The Knowledge Engineering Review, Vol 7:3, 1992, 215–249

Knowledge systems for financial advice

V. MITAL

Centre for Computers in Law and Finance, Brunel University, Uxbridge UB8 3PH, UK

Abstract

The financial domain has seen considerable knowledge systems activity. Some of the efforts have resulted in deployable systems, but many others have met with much less success. Commentators have attempted to discern broad, universal indicators which explain or predict success. More recently, however, it has become clearer that a principal cause of the difficulties encountered is the incongruence between the real world task and the system architecture employed by developers. Proceeding from the latter perspective, this paper concerns itself with knowledge systems for the provision of investment related financial advice. The various tasks involved are differentiated and correlated to the system architectures employed by developers. It is seen that while simple rule-based systems suffice for some tasks, others indicate the need for techniques such as Case-Based Design.

1 Introduction

From the knowledge system developer's point of view, the financial domain comprises a variety of activities, some of which have only a tenuous relationship with others. Thus, the domain encompasses brokerage and trading of financial instruments; short- and medium-term prediction of the movement of market indices; monitoring of transactions for compliance with the regulatory framework; insurance underwriting; loan authorization; handling personal and corporate financial affairs, bearing in mind tax implications; decision-making within corporations in respect of operational objectives; and managing and optimizing of investment portfolios. Very many financial knowledge systems have been prototyped, a substantial number have been deployed in the field, and several have resulted in sustained use by a significant number of professionals.

Perhaps the best example of widespread knowledge systems activity involves risk analysis, principally insurance underwriting and assessment of loan applications. Some such applications are fairly routine, or have been essentially reduced to statistical tables, e.g. life insurance for the general public and applications for loans for household items. However, the situation is rather different at the top end of the market. Making the decision on insuring a timber factory or giving a loan for the expansion of a sizeable business involves highly skilled experts. A very large number of factors is potentially relevant, though the expert can usually reduce this fairly quickly to a salient few by the application of specialized knowledge. It is possible to increase profit margins, or reduce the possibility of risk, by providing knowledge systems to assist or, sometimes, take over part of the manual analysis and decision-making.

Of the risk analysis systems that have been developed for insurance companies and banks, several have been deployed in the field and used by experts as tools during the performance of their professional tasks. The best-known systems in this category are perhaps those based on the non-procedural dataflow language, SYNTEL (Duda et al., 1987). Some systems, on the other hand, have not quite made it to line tasks, and have either been abandoned or have metamorphosed into training aids (Butler & Chamberlin, 1988).

Given the level of resources committed to the developments, the failure rate may be considered to be high. This has led some commentators to regard the whole field with great scepticism

(Stevenson, 1989), while others have tried to discern universally applicable factors differentiating failures from successes. In the latter vein, Coats (1988) looked at the problem types tackled by financial knowledge systems, and concluded that success requires that there be a single low-level closed domain with static rules, a task that takes an expert from a few minutes to a few hours, and that experts are routinely able to teach the relevant skill to neophytes, thereby indicating that they are accustomed to explaining themselves. There is no need to challenge such hypotheses, because an application of the type indicated by Coats will very likely have relatively easy knowledge acquisition and verification and, possibly, require the simpler of knowledge-base design techniques.

More recently, it has been realized that universal problem descriptors are predictive of performance only when it is assumed that the architectural tools and techniques available to the developer are limited to, say, the basic production-rule based type. With the richer toolset now in the armoury of developers, it is necessary to go behind the broad problem characteristics and look at the precise nature and structure of the task(s) tackled by a knowledge system. It is realized that many of the problems with systems arise due to the incongruence or mismatch between the task structure on the one hand, and the knowledge-base architecture and representational formalism employed on the other.

Consequently, more recent analyses of financial knowledge systems have given greater emphasis to detailing the nature of the application area and, more specifically, the particular tasks in question. This is the approach taken by Freedman (1991) for data-intensive systems for trading in financial products and commodities; Denna et al. (1991) do the same for systems for auditing.

This paper will concentrate on knowledge systems in an area very different to those mentioned above, though no less commercially important. The area covers various kinds of investment related financial advice and planning, often with important tax implications. The basis for evaluation and assessment will be presented by considering the different kinds of tasks which the various systems seek to assist or carry out. The nature of the tasks will be examined in detail, and it will be seen how the system architecture varies correspondingly.

2 Financial advice: the commercial context

Advice is available to the non-professional investor from a number of sources and in various forms. Service is offered under many labels such as financial planning, estate planning, investment advice and tax advice. The labels are ill defined and the more prestigious among them, for example finanical planning, are often used to give an added cachet to a relatively rudimentary service[1].

To consider the application of knowledge systems to the financial services industry, and by the traditional advisory professions, including accountants, attorneys, solicitors and actuaries, we have to go beyond the labels and be aware that the tasks to be supported differ considerably. Take the following situations:

- An advisor recommends to an individual or business organization the purchase of a particular financial service ('product') as being tax-efficient, capable of providing a desired level of income, growth or protection, or otherwise suitable in the client's circumstances. We will call this task *product selection*.
- An adviser assesses some of the circumstances of a client who has either a lump sum or a recurrent amount to invest, with a view to finding out what general types of transactions are most suitable. If one of these transactions involves the purchase of a product available from a financial institution, the advice may be made more specific (the adviser may or may not be tied to a

[1]The labels can occasionally have legal or regulatory significance. Some states in the USA have considered extending the licensing requirements for investment advisers and the regulatory regime of securities commissioners to all persons calling themselves 'financial planners', irrespective of the service offered. The Investment Advisers Disclosure and Enforcement Act of 1990 proposed to regulate the titles 'financial planner', 'financial consultant' and 'financial adviser'.

Table 1 Premium on expertise

Market change	*Effect on organizations*	*Result*
Competitive deregulation	Cover more market sectors	Shortage of experienced staff
More demanding clients	Differentiate self by offering more varied (hence complex) financial services	Professionals need better training
Regulation of conduct of business	Guard against legal liability and regulatory curtailment of business due to substandard conduct	Professionals need aids and assistance

particular institution[2]). Acceptance of the advice could lead to the sale of the relevant products. We will call this *investment advice*[3].

- A professional undertakes a review of the entire financial and relevant personal circumstances of a client. He identifies needs such as minimizing the tax liability of the estate on death, differing the payment of tax on capital gains or setting up a compensation scheme for the principals of a small business. To achieve the objectives, the adviser produces a coordinated scheme of temporally ordered transactions. Where most of the required transactions can be satisfied by using standard financial–legal devices and off-the-shelf financial products, we will refer to the service as *financial planning*[4].

The product selector and the investment adviser may, as mentioned, either be tied to a particular financial services organization or hold oneself out as being independent, able to offer products from multiple sources. In either event, the execution or implementation of transactions pursuant to recommendations is likely to be effected by the recommender. The financial planner is less likely also to purvey the products associated with his advice: a broker or other specialist would be commissioned, instead. However, commercial pressures are bringing into the financial planning arena those who are ultimately interested in vending products, rather than deriving an income directly from the provision of a planning service. This is naturally affecting the traditional advisory professions such as lawyers, accountants and actuaries.

The financial services industry itself has undergone a major deregulation, aimed at removing minimum prices, eliminating controls on the kinds of services that a financial services organization may provide, and lifting restrictions on where an organization may operate[5]. Personal and corporate incomes have risen increasing the potential clientele, but steep fluctuations in the financial markets and volatility in economic confidence has made clients more discriminating. As a counterweight to competition-related deregulation, there is increasingly stringent and interventionist regulation, or self-regulation, of the manner in which service is provided. The net result is summarized in Table 1.

It is possible to obtain a commercial edge in a crowded market by offering a more competitively priced or sophisticated (read 'more complex') financial product. Another route to the same end is more indirect: attracting and keeping clients by increasing the standard of the *collateral* service or advice offered by the organization's representatives or employees when vending the products. As

[2]Under the (self) regulatory regime set up pursuant to the Financial Services Act 1986 in UK, a tied agent normally has to inform a prospective client of the fact that he is authorized to advise in respect only of the products of the relevant organization (and closely related organizations).

[3]On occasions, the recommended transactions may have no direct investment element, e.g. term assurance policies, though they may indirectly free other asets or income for investment.

[4]There are many associated or specialized aspects of this service: 'estate planning', 'investment planning', 'educational fees planning', etc.

[5]Perhaps, one should not talk of deregulation in the past tense: it continues apace. For instance, the latest proposals from the US administration may lead to commercial banks moving into areas previously the preserve of investment and insurance organizations.

financial institutions move into areas previously dominated by the advisory professions, the latter too must improve or widen their service to stay competitive.

One outcome of the commercial pressures has been a resort to knowledge systems that aim to increase the effective skill level at the point of service. In fact, of all the areas in finance, this is where knowledge systems have been most notably successful[6]. The current commercial emphasis is on overcoming many of the limitations inherent in the designs of existing systems and obtaining greater functionality, the byproduct being a significant commercial advantage for organizations which are then able to exploit the under-served sectors of the market.

In subsequent sections we will look at the principles involved in developing knowledge systems for providing or assisting in the provision of some of the services mentioned above, namely, product selection, investment advice and financial planning[7]. We will look in some detail at the relevant tasks, pointing out the commercial significance and technical implications of the distinctions between product selection and investment advice *inter se* and between these and financial planning. This will allow a better appreciation of the various approaches taken and make starker the limitations.

