = 200 \\ \epsilon(\{\text{Cost of Goods}\}, \text{Profit}) &= |P_{t_1} - (R_{t_1} - \text{CoG}_{t_0} - \text{OE}_{t_1})| = 181 \\ \epsilon(\{\text{Other Expenses}\}, \text{Profit}) &= |P_{t_1} - (R_{t_1} - \text{CoG}_{t_1} - \text{OE}_{t_0})| = 20 \end{aligned}$$

which shows that both Revenues and Cost of Goods had major effects on Profit, while Other Expenses were relatively insignificant. Kosy and Wise (1984) and Kosy (1989) discuss the effective use of this measure and a variety of explanation related issues including natural language processing and generation.

In the above example, ROME produced its explanations of the behaviour of the entity represented by the accounting model solely in terms of financial quantities already available (Revenues, Cost of Goods, and so forth). A more satisfying notion of explanation invokes hidden exogenous variables—variables which have not been or cannot be measured directly, but that are related to and influence observed variables. For example, in a more elaborate model of a firm, demand for products monotonically influences profit margins, so that a good explanation to propose for the above example is that “demand has decreased”. This explanation is also intuitively satisfying because there is a sense in which the decreasing demand could *cause* the decline in profits. Demand, however, is difficult to measure directly, and in any case is not part of an accounting model as defined earlier. Thus, a prerequisite to producing such explanations is thus to represent what causes the accounting numbers to be what they are.

### 6 Explanations of accounting numbers

Accounting numbers are only the end product of exchanges (buying, selling, and other commitments) between economic entities. Exchanges are recorded in financial records, and these financial records are then summarized and reported as accounting numbers. Figure 2 lays out these relationships. Different vocabularies apply to each category of information; for example, economic forces include Demand, Elasticity, Barriers to Entry, and so forth, while operational records refer to the actual business of a firm such as Volume or Units of goods produced, their Cost and selling Price per unit, and so forth. This is in contrast to accounting models, which measure only dollars. The *cause* relation in Figure 2 subsumes many factors, such as the pricing policies of management (influencing what Exchanges takes place) or the nature of the Reporting system (for




Figure 2 Some casual relationships among types of financial information

$$\begin{aligned}
 \text{Gross Profit} &= \text{Revenues} - \text{Cost of Goods} \\
 \text{Revenues} &= \textit{Volume} \cdot \textit{Unit Price} \\
 \text{Variable Cost} &= \textit{Volume} \cdot \textit{Unit Cost} \\
 \text{Cost of Goods} &= \text{Variable Cost} + \text{Indirect Cost} \\
 \text{Indirect Cost} &= 15\% \cdot \text{Variable Cost}
 \end{aligned}$$

**Figure 3** Sample financial model. Operational variables are italicized (after Kosy and Wise, 1984).

example, in a typical company, shipping an order of goods to a customer simultaneously increments both the Revenues and Receivables accounting numbers).

Unlike accounting models, models that include information about economic forces and financial records appeal implicitly or explicitly to causality to support what-if analysis and explanation. Kosy and Wise (1984) present an example of ROME producing explanations for a model from which the example in Figure 3 is taken. The intermediate variables Variable Cost and Indirect Cost relate accounting and operational variables to each other. The intent is that accounting figures such as Gross Profit are to be explained by perturbations in Volume, and so forth.

Daniels and Feelders (1991) use a similar approach to explanation as ROME, using models with explicit qualitative causal relations in addition to quantitative equations. For example, in their notation the relation

$$\text{Volume} \stackrel{c}{\leftarrow} \{(\text{Advertising}, +), (\text{Credit period}, +), (\text{Unit Price}, -)\}$$

means that the number of units sold is positively influenced by advertising and the length of the grace period granted to credit customers, while negatively influenced by price. In their system, if both Advertising and Volume increased, the increase in Advertising would be said to explain the increase in Volume.

Although quantities such as Volume and Advertising are rarely found in accounts prepared according to standard conventions, in ROME and related systems they are assumed to be observable. However, the nature of the explanation problem changes significantly when observables are to be explained in terms of hidden exogenous variables.

## 7 Diagnosis

For the purposes of this article, a *diagnosis* is an explanation involving hidden exogenous variables—variables which have not been or cannot be measured directly, but that are related to and influence observed variables. For example, using the model from Figure 3, suppose we observe that Gross Profit has decreased since last period, but that only accounting numbers are observable. Figure 4 shows three diagnoses. All three of the diagnoses hypothesize disturbances in the

$$\begin{aligned}
 &\text{Unit Price decreased} \\
 &\quad \longrightarrow \text{Revenues decreased} \\
 &\quad \quad \longrightarrow \text{Gross Profit decreased} \\
 &\text{Unit Cost increased} \\
 &\quad \longrightarrow \text{Variable Cost increased} \\
 &\quad \quad \longrightarrow \text{Cost of Goods increased} \\
 &\quad \quad \quad \longrightarrow \text{Gross Profit decreased} \\
 &* \text{Volume increased} \\
 &\quad \longrightarrow \text{Revenues increased} \\
 &\quad \longrightarrow \text{Variable Cost increased} \\
 &\quad \quad \longrightarrow \text{Indirect Cost increased} \\
 &\quad \quad \quad \longrightarrow \text{Cost of Goods increased}
 \end{aligned}$$

**Figure 4** Potential diagnoses from Figure 3

exogenous variables Unit Price, Unit Cost, and Volume. Because we assume that these hypotheses cannot be confirmed by direct observation, the consistency, parsimony and likelihood of the diagnoses are important criteria. In this case, all three are consistent, and all are parsimonious with only one disturbance to each exogenous variable is hypothesized in each. The first two are better explanations than the third (\*), because for the first two the hypothesized disturbance entails the observation that we wanted to explain, while the third does not: if both Revenues and Cost of Goods increased, then Gross Profit might have either increased, decreased, or stayed the same.

### 7.1 Diagnosis with qualitative financial models

The CROSBY system (Hamscher, 1990, 1991 (b)) addresses the diagnosis problem for accounting models augmented with hidden operational variables. CROSBY produces explanations like those in Figure 4 by applying the techniques of model-based diagnosis (Davis & Hamscher, 1988; Hamscher *et al.*, 1992), in particular the heuristic search technology in SHERLOCK<sup>1</sup> (de Kleer & Williams, 1989). The space of potential diagnoses is defined by the cross product of values that can be taken on by selected exogenous variables. In physical domains these variables would be boolean-valued and refer to diseases, states of components in designed artifacts, or multiple-valued and refer to behavioural modes of components. One version of CROSBY (Hamscher, 1990) uses a qualitative model with propagation of constraints as described in section 4; in this version the exogenous variables are the signs of the exogenous parameters of the financial model and their first derivatives.

