SELF-EXPLANATORY FINANCIAL PLANNING MODELS

Donald W. Kosy and Ben P. Wise

The Robotics Institute, Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

The purpose of computing is insight, not numbers. -- R. W. Hamming

ABSTRACT

A financial model is a representation of the activities of a business in terms of quantitative relationships among variables that can help an analyst understand the financial consequences of past activities or assumed future activities. The equations comprising such models form a kind of knowledge base which can be used to generate explanations. In this paper we give some background on financial models, discuss two sorts of explanations in this domain, and present a procedure for explaining model results.

Introduction

"It is February 1974 and as President of the Battery Company you are a little concerned at the results for 1973 that you have just received. Despite a 20% increase in sales over 1972, profits have decreased by 1%.

You feel that the decrease in profit could be due to a combination of three causes: increase in overhead expenses, decrease in contribution (or profit) margins (difference between selling price and direct manufacturing cost) or a change in product mix toward less profitable units. Alternatively, you would like to know how the additional revenues from increased sales were spent. You would like to investigate the cause of the decreased profit using The Information System."

Thus began the statement of a problem that Malhotra gave to a number of managers and management students as part of his investigation into the utility and feasibility of an English language question-answering system to support management [6,7]. In order to determine the design specifications for such a system, e.g., the vocabulary, grammar, and types of questions it would have to deal with, an "ideal" system was simulated that was capable of "perfect" interpretation and response to naturally occurring questions and commands. Users could ask about what the system could do, what kinds of data it had, how computed values were derived, and what the data values were, either for a particular plant, product, customer and year, or aggregated over subsets of these, as the user's question required. The simulation was conducted by sending user inputs to another terminal where a human experimenter would interpret it and create responses on the user's terminal. The responses provided were those that Malhotra felt could be reasonably produced by a computer system, either because a simplified prototype he had developed could produce them or because they seemed to require only straightforward extensions to that prototype.

Malhotra's prototype embodied an early version of what have come to be called "financial modeling languages" [11] or "decision support system generators" [8]. Spreadsheet calculators, such as Visicalc [2], are simpler systems that also fall into this class. Although they lack a natural language interface, these systems allow users to interactively display data, aggregate it, compute functions of it (e.g. averages, percentages, ratios, etc.) and to define algebraic models that assist in business decision-making. Given historical data, the results they produce are similar to the figures that appear on financial reports. An example of a report generated for the Battery Company is shown in Table 1.*

These systems are not, of course, limited to only historical data. They can also generate hypothetical data, or projections, based on assumed data and expectations about the future. The first two columns in Table 1, for example, show historical data on Battery Company operations and the last three show projections.

However, neither Malhotra's natural language prototype nor more recent systems allow our president's question to be asked directly, to wit:

> *Why did profit go down in 73 even though gross sales went up?*

A little reflection on Table 1 may suggest other similar questions, such as,

> *Why do gross sales go up? in 75? in 76?*
> *Why does gross margin go up so little in 76?*
> *Why is there a peak in profit in 75?*
> *Why does unit cost go down in 74?*

These questions call for an explanation of results, not just a presentation of them, and the task of explaining results has traditionally been left to human analysts. The purpose of this paper is to show that, with suitable underlying models, generating such explanations by machine is not difficult and can be quite useful. The technique to be presented has been developed for use in the ROME system, a Reason-Oriented Modeling Environment for business planning managers [5].

The Explanation Problem for Financial Models

Financial Models

A financial model is a representation of the activities of a business in terms of quantitative relationships among financial variables. Financial variables are variables that have some economic or accounting significance and the relationships among them can generally be expressed by formulas and conditional statements.

*Due to space limitations, the data and model presented in this paper represent only a one-plant, one-product, one-customer version of the original Battery Company.