We start by characterizing the concept of 'reasonableness', as this not only affects how the efficacy of task performance must be judged, but also places in perspective the kind of knowledge that is obtainable from experts.

3 The concept of reasonableness

In this section we refer to financial planning, investment advice and product selection collectively as *financial advice*. The financial adviser's decision-making aims to achieve some real world qualitative goal, for example:

- suggest an investment which minimizes potential capital gains tax liability while achieving moderate security of investment, or
- identify a product that will result in sufficient liquidity to pay off the liability in respect of estate taxes in the event of the death of the client at any time within the next 15 years.

If the result of the decision-making process does not adequately achieve the qualitative goal, loss may be caused to a client. If the client challenges the decision on the grounds of inefficacious performance of service—whether under the law relating to negligent performance of service, contract, or breach of some statute—the fact of loss being caused is generally insufficient to conclude liability (Mital & Johnson, 1990). The path taken to make the decision has to be investigated to see if professional skills corresponding to a certain standard were employed during the decision-making.

As no direct test of financial advisory skills is generally available, a *notional* peer review (termed 'the reasonableness test') is applied to see whether the decision-maker did what reasonably competent professionals of equal standing would do in similar circumstances. Admittedly, there is some doubt as to whether the comparison is with what a professional would actually do in practice, or with what the court thinks he ought to do. For practical purposes, however, the distinction is frequently blurred, as it is the court itself that decides who is reasonably competent, and that decision may or may not select practioners of average competence. The reasonableness test is a

[6]An example is the Advisa system, available from Datasure Ltd., Cheltenham, Glos. in various versions, which is used by more than one hundred financial services organizations. It was conceived and specified by Mital jointly with Nigel Chambers, a financial product development actuary and the then technical director of The Moorgate Group plc; Mital also led the design and prototyping. The system covers the full range of products likely to be offered by insurance companies and mutual fund organizations, and helps ensure that advisers act in accordance with the self-regulatory regime of the Financial Services Act 1986.

[7]We will not consider the open-ended services that a tax attorney or accountant may provide to a very wealthy client or a large business organization. The exclusion is made because it is extremely difficult, if not impossible in the state of the art, to build knowledge systems that can cope with the wide range of circumstances faced in these situations and the variety of solution options and room for creativity available to these professionals.

complex concept, varying substantially in its scope and stringency according to the legal and factual context in which fault-based liability is sought to be established. Nevertheless, we need to characterize certain notions and draw out some general principles, albeit at the expense of legal rigour. The decision-making typically involves reasoning about:

(i) real world facts, including personal or corporate affairs;
(ii) the relevant law and its applications to the current circumstances; and
(iii) financial matters including properties of different financial instruments and macroeconomic factors such as inflation and prospects in the concerned sectors.

The reasonableness test is stringent in respect of skills employed in obtaining and analyzing the real world facts of the situation. Widespread disagreement between experts is not really countenanced. For instance, in a case certain written instructions were given to the policy broker by a client regarding where his ship would be calling. The question was whether the broker should have known from the routing instructions that it was necessary to obtain from the insurer a specific endorsement covering port calls at all the Canary Islands rather than just Tenerife. It was said[8]: "If nine brokers of experience out of ten would have done the same as the Defendant under the same circumstances, or even if as many out of a given number would have been of his opinion as against it, he who only stipulates to bring a reasonable degree of skill to the performance of his duty, would be entitled to a verdict in his favour."

The test for reasonableness is much less strict when judging skills in applying legal knowledge to world facts. An investment adviser is not usually a lawyer or an accountant and is not expected to know all the intricacies of the debatable points of law. While he is expected to be familiar with the well-settled law, where there is genuine doubt it is normally not unreasonable for him to choose a plausible interpretation.

Much has been written about knowledge systems for legal reasoning and the great difficulties encountered en route to building practical systems (Gardner, 1987). However, what we have here is *conduct-oriented* reasoning, or reasoning aimed at determining the course of conduct of a financial professional, rather than contentious or adversarial reasoning, which is what is more usually meant by legal reasoning (Mital & Johnson, 1992).

In the present context, the concept of reasonableness is the defining characterizer of conduct-oriented legal reasoning. The effect of reasonableness on conduct-oriented legal reasoning is profound. No doubt, there are 'hard cases' (Gardner, *op cit*), but the imperative need to make a commercial decision is recognized by the law. Therefore, professionals are allowed, within limits, to choose any one of several plausible interpretations of legal rules. They are also permitted to abstract plausible 'rules' from cases for daily commercial application. The heuristics in the form of interpreted legal rules, and the defeasible 'rules' abstracted from cases can be supplemented by hypothetical examples and illustrations to fill the gap in respect of those situations which have not previously been authoritatively interpreted by courts. As such, if one can anticipate a potentially 'hard case', which admittedly is not easy to do in an exhaustive manner, one can generally obtain heuristics to deal with it.

As to knowledge in respect of financial matters, organizations have different market stances and perspectives, and their interest in the outcome of problem solving in identical factual situations can vary fundamentally. For instance, both collective investment houses and life assurance organizations may offer the same type of financial services, but their underlying investment spreads and policies may differ, colouring their view of what is better. The law recognizes commercial realities to an extent and hence, ordinarily, it is sufficient that a professional's view of the financial environment conforms to that which may be held by a *significant but 'respectable' minority* of professionals.

Certainly, the reasonableness test gives considerable room for manoeuvre to the expert. It is not necessary to find the 'most optimal' solution, whatever that might mean, nor always to reconcile the

[8]Chapman v. Walton (1833) 10 Bingham 57.

views of multiple experts. This is undoubtedly one of the causes of the relative computational tractability of financial advisory knowledge. However, the reasonableness test should not be thought of as a licence to err.

Around the solutions that are considered to be acceptably achieving some qualitative goal, there lies a class of purported solutions that are legally defensible, but only in the sense that it is difficult to conclude that they are actually *unreasonable*. This latter class can be called the 'defensible penumbra'. But it is not always possible to avail of the defensible penumbra. It is essential that a financial advisory professional is not biased and considers all the relevant facts of the situation before him.

In the current state of the art, knowledge engineers are constrained into representing rule-like heuristics in such a manner that a *subset* of facts is sufficient to lead to the selection of a certain solution, or at least suggest a particular hypothesis for further exploration. It is obvious that, in the case of a knowledge system, it is trivial to show that the embodied heuristics are such that whenever a particular *subset* of potentially relevant facts in encountered, it invariably leads to a certain solution which does not achieve the qualitative goal but is within the defensible penumbra. However, if such a challenge were to be made in the case of a human professional acting without the aid of a knowledge system, it would be necessary to give evidence of the pattern of the professional's behaviour over a number of transactions—courts cannot access the supposed internal workings of the human mind (Mital & Johnson, 1990).

Having characterized the concept of reasonableness, one can appreciate the significance of the knowledge and reasoning involved in specific applications. Thus, we have laid the foundations on which details may be built of knowledge systems for the various tasks referred to above. Next we elaborate on the relevant distinctions between the different tasks and how the distinctions have not only commercial but also technical implications.

4 Distinguishing product selection from investment advice

Differentiating investment advice from product selection can be quite difficult. At one level, the two are not fundamentally different. The ultimate result can be the same: the client deciding whether or not to carry out certain transactions involving financial products. The motive of the adviser may also be the same: receiving a commission from the product-vending institution. It is the scope of the reasoning involved in the two processes that separates one from the other.

The investment adviser analyzes the client's situation, formulates or acquires some goals, and then recommends whatever transactions 'appropriately' achieve those goals; further, under certain circumstances, the adviser may be undertaking to recommend those transactions which '*most* appropriately' achieve the goal, increasing the scope of the task. It may be that no transaction which provides a commission to the adviser is among those which should be recommended. Aternatively, if the adviser limits his repertoire to certain products, there is a strong probability that a 'no purchase' is recommended.

The product selector works from a different perspective. His paramount aim is to find whether any, and if so which, out of the list of products he vends is appropriate (or most appropriate, if that be the task undertaken) in the client's circumstances. Essentially, goals provide the adviser with something to pin a recommendation on. It may be that a 'no purchase' recommendation is made, but this is because no helpful goal can be found to apply.

Certainly, the distinction is subtle. But it is also a crucial one to appreciate. Let us elucidate by means of a dispute reported in the newspapers.

4.1 An illustration of the distinction

Keith Douglas had saved £6687 towards his 11 year old daughter's school fees when, in early 1987, he approached James Petter, a financial adviser tied to the renowned financial institution, Hill Samuel (Hunter, 1991). Douglas' money was lying in a building society—similar to Savings & Loan

institutions in USA, but perceived as rather more secure. The payment of fees was to start one year hence and was required for a total of four years. Petter suggested that a better return would be achieved through a Hill Samuel school fees plan. The underlying portfolio of the plan involved investing in the equity of small companies. Douglas followed the advice. Then came the October 1987 stock market crash which reduced the value of the investment considerably, so much so that the amount in 1991 stood at less than the £5000 even though not a single payment of fees had been made. There were a number of other alternative transactions which could have been recommended. For instance, the following, neither of which were among the products apparently available from Hill Samuel:

- As the school had already been chosen, the money could have been given to the school as a 'composition fee'. The school would have bought an annuity with the money, designed to offset the fees. The school's charitable status would have obtained a tax break on the income element of the annuity.
- Investment in an education trust would have led to a somewhat lower return than the then bull market was forecast to provide. However, the income would have been tax free and guaranteed.