In essence, CROSBY computes the probability of possible diagnoses by using Bayes' rule in the form  $p(D|O) = p(O|D)p(D)/p(O)$ , where  $O$  is the set of observations and  $D$  is the diagnosis—a set of hypotheses about the values of exogenous variables. Although, in general, a probabilistic formulation of diagnosis with  $n$  hypotheses and  $m$  observations could require  $2^{n+m} - 1$  joint probabilities with intractable computations, CROSBY exploits three key ideas to make this practical.

- (1) All of the probabilities of the form  $p(O|D)$  are computed directly from the model, dramatically reducing the number of probabilities required. An important case is  $p(O|D) = 1$ , which occurs when the observations are actually entailed by the hypotheses as in the first two examples of Figure 4.

When constraint propagation over the signs of variables deduces an assignment for a variable, it is logically entailed; this provides  $p(O|D) = 1$  for many cases. As a result, CROSBY only needs estimates of the prior probability, and the resulting distribution sums to unity. For example, the variable [ $d/dt$  Unit Cost] might take on the mutually exclusive values [0], [+], or [-] with prior probabilities 0.95, 0.03, and 0.02, respectively. Independence is assumed among all variables.

- (2) Many of the intermediate results in computing  $p(O|D') = 0$  arising from a logical inconsistency between the observations and some set  $D'$  of hypothesis is called a *conflict*. Conflict can significantly prune the search space of diagnoses.

Each value in the domain of each key variable is associated with an ATMS assumption (de Kleer, 1986). Each new deduction is given a label corresponding to the sets of minimal sets of assumptions needed to deduce that value. By propagating labels through a network of Horn clauses resulting from inferences made by the constraint propagator, explanations are constructed for every assignment of a value to a variable. When values deduced using different assumptions disagree, the set of underlying assumptions is declared to be in conflict. All supersets of that conflict are inconsistent.

<sup>1</sup> Hence the name: Crosby was a banker mentioned in the Sherlock Holmes "Adventure of the Golden Pince-nez."

- (3) We need only compute the *relative* likelihoods of different diagnoses. This allows CROSBY to virtually ignore the term  $p(O)$  in Bayes' rule since it is a normalization factor affecting all diagnoses equally. Moreover, this supports a heuristic search strategy that terminates after finding the few most likely diagnoses.

In CROSBY a diagnosis is a consistent set of assumptions that is maximal in the sense that no assumption can be added to it without making it inconsistent. CROSBY performs heuristic best-first search through the space of diagnoses. The evaluation function for each diagnosis is the upper bound of its probability, conditioned on any observation made so far. When an observation is neither entailed by, nor in conflict with a diagnosis  $D$ ,  $p(O|D)$  is estimated as  $1/m$ , where  $m$  is the number of possible outcomes.

### 7.1.1 Example

In the example from Figure 3, the observation is  $[d/dt \text{ Gross Profit}] = [-]$ . Suppose that the three exogenous variables [Unit Price], [Unit Price] and [Volume] and their derivatives have the same probability distributions:

$$\begin{aligned}
 p([\text{Unit Cost}] = [+]) &= p([\text{Unit Price}] = [+]) = p([\text{Volume}] = [+]) = 1.0 \\
 p\left(\left[\frac{d}{dt} \text{ Unit Cost}\right] = [+]\right) &= p\left(\left[\frac{d}{dt} \text{ Unit Price}\right] = [+]\right) = p\left(\left[\frac{d}{dt} \text{ Volume}\right] = [+]\right) = 0.03 \\
 p\left(\left[\frac{d}{dt} \text{ Unit Cost}\right] = [0]\right) &= p\left(\left[\frac{d}{dt} \text{ Unit Price}\right] = [0]\right) = p\left(\left[\frac{d}{dt} \text{ Volume}\right] = [0]\right) = 0.95 \\
 p\left(\left[\frac{d}{dt} \text{ Unit Cost}\right] = [-]\right) &= p\left(\left[\frac{d}{dt} \text{ Unit Price}\right] = [-]\right) = p\left(\left[\frac{d}{dt} \text{ Volume}\right] = [-]\right) = 0.02
 \end{aligned}$$

Letting “[ $d/dt$  Unit Cost] = [+], . . .” denote the diagnosis in which Unit Cost increased and all other exogenous variables remain constant, constraint propagation results in the observation [ $d/dt$  Gross Margin] = [-] being entailed by increasing Unit Cost or decreasing Unit Price, while merely consistent with either increasing or decreasing Volume. Letting  $m = 3 = |\{-, [0], [+]\}|$  and hence  $1/m = 0.33$  this yields

---

$p(O D)$
$p\left(\left[\frac{d}{dt} \text{ Gross Profit}\right] = [-] \mid \left[\frac{d}{dt} \text{ Unit Cost}\right] = [+], \dots\right) = 1.0$
$p\left(\left[\frac{d}{dt} \text{ Gross Profit}\right] = [-] \mid \left[\frac{d}{dt} \text{ Unit Price}\right] = [-], \dots\right) = 1.0$
$p\left(\left[\frac{d}{dt} \text{ Gross Profit}\right] = [-] \mid \left[\frac{d}{dt} \text{ Volume}\right] = [+], \dots\right) = 0.33$
$p\left(\left[\frac{d}{dt} \text{ Gross Profit}\right] = [-] \mid \left[\frac{d}{dt} \text{ Unit Cost}\right] = [-], \dots\right) = 0.33$

---

Using Bayes' rule. CROSBY finds the following four top diagnoses, with other diagnoses—for example those involving more than one variable increasing or decreasing—having lower probabilities,  $\alpha$  denotes the normalization factor  $1/p([d/dt \text{ Gross Profit}] = [-])$ .