	1972	1973	1974	1975	1976
Volume	100.00	120.00	132.00	145.20	145.20
Selling price	35.00	35.00	36.40	37.86	39.37
Gross sales	3500.00	4200.00	4804.80	5496.69	5716.56
Labor/unit	9.00	9.00	9.36	9.73	10.12
Matl. price/unit	8.00	8.00	8.64	9.33	10.08
Material/unit	8.00	8.00	7.34	7.93	8.57
Shipping/unit	2.00	2.00	2.08	2.16	2.25
Unit cost	19.00	19.00	18.78	19.83	20.94
Variable cost	1900.00	2280.00	2479.49	2879.19	3040.42
Indirect cost	285.00	342.00	371.92	431.88	456.06
Production cost	2185.00	2622.00	2851.41	3311.07	3496.49
Gross margin	1315.00	1578.00	1953.39	2185.62	2220.07
Operating exp.	415.00	630.00	720.72	824.50	857.48
Interest exp.	0.00	0.00	0.00	0.00	0.00
Depreciation	35.00	35.00	35.00	29.00	29.00
Mgmt. salary	182.00	236.60	246.06	255.91	266.14
Overhead cost	632.00	901.60	1001.78	1109.41	1152.63
Profit	683.00	676.40	951.60	1076.21	1067.45
Profit margin	16.00	16.00	17.62	18.03	18.43

Table 1: Financial Model Results for the Battery Company

The time span encompassed by the model is normally divided into time periods and output is generated by computing values for each variable for each period and displaying the values of selected variables on a report.

There are three categories of formulas in a typical model. *Exact formulas* correspond to definitions and equivalences, e.g. "sales = volume * selling price" and "beginning inventory(period) = ending inventory(period - 1)". *Approximations* are essentially estimating relationships for endogenous variables, i.e. variables taken to be "internal" to the system of activities being modeled. These formulas are intended to yield the aggregate effect of (very) complex causal processes without actually simulating or even defining those processes. Examples include the use of historically-derived ratios to estimate one value from another and the use of cross-sectional regression equations. Finally, *predictions* are formulas used to estimate values for exogenous (external) variables, such as the price a firm must pay for its raw materials. All the numerous forecasting methods, such as growth rate factors, trend extrapolation, exponential smoothing, and the like, fall into this category. Table 2 shows the formulas used to generate the numbers in Table 1, grouped into the three categories.

Similarly, there are three kinds of input data to a financial model: *actual data*, *approximation parameters*, and *prediction parameters*. Actual data are historical, factual, non-negotiable numbers while the parameters are negotiable numbers, estimates, and assumptions. Approximation parameters appear in the approximation formulas and prediction parameters appear in the prediction formulas. Parameters for the Battery Company model are the constants that appear in Table 2.

If we think of a financial model as a kind of "knowledge base" from which we can "infer" (numerical) properties of business activities, we can make an analogy here with backward-chaining rule-based systems like Mycin. The formulas in a model correspond to rules and evaluating formulas corresponds to drawing conclusions. Rules change the degree of belief in propositions while formulas change the values of variables. The derivation of a value spawns a directed acyclic graph of subderivations much like the goal tree generated by backward chaining. The amount by which belief in a proposition changes can depend on judgmental factors and the amount of change in a variable can depend on judgmental parameters. Not to push the analogy too far, a rule-based system is is much more complicated since it depends on pattern matching and allows for more than one rule to contribute to the degree of belief in a conclusion. Nevertheless, the analogy suggests that the same explanation techniques that are used in Mycin [4] might also work for financial models. The next section shows why these techniques are inadequate for our problem.

Explanations

The purpose of an explanation is to make clear what is not understood. Depending on their initial level of understanding, users of financial models can benefit from two sorts of explanations. The first sort deals with the model itself and involves showing how it corresponds to reality and why that correspondence is justified. Such an explanation might include, for example, a description of what financial entity some variable represents and a justification for why some approximation was chosen to assign a value to it. The second sort deals with the results of the model and involves showing how those results were derived and why the derivation produces the results observed. In this paper, we focus on explaining results rather than explaining the model.