Douglas has sued Petter for professional negligence. It is argued for the defence that Douglas knew of the nature of associated risks, and the recommendation was appropriate because the investment period was five years and the stocks were forecast to make overall gains in that period.

Even without knowing the outcome of the case, the matter is of illustrative value to us because the distinction between investment advice and product selection becomes clear when one looks at the defence put forward by Petter, one who is, arguably, providing the latter service. Clearly, the defence does not appear to say that the client's circumstances could best be satisfied by a small companies fund. There does not appear to be any attempt at comparing the selected product with the available alternatives (Hunter *op cit*). Instead, the contention is that the particular product could (reasonably) be classified as being suitable or appropriate. Which, one could say, is rather a different matter. In product selection, showing reasonableness of the decision process requires little more than the selected product is itself appropriate; the merits may be considered virtually without reference to other competing products.

This has implications for the development and utilization of knowledge systems. It makes system development relatively easy, but at the expense of limiting the commercial utility or advantage the system otherwise provides.

4.2 Commercial and technical significance of distinction

A financial services institution with a large, motivated, direct sale force and a strong, distinctive, brand image can possibly rely on offering product selection alone. Such an organization finds it easy to attract or otherwise approach prospective customers. Sometimes, a client already knows the broad nature of the transaction desired. It is then only a question of verifying that a product is available from the institution which will satisfy the requirements. At other times, a fact finding and needs analysis session has to be conducted by the institution's representative before this can be done. The crux of the matter is that the prospective customer knows—or, strictly speaking, should be made aware—that he is being *sold* one or more of a limited range of products.

Even if the adviser does not represent a single product-vendor and purports to provide independent advice, it is difficult to break out of the buyer–seller relationship with its associated, constraining, ethos. Things take a quite different perspective if the potential investor understands that the advice will be open ended and will consider all his needs, whether or not they can be satisfied from inventory. Furthermore, that the advice will be portable, in that the information given will be useful even out of the adviser's office and even in the context of somebody else's products.

The competitive market is making many organizations which habitually relied on product selection to move to offering investment advice. Where the adviser is able to offer only one

organization's products, it is true that, by its very nature, investment advice in the above manner will lead to more 'no purchase' recommendations. However, it is becoming increasingly clear that offering the more open ended service *eventually* results in the advice being accepted by a higher proportion of those advised and, probably, in attracting other potential customers.

Unfortunately, not all product sellers have the skills to make the assessments required for competent investment advice. High staff turnover and other factors may mean that adequate training is not always available before the persons have to start dealing with customers. Further, the standard of service that needs to be offered under the label of investment advice is becoming ever higher as expectations rise. There can be little doubt that knowledge systems have a most important role to play.

5 Distinguishing investment advice from financial planning

Most people will be familiar with some aspects of financial planning, at least personal financial planning or its specialist areas such as estate planning. The constant entreaties in advertisements coupled with the exhortations in the financial pages of newspapers have served to raise awareness of the benefits of planning for the future: for retirement, for death, for children's education, or simply to avoid paying more tax than is necessary within the letter of the law. Businesses also carry out financial planning: to minimize the tax on income from the return on investments or capital gains, to ensure that money is available for new ventures, to plan affairs so that foreseeable liability may be paid off, and so on.

The term 'financial planning' is, unfortunately, also often used in a loose sense to add status to advice concerning a narrow aspect of a person's or business's financial situation. Even more confusingly, it is often adopted by those who are doing little more than selling one or more of a limited collection of financial products. Our use of the term is in a more specialized sense. We are concerned with the activity of assembling or synthesizing a comprehensive solution to meet the whole of, or a distinct but substantial part of, a person or business's financial objectives; with reorganization of existing affairs as necessary.

The principal difference between investment advice and financial planning, for our purposes, is that, in the former, the whole of the existing situation of the client is not sought to be varied. Generally, attention is confined to only that element of the client's finances which is currently available for investment. The line is blurred when a client is advised to, say, liquidate an equity holding or terminate a low yielding insurance policy with a view to increasing the amount available for investment in the currently recommended transaction. It is a matter of degree: a 'minimal' amount of interference with the existing situation, particularly when the change is isolated in its effect, does not take the recommendation out of the ambit of our delineation of investment advice. Nor does substituting one component of the existing financial arrangement with another of the same broad type change the character of the service.

In contrast, producing a coordinated financial plan involves synthesizing a solution from multiple components, where there are significant interactions present. The components are associated with various constraints, fixing the violation of a constraint for one component may lead to another's being violated. A coordinated financial plan achieves an equilibrium. This task is a difficult one and is at the frontiers of what can be tackled by knowledge systems. More tractable is the task of assessing a situation and making recommendations based on the classification of the situational facts, with the helpful assumption that the recommendations do not materially alter those facts upon which the classification is itself based.

5.1 Commercial significance of distinction

Traditionally, financial planning and variations like as estate planning, capital gains tax planning and portfolio construction, are available from sources such as accountants, attorneys, solicitors and

specialist financial planners. Some commercial banks also provide this service, in the contemplation of obtaining (or keeping) investment related business for themselves or trusteeship of an estate for an associated corporation.

Wealthy individuals and sizable businesses have long accepted the need to pay for advice aimed at arranging their affairs to achieve certain goals. They do not expect the adviser to rely on receiving their remuneration from the vending of products consequent to the execution or implementation of the advice. Most such financial planning is provided by highly skilled persons who charge accordingly for their services. It is fair to say that the use of knowledge systems is not really contemplated by these experts.

Now, however, financial planning is increasingly being sought by middle-income groups. While there is, particularly in the USA, a substantial body of clients who manifest this recognition by their willingness to pay anything from $250 to $5000 for a planner's services, fee-based financial advice is still an irksome concept.

This has brought into the market those for whom the primary commercial advantage is obtained from revenue heads other than fee-income for the planning service. Brokers and mutual fund companies (such as The Fidelity Group and Franklin Funds in the USA) offer a comprehensive employee retirement planning service for small- and medium-sized businesses. Often no charge is made for the planning, the profits being generated from the associated management of the retirement portfolio. Other financial services institutions, such as insurance companies and mutual fund organizations, are also becoming involved in this area. Generally, their aim is to sell their own products on the back of any (re)arrangements recommended.

Insurance policies in trust or as gifts have long been used to mitigate estate taxes liability. Investments in bonds linked to mutual funds or unit trusts are currently being recommended as suitable investments for trustees of an estate. It is suggested that the conflicting interests of the testator's spouse in receiving income, and of the children in receiving accumulated capital, may be reconciled by a careful mix of income producing and capital growth bonds. By offering an estate planning service and, perhaps, arranging for the will to be written at a nominal cost, financial institutions hope to have influence in the post-death arrangements, whilst also obtaining a marketing advantage in the interim.

A client readily reveals a great deal of information when the aim is to identify ways of achieving and maintaining a desired life style, or to ensure that, in the event of death, the distribution of the estate will be smooth and in accordance with prudent principles. The information is itself valuable from a marketing point of view. Further, the relationship between client and adviser is closer and longer lasting because of the need to make periodic reassessment of any arrangements.

Inevitably, the effort expended in providing financial planning is greater than that in giving investment advice. We are referring not merely to the time taken to provide the service. The training level of personnel in the former case also needs to be far higher. As such, financial services organizations are interested in knowledge systems that can either reduce the time taken in dealing with individual planning cases, or enable less well trained or lesser experienced persons to become as qualitatively effective as their more skilled counterparts. It is hoped that computers will help to ensure that the standard of service does not fall below the requisite level, for legal liability for performing such an undertaking in a negligent manner can be substantial; also, regulators can suspend authorizations, with catastrophic consequences.

However, as shall be seen below, systems for financial planning in the sense understood here are complex and difficult to build. Moreover, the risks associated with providing the more sophisticated and onerous service in a negligent or inefficacious manner are also greater. Consequently, most institutions continue to concentrate on the development of investment advisory systems, leaving financial planning to researchers.

6 Analyzing the tasks

6.1 The product selection task

Product selectors are able to provide simple heuristics to knowledge engineers[9]. Whether it is their training or the way they communicate, the experts seem most keen to express their knowledge in terms of a few basic facts which clearly point to certain product lines as being promising, for example:

- If the client is an unmarried working woman above 25 but with at least 10 years to go to retirement, and has a dependent child, explore disability insurance.
- If the business has fewer than 10 employees, the company wishes to set up a retirement scheme for employees and it is desired that more than 15% of income be sheltered from tax, explore a Keogh plan.
- If the client is 'resident' or 'ordinarily resident' for capital gains tax purposes, age is less than 60 and investment in equity funds is being considered, rule out off-shore non-distributor funds.

Most developers apply heuristics such as the above to narrow down the search to a single or a small number of product lines. Once this is done, further rules are applied to narrow down the choice, but only within the areas suggested in the initial screening process.

The heuristics obtained from the experts can readily be decomposed into chunks or modules that are isolated for technical purposes, in the sense that the outcome of reasoning with one module has little influence on and does not need direct reference to the decision-making in another area. It is largely possible to treat singly product 'tracks' such as the following:

- contingent protection for dependants
- single premium or lump sum investment
- recurrent contribution investment
- investment maturing at death for estate taxes provision
- investment to reduce size of estate
- group retirement plans for employees.

Each track can be decomposed into those products or product types which fall within it, for instance:

- Group retirement plans for employees may be decomposed according to the number of employees, e.g. less than 10, between 11 and 2000, and more than 2000 employees.
- In the USA, plans for medium-sized businesses might include SEPs (Simplified Employee Retirement Plans), Keogh plans, 401-K plans amd IRAs (Individual Retirement Accounts).