Rank	$p(D O)$
(1)	$p\left(\left[\frac{d}{dt} \text{Unit Cost}\right] = [+], \dots \left[\frac{d}{dt} \text{Gross Profit}\right] = [-]\right) = 0.029\alpha$
(2)	$p\left(\left[\frac{d}{dt} \text{Unit Price}\right] = [-], \dots \left[\frac{d}{dt} \text{Gross Profit}\right] = [-]\right) = 0.018\alpha$
(3)	$p\left(\left[\frac{d}{dt} \text{Volume}\right] = [+], \dots \left[\frac{d}{dt} \text{Gross Profit}\right] = [-]\right) = 0.009\alpha$
(4)	$p\left(\left[\frac{d}{dt} \text{Unit Cost}\right] = [-], \dots \left[\frac{d}{dt} \text{Gross Profit}\right] = [-]\right) = 0.006\alpha$

CROSBY thus finds the two relatively most likely diagnoses for declining Gross Profit to be that either Unit Cost has increased or Unit Price has decreased, which is in accord with intuition. These two can be discriminated with additional observations, chosen using techniques detailed elsewhere (de Kleer & Williams, 1987; Hamscher, 1990). For example, the observation that revenues were steady  $[d/dt \text{Revenues}] = [0]$  would result in further constraint propagation that would show diagnoses (2), (3) and (4) to be inconsistent.

In this ranking, changes in Volume as represented by diagnoses (3) and (4) are possible but relatively less likely, in part because they do not entail the observation. To rule out such diagnoses entirely, it is only necessary to replace  $1/m$  with 0; this leaves only the *abductive* diagnoses (Console *et al.*, 1991; Poole, 1991), those that fully explain all observations. One problem with the abductive strategy is that if there is even one observation not entailed by any diagnosis, there are no diagnoses at all.

### 7.1.2 Explanations

The diagnoses that CROSBY finds using the qualitative financial model form the basis of explanations that are comprehensible to human users. The strategy for doing this relies on the results of constraint propagation. For each set of assumptions, termed an *environment*, there is a network of Horn clauses whose conclusions are supported in that environment. These are called the *active clauses* of that environment. Each network of active clauses supports some observations that the user may wish to see an explanation for given a certain environment. Each fact, such as  $[d/dt \text{Gross Profit}] = [-]$ , may be supported by several active clauses, and for clarity it is best to select just one to display to the user. Three local criteria are used to make this choice among the supporting clauses of the fact: (i) the clause chosen must be active in the diagnosis that the user is examining; (ii) the clause must be that which is active in the environment of smallest cardinality; and (iii) the clause must have the fewest antecedent facts. Recursive application of these local criteria within a diagnosis yields a directed graph of active clauses that forms an explanation structure. There is no guarantee of global economy in terms of facts or clauses, but the results are easily comprehensible in practice. The first two explanations in Figure 4 are examples. Figure 5 shows the first explanation in Figure 4 as drawn by CROSBY using the Symbolics GENERA grapher.

### 7.1.3 Limitations

The qualitative version of CROSBY exhibits at least two limitations, one at a technical level and one at the knowledge level.

The technical difficulty is that the qualitative values do not permit appropriate levels of ambiguity; for example, in any real company Revenues *always* change from period to period, but it is important to distinguish relatively large changes from small changes, the latter of which can only be represented as  $[d/dt \text{Revenues}] = [0]$ . This also interacts with the probability calculations in that

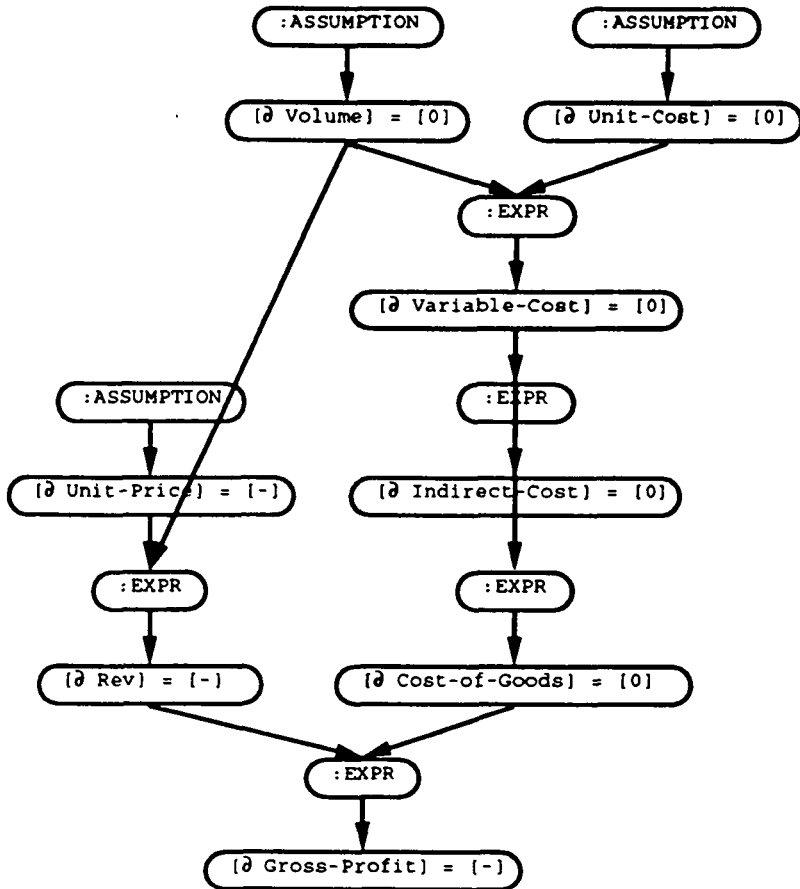


Figure 5 A CROSBY explanation. Each node refers to an equation from the model in Figure 3; ∂ means d/dt

generally the larger the change, the less likely it is *a priori*, but its simple sign representation does not provide a way to formulate this.

The knowledge level limitation arises from the fact that equations like Revenues=Volume · Unit Price (Figure 3) implicitly represent causal relationships—in this case, that revenues result from selling goods. Using equations as proxies for causal relationships in this context is a reasonable strategy because there is a clear distinction between unobservable (operational) variables on the one hand, and observable (accounting) variables on the other; this ensures that the hypotheses predict the observations and not the reverse. But while CROSBY presumes that experts can estimate the relative likelihoods of perturbations of these operational variables, in fact there is no reason to believe that this is any easier to do than for accounting variables. It may be just as hard if not harder for an expert to make judgements about the relative likelihood of changes in (say) Volume as it is to judge likelihoods of changes in (say) Revenues. A more practical strategy is to represent exogenous variables for which analysts are familiar with and comfortable in making judgements about even if the resulting equations are more remote proxies for causal relations. Financial ratios like those introduced in section 3 can fill this role:

$$\text{Return on Equity} = \frac{\text{Income}}{\text{Equity}}$$

$$\text{Gross Profit Margin} = \frac{\text{Gross Profit}}{\text{Revenues}}$$

$$\text{Inventory Turnover} = \frac{\text{Cost of Goods} \times 2}{\text{Inventory}_{\text{start}} + \text{Inventory}_{\text{end}}}$$

Ratios are used because analysts may expect that most things being equal, they remain approximately constant over time for a given company. Cross sectional parameters such as Market Share are technically similar, and provide a basis for comparison between companies of differing sizes. The definitions of these ratios and cross sectional parameters have a basis in causality (for example, the notion that physical goods spend a certain amount of time in inventory supports the *Inventory Turnover* ratio), although the relationship is neither as immediate nor obvious as with operational models.