There are several kinds of results that we might want explained. First of all, there are the results that are explicit in the output report and are produced directly by formulas. It seems to us that explaining these is simple. To answer a question like *Why is operating expense equal to 724.84 in 74*, for example, we can imagine nothing better than a display of the associated formula and the the values it was used with. In other words, we interpret a *why* question about the value of a variable as a *how* question about it's derivation and show the derivation. A more difficult problem arises, however, if the user questions the formula, e.g. *But why does operating expense = .15 * gross sales?* Clearly, such questions should be answerable by giving the justification for the formula, or for its parameter values. But notice that this really calls for an explanation of the model: *why did the model*

Definitions

gross sales = volume * selling price	production cost = variable cost + indirect cost
gross margin = gross sales - production cost	profit = gross margin - overhead cost
variable cost = unit cost * volume	profit margin = selling price - unit cost
unit cost = labor/unit + material/unit + shipping/unit	
material/unit = matl. price/unit * (1 - volume discount)	
overhead cost = operating exp. + interest exp. + depreciation + mgmt. salaries	

Approximations

operating expense = .15 * gross sales	indirect cost = .15 * variable cost

Predictions

inflation = .04	interest expense = 0
depreciation(74) = 35	volume(74) = volume(73) * 1.1
depreciation(75) = 29	volume(75) = volume(74) * 1.1
depreciation(76) = 29	volume(76) = volume(75)
selling price(y) = selling price(y-1) * (1 + inflation)	
mgmt. salaries(y) = management salaries(y-1) * (1 + inflation)	
labor/unit(y) = labor/unit(y-1) * (1 + inflation)	
matl. price/unit(y) = matl. price/unit(y-1) * 1.08	
shipping/unit(y) = shipping/unit(y-1) * (1 + inflation)	
volume discount = 0 if volume <= 130, .15 if volume > 130	

Table 2: Battery Company Model

builder choose this formula/this parameter value to compute that variable? How to do that goes beyond our present focus.**

The other kinds of results are all implicit in the report and hinge on comparisons the user makes between values. The questions posed in the introduction ask about results of this kind and answering them involves explaining the difference. We can classify these kinds of results and their associated explanations along several dimensions.

1. **Referent of comparison.** All questions focus on a particular variable, which is the subject of the question sentence, but the referent it is compared to depends on the question. In a question about change, e.g. *Why did gross sales go up in 73?*, the referent is the value the subject variable had in a previous period. In questions of relative magnitude, e.g., *Why is depreciation so small?*, the referent is the user's expectation for the value of the focus variable. Otherwise, the referent is explicit, e.g. *Why is sales of product A > sales of product B?* In any case, the result to be explained is the difference between the focus value and the referent.

2. **Implicit referents.** There are two sources for a user's expectations about values which we will call "local" and "external". Local expectations come from the set of values observed on the report and are essentially local averages. So, for example, we interpret a question like *Why did gross margin go up so little in 76* as *Why is the change in gross margin small in 76 compared to the average of the changes in other periods?* External expectations come from a user's pre-existing knowledge of either analogous or prescriptive values. Analogous values include historical norms, industry averages, values observed for competing firms, and the like, while prescriptive values are goals (target values) the user knows to have been set.

** but see [9] for a technique that ought to apply if a financial modeler's knowledge could be suitably represented.

3. **Level of specificity.** A user may phrase his question in terms of mere difference (*Why did x change?*), direction of difference (*Why did x go up?*), or magnitude of difference (*Why did x go up so much?*). An explanation should take these different levels of specificity into account by referring to directions or magnitudes when the user implies he desires it.

4. **Interval to be covered.** A question may ask about a single difference (*Why does x go up in 74*), several differences (*Why does x go up in 74-76*), or all the differences (*Why does x go up?*). We interpret the latter questions as calling for a summary explanation that attempts to account for all the differences in the interval using the same factors. If that is not possible, we would like an answer to at least group similar explanations of individual differences into subinterval explanations and to indicate the contrast among the members of the set. Along the same lines, questions about peaks and dips seem to demand an explanation which covers the interval of inflection (at least two time periods) and accounts for the inflection by a single set of factors, or by a contrasting set.

5. **Violated presuppositions.** In general, a user may ask for an explanation of a result either because he simply wants to obtain the reason or because he can think of a reason to believe the contrary and wants to resolve the conflict. He can highlight the second case, however, by asking a *why not* question or using a contrastive subordinate clause, e.g. *Why did profit go down in 73 even though sales went up?*. It is then necessary to infer the presupposed relationship and to show in the answer why it does not hold for the situation at hand.

