



Report 83-04
Stanford -- KSL

Scientific DataLink

A Report on FOLIO: An Expert Assistant
for Portfolio Managers.
Paul R. Cohen, Mark D. Lieberman,
May 1983

card 1 of 1

**Heuristic Programming Project
Report No. HPP-83-4**

May 1983

**A Report on FOLIO:
An Expert Assistant for Portfolio Managers**

Paul Cohen and Mark D. Lieberman

Department of Computer Science
Stanford University

Abstract

FOLIO is an expert system to assist portfolio managers. It interviews a client and, on the basis of expert knowledge, determines the client's investment goals and the portfolio that best meets them. For example, FOLIO may determine that one client requires a preponderance of tax-free investments and a substantial hedge against rising short-term interest rates, while another is best served by a mix of low-risk dividend-oriented stocks and intermediate-term bonds. FOLIO is a test bed for a theory of heuristic reasoning about uncertainty (Cohen and Grinberg, 1983), and its task has many parallels to established AI paradigms such as diagnosis in medicine and construction of a student model in ICAI domains (Barr and Feigenbaum, 1982). FOLIO uses a goal programming algorithm (Hillier and Lieberman, 1980) as a relaxation method for resolving the client's multiple goals into a portfolio that fits them optimally. This paper discusses the motivations for building FOLIO; its task, design, and operation; and examples of its recommendations.

Motivations for FOLIO

Modern portfolio theory (Rudd and Clasing, 1982; Sharpe, 1981) is currently applied by large institutions to problems of *asset allocation* for large clients, such as pension funds, as well as for small, individual clients. The central idea of these applications is to divide up the client's assets between several *funds* that have different characteristics. The relevant characteristics of the client are the client's desired risk and return; given these, an optimum allocation of assets can be found with a quadratic programming algorithm (Sharpe, 1970). Other characteristics of the client, besides those associated with risk and return, can enter into the problem but in practice are rarely considered. This may be because it is difficult to assess these factors accurately. The client's risk tolerance is uncertain, as is his or her tax status, asset structure, and investment goals. The information provided by the client is often incomplete and inaccurate, and sometimes it is even contradictory. Thus, no matter how sophisticated an asset allocation algorithm is, its recommendations are subject to the same suspicion as the data on which they are based. Garbage in, garbage out.

Our goal in the FOLIO project has been to develop an asset allocation program that is sensitive to the uncertainty implicit in client data. Our research takes two paths: first, our work on an expert system to interview a client and recommend a portfolio and, second, our development of a *theory of endorsers*, (Cohen and Grinberg, 1983) to facilitate reasoning about the uncertainty that arises during the interview. This paper discusses the expert system research.

We will digress briefly to discuss the theory of endorsers. The theory includes a representation of states of certainty that supports sophisticated reasoning about uncertainty, and a body of heuristic

knowledge to effect this reasoning. It is fundamentally different from numerical methods of reasoning under uncertainty, such as subjective Bayesian methods (Duda, Hart, and Nilsson, 1976), because it makes explicit the reasoning that is implicit (and hidden) in numerical degrees of belief and their combining functions (Shortliffe and Buchanan, 1975). For example, two pieces of evidence for a client's risk tolerance are age and the stock/bond ratio of the client's current portfolio. Age is but weakly correlated with risk tolerance while the stock/bond ratio is *direct* evidence for risk tolerance. Thus, age weakly suggests risk tolerance and can be contradicted by the stock/bond ratio, and the stock/bond ratio strongly suggests risk tolerance and *cannot* be contradicted by age. This asymmetry is important. It is maintained, in the theory of endorsers, by recording what kind of evidence supports a hypothesis about risk tolerance. If the evidence is age, and a contrary hypothesis is supported by the stock/bond ratio, then the latter hypothesis is preferred. This reasoning can be summarized by assigning a relatively small number—a *degree of belief*—to conclusions based on the client's age, and a larger number to those based on the composition of the client's portfolio. But if these numbers are to be used in lieu of the reasoning on which they are based—the reasoning that the stock/bond ratio contradicts age but not vice versa—then they can only result in less interesting reasoning; and if they are maintained in addition to the original reasoning, then they seem redundant. The theory of endorsers is based on the idea that we have reasons for preferring some kinds of evidence to others, and it provides a representation and heuristic methods for making the reasoning explicit.