7.2 *Diagnosis with quantitative financial models*

The current version of CROSBY (Hamscher, 1991 (b)) uses a model of companies that incorporates quantitative financial analysis techniques (Helfet, 1986). Figure 6 shows an overview of the data used, and makes explicit the categories of hypothesis involved. For example, in financial analysis it is common to use historical data to make projections, with the crucial implicit hypothesis that historical trends will continue. Similarly, reported current period data can be used in analyses only under some hypothesis about the magnitude and direction of errors in the data. CROSBY makes these hypotheses explicit, embodying them as ATMS assumptions, and finds diagnoses essentially as described in section 7.1.

As before, the model consists of a set of equational constraints. Each variable has a value that is expressed as an interval of the real numbers—the interval  $[0, +\infty)$ , for example, denotes all nonnegative numbers. The constraint language ACP (Hamscher, 1991 (a)) propagates intervals through constraints and records justifications and ATMS labels for every interval value predicted. Under a particular set of assumptions, only the most specific interval is visible at any given time.

Each financial number has several attributes: its reported value (i.e., what the company management says it is), its actual value (which is never directly observable) and finally its error,

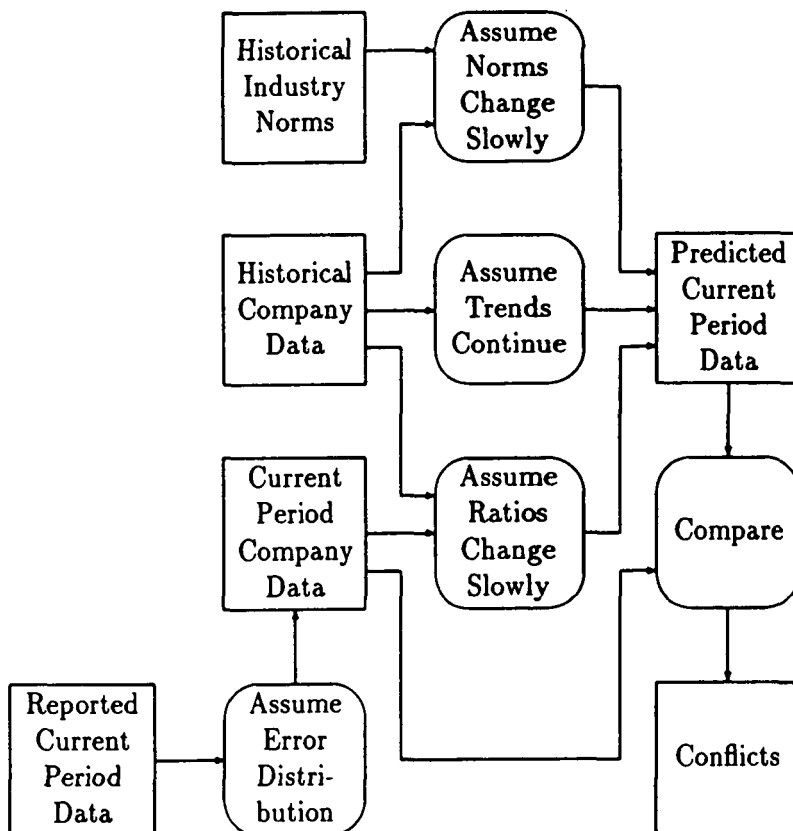


Figure 6 Information flow in CROSBY. Boxes represent data assets; ovals represent sets of model constraints and the hypothesis that they embody

which is the percentage difference between the reported and actual. Errors are modelled explicitly in this way because CROSBY is intended to be used in an audit planning context, in which the auditor is interested in finding anomalies in the accounting numbers that might be best explained by reporting errors. Every reported financial number has five possible error sizes: badly understated ( $-1.0, -0.10$ ), somewhat understated [ $-0.01, -0.10$ ], approximately correct [ $-0.01, +0.01$ ], somewhat overstated ( $+0.01, +0.10$ ) and badly overstated ( $+0.10, +\infty$ ).

Suppose, for example, that a company reports its sales in \$000s as 555. The reported sales is recorded as the interval  $[554.4, 555.5)$ , since we know that the figure 555 has been rounded off. Under the assumption that the error is  $[-.01, +.01]$ , the actual sales is the interval  $(548.96, 561.05)$ . It is this latter interval that participates in constraints that relate, for example, Revenues to Receivables. Under the assumption that the error was  $(+0.10, +\infty)$ , the actual revenues would be  $(609.95, +\infty)$ . There is a probability distribution for the error of any financial number; for example, understatement of Revenues or Assets is typically far less likely than overstatement.

ACP constraints also define the ratios and cross sectional parameters used by financial analysts, as shown earlier. These variables are constrained to meaningful intervals, for example, Gross Profit Margin must fall within  $[0,1]$ . These ratios and their movement (or lack of movement) over time are the basis for predictions by the model. Each exogenous variable has *extrapolation methods* for generating predictions of its value based on historical trends. For example, ratios such as Receivables Turnover, or cross sectional parameters such as Market Share, may be predicted as:

- Constant. The actual is within an interval around the previous actual. The size of the interval depends on the historical regularity of the number.
- Increase (Decrease). The actual falls in an interval above (below) the interval predicted by the Constant method.

As with errors, there is a probability distribution over alternative extrapolation methods. Other extrapolation techniques used in CROSBY include linear weighted averaging and univariate regression.