It may be seen that the major problem in explaining a difference does not lie in determining the difference of interest. Although a small amount of inference may be required to choose an implicit referent, and perhaps somewhat more to determine a presupposition, if these were problematic, one could simply ask the user to select among the possible interpretations. Nor is there a problem in showing the mathematical derivation of a difference. Rather, the problem lies in clarifying that derivation, which is the topic of the next section.

An Explanation Procedure

While it would be truthful to explain a model's results by exhibiting the formulas, the input data, and exclaiming "The math works out that way", it would not be clear. When we asked human analysts to explain model results, they tended to cite only the most important factors involved. What they did in answering specific questions gave us a set of goals for artificial explanations:

- distinguish the relevant parts of the model from the irrelevant
- distinguish the significant effects from the insignificant
- translate quantitative information into a qualitative characterization
- summarize if the same reason accounts for more than one result

General Strategy

To explain a difference, Δy, our general strategy is to first find a set of variables, A, which "account" for it and then to express that information to the user. Suppose, for the sake of simplicity, that we have a direction question -- *Why did y go up?* -- so that Δy is the change in variable y. The relevant part of the model is then the formula that computes y, say f, $\Delta y = f(a_2, b_2, c_2, ...) - f(a_1, b_1, c_1, ...)$ where the subscripts on the arguments denote the two different time periods, and $A \subseteq \{a, b, c, ...\}$. We first delete from S all variables that didn't change, since they clearly have no effect on Δy. Call the reduced set S*.

To determine A, we need to determine the "significance" of each variable in S* and collect the the smallest subset whose joint significance is sufficient to account for Δy. Our initial approach (the obvious one) was to loop through all possible sums of partial derivatives until nearly all of the difference had been accounted for. For example, we would stop with the single variable a if $(\partial f/\partial a)\Delta a \sim \Delta y$. This method turned out not to work because of two fundamental flaws. First, it assumes that the value of $\partial f/\partial a$ is nearly the same at both time points and this was not always true. When $\partial f/\partial a$ changes markedly from period 1 to period 2, there is no clear way of deciding whether it should be evaluated at a_1, or a_2, or perhaps some value in between. Second, it assumes that all the other variables in S remain constant, and this was rarely true. The result was that the above test would often fail on a variable that was significant and succeed on one that wasn't.

So we defined a new measure, called $\varepsilon(X,y)$, to indicate the *effect* of the set of variables in X on y in one context, such as one time period, relative to another. The general definition is

$$\varepsilon(X,y) = y_2 - f(Z)$$

where the vector Z contains the values of variables in X evaluated in context 1 and values for the other variables in S evaluated in context 2. If X contains just the variable a, for example, $\varepsilon(X,y) = y_2 - f(a_1, b_2, c_2, ...)$. Thus, $f(Z)$ gives the value y_2 would have had if all other variables had changed except those in X, and $\varepsilon(X,y)$ gives the amount of y_2 contributed by the change in the X variables. Restating this in words, we measure the effect of a variable by what the result would have been without the influence of that variable, leaving all other influences intact.

If the total effect is large enough for some X, we conclude that X = A and Δy is accounted for. The test we use is

$$1/\theta > \varepsilon(X,y)/\Delta y > \theta$$

where θ is the fraction of the difference considered large enough to be sufficient. The bound on the high side is needed when variables not in X counteract the effects of those in X. If the former effects are large enough, they should be included in X and so the test should fail. The value of θ was set empirically to be .75. We also associate with each variable x_i in A its relative effect on y, $\alpha_i(y)$, where $\alpha_i(y) = \varepsilon(x_i,y) / \Sigma |\varepsilon(x_i,y)|$.

When A is found, we can answer the original question with an explanation. In general, the answer given includes (1) the differences that account for Δy, (2) the formula f, (3) the primary explanatory variable, and a qualification, which expresses counteracting or reinforcing effects. What we do then depends on the specific form of the question and the contents of A, so before discussing that, it will be helpful to look at a specific example.

Details

Let us consider the first example from the introduction, *Why did profit go down in 73 even though gross sales went up?* The following describes the processing steps.