Many of the products have multiple purposes, e.g. whole life and universal insurance policies are useful for pure investment, even where the risk protection element is not particularly relevant. Therefore, there is some interaction due to either the same product serving dissimilar purposes in different tracks, or different products serving the same (secondary) purpose in several tracks. Such interaction, while encountered, is at a minimal level and it is possible to let the repetition occur during reasoning, while treating each category separately. Duplicate answers can be consolidated at the end, when a spreadsheet-like exercise is carried out to see that available funds are spread appropriately between the various selected products. The spread is appropriate when projections based on estimates of rates of return, inflation, etc. give the desired end-values.

The facility with which chunks of knowledge may be isolated makes it relatively easy to design several small knowledge-bases or coherent rule-sets which can then be coordinated externally. In a simple production-rule representation, the above knowledge is turned into rule sets, each of which can be reasoned with in a largely self-contained manner. The systems use both backward and

[9]As mentioned before, some of the persons selecting products are attached or tied to a single product providing institution. Others are independent, though how wide their net is cast varies. Consequently, different heuristics are avilable from experts depending on the experts' role and market stance.

forward chaining through rule sets to make decisions. Backward chaining with a goal such as 'Far Eastern Utilities Fund-linked capital growth bonds are appropriate' can lead the user–client pair (or the user alone, when relying on pre-filled forms) being asked for all that information which is relevant to the choice. Any communication between rule sets can be done by means of pre-specified attribute values or variables. In more sophisticated rule-based systems, strategic knowledge is distinguished from other rules. Where structured knowledge representation is also available, inter-concept relationships, such as IS-A and PART-OF are obtained. For instance, the relationship that 'a Keogh plan IS-A retirement plan' can be shown by making the former a specialization of the latter. The relationships may be more complex where one concept shares more than one superconcept or generalization, for example: a with-profit whole life insurance policy IS-A protection scheme as well as a long-term investment.

Simple architectural principles can suffice, even where strategic knowledge is not separated fully from the reasoning knowledge and where the task structure is not as explicit as might be. Errors in product selection systems are less of a problem than in knowledge systems generally, because small portions can be separately tested; also, solutions based on any errors that do remain are less likely to be transmitted to clients because the user knows roughly what to expect from each track. The general dearth of interaction between knowledge-bases or rule sets means that an explanation based on canned text attached to the proof trace of each component of the solution is relatively coherent. Therefore, there is a high probability that the user will be able to spot erroneous reasoning by looking at the explanation for each component, largely separately. The explanation facility of the system, usually based on canned text attached to the rules, will give information on why the recommendation is thought to be suitable. All the client has to do is to make up his mind whether or not to accept the recommendation. Of course, it is a time-honoured marketing adage that the client should never be 'confused' with too many options. It is only when one recommendation has been rejected that the next avenue is pursued. A wider ranging product selection session needs to do a somewhat more open-ended evaluation of possibilities. This is a task for a decision-tree, which can readily be implemented in a rule-based representation. Finally, the task that a product selector undertakes to perform for a client, particularly when the client is fully aware that, essentially, he is being *sold* one or more of several products, is not nearly as onerous in terms of potential legal liability.

The way that the systems are used, i.e. with the client usually present and affirming or disaffirming partial results, also goes towards easing the burden on the knowledge engineer. Often the concepts referred to in premises and conclusions of rules—'resident', 'age', 'number of employees', 'marital status', etc.—are not themselves inevitably basic facts or primitives. Fortunately, the conditions under which the derivable concepts can be ascertained directly by asking the user can be specified by the expert. The possibility of asking the user for information is extensively used to reduce the complexity of the system—many of the askable concepts are treated as primitives, with no further rules needed to infer them.

A large number of knowledge systems for product selection exist. The systems are generally integrated with standard projection and illustration software which would tell a client the estimated monetary benefits to be obtained if the client should, for instance, purchase a selected retirement plan and follow the recommendations as to the monthly amount to be contributed. By and large, they are successful in their primary role: supporting the selling of products whilst constraining over-enthusiastic or inexperienced representatives.

The process of development of product selection systems, as described above, sounds mundane, almost routine. This is how it is, at the *technical* level. However, short cuts cannot be taken. The biggest problems arise from the shortage of competent experts, or at least those who are verifiably competent. There is the temptation, in the face of the simplicity of the task, for knowledge engineers to bypass the experts. Most organizations have enormous volumes of material by way of illustrations and guidelines for their representatives. However, while this material is fine for familiarizing a knowledge engineer with the concepts involved, it is too sales-oriented. If one studies the illustrations, almost every product is suitable for every prospective customer. Quite

understandably, there is too much concentration in such material on the benefits to be obtained and not enough on the potentially detrimental side-effects.

Indeed, illustrations and guidelines which may be adequate to train company representatives are certainly not a substitute for a full-scale knowledge acquisition exercise with the co-operation of one or more *experts*.

In spite of this caveat, we may have adopted a rather dismissive tone towards the product selection task; therefore, it would not be right to discuss below specific systems that fall in this category.

6.2 The investment advisory task

Let us now look at in more detail the provision of investment advice. A mentioned above, we use the personal investment domain for illustration. The essential principles are fully applicable to advice that may be given to a firm or corporation, say, on the placement of its liquidity. In fact, in the case of investment advice to businesses, the complexity is likely to be *lower*, because the time frames for investment are much shorter and there will be fewer unpredictable considerations, such as whether the client trusts his spouse with absolute ownership of some shares, or whether one dependant is to be favoured over another in some subtle manner. An example from personal investment follows.

Consider the situation where the client is a married man of 45 who seeks open-ended advice about how to invest a sum of £50,000 recently inherited from a relative. The man is married and has a daughter aged 20 who is likely to get married in the near future. Some of the relevant features of the client's financial position are as follows:

- Client's income: £15,000 per annum.
- Spouse's income: £2500 per annum.
- House worth £75,000, in client's name, with outstanding mortgage of £40,000.
- Client's retirement provisions: employment related pension will provide equivalent of about 20% of final salary; state pension entitlements are additional. Client's spouse, who works part-time, has no retirement provision.
- Couple's savings in a joint account: £5000.

Different advisers would treat the above situation differently. Much depends on the kind of market stance the organization they represent takes. If they are independent advisers, much would hang on their particular view of, say, the relative standing of different product vendors, possibly influenced by the service they themselves have previously received from the organizations. The advice might be:

(a) Make additional voluntary contributions to the company retirement scheme—perhaps opting out of the State scheme and applying that amount to the company scheme.
(b) Take out insurance against sickness and disability.
(c) Take advantage of the spouse's separate income tax assessment by putting sufficient money in her name in an account that pays interest gross, thus utilizing her personal allowance.
(d) Put sufficient money to cover marriage expenses in a three month access account.
(e) Put the balance in a mix of capital growth and income bonds linked to equity funds.

Whilst providing the advice, a number of other options were open to the professional, including:

(i) Paying off that part of the mortgage which does not attract tax relief, thereby reducing the amount invested in equity related mutual funds.
(ii) In addition to contributions to pension schemes, buying inflation-linked gilt-edged (government) securities maturing at about the time of retirement.

Even more options would have been open if the couple had been other than low-rate income tax payers, or if the sum inherited was sufficiently large to require thinking about sheltering the investment income from tax and any capital gains from liability.

Such added factors might suggest that the task is completely unbounded and that a plan is being synthesized, with all the difficulties that entails. Fortunately, the superficial complexity is deceptive. What we have actually is a *task tree* in which the ordering of tasks is pre-ordained: look at protection before capital growth; at school fees before providing for marriage; and so on. Sometimes, it may be subject to alteration within the scope of well defined task groups: within the task group that chooses investment, long-term investments may be considered before or after short-term ones, depending on local conditions.

One may think of the above as planning, if that helps, but it is planning with some extremely facilitative features. Imagine designing a car (just the external appearance of the carriage, rather than the chassis and the mechanics), assuming that the overall length and width are given. Perhaps, one could start by deciding whether a boot or a tail gate is required, whether two or four entrance doors are needed, then fix the size of the doors, the spacing between them, then the glass area, and so on. It does not sound too difficult. But how does one know that everything will fit together? What tells the designer that the entrance doors should be symmetrically placed, that they should swing open laterally—rather than upwards (*a la* Delorean)—without interfering with each other, that the tailgate should be hinged at the top rather than the bottom?

However, these problems do not confront the armchair designer if he is able to find from his own knowledge, or is informed of, the 'correct' order in which to approach the lower level decisions, indeed, *which* lower level decisions to perform, without getting involved in possibly endless search and revision. This order may leave numerous gaps, two hinges or three, the hinges open towards the front of the car or back, etc., but these matters can be dealt with after the abstract plan is assembled, without unduly disturbing the earlier decisions.

In short, major decisions can be taken and choices made at an early stage, in the knowledge that they will not be subject to significant revision.

A knowledge system does not generally have recourse to most of the knowledge which enables a designer quickly to decide the order in which to approach a job. But there are ways in which the system can be informed of the so-called correct order of task performance. This matter will be discussed in more detail when financial planning is addressed. It is sufficient here to note that when the planning is quite *routine*, in the sense that the components of the plan do not themselves need to be redesigned, it is possible to obtain from experts a small number of skeletal plans which need only be instantiated or filled in by making decisions indicated by the skeleton (Brown & Chandrasekaran, 1989). There is then still the difficulty of finding the right skeletal plan(s) to deal with a particular problem. Fortunately, this difficulty is almost completely absent in the investment advisory application because a *single* skeletal plan is thought to be sufficient. This skeletal plan is implicit in the structure of the task tree. Moreover, the gaps that are inevitably there when a plan is followed, are largely dealt with by the user (investment adviser).