CROSBY has been tested on a variety of examples. One of these examples concerns MiniScribe, a maker of disk drives. From 1986 through to 1988 MiniScribe management overstated revenues using such techniques as shipping boxes full of bricks and booking shipments to their own warehouses as sales (Kaufman, 1989; MiniScribe, 1989; Zipser, 1989). The financial data was suggestive of overstatement at least as early as the first quarter of 1988, since MiniScribe appeared to be simultaneously increasing its margins as well as its revenues during a period when comparable disk drive manufacturers had poor sales and accumulating inventory. Given data on various combinations of Revenues, Inventory, Receivables, Cost of Goods, and the comparable figures for six close competitors, CROSBY usually concludes with high confidence that MiniScribe Q1/88 Revenues are overstated.

One of the deficiencies of this version of CROSBY is that it builds in unwarranted independence assumptions. For example, CROSBY treats errors in Revenues and Receivables as two independent exogenous variables. In reality they are closely linked; in a typical company, shipping an order of goods to a customer simultaneously increments both numbers, so it would, for example, be difficult to increase Revenues in the short term without also increasing Receivables. The remainder of this section considers ways of modelling some of these interdependencies.

### 7.3 Diagnosis with model of incentives and transactions

In real companies, the management has many ways of influencing the accounting numbers reported, in ways ranging all the way from legitimate but indirect—for example, reducing selected prices to stimulate sales volume, decrease inventory and increase revenues in the short term—to rather blunter methods such as those practised at Miniscribe. Representing and reasoning about the *incentives* influencing the actions of management in presenting financial results is the domain of IRE (Dhar et al., 1988).

IRE generates expected financial results using a simple trend extrapolation technique, then constructs explanations for deviations between expected and reported financial numbers by using these management incentives among other hidden exogenous variables. For example, suppose there is an incentive for management to maintain the Current Ratio (Current Assets/Current Liabilities) greater than 2.0. This incentive is a possible explanation for any unexpected increase in Current Assets or one of its constituents such as Cash or Receivables; similarly for an unexpected decrease in Current Liabilities or its constituents. This is taken to be a particularly likely explanation if the ratio were reported to be very close to the specified value of 2.0. For each reported account IRE assembles these incentives to explain unexpected deviations. The control strategy for choosing among alternative explanations is based on a cognitive model derived from numerous expert protocols and implemented as pattern-action rules (Peters, 1990).

Incentives can be viewed as additional exogenous variables that remove some of the unwarranted independence assumptions as discussed above. Continuing the example above, the incentive to maintain the Current Ratio above 2.0 is a common underlying cause for variations in both Assets and Liabilities. Likewise, IRE represents some of the accounting transactions that underly the financial numbers, and could propose a single perturbation that explains deviations in multiple accounts. For example, if shipping an order of goods to a customer simultaneously increments both Revenues and Receivables, then a phoney shipment would result in overstatement of both numbers.

#### 7.4 Diagnosis with microeconomic models

Another class of relationships among apparently independent financial variables is the area of microeconomics, in the interaction between the actions of an individual company and other individuals via markets. In CROSBY, for example, Gross Margin and Market Share are treated as independent variables. In competitive markets, companies acquire Market Share by reducing their prices and hence their margins. In that case the variables are inversely correlated. On the other hand, this relationship need not hold in monopolistic or oligopolistic markets.

In contrast to management incentives and accounting transactions that influence how operational records are reported (Figure 2), these microeconomic relationships influence the operational records themselves. A microeconomic model would use concepts from the *theory of the firm* such as market demand, barriers to entry, fixed and variable costs of production, production capacity, and derived concepts such as utilization, marginal cost of production, elasticity of demand, and so forth, to explain financial performance Bouwman (1983) describes a program that uses some of the concepts—none of which correspond to directly observable quantities—to produce diagnoses of financial performance. Figure 7 shows one example. The program used qualitative characterizations of variables, such as “too high” and “too low”, and assembles relationships between them to account for one or more observations, the latter of which are typically accounting numbers. Like IRE, the control strategy for choosing among alternative explanations is based on a cognitive model of expert behaviour.

*Capacity*: too low  
 → Units Produced: too low  
   → Units Sold: too low  
     → Finished Goods Inventory: too low  
 → *Utilization*: too high

**Figure 7** A financial diagnosis involving microeconomic variables (italicized). (From Bouwman, 1983)

### 7.5 Summary—Diagnosis

This section began by contrasting *diagnosis* with *explanation*. Both tasks are concerned with constructing arguments to support observable variables, and one of the key preference criteria distinguishing good explanations and diagnoses from poor ones are that they should appeal to causal relationships. This holds even more when the causal relationships are represented indirectly by equations that represent features of companies such as their performance ratios. The crucial distinguishing feature is the observability of the exogenous variables. In diagnosis, the *hypotheses* about the values of exogenous variables are not observable, and so in addition to the preference for causal arguments, a variety of other preference criteria come into play, such as consistency (are the hypotheses and their consequences mutually consistent?), adequacy (does the diagnosis account for all observations?), parsimony (does it propose more hypotheses than necessary?), and likelihood (are the hypotheses themselves plausible, and are the observations plausible given the hypotheses?). CROSBY (Hamscher, 1990, 1991b) takes the approach of embedding all these criteria into a single quantitative measure, the posterior probability of a diagnosis given the observations. Other systems (Peters, 1990; Bouwman, 1983) rely on expert problem solving heuristics to control the process of proposing hypotheses and resolving inconsistencies among them. Despite these process differences, all share the distinguishing features of model-based reasoning: an emphasis on categorical reasoning, declarative representation of relationships between qualitative and quantitative features of business entities, and idealization of relationships to permit computational tractability.

## 8 Design

The previous sections have treated the tasks of analysis, explanation, and diagnosis and have discussed how ideas useful in physical domains can be carried over into financial ones. This final section treats the task of *design*, which in this article means assembling simple structural elements into a complex artifact whose behaviour achieves specified functional requirements. One key difference between diagnosis and design is that rather than analysing a fixed structure or set of variables, the structure itself is to be determined by the task. Roughly speaking, design is more than assigning values to variables; it is an exploration of what variables would be appropriate.