Figuring out what the client needs and wants—what we call *client analysis*—has much in common with other AI problems, such as medical diagnosis and developing a student model for an ICAI program. In these cases, we seek an analysis of the subject's state of health or state of mind, given incomplete and uncertain data and inference rules. We believe that the theory of uncertainty developed in the FOLIO project and described in (Cohen and Grinberg, 1983) will be applicable to these general problems.

The Asset Allocation Problem

FOLIO's tasks are client analysis and asset allocation for the small investor. Small, in this context, means less than a million dollars. Asset allocation is a simplification of the more general problem of portfolio management, because the goal is to allocate a client's assets to one or more *funds* of securities, rather than to individual securities. FOLIO recognizes nine funds (or classes) of securities, listed in Figure 1. It infers the client's goals during an interview and, at the end, recommends percentages of each of the security classes in Figure 1. This composition is guaranteed to provide an optimum fit to the client's goals. FOLIO leaves the investment advisor free to instantiate its recommendations with particular securities. Since asset allocation involves shares in funds of

1. Dividend-oriented, lower risk stocks
2. Growth-and-yield, lower risk stocks
3. Growth-and-yield, higher risk stocks
4. Commodity-sensitive, lower risk stocks
5. Commodity-sensitive, higher risk stocks
6. A mixture of maturities of government and corporate bonds
7. A mixture of maturities of municipal bonds
8. Discount bonds
9. Money-market funds and other cash-equivalents

Figure 1: The classes of securities used by FOLIO

securities, FOLIO requires only aggregate knowledge about the properties of the securities in each fund; for example, it knows the average riskiness and rate of return for the entire fund, not for individual securities within it. This has three important consequences. First, the program need not know thousands of stocks and bonds individually and intimately. Second, since aggregate figures change more slowly than those for individual securities, FOLIO can be kept current without perpetually changing parameter values for large numbers of securities. Third, the responsibility for security analysis rests with the investment advisor, not with FOLIO. This is consistent with the investment philosophy of FOLIO's consulting expert,¹ who holds that proper diversification over securities within a class is less risky (preferable) to holding too few securities. The choice of which securities to hold, say, from the class of high risk growth stocks, is judged to be less important than the choice of that class for a client.

The Structure of FOLIO

FOLIO has three main components: a set of interview functions, a forward-chaining production system for inferring the client's goals, and a *goal programming* algorithm to maximize the fit of the client's portfolio to his goals. The interview functions are very simple, distinguished only by an editor to allow the client to change his answers. It was decided that, since most clients are asked the same questions, each should be asked all the questions at the beginning of the consultation. No effort has been made to ask questions on a goal-directed, "as-needed" basis. After questions are answered, and the answers edited, FOLIO derives several dozen numerical parameters and infers the client's goals. But we defer discussion of this stage until we have described what we mean by client goals and how they are satisfied.

¹Mark Lieberman

Maximizing the Fit to the Client's Goals

FOLIO uses a goal programming algorithm to maximize the fit of a portfolio to the goals that the portfolio is supposed to satisfy. Goal programming is a kind of linear programming in which the objective function is made up of terms describing the deviations from the target values of each goal. The result is a solution that minimizes the summed deviation from all the goals.