It is an essential part of investment advice to produce a written statement detailing the assessment of the client's situation and the reasons for making particular recommendations. One of the attractions of knowledge systems is the partial automation of the laborious task of producing these reports. Virtually all the systems that have been fielded rely on canned explanatory text attached to rules or frame slots. This text is progressively inserted into the report as the rules are fired or slots accessed. It is up to the user to word-process the accumulated text. This is where an explicitly represented task tree proves useful. A well-defined problem solving framework, with separation of strategic knowledge, results in coherent and orderly explanations of the reasoning beig produced. From this it follows that the insertion of text in the report file will also be better. By following such more sophisticated architectural principles, it is possible to produce more powerful and robust devices.

6.3 The financial planning task

It must be noted at the outset that financial planning has proved a difficult area for knowledge systems to tackle, unlike the case for applications involving financial product selection or

investment advice. Below we will look more closely at the reasons why this should be so. It will be seen that, notwithstanding the problems, it is possible to go some way towards commercially useful financial planning systems by limiting the functionality aimed for, and by utilizing some advanced techniques.

At the technical level, it is generally true to say that financial planning for individuals or families is more complex than doing likewise for businesses. There can be very complex tax considerations when assets are being allocated among the components of a business's portfolio. As large sums may be involved, a wide range of financial instruments may be relevant. However, the time frame which must be considered for personal financial planning is usually much longer, which exacerbates the uncertainties that must be confronted and the number of factors which may affect the outcome is larger. For instance, emotional preferences are more likely to carry weight in a family situation than in a business; often personal financial planners find that their recommendations are rejected by the client for reasons hard to articulate. As our purpose is to illustrate the principles of developing and using systems of a particular type, we lose nothing of essence by confining the discussion to the personal finances domain.

The *objectives* which must be achieved in personal financial planning are similar to those indicated above in the context of investment advice. For example:

- Estate insulation by insurance against disability, illness, etc.
- Provision for dependants in case of death.
- Saving for educational fees.
- Generation of income from assets to cover normal expenditures or for specific purposes.
- Differing the imposition of tax on capital gains.
- Minimizing estate or inheritance taxes by arranging affairs during the lifetime of the client or by testamentary dispositions and provisions.

A main difference with investment advice is the degree to which objectives may interact or conflict. The fact that existing affairs could not be significantly rearranged attenuated the complexity when service was confined to investment advice. The financial planner is charged with not only allocating the fresh assets to be invested but, potentially, also altering the existing state so as to contribute to the achieving of the objectives. The planner cannot ignore a less than satisfactory arrangement resulting from a clash of objectives, saying that it is out of his brief.

When, as is likely, all objectives cannot be achieved, there has to be some mechanism for resolving competing claims, for example, by weighting or prioritizing. A similar difficulty was identified above in relation to investment advisory systems. There, the device for dealing with the difficulty involved deciding in advance that, say, insurance was more important than capital accumulation; then the need for insurance was first completely fulfilled before the available investment amount was applied towards the other need; alternatively, the amount was split between the various purposes in some prespecified or easily determined proportion. It was left to the user to process these *suggestions* and alter them by removing some of the components, or reallocating the funds, with the help of a spread sheet. Nevertheless, while these simple resolutions of conflicts and interactions may be adequate and acceptable when a small part of the client's life is being arranged, a financial planner has to take a more sophisticated approach:

(a) It has to be realized that articulated objectives are generally tentative, and not sacrosanct. The client may agree that supporting a continuation of the present life style by insulating the estate over the next 15 years is crucial but may retreat when it becomes clear that it is at the expense of, say, lowering of post-retirement living standard.

(b) One financial product or instrument or legal device may serve multiple purposes, but with different degrees of effectiveness. A whole-life insurance policy may provide both capital accumulation and an element of contingent protection. Under some economic circumstances, it may be the preferred option in the first role. When considered without knowledge of its first role, for some age groups or persons with certain occupations, it may not be the most cost

effective choice to perform its second role. Yet, its dual purpose may make it better than having two different transactions. In other words, the product may be the best choice in a dual role, but not when either role is considered in isolation. The planner, therefore, cannot look at each area separately.

(c) A solution to achieve one objective may conflict with the achieving of another. It may be desirable, from the point of view of reducing the taxable size of the estate, to transfer some assets into a trust. This may frustrate another objective, that the client should continue to receive income from assets.

Clearly, a complex, open-ended task is involved.

6.3.1 Essentially a design task

In knowledge systems terms, the problem as characterized above is an exercise aimed at *transforming* one state of affairs into another, more appropriate one, by searching through a large space of possible solutions, only some of which are efficacious. This is not quite 'planning' as understood in traditional artificial intelligence domains such as planning a machining process to achieve maximum utilization of equipment, or planning a bombing mission so that sorties conducted from multiple bases do not interfere with each other and have maximum impact.

In AI planning systems, the current state has to be transformed into the 'goal' state in a number of steps, some serial and others being executed in parallel. The crucial point is that the transformation steps, and their sequence, if any, form an essential part of the solution. The solution is valueless without a full specification of the steps.

In financial planning, there certainly are some temporal or sequential considerations: in a long-term plan, the client may be advised to put money in insurance bonds for ten years and then withdraw just over 5% of the original amount per year so as to keep the tax liability low. Or he may be advised to sell shares one day and buy them back the next day in a 'bed-and-breakfast' transaction, so as to elevate the base from which capital gains are calculated. However, these are essentially stereotypical elements of the solution and are *not to be thought of as 'steps'*. The 'bed-and-breakfast' comes as a whole, not as a number of steps which must then be put in the right order. Financial planners recommend one of only a few, stereotypical ways of withdrawing the growth element from the bond: for instance, the withdrawal from the bond can be built-in by creating ten separate bonds, each of which may be surendered in a separate year.

Clearly, the process is not so much like planning as *design* by assembling standard components or components that themselves require minimal design[10]. The important difference is that, *in practice*, the final design is quite useful even without explicitly ascertaining the interim steps taken to get there. The users can be expected to fill the gaps themselves. As such, a design system, at least in the financial advice domain, can virually ignore the temporal or ordering element. This fact will assume significance when we discuss designing by using the knowledge contained in cases which are records of specific (or hypothetical) instances of problem solving. It must be noted , however, that for all the diffeences between design and planning brought out here, they are similar in many other aspects and will need below to be considered together for some purposes.

Having laid the foundations by describing the relevant tasks, next we will study the architectural principles behind a variety of systems encompassing a range of investment advisory and financial planning tasks.

7 Advisa: a large, integrated, rule-based system

The Advisa system from Datasure Ltd. is a substantial system which was developed between early 1986 and mid-1987 by a team of four knowledge engineers/designers (in addition to several systems

[10]Variation in the allocation of the amount to different products in the solution set is, strictly speaking, not design, provided it is kept within the governing limits. Violating the limits, if any, would change the character of the instrument. For example, when the contributions to a retirement plan exceed a certain proprtion of the income, tax relief is no longer available. As stated above, we are not concerned with the very top end of clients for whom attorneys and accountants create novel tax shelters.

analysts and programmers). Around six experts, including an actuary, contributed their knowledge over hundreds of hours.

The design was influenced by the multi-million pound cost of development which was partly met by venture capitalists. This association meant that the project needed to have plans and follow them closely. Consequently, an entirely prototype drive evolvement of the system was frowned upon. Prototyping was restricted to proving carefully chosen portions of the design that were felt to be representative of the whole in complexity, while being of a significant size. Based on this experience, functional and design specifications were drawn up in the manner normally adopted in large information systems development. As, in addition to knowledge-bases, the system included conventional programming for database management and access to value added networks, the reliance on conventional specifications was most useful. Also, this approach allowed the employment of systems analysts for a part of the task that would otherwise have to be done by knowledge engineers—the latter type of personnel were, and remain, more difficult to recruit.

Advisa was originally meant for independent financial advisers who, by virtue of the regulatory guidelines, are obliged to consider the entire range of products on the market. The system conducts an initial fact-find of the core information such as age, marital status, income, assets, aversion to risk, likely expenditure patterns, any particular needs for lump sums or income, etc. Other information is asked for, if needed, when a particular area of investment (or insurance) is being considered. From this core information certain *needs* are identified and are refined by presenting to the user (who might be accompanied by the client) projections based on parameters including the likely inflation rate and expected returns on investment.

The needs relate to stereotypical investment goals such as obtaining a tax free lump sum on early retirement, providing a guaranteed income over the next few years, sheltering from tax the income produced from assets, providing health insurance, and so on. The goals are then attempted to be satisfied in a pre-ordained ranking order. It is possible to amend or re-specify the goals and cause the system to search on an altered basis. Thus, the user is able to have some influence over the eventual outcome. The accepted decision is accompanied by a report that has been automatically produced from the text attached to the relevant rules.

The system is fully integrated with a database which holds, amongst other items, large quantities of historical performance data on mutual funds and insurance policies. This allows comparison between similar types of products from a number of sources. Data showing dependability or stature of the various institutions is also held. Once specific products are selected, the system can obtain automatic quotations and proposals from the concerned institutions by dialling-up value added networks. Another feature is integration with back-office software such as file-management systems which hold information on the financial arrangements of clients.