There are several alternative characterizations of the design process (Maher, 1990), but most presume detailed behavioural models of structural elements, along with distinct phases of proposals, analysis/verification, and refinement of structures composed from these elements. One of the many examples of model-based approaches to design is the IBIS system (Williams, 1990) that designs hydromechanical regulators. The requirements of these regulators are typically to maintain a constant flow of fluid to maintain a constant level in containers from which fluid is being drawn, or to perform similar functions. The structural elements of these designs are containers, pipes, valves, levers, and so forth. IBIS has a set of qualitative equations (section 4) that model the behaviour of each structural element, each of which can be viewed as establishing some interactions between different variables; a pipe, for example, establishes interactions between otherwise independent flow and pressure variables. Given a specification of a desired behaviour, IBIS (i) proposes a set of interactions that would achieve the desired behaviour, (ii) augments the device structure with structural elements that establish those interactions, and (iii) verifies through mixed qualitative and quantitative algebraic manipulation that the path produces the desired behaviour. If step (iii) (verification) fails, IBIS can either backtrack to step (ii) (refine the device structure) or step (i) (propose a different path).

The financial design task performed by the INTELLIGENT-HEDGER system (Benaroch & Dhar, 1991a,b) bears important similarities to the model-based physical design of work of which IBIS was an example. INTELLIGENT-HEDGER designs financial instruments whose function is to hedge against anticipated risks. The structural elements of these financial instruments are objects such as stocks and bonds as well as instruments derived from these, such as options on bonds and combinations of

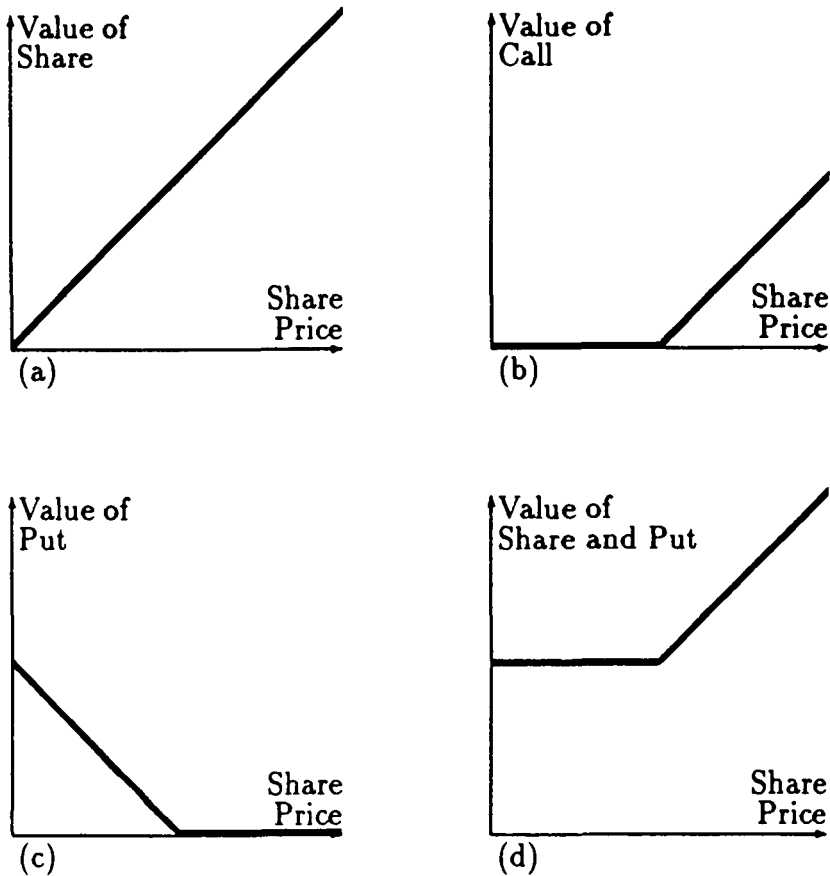


Figure 8 Examples of payoff profiles (after Brealey & Myers, 1981)

options. The behaviour of any instrument can be represented as a *payoff profile* based on possible future events. A simple example is the purchase of a stock (Figure 8(a)): the value of the instrument in the future (vertical axis) is simply its market price (horizontal axis). A *call* option is the option to buy a stock (or any instrument) in the future at a specified *exercise* price. A call option is worthless if the stock price is below the exercise price; otherwise its value is the difference between the stock price and exercise price (Figure 8(b)). A *put* option is the option to sell the instrument at an exercise price, so its payoff profile is symmetric around the exercise price (Figure 8(c)).

The opportunity for *hedging* arises both from the variety of instruments, with their varying sensitivities to external events (overall market movements), changes in interest and inflation rates, and so forth) and from the ability to combine them. Figure 8(d) shows the result of buying both a share and a put option on that stock; it is the vertical sum of (a) and (c). Buying the option hedges against the risk of a large drop in the share price.

INTELLIGENT-HEDGER takes a model-based approach to the problem of designing a hedge against some specific task. For example, a company that wishes to borrow money in a few months' time would wish to design a hedge against the risk of an imminent increase in prevailing interest rates. INTELLIGENT-HEDGER begins with an initial knowledge base of available basic instruments; the user provides both a goal profile and a qualitative set of expectations about macroeconomic indicators (such as the short and medium term movement of interest rates). INTELLIGENT-HEDGER generates an initial set of possible instruments via the principle of matching risks; in this example this would filter out candidates such as "buy stock" that are only indirectly sensitive to the key risk of interest rate movements. The analysis/verification step of the design is done via a qualitative analysis of each instrument, using a kind of simulation to generate the payoff profile for each candidate under each macroeconomic scenario. Some payoff profiles thus generated may only partially match the goal profile; INTELLIGENT-HEDGER can synthesize new instruments by combining these profiles

(Figure 8(d)). Finally, all candidate instruments are ranked according to some qualitative preferences (such as preferring to build the hedge out of widely traded rather than custom instruments).

As a prototype, *INTELLIGENT-HEDGER* exhibits all the key advantages of a model-based system. It is compositional, because it can assimilate rapid changes in the knowledge base; the latter property is exceptionally important in this domain because the specific instruments available for trade change daily. It is robust, because it is based on fundamental principles of hedging (matching risks, for example) rather than on collection of examples. Finally, it is reusable, in the sense that it derives qualitatively plausible hedges with enough specificity as to allow the use of existing quantitative analysis programs for subtasks such as identifying underpriced instruments.