1. **Interpreting the question**. As outlined above, it is necessary to determine the focus of the question, referent of comparison, level of specificity, interval to be covered, and presuppositions. The ROME system uses a pattern-matching parser [3] to extract and label the parts of the input sentence, and a straightforward set of linguistic tests to make the determinations. For example, the verb or complement of the main clause establishes the type of comparison, and use of a time modifier indicator sets the interval to be covered. If the referent is implicit, we assume the expectation is local unless it is not satisfied by the data displayed, in which case we look for a global expectation. ROME allows the specification of external expectations for values, and their sources, and we use the first expectation found (if there is one) that has the right relation to the focus. In the question at hand, the focus is *profit* for the period *73*, the referent is *profit* for period *72*, and the level of specificity is *direction*.

To apply the explanation procedure, the focus and the referent must be comparable. In the present system, this means they must be computed by the same formula so that the difference in value arises from different contexts of evaluation. The contexts allowed are set by internal indices on the variables (e.g. time, plant, product, etc.) which range over different instances of entities of the same semantic type. The types are represented as elements in a semantic network using the frame-style language SRL [10]. If the variables are for some reason not comparable, a message is produced giving the reason.

Our treatment of presuppositions has not gone beyond the *ad hoc* stage. Currently, we just save the variables involved for later use in deciding when to stop the explanation, as described below.

2. **Identifying significant effects**. Since both gross margin and overhead cost change in the formula for profit, S* = {margin, overhead}. Working out the calculations gives

$$\varepsilon(\{\text{margin}\},\text{profit})/\Delta\text{profit} = 257.27/-6.6 = -38.98$$

$\varepsilon(\{\text{overhead}\},\text{profit})/\Delta\text{profit} = -264.6/-6.6 = 40$
$\alpha_{\text{margin}} = -.493$
$\alpha_{\text{overhead}} = .507.$

Since neither value of ε/Δ passes the significance test, both are needed to explain the difference (which the procedure discovers when it considers X = {margin, overhead}).***

3. **Characterizing effects qualitatively**. All the differences for variables in A are translated into direction and magnitude descriptors. The magnitude descriptors are normally percents since these are more familiar to users than our α values.

4. **Expressing the answer.** The answer generator is template-driven, where the templates are just those needed to express formulas, simple comparisons between variables, change, relative significance, reference to change, conjoined noun phrases, and contrast between propositions. The first sentence states either the most significant cause of the difference, based on the α values, or all the causes if they are positively correlated with the difference. For the question at hand, the first sentence is:

> *Profit went down in 73 primarily because overhead cost went up and profit = gross margin - overhead cost.*

The next sentence expresses the qualification, if any, such as a contrast among counteracting factors, a statment of primary cause, or a statement of additional cause. In this case, the qualification is:

> *Although gross margin went up by 19%, overhead cost went up by 41% and the latter outweighed the former.*

5. **Continuing the explanation.** Without the presupposition, the explanation would normally stop at this point with the message *Would you like me to continue?* However, since the answer has not yet mentioned gross sales, it has not yet been related to the presupposition, so we continue down the derivation path that leads to gross sales. The four previous steps yields the following continuation:

> *Gross margin went up in 73 primarily because gross sales went up and gross margin = gross sales - production cost. However, the increase in gross sales was not enough to affect the change in profit. Would you like me to continue?*

Notice that the qualification is one relevant to the violated presupposition, not the contrast between sales and production cost that would otherwise be generated. It is known that the increase is not enough because gross sales has a positive influence on gross margin and hence on profit but the change in profit was negative.

Continuing one step further will illustrate two final points. Since the presupposition variable has been mentioned, the system returns to the primary path:

> *Overhead cost went up in 73 because operating exp and mgmt salaries went up and overhead cost = operating exp. + interest exp. + depreciation + mgmt. salaries. The increase in overhead cost was due primarily to the increase in operating exp. However, the effect of mgmt. salaries was also significant to the change in profit.*

The first thing to notice is that the last sentence mentions an effect on the initial difference to be explained. It can happen that, while sufficient to explain a local difference, a particular set of variables A is not sufficient to account for a difference higher up in the derivation tree. That is, the higher level difference would not have been observed without the effect of variables that happen not to be significant to the local difference. Hence, the general strategy described above also includes a test for significance with respect to higher level variables and adds variables to A as required. These secondary "long distance" effects are important to an accurate explanation and would be missed by a purely local analysis.