Fourteen goals, listed in Figure 2, received much discussion. Each goal is represented by five parameters: a *target* value that is more desirable than any other; a *penalty for exceeding* the target value, which is a monotonically nondecreasing function of the difference between the target value and the actual value; a similar *penalty for falling short* of the target value; a *lower bound*, below which the penalty becomes infinite; and an *upper bound*, above which the penalty becomes infinite. These parameters are described by four linear constraints on each goal:

$$F_1 + F_2 + \dots + F_g \geq \text{lowerbound}_{g1} \quad (1)$$

$$F_1 + F_2 + \dots + F_g \leq \text{upperbound}_{g1}$$

$$F_1 + F_2 + \dots + F_g \geq \text{target}_{g1} - \text{diff}^-_{g1}$$

$$F_1 + F_2 + \dots + F_g \leq \text{target}_{g1} + \text{diff}^+_{g1}$$

Where F_i is a number that indicates whether and to what extent fund i satisfies a given goal, in this case, $g1$.

Assuming that $\text{lowerbound} \leq \text{target} \leq \text{upperbound}$ (which is always the case), if a portfolio could be found such that diff^- and diff^+ were equal to zero for every goal, that portfolio would be a perfect fit to its goals. This rarely happens, however, so, instead, one might try to find a portfolio that minimizes the sum of diff^+_g and diff^-_g for all goals g . In fact, an additional embellishment allows us to emphasize some goals over others:² We can multiply diff^+_g and/or diff^-_g by a number that indicates the importance of achieving goal g , and try to minimize the following sum:

$$(\text{diff}^+_{g1} * p^+_{g1}) + (\text{diff}^-_{g1} * p^-_{g1}) + \dots + (\text{diff}^-_{g14} * p^-_{g14}) \quad (2)$$

In fact, p^+_{gi} and p^-_{gi} are just the respective penalties for exceeding and falling short of the target value for goal i . Equation 2 is minimized with a modified, two-phase simplex algorithm (Hillier and Lieberman, 1980; Dantzig, 1963). Parts of the algorithm are written in Franz Lisp (FOLIO's "native

²We are grateful to Professor William F. Sharpe, of Stanford University, for suggesting this model for FOLIO

language"), although the array pivoting procedures are Pascal "foreign functions." The algorithm takes about five minutes to produce a portfolio for a client with three or four accounts on a lightly loaded DEC 11/780.

1. Hedge against rising short term interest rates
2. Hedge against falling short term interest rates
3. Hedge against rising long term interest rates
4. Hedge against falling long term interest rates
5. Hedge against rising inflation
6. Hedge against rising deflation
7. Tax free interest
8. Taxable interest
9. Dividends
10. Capital gains
11. Capital preservation
12. Risk to interest
13. Risk to principal
14. Rate of return

Figure 2: The goals that FOLIO considers for each client

Having discussed how FOLIO derives a portfolio from a description of the client's goals, we now turn to the problem of inferring those goals from data collected during the interview. Note that each of the 14 goals is described by 5 parameters, for a total of 70 goal-statements. Examples are:

Goal: Interest

lower bound = 3%
 upper bound = 7%
 target = 5%
 $p^+ = 100$
 $p^- = 0$

Goal: Hedge against increasing long term interest rates

target = 30%
 $p^- = 300$

The first case illustrates a complete specification of a goal for interest income. Evidently, the emphasis is on achieving a relatively low rate of interest, although less than 3% is unacceptable. The target is 5%, and the penalty for exceeding it is equal to 100 times the difference between the target value and the actual value. There is no penalty for failing to achieve 5%. The second example illustrates a partial specification of the a goal for hedging against increasing long term interest rates. In this case, 30% of the positions in the portfolio should comprise such a hedge, and the penalty for

failing to achieve this is fairly stiff: 300 times the difference between the target and actual values. (Penalty coefficients run from 0.0 to 1000.0 in FOLIO.)

Heuristic Rules for Inferring Goals

FOLIO infers a client's goals with heuristic rules. These are represented in MRS (Genesereth and Smith, 1982), a predicate-calculus-like language that allows for multiple ways to store and infer propositions. MRS also has some default inference methods, such as forward- and backward-chaining; and FOLIO uses the former. Figure 3 shows some examples of these rules ineptly translated into English.