## 9 Observations and future directions

The set of expert systems demonstrating the applicability of model-based reasoning methods to financial domains is currently quite small, so this article has by necessity been a survey of a sparsely populated area. However, at least one observation can be made about them as a group. While these programs all represent intriguing explorations, there appears to be further potential leverage provided by the model-based approach that they do not yet exploit and, conversely, there is a sense in which they do not yet profoundly challenge model-based reasoning techniques. Each exploits only simple behaviour prediction techniques from physical problem solving; for example, the use of qualitative simulation in *INTELLIGENT-HEDGER* (Benaroch & Dhar, 1991b), while an excellent example of how this transfer contributes to the generality and robustness of the overall system is insufficient for actually finding the best tradeoff, an inherently quantitative question. Of the programs surveyed, only *IRE* (Dhar et al., 1988) exploits decomposability in an explicit way, using the decomposition of financial statements into accounts to focus its attention. Similarly, although all the programs have implicit idealizations and assumptions in their models, only *CROSBY* attempts to manipulate them explicitly, and its alternative extrapolation methods are only a shadow of the variety of modelling assumptions manipulated by other systems (Addanki et al., 1991; Falkenhainer & Forbus; 1991; Weld & Addanki, 1990).

The key reason for this lack of aggressiveness appears to be the limitations of the models used. Accounting models, in particular, are a reasonable place to start with this line of research—accounting data is quite accessible in this domain and accounting relations are computationally simple and easy to acquire—but financial information is only the end product of many more complex and subtle relationships (Figure 2). From a practical viewpoint, supporting the problem solving of real experts in financial domains will require those more complex models; from a scientific viewpoint, nothing will be learned about model-based reasoning in general through studying this domain until the models either (i) explicitly represent such intuitively *non*-physical phenomena such as agency, intention, preferences, and so forth, or (ii) match the complexity of those in physical domains.

A representative of research in category (i) is represented by the work of Bridgeland (1990) on extending qualitative simulation to support microeconomic reasoning. Briefly, some economics models presume an agent who is actively seeking to maximize utility. The preferences of the agent—and hence the dynamic evolution of a system of which that agent is a part—impose a partial order on alternative states, such that the system will evolve toward states preferred by the agent. In a qualitative simulation framework, there is no way to formulate such constraints among alternative states. Bridgeland (1990) describes how to extend qualitative simulation with a language capturing the preferences for one or more agents.

Research in category (ii), in which the models match or exceed the complexity of models in physical domains, is exemplified by work in progress on the *SAVILE* project (Hamscher, 1992). A domain model in *SAVILE* represents the steps involved in producing and processing financial data in a company, and is based on earlier accounting systems research (Bailey et al., 1985, 1989). In terms of Figure 2, *SAVILE* models in detail the causal relationships among operational records, and how

the reporting system produces the accounting numbers. The goal is a modelling framework that can support a variety of tasks related to audit planning such as evaluating internal controls, designing tests, and assessing inherent risk. A key issue in this early stage of effort has been to develop a foundational ontology for data processing systems, broadly construed to include both manual and computerized processing steps. A typical information system is an order of magnitude more complex than the most complex digital system tackled by any model-based reasoning prototype to date, namely, a single computer board (Hamscher, 1990c), hence, as with any program designed to reason about such a complex system, the complementary issue is to design the representation so as to allow the suppression of details that are not significant for the problem at hand. In this sense, SAVILE tries to address issues relevant to model-based reasoning in general.

A principal lesson of this survey is that although there are ample opportunities for useful systems and significant research in this area, accomplishing either goal depends crucially upon moving beyond qualitative accounting models to exploit the true richness of the underlying domain.

### Acknowledgements

Beau Sheil provided valuable feedback on early drafts. Michael Benaroch, Jim Peters and Brian Williams helped to clarify aspects of their work.

### References

- Addanki, S, Cremonini, R and Penberthy, JS, 1991. "Reasoning about assumptions in graphs of models" *Artificial Intelligence* 51 (1-3) 145-177. (Also in: J de Kleer and B Williams (eds), *Qualitative Reasoning about Physical Systems II* North-Holland MIT Press.
- Bailey, AD, Duke, GL, Gerlach, J, Ko, C, Meservy, RD and Whinston, 1985. TICOM and the analysis of internal controls "The Accounting Review" LX (2) 186-201.
- Bailey, AD, Han, KS, Stansifer, RD and Whinston, AB, 1989. "The advanced internal accounting control model using a logic programming approach" Working Paper.
- Bailey, AD, Kiang, Y, Kuipers, B and Whinston, AB, 1990. "Analytical review and qualitative and causal reasoning in auditing" Draft.
- Benaroch, M and Dhar, V, 1996. "An intelligent assistant for financial hedging". In: *Proc. 7th IEEE Conf. on AI Applications* 168-174, Miami, FL.
- Benaroch, M and Dhar, V, 1996. "A knowledge-based approach to solving hedge design problems". In: *1st Int. Conf. on AI Applications on Wall Street* 140-145, New York.
- Bouwman, MJ, 1983. "Human diagnostic reasoning by computer: An illustration from financial analysis" *Management Science* 29 (6) 653-672, June.
- Brealey, RA and Myers, SC, 1981. *Principles of Corporate Finance*. McGraw-Hill.
- Bridgeland, DM, 1990. "Three qualitative simulation extensions for supporting economics models". In: *Proc. 6th IEEE Conf. on AI Applications* 266-273, Santa Barbara, CA.
- Console, L, Theseider Dupré, D and Torasso, P, 1991. "On the relationship between abduction and deduction". *J. Logic and Computation* 1 (5) 661-690.
- Daniels, HAM and Feelders, AJ, 1991. "Combining qualitative and quantitative methods for model-based diagnosis of firms" In: Singh, M and Trave-Massuyes, L (ed), *Decision Support Systems and Qualitative Reasoning*. North-Holland.
- Davis, R and Hamscher, WC, 1988. "Model based reasoning: Troubleshooting". In: HE Shrobe (ed), *Exploring Artificial Intelligence: Survey Talks from the National Conferences of Artificial Intelligence* 297-346. Morgan Kaufmann. (Also in: PH Winston and SA Shellard (eds), 1990. *Artificial Intelligence at MIT: Expanding Frontiers* MIT Press. And in WC Hamscher, J de Kleer and L Console (eds), 1992, *Readings in Model-based Diagnosis* Morgan Kaufmann.)
- de Kleer, J and Williams, BC, 1987. "Diagnosing multiple faults", *Artificial Intelligence*, 32 (1) 97-130, April. (Also in: Ginsberg, ML (ed), 1987, *Readings in Nonmonotonic Reasoning* Morgan Kaufmann. And in: WC Hamscher, J de Kleer and L Console (eds), *Readings in Model-based Diagnosis* Morgan Kaufmann.)
- de Kleer, J and Williams BC, 1989. "Diagnosis with behavioral modes". In: *Proc. 11th Int. Joint Conf. on Artificial Intelligence* 1324-1330, Detroit, MI. (Also in WC Hamscher, J de Kleer and L Console (eds), 1992, *Model-based Diagnosis* Morgan Kaufmann.)
- de Kleer, J and Williams BC, (eds), 1991. *Qualitative Reasoning about Physical Systems II* Elsevier.
- de Kleer, J, 1986. "An assumption-based TMS" *Artificial Intelligence* 28 (2) 127-162. (Also in: Ginsberg, ML (ed), 1987, *Readings in Nonmonotonic Reasoning* Morgan Kaufmann.)