One of the parties associated with the development of Advisa was a major financial information vendor/publisher. The main business of this organization is to collate historical data on investment performance of various insurance products, mutual trust funds, shares of investment trust companies, etc., and distribute it in the form of periodical publications such as the well-known magazine, *Planned Savings*. The year, 1986, saw the start of the regulatory regime which required those professing to be independent investment advisers to consider all the available products on the market before making recommendations. Not all these advisers were likely to be capable of utilizing the copious, printed data in an effective manner. Business directions made it essential that means be provided to distribute the data in a computer format, with aids to allow retrieval in a context sensitive manner.

The Advisa system was designed, in part, to cater for this strategic requirement. Accessing information so that product types may be compared, whether during an advisory session or at another time, the system has an independent utility as a powerful database. The knowledge-based interface is more easily used by those not overly familiar with computers in comparison to most traditional database management interfaces.

A system like Advisa can be sold or rented to advisers—including, in appropriate versions, tied representatives of financial institutions—for use on their personal or office computers. Alterna-

tively, it is possible to run it on a value added network to which persons can have access over a telephone line, making payment by subscription.

It must be appreciated that a tremendous effort went into analyzing the business environment before the design was embarked upon. The participants appreciated throughout that knowledge systems technology can give them a major edge over rivals, even in the most conservative of areas. The faith was vindicated: the system is successful in commercial terms. It is in use with over 100 organizations and numerous manufacturers have taken steps to make their finance office management system integratable with Advisa.

8 Taxadvisor: an early rule-based system

On the subject of basic rule-based systems, for the sake of historical accuracy, and so as not to suggest that Advisa's is a pioneering approach in a fundamental technical sense (as opposed to the sheer scale of the task undertaken, with its associated methodological innovations), it is necessary to mention an early system called Taxadvisor. This system was developed in the early 1980s (Michaelson, 1984). It used the Emycin shell, which is a derivation from the classic system, Mycin (Shortcliffe, 1976). The underlying approach of such systems, including unsegregated rules, backward chaining, propagation of confidence factors as measures of belief in conclusions, is well known.

The system advises corporate executives and passive investors. The advice is meant to include some aspects of estate planning, but only at a relatively simple level, and certainly not *planning* in the sense in which we are using the term here[11]. The order of exploring the tax situation is essentially predetermined. Once the identifying data has been collected, the sequence in which goals are considered is as follows:

1 Insurance to protect client's (and spouse's) assets in case of accident, disability, illness, etc.
2 Life insurance to provide for client/spouse and family in case of premature death.
3 Retirement provisions.
4 Estate planning not involving permanent transfer of wealth—i.e. tax shelters, revocable trusts and short-term trusts.
5 Estate planning by means of transfers by way of gifts and testamentary dispositions.

Recommendations are produced in each area. The solution for items (4) and (5) is expressed in extremely general terms, owing, no doubt, to the complex interactions extant in these areas. The generality means that the advice produced need be no more specific than, say, the following:

- Life-time planning:
  gift of life insurance policies
  gift of assets expected to appreciate in value.
- Testamentary provisions:
  generation-skipping transfers of remaining property
  tangible personal-use property to spouse
  stock to son in outright form
  home to charity in outright form.

The Taxadvisor model is simple, yet powerful. However, it must be noted that the advice produced by a system based on the Taxadvisor model (as Advisa and others are, to some extent) cannot be equated with estate planning because the advice is not detailed or specific enough. It is only when one gets to particulars of recommended steps that the possibility of interacting constraint violations becomes apparent. This includes taking corrective action when one choice leads to the violation of some restriction in respect of another, earlier, choice. Further, the system does not itself assist in

[11]Further, being experimental, the system excludes from its consideration complicating matters such as the following: (a) client owns a controlling interest in a business; (b) the combined estate of client and spouse is lower in value than the single unified credit equivalent for purpose of estate taxes; (c) a substantial part of the estate consists of assets held as community property of the client and spouse.

refining or revising the advice because the consequences of altering one element of the solution on the other elements are not catered for. If the user himself were to alter part of the system's recommendation, he would need to have a comprehensive and in-depth knowledge of the domain, for the implications are subtle and potentially far reaching.

Such systems are, therefore, most useful in areas where solution elements can be made specific by largely procedural or numerical processing. Examples of such areas are advising in respect of items 1 to 3 above; then, the solutions can be refined by means of a spreadsheet which allows sums allocated to different elements to be altered and the effect to be seen. Restricting the scope of the advice to such matters does not limit the functionality of the system; the system is still valuable to moderately experienced personnel, who happen to be the ones most to benefit from being aided by knowledge systems.

Advice on testamentary matters requires extension of model

It is interesting to note that there are occasions where even minimal advice in relation to life-time estate planning and testamentary provisions can be commercially useful. For example, there is currently a strong trend in the UK for insurance companies to set up will-writing services. Under the schemes, company representatives obtain the relevant information from (mainly middle income) clients and pass it on to experts qualified to draft the will. The idea is to offer the service for a low cost, the benefit is expected in terms of increased access to clients and, dare one say, respectability.

Several systems are currently under development to provide general advice about the form of a will that is suitable in the circumstances. This information is presented to the client for preliminary evaluation, so that any changes can be made whilst the representative is still with the client. However, because such advice will be somewhat more detailed than that produced by Taxadvisor, the architecture underlying the model has to be extended. In particular, it is necessary to allow the user (at the prompting of the client) readily to make revisions in the advice by changing constraints and preferences. This is being addressed in current developments through a more explicit representation of the task structure.

8.1 Over-the-counter advice in a bank

A system that takes the Taxadvisor and Advisa approach further is PORTAFOGLIO (Chierici et al., 1989). The system is meant to provide over-the-counter investment advice in a bank: Cassa di Risparmio di Parma, or CRParma. The variety of products dealt with is limited compared to Advisa, primarily bonds and securities, but the goals of the client are similarly wide-ranging.

In Advisa, low level tasks are performed in an order or rank that is embedded in the organization of the rules in the knowledge-base. The uniformity of the rule-representation means that changing the order would require alteration to a large number of rules. The lack of explicitness of the structure makes it difficult to correlate the system's knowledge to that of the expert—the explanations do not reveal enough. Also, the quality and coherence of reports recording and justifying decisions made during an advisory session is limited.

PORTAFOGLIO, therefore, provides a more explicit task structure to obtain enhanced performance: the quality of service is particularly important to a major bank. What may be tolerable for an insurance comany's salesman will not do in the case of an assistant bank manager.

8.1.1 The architecture

In the task tackled by PORTAFOGLIO, the purposes (or sterotypical goals) of the client include the following:

- capital gain
- gradual accumulation
- periodic income
- liquidity.

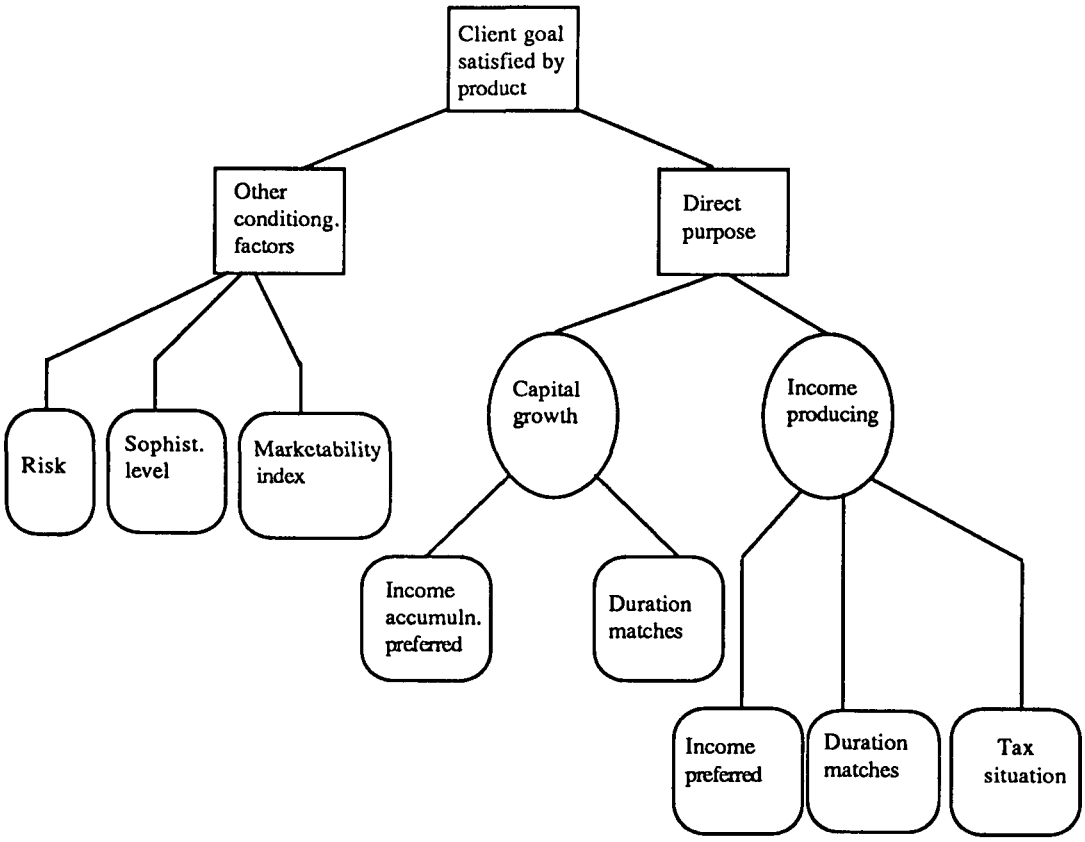

Figure 1 Client's purpose as a decision tree

The financial products are organized in around 15 classes, including:

- cash
- short-term bonds
- long-term bonds
- shares.