Second, the explanation halts at this point since it has reached the leaves of the derivation tree. We define a leaf to be either an input value to the model or a result produced by a prediction formula. The latter case reflects our view that explanations are relative to models and that the representation of a model must distinguish the endogenous from the exogenous. This does not preclude explaining results from interrelated models, but simply breaks off explanations at model boundaries.

Discussion and Conclusions

We have presented an explanation procedure which couples an analytic technique with a natural language facility in order to explain differences between values of variables. The procedure seems to work well when:

- only a few variables out of many account for the difference to be explained
- the variables form a natural hierarchy via their formulas
- lower level variables and their values have a priori meaning to the user
- the complexity of the model comes from the depth of the derivation trees, not the complexity of the computations.

These conditions are well met in many financial models, and the procedure can be applied to a number of contrastable situations, such as actual vs historical comparisons, budgeted vs actual, and scenario vs scenario. However, the procedure offers no direct help in explaining iterative computations, such as those used in probabilistic models, discrete event simulations, or optimization algorithms.

Financial modeling is not the only domain in which this procedure might be used since there are many domains where at least some knowledge is encoded in quantitative relationships. The QBKG program [1], for example, uses a similar sort of procedure to explain the reasons for the backgammon moves it selects via a quantitative evaluation function. It may be seen in [1] that the form of this function and its use in the selection task approximately satisfy the above criteria. However, the test of significance used is very much different and the details of expressing the explanation are very specific to QBKG's particular function. Our procedure is more general but it does not incorporate any knowledge of what the value of a formula will be used for. The difficulty we see in using our technique as it stands to explain a heuristic selection lies in making all the terms and coefficients in the evalutation function meaningful to the user and in generating a meaningful characterization of the degree of difference in worth of different alternatives.

Financial models are intended to provide their users with insight

*** In the case of a no-change question -- *Why did y remain constant* -- a specialist procedure is invoked which looks for cancelling effects or the complete absence of change in the terms of the formula.

into the consequences of financial activities. It appears that automated explanation of the results can enhance that insight by focusing the user's attention on the major reasons for those consequences.

References

[1] Ackley, D.H. and Berliner, H.J., "The QBKG System: Knowldege Representation for Producing and Explaining Judgements," Computer Science Department, Carnegie-Mellon University, March 1983 (abridged in Berliner and Ackley, "The QBKG System: Generating Explanations from a Non-Discrete Knowledge Representation," *Proceedings of AAAI-82*, August 1982, pp 213-216).

[2] Beil, D.H., *The Visicalc Book*, Reston, Va.: The Reston Publishing Co., 1983.

[3] Boggs, W.M, Carbonell, J.G., and Monarch, I., "The Dypar-I Tutorial and Reference Manual," Computer Science Department, Carnegie-Mellon University, 1984.

[4] Davis, R., *Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases*, PhD Thesis, Computer Science Department, Stanford University, July 1976.

[5] Kosy, D.W., and Dahr, V., "Knowledge-Based Support Systems for Long Range Planning," The Robotics Institute, Carnegie-Mellon University, 1983.

[6] Malhotra, A., *Design Criteria for a Knowledge-Based English Language System for Management: An Experimental Analysis*, PhD Thesis, Sloan School of Management, MIT, 1975.

[7] Malhotra, A., "Knowledge-Based English Language Systems for Management Support: An Analysis of Requirements," *Advance Papers of IJCAI-4*, September 1975, pp 842-847.

[8] Sprague, R.H., Jr., "A Framework for the Development of Decision Support Systems," *MIS Quarterly* **4**, December 1980, pp 1-26.

[9] Swartout, W.R., "XPLAIN: A System for Creating and Explaining Expert Consulting Programs." *Artificial Intelligence* **21**, 1983, pp 285-325.

[10] Wright, M., and Fox, M.S., "SRL/1.5 User Manual", The Robotics Institute, Carnegie-Mellon University, 1982.

[11] *An Introduction to Computer-Assisted Planning Using the Interactive Financial Planning System*, EXECUCOM Systems Corp., Austin, Tx., 1980.