- Dhar, V, Lewis, B and Peters, J, 1988. "A knowledge based model of audit risk" *AI Magazine* 9 (3) 57–63, Fall.
- Elliott, JW and Uphoff HL, 1972. "Predicting the near term profit and loss statement with an econometric model: A feasibility study" *J. Accounting Research* 259–274, Autumn.
- Falkenhainer B and Forbus, KD, 1991. "Compositional modeling: Finding the right model for the job" *Artificial Intelligence* 51 (1–3) 95–143. (Also in J de Kleer and B Williams (eds), 1992, *Qualitative Reasoning about Physical Systems II* North Holland MIT Press.)
- Feigenbaum, EA, McCorduck, P and Nii, HP, 1988. *The Rise of the Expert Company*. Times Books.
- Hamscher, WC, de Kleer, J and Console, L, (ed), 1992. *Readings in Model-based Diagnosis* Morgan Kaufmann.
- Hamscher, WC, 1990. "Explaining unexpected financial results" Technical Report 11, Price Waterhouse Technology Centre, Menlo Park, CA 94025. (Also in *Working Notes of the Spring Symposium on Automated Abduction*, 96–100, AAAI Press, 1990.)
- Hamscher, WC, 1991a. "ACP: Reason maintenance and inference control for constraint propagation over intervals" In: *Proc. 9th National Conf. on Artificial Intelligence* 506–511, Anaheim, CA, July. (Also in WC Hamscher, J de Kleer and L Console (eds), 1992, *Readings in Model-based Diagnosis* Morgan Kaufmann.
- Hamscher, WC, 1991b. "Model-based financial data interpretation" In: *1st Int. Conf. on AI Applications on Wall Street* IEEE Press. (Also in: *Working Notes of the Model-based Reasoning Workshop* AAAI Press, 1991. And in: *Working Notes of the 2nd International Workshop on Principles of Diagnosis* Technical Report RT/DI/91-10-7, Dipartimento di Informatica, Università di Torino, 1991.)
- Hamscher, WC, 1991c. "Modeling digital circuits for trouble shooting" *Artificial Intelligence* 51 (1–3) 223–271, October. (Also in: WC Hamscher, J de Kleer and L Console, (eds), 1992, *Readings in Model-based Diagnosis* Morgan Kaufmann.)
- Hamscher, WC, 1992. "Modeling accounting processes to support multiple tasks: A progress report" In: *Proc. 10th National Conf. on Artificial Intelligence*, San Jose, CA, July.
- Hart, PE, Barzilay, A and Duda, RO, 1986. "Qualitative reasoning for financial assessments: A prospectus" *AI Magazine* 7 (1) 62–68, Winter.
- Helfert, EA, 1986. *Techniques of Financial Analysis*. Homewood.
- Hillier, FS and Lieberman, GJ, 1980. *Introduction to Operations Research* Holden-Day.
- Hyvönen, E, 1991. "Interval constraint spreadsheets for financial planning." In: *1st Int. Conf. on AI Applications on Wall Street* 302–311, IEEE Press.
- Karp, PD and Wilkins, DC, 1989. "An analysis of the distinction between deep and shallow expert systems" *Int. J. Expert Systems Research and Applications* 2 (1) 1–32.
- Kaufman, S, 1989. "Report Finds Massive Fraud at Miniscribe" *San Jose Mercury News*, 4D, September 12.
- Kosy, DW and Wise, BP, 1984. "Self explanatory financial models." In: *Proc. 4th National Conf. on Artificial Intelligence* 176–181, Austin, TX, August.
- Kosy, DW, 1989. "Applications of explanation in financial modeling." In: LE Widman, KA Loparo and NR Nielsen, (eds), *Artificial Intelligence, Simulation, and Modeling* 487–509. Wiley-Interscience.
- Kuipers, BJ, 1987. "Abstraction by time-scale in qualitative simulation." In: *Proc. 6th National Conf. on Artificial Intelligence* 621–625, Seattle, WA, August. (Also in DS Weld and J de Kleer (eds), 1990, *Readings in Qualitative Reasoning about Physical Systems* Morgan Kaufmann.)
- Maher, ML, 1990. "Process models for design synthesis." *AI Magazine*, 11 (4) 49–58, Winter.
- MiniScribe, 1989. Summary of Report of the Independent Evaluation Committee of the Board of Directors of MiniScribe Corporation. Exhibit provided in 8-K Disclosure, September 12.
- Peters, JM, 1990. "A cognitive computational model of risk hypothesis generation." *J. Accounting Research* 28 (Supplement) 83–103.
- Poole, D, 1991. "Representing diagnostic knowledge for probabilistic Horn abduction." In: *Proc. 12th Int. Joint Conf. on Artificial Intelligence* 1129–1135, Sydney, Australia, August. (Also in: WC Hamscher, J de Kleer and L Console (eds), 1992, *Readings in Model-based Diagnosis* Morgan Kaufmann.
- Saltzman, S, 1967. "An econometric model of a firm." *The Review of Economics and Statistics*, 49 332–342, August.
- Weld, DS and Addanki, S, 1990. "Query-directed approximation." Technical Report 90-12-02, Department of Computer Science and Engineering, FR-35, University of Washington.
- Weld, DS and de Kleer, J, 1990. *Readings in Qualitative Reasoning about Physical Systems* Morgan Kaufmann.
- Wild, JJ, 1987. "The prediction performance of a structural model of accounting numbers" *J. Accounting Research* 25 (1) 139–160, Spring.
- Williams, BC, 1990. "Interaction-based invention: Designing novel devices from first principles" In: *Proc. 8th National Conf. on Artificial Intelligence*, 349–356, Boston, MA, August.
- Zipser, A, 1989. "How pressure to Raise Sales Led MiniScribe to Falsify Numbers" *Wall Street Journal*, 1 September 11.