For each product class there exists a tree that shows the relationships between the class's characteristics and the client's purposes (see Figure 1 for a portion of the tree).

The tree takes into account 'other conditioning factors' such as preparedness to bear risk, investment culture of the client and various rates, all of which are derived from the information supplied by the client in the fact-finding session. At the head of the tree is the goal that the product is satisfactory. Its sub-goals are the (direct) purpose of the client and the relevant 'other conditioning factors'. The client's purpose is to make capital gains, rather than receive a periodic income from investments. The purpose is decomposed into lower level tasks that involve ascertaining, for instance, whether the product favours accumulation, whether the duration of accumulation is in accordance with the client's requirements, and the acceptability of the tax position.

An evaluation of products against the client's goals and related information is carried out by trying to satisfy the decision tree associated with each product. This produces a list of a number of potentially suitable products. In the next phase of analysis—the *configuration* stage—rules are applied to see how the available investment amount is to be distributed among the products. At this stage many of the candidates on the list of suitable products are eliminated by application of further rules. For example, if 'Bond (Closed-end) Investment Trusts' and 'Capital (Closed-end) Investment Trusts' were both on the list, the rules might abstract the features of the two recommendations and suggest 'Balanced Investment Trusts' as a more general solution which can perform the desired functions. It will be noted that configuration is tantamount to filling in the details in the gaps left by the 'skeletal plan' in the pre-specified task structure; decisions at this stage

are not so far-reaching as to render invalid or cause the revision of the choices implicit in the structure of the skeletal plan.

As in Advisa, the system accesses product databases to best execute the configured solution by means of the bonds and securities that the bank wishes to sell. The client is allowed to state constraints on individual bonds and securities.

A feature which is most interesting, commercially, is the product database which contains a parameter called the 'marketability index'. This parameter is determined by the bank and indicated its relative willingness to sell particular securities. The higher the index, the more the bank would prefer to sell the relevant item, provided of course, they satisfy the purpose of the client.

9 Classifying the investor, not the product

The INVEST system has been developed to give advice to clients of the Volksbank Münster: a bank with many branches in rural areas (Heuer et al., 1988). In spite of its purchasing power, the rural population is underserved in terms of competent investment advice. The idea is to make such advice available in even small branches. The scope of the advice is initially confined to securities available on the stock market.

The architecture of the system is such that the task tree is more flexible even than that of PORTAFOGLIO, with less need to depend on pre-specifying the order of evaluating hypotheses in respect of components of the ultimate solution set. Therefore, the solution set is less biased towards products that happen to have a higher priority on the task structure. In other words, those clients who have complex requirements and for whom product lines cannot be decided at an early stage will receive a higher quality and more sophisticated advice.

The central idea of the INVEST architecture is the existence of a number of *client* profiles which are meant to be representative of a broad range of clients. These are called *hypothesis prototypes* and are represented by means of frames. A hypothesis frame can have slots leading to other, more specialized hypotheses. The client's situation is also represented by means of a frame. A match is carried between the client and each high level hypothesis prototype. If there is sufficient degree of match with a hypothesis, its lower level hypotheses, if any, are considered in order to make the advice more specific.

It may be noted that reasoning from client profiles is rather different from having the *product* profiles expressed, as they are, implicitly, in the systems described in previous sections. The service is client-oriented, and not based on finding the first product that is capable of being recommended.

9.1 Illustrating the matching operation

The hypothesis prototype called 'Shareholder' asserts that any client with sufficiently close or similar characteristics to those in the frame should be recommended shares. The specialized hypotheses of this prototype classify the profile more exactly so that particular types of shares or the share mix can be identified. Part of an instance of the Shareholder frame is as follows:

```
(Instance Shareholder
  ...
    (Acceptable Risk :=
      (HList (low . −1000)
             (average . 500)
             (high . 1000)
    weight := 3)
...))
```

The HList contains a list of pairs in which the first element is a possible value for the 'Acceptable Risk' slot, the second element is the corresponding numerical representation which may be used when calculating the degree of match. The numerical value −1000 means low risk, 500 is average

risk, 1000 is high risk. When matching, the value of the Acceptable Risk slot of the client profile frame is seen. If, say, the value is 'high', the number 1000 is assigned to it. Then the number is multiplied by the weight, 3, to get an idea of the degree of match. This kind of matching is carried out for all the relevant slots and a cumulative match number is arrived at. The developers have roughly associated some semantics with the numbers. In matching, 1000 is 'fits very well', −1000 denotes 'doesn't fit at all'. The gradation between these two end points is marked by steps of 250.

It is reported that the experts were quite readily able to understand the significance of the grading scheme, and there was little need to revise the numbers assigned in the HLists (Heuer et al., *op cit*).

The task structure is quite sophisticated because hypotheses are not considered in a pre-ordained order or partial order. Instead, they are individually suggested or *triggered* by the presence of certain facts in the client's circumstances. For example, the desire to reduce tax may trigger municipal bonds. In the first pass through these facts, some rules with this limited purpose (i.e. selection of municipal bonds) are applied and a list is created of all the hypotheses that have received triggers.

If a hypothesis has multiple triggers, it does not mean that it is more likely to be true. Instead, the inference made from the number of triggers is that the hypothesis with a higher number of triggers should be considered *before* one with fewer triggers. This is part of the meta-knowledge and is kept distinct from the judgemental knowledge.

Once a hypothesis, or one of its specializations, if present, is confirmed, it is added to the investment proposal list. When there is a problem with the execution of the product suggested by a confirmed hypothesis, e.g. if foreign security is suggested but is not available from the bank, a parallel proposal including foreign securities is created. Several alternative proposals can thus be produced.

However, we believe that there is a limit to the usefulness of relying solely on matching clients with typical investor profiles. A very large number of profiles would be needed in a system with a purview broader than that of INVEST. If it is required to deal with complex products with varied but partially overlapping features, product profiles must also be explicitly held in the knowledge-base. The reasoning must include looking at both. Although the client profile might be used to determine the initial selection, a triggered hypothesis should be refined against a more detailed representation of the characteristics of the products themselves.

It may also be noted that the system, powerful as it is, still does not attempt financial planning. There is some provision for revising a proposal when it appears inappropriate in the light of facts discovered later. Essentially, however, components are merely added to the investment proposal, with the order of addition having little intrinsic significance. It is implicitly assumed that there is little or no interaction between the various components. This assumption crucially differentiates the exercise from financial planning (or design).

10 Design from first principles and abstract 'plans'

The basic problem faced by a computer system when creating a design is that, having proceeded down a particular track, it may find that a constraint has been violated, causing it to backtrack and try another path to search[12]. The problem, when approached as a first-principles state space search, is thought to be computationally intractable. The system may well have knowledge regarding how to correct or *fix* the violation of any constraint. But the fix may cause some other restriction to be breached, leading to backtracking again.

To counter the above problem, domain-specified knowledge is acquired to provide the system with a narrower search space. With much effort during the knowledge engineering stage, it is

[12]The process of design is not very well understood in the context of knowledge systems. There are several theories, but few are complete enough for application. We certainly do not propose to try and correct this here and will limit ourselves to those principles which are pertinent to the present application and are necessary to appreciate the architecture to be described.

possible to acquire heuristics from experts which allow shortcuts to be taken. For example, knowledge can be elicited showing how goals may be prioritized and how competing goals may be resolved without affecting other parts of the environment. A system that successfully relies on such heuristic knowledge is PlanPower, which was produced by Applied Expert Systems, Inc. (Stansfield & Greenfield, 1987). Obviously, full details of the architecture are not published, but an explanation of the principles can be given here.

In PlanPower, the pruning of options begins right at the first stage, when the raw information about the client is analysed and reduced to a 'diagnosis' of the situation of the client. In other words, the client is *classified* by means of back-chaining through rules and the client's problems are identified as, what are termed, 'observations':

- debt payments are high proportional to cash flow
- particular aspect of the tax position is ill-planned
- client will not achieve certain objectives if the status quo is maintained.

Based on the observations, further rules are applied to evaluate all available investment vehicles and to see which of them should be further explored. A list of possibilities is produced. There are heuristics such as that one should always handle insurance before allocating discretionary assets or that, when some existing position has to be changed, certain assets within an investment category should be sold before others. Using these, the possibilities are then *partially prioritized* through considerations such as client preferences and tax implications. This is clearly a crucial stage because the prioritized list forms an agenda for the planning stage by keeping a check on the scope of the search. The prioritized possibilities list may be viewed as an abstraction of the final actions to be recommended. From there onwards, planning is a refinement process in which details in each area of recommendations are filled in. There is still considerable interaction between the different investment categories and an *asset allocation model* is used to minimize the need for backtracking during the construction of a portfolio. There is also much involvement from the user who can use 'what-if' type spread sheets to see the effect of changes, and can additionally influence the priorities on the possibilities list.