IF	the client needs a relatively large sum of money in less than a year,
THEN	set the lower bound for the goal of "preserve capital" to produce the needed amount.
IF	the client needs a relatively large sum of money in more than a year and the client's risk-to-interest measure is greater than or equal to 2.0,
THEN	set the target value for "dividends" to produce part of the needed amount and set the target value for "interest" to produce the rest.
IF	the client's tax bracket is over 30%,
THEN	set the upper bound on "interest" to produce the minimum needed income and set the target value on "interest" to zero.
ALWAYS	set the target value on hedges against inflation to its maximum.
IF	a client is on a fixed income
THEN	set the penalty for not achieving the maximum hedge against inflation to be 500

Figure 3: A few of FOLIO's heuristic rules for inferring goals

In general, the conditions of FOLIO's rules test some aspect of the client's assets structure, risk status, or tax bracket. Most conclusions specify one or more goal statements. One of the rules in Figure 3 has no conditions. FOLIO has eight such rules. Six of the rules set the target values for each of six hedges to the maximum, one sets the target value for total rate of return to the maximum, and one ensures that every client gets at least enough interest income to exploit the dividend-exclusion clause in the federal income tax.

The first of the rules in Figure 3 asks whether the client needs a chunk of cash in less than a year. If so, it is the judgment of FOLIO's consulting expert that the necessary sum should be invested in a

fund that preserves capital (in practice, a money market fund). Note that, because a lower bound is set in this example, the program will try to guarantee that at least the needed amount is invested in money market funds.³ This rule is thus quite conservative: The client is virtually guaranteed that the necessary amount of money will be protected against capital losses. The next rule is less conservative and is not used unless the client is willing to accept some risk. It says that, if the sum in question is needed in more than a year and the client is moderately risk tolerant, then some of the assets should produce dividends, and the rest should produce income. FOLIO uses a simple method (not illustrated in the English translation of the rule) to determine how much of each should be produced: In this case, the amount of assets that should be invested in dividend-producing positions should be proportional to the client's tolerance of risk to principal. If the client will tolerate a lot of risk, then most of the chunk of money that was mentioned in the condition of the rule will be invested in stocks. If the client is intolerant to risk, most of the money will be invested in bonds.

The third rule captures the concept that a person's portfolio should not generate more interest than he (or she) needs. Interest is taxed at income-tax rates, while capital gains are taxed more favorably, thus, the portfolio of a client in a tax bracket over 30% must never produce more interest than is absolutely necessary, and should ideally produce very little interest. This is an appropriate time to mention that FOLIO works with several different accounts simultaneously, if the client has more than one account. FOLIO would recommend that a client's portfolio produce interest in the account that is taxed the least; indeed, it encourages the client by its recommendations to seek as much interest as possible in tax-free and tax-deferred accounts. By similar reasoning, FOLIO likes to buy preferred or dividend-producing stocks with the assets in a corporate account, to take advantage of the fact that U.S. corporate accounts only pay tax on 15% of the dividends they receive from other corporations.

The fifth rule notes that a client on a fixed income is in special need of hedges against inflation. The penalty for failing to achieve the maximum possible is quite stiff.

The current version of FOLIO uses about 55 of these rules to infer one or more of the five parameters of each of 14 goals. Many rules conclude about all five parameters of a goal. It is possible to have multiple values for a goal parameter; for example, the target value for the "preserve-capital" goal might be set to both 2.0% and 7.0% by different inference rules. Currently, such conflicts are

³In practice, if contradictory bounds are set, the linear programming algorithm will be unable to find a feasible solution. For example, if we require that 80% of a client's assets produce tax-free interest, and 60% produce capital gains, and 40% produce taxable interest, the LP will blow up because 180% of the client's assets are being required for three mutually exclusive goals. Of course, it is perfectly permissible to want 80% of the assets to produce dividends and 80% to produce capital gains, because these aren't mutually exclusive. In general, FOLIO's rules set target values and penalty functions instead of bounds to avoid the problems inherent in setting inflexible bounds.

resolved by some 20 heuristic rules; for example, if there are multiple values for p^+ for a goal, the maximum of the values should be chosen, we anticipate the conflicts to be resolved by selecting the value in which FOLIO has the certainty, but this awaits the implementation of the theory of endusers.