The implementation of PlanPower in such a difficult domain is an impressive feat. It required a team of ten knowledge engineers and 12 systems people for its development, totalling about 50 man years. The completed system is meant for banks and other institutions to prepare plans for major clients. The skill level of the user needs to be very high as much intervention is needed. The route followed by the developers of PlanPower is certainly one which is open to developers, provided they are extremely persistent and are confident that the heuristics obtained, at great cost, do not become inapplicable or out of date before the investment can be recouped.

11 More structured approaches to design

More recently, developers of design systems have taken approaches where knowledge structures at a higher level of abstraction than rule-like heuristics are employed to guide the search. Chunking and structuring knowledge at this level of granularity means that it can be obtained more reaily, particularly if multiple experts are involved. Each expert can be asked to give knowledge about his own area. Unlike this, in the PlanPower model, any differences in search reducing heuristics would generally have to be reconciled and then applied uniformly over large areas of design. The knowledge structures can be explicitly represented with all the benefits that brings, including the ability to make local changes without affecting everything else.

11.1 'Routine' design

One such structured approach to design is taken by Brown and Chandrasekaran (1989) in the context of what they term 'routine design'. In such a task:

(a) The *structure* of the entity to be designed is already known in that the *nature* of the parts or components comprised within the entity is known and the entity can be hierarchically decomposed in terms of its components.
(b) The knowledge for designing individual components is available as *design-plans*.

It must be clarified here that the term 'design-plan' is used to denote the knowledge structure which provides knowledge for the *design* task. It, effectively, contains a sequence of actions to be performed and the constraints to be checked, representing *one* method to design the component. It must not be confused with a *financial plan*, which is the final artifact or solution produced by the design activity.

The system architecture adopted for routine design has a task-tree corresponding to the hierarchy of parts. At each node in the tree there is placed a program called a 'specialist', which has access to a number of design-plans. A specialist chooses an applicable design-plan, uses it to fill out the design and passes control to specialists lower down in the tree. When the lower level specialists have done their part of the design, control returns to the parent which then fills those details which could not be ascertained without the lower level decisions. In this way, the search is focused by means of localized knowledge.

This method has been tested in a system called AIR-CYL meant to design an air cylinder. Those who have used or, better, taken apart a bicycle pump will realize that the components of a functioning air cylinder include a piston or plunger, the tube inside which the piston moves, a rod connecting the piston to the outside and a cap and a spring at the end of the tube. The component hierarchy is known. Moreover, the domain has a strong theory. For example, not only are fixes for constraint violations fully known, they are actually predictable. One knows the *cause* of every manifestation.

11.2 Financial planning is not 'routine' design

The financial planning application is rather different to that described in the previous sub-section. All the components of a financial plan are not known, nor is it possible to decompose the task into a hierarchical form. Even if one equates purposes or objectives with high level components in a mechanical design, there is the difficulty that it is possible to omit some objectives altogether. Even if it were possible to structure the task hierarchically, it would not be possible to say that a design-plan stays valid or applicable for one node in the component hierarchy irrespective of what goes on elsewhere. In other words, any design-plans are too fine or low level to be isolatable. This is not surprising and follows from the nature of the reasoning involved here.

An expert financial planner's knowledge is usually articulated in the context of 'cases' which are specific (or hypothetical) instances of problem solving behaviour[13]. It is possible for an expert to point out the merits or demerits of choices in a financial plan produced in a specific case, but it is very difficult to make generalizations. For example, it may be possible to obtain reasons from experts why a choice in a case is better than a discarded alternative. However, it is not always possible to say, outside the context of the case, *why* one product is more suitable than another. If an explanation is given, it is likely that it is essentially aimed at demonstrating the *reasonableness* of the solution, rather than at explicating the relationships and dependencies between individual components of the solution and particular aspects of the client's situation. In fact, the presence of the reasonableness test itself often breaks the chain of dependencies: within the defendable penumbra of reasonableness a choice may not actually be right but neither, provided certain conditions are met, may it be wrong.

[13]Design-plans are not the same as 'cases'. Design-plans may be thought of as chunks of frozen control knowledge that are able to create a small part of the overall design. They do not tell us what design succeeded in the past.

Consequently, perhaps the lowest level knowledge structure that ought to be used in this domain is one which has a high *normative status*: namely, specific cases or instances of actual (or hypothetical) problem solving. The reference is to Case-Based Design (CBD).

11.3 Case-based design

In spite of the differences between planning and design mentioned above, some of the concerns of CBD are closely related to those of Case-Based Planning (CBP). The present sub-section will consider them together.

CBD is founded on the premise that by working from existing designs a designer is able to work in the absence of a complete model of the domain. It has been said that when cases are relied upon:

> "[T]he system needs only enough knowledge about the purpose of the case and the causal relationships in the domain to explain why this sequence of actions was performed. The domain theory must be much more complete if the goal is to perform validations ... However, for the domains of interest to [Case-based reasoning], it is extremely unlikely that there will be a domain theory complete enough to allow even this justification reasoning. In a real domain like robotic cell fault recovery, we do not know enough—and we cetainly do not know how to represent enough—domain theory to relate each action in each case to what we know about the state of the world when the case was executed". (Barletta & Mark, 1988).

Strong claims, based on psychological studies, have been made about how designers or planners use cases that have worked before. For instance, it is said that they use cases *directly* rather than trying out combinations of sub-designs or sub-plans (Hammond, 1988). We need not, however, be concerned with this strong view because we can proceed by taking CBD to be no more than a practical way of adapting the technology to the peculiarities of the domain.

Proceeding at this pragmatic level, a typical CBD system carries out the following steps (no sequence is implied in this order of presentation):

- Represent the past cases in a form which enables easy acquisition of the knowledge required to build the case-base.
- Classify and index the cases for efficient retrieval on some assumption of the likely problem cases which will be encountered. Cross-index to allow directed search and enable inferencing.
- Process the problem description to match the vocabulary and representation formalism in which the cases and the indices are defined.
- Retrieve the candidate or target cases based on the classification and the retrieval indices. Evaluate for similarity with the test case. Move within the case-base using cross-indices until sufficiently similar cases are found.
- Modify a target case to fit the situation described in the test case, and use others of the target cases to assist in the adaptation.

Most of the work on CBD has been done in applications where there is a relatively small number of cases to be handled: generally, less than a hundred. It is feasible to represent these cases in a sophisticated manner, for example by using the conceptual dependency formalism (Schank, 1982). Such a formalism allows the situations and solutions found in cases to be expressed comprehensively and explicitly by reducing the domain concepts to primitive relationships. There are, however, some problems with such an approach. Certainly, a conceptual dependency type formalism is designed to be computationally efficient in handling concepts expressed in terms of primitives, but a complex domain cannot readily be reduced to a small number of primitives. Therefore, a representation at such a low level of granularity can become unwieldy when faced with even a moderate number of cases. It is expensive to maintain a case-base in this manner, both in terms of the coding time involved and the computational power required to deal with the mass of information.

12 Practical financial planning using a variation of CBD

The requirements of CBD in the financial planning are such that large amounts of information, such as that which is contained in several hundred cases, must be represented, maintained and updated. Yet the system must be able to reason sufficiently with the cases to be able to judge the usefulness of a case or set of cases in satisfying the conjunctive design objectives. Several compromises are, therefore, necessary.

Once cases which are, potentially, capable of providing the solution to the problem situation, the question is whether the system itself should do any adaptation or modification of the solutions found therein and produce the solution to the problem situation. While such a capability is certainly desirable, incorporating it in a financial planning system is extremely difficult. The ability to transport knowledge from one case to another, newer one requires the assumption that there exists generalized knowledge which allows alterations to be made to a case without affecting its validity. This is a dangerous assumption in a domain with a theory as weak as the one we are faced with. Moreover, as mentioned earlier, cases have a special normative status in this domain because of the (usual) acceptance of the practice of 'respectable' professionals as by itself being reasonable[14]. The practice is best evidenced by what is done in specific cases, rather than through abstract generalizations.

The experimental system which we will describe next does not attempt to modify or amalgamate cases to provide a fully developed solution. It is felt that there is commercial utility in having before the user those few out of thousands of cases which potentially provide an answer to the current problem. The user (the person providing a financial planning service) can readily see those objectives which the retrieved cases do satisfy, and those which they do not.

Whatever alterations are made by the user, they are likely to take far less time than if he had started from scratch. Also, if he is relatively inexperienced, he is more likely to be able to arrive at a high quality plan when being guided by having before him relevant precedents than otherwise. The commercial usefulness of the approach is predicated on the ready availability of an adequate population of cases, perhaps in the form of records of financial plans produced for clients with a variety of requirements.

A sizable and busy financial planning practice usually keeps records of instances where a recommendation was made to a client or when the composition of a client's portfolio was rearranged. The regulatory regimes in most jurisdictions now make it imperative to keep skeletal records of transactions, many organizations insist on more detail. Basically, the record of a case can be assumed to include the following information.

The personal and financial details of the clients (the 'situational facts')

Often this is in the form of a pro forma questionnaire which would have been filled up by the client alone or with the assistance of a financial planning professional. Some information in it might be missing, the professional might have noted down assumptions where necessary.

The attitude, financial culture and objectives of the client

A few of these parameters may be directly adduced from the client, others would be derived by the financial planner by application of his expertise. While there is no standard lexicon for the definition of this information, generally a few organization-specific stereotypical terms and headings would be used. It may be noted that this vocabulary is becoming increasingly normalized

[14]An exposition of the law relating to professional negligence notes: "Where the matter is regarded as within the common knowledge of laymen, as where the surgeon saws off the wrong leg . . . the jury may infer negligence without the aid of any expert. [