Note that the rules' conclusions do not specify any of the funds in Figure 1, but, rather, they refer to the goals in Figure 2. The goal programming algorithm then makes a selection of funds that best satisfy the goals. For example, in the second rule in Figure 3, we noted that the client was willing to tolerate a certain level of risk. But rather than specifying a proportion of stocks from one of the first five rows of Figure 1, we simply asserted the goal of producing dividends. The linear programming algorithm is then responsible for deciding which kinds of stocks are both sufficiently risky and produce the appropriate amount of dividends. Depending on how much risk the client can stand, he would probably end up with yield-oriented stocks or less risky growth-and-yield or commodity-sensitive stocks. But the decision is made automatically.

Summary and Future Research

FOLIO is an expert assistant for portfolio managers. It performs client analysis and suggests allocations of clients' assets to a small number of asset classes. It has three conceptually distinct parts: The first conducts an interview and produces many useful parameters (such as tax bracket, and proportion of current investments in equity). The second part uses expert, heuristic inference rules to infer the client's goals. Lastly, FOLIO uses an optimizing technique to ensure that the final portfolio comes as close as possible to satisfying the client's goals. FOLIO's recommendations are in terms of the classes of investments that were shown in Figure 1. The investment advisor who uses FOLIO then implements these recommendations by buying securities that belong in the recommended classes.

FOLIO is as yet untested, although it is currently being run on extant accounts. The program is immune to a problem that we believe troubles human advisors, namely, forming a preconception about the client and then ignoring interview evidence that is contrary to the preconception. We expect that a program like FOLIO could significantly increase an advisor's efficiency, and, of course, it is a natural front-end for an electronic filing system for storing client information. Running alone, overnight perhaps, it could review old accounts in the light of current economic circumstances and alert the advisor to accounts that may need adjustments in their portfolios.

References

- Barr, A. B., and Feigenbaum, E. A. 1982. *The Handbook of Artificial Intelligence, Volume II*.
Los Altos CA.: William Kaufmann, Inc.

- Cohen, P. R., and Grinberg, M. R. 1983. A Theory of Heuristic Reasoning about Uncertainty. HPP Memo 83-5. Department of Computer Science, Stanford University.
- Dantzig, G. B. 1963 Linear programming and extensions. Princeton University Press. Princeton, NJ.
- Duda, R. O., Hart, P. E., and Nilsson, N. 1976. Subjective Bayesian Methods for Rule-Based Inference Systems. Technical Note 124. Artificial Intelligence Center, SRI International.
- Genesereth, M. R. and Smith, D. E. 1982. Meta-Level Architecture. HPP Report 81-6. Department of Computer Science. Stanford University. Stanford, CA.
- Hillier, F. S., and Lieberman, G. J. 1980. Introduction to Operations Research. San Francisco, CA.: Holden-Day, Inc.
- Rudd, A. and Clasing, H. J. 1982. Modern Portfolio Theory. Homewood, Illinois: Dow Jones-Irwin.
- Sharpe, W. F. 1970. Portfolio theory and capital markets. New York: McGraw-Hill
- Sharpe, W. F. 1981. Investments. Englewood Cliffs, NJ.: Prentice-Hall.
- Shortliffe, E. H. and Buchanan, B. G. 1975. A model of inexact reasoning in medicine
Mathematical Biosciences 23, 351-379

Copyright © 1985 by KSL and
Comtex Scientific Corporation

FILMED FROM BEST AVAILABLE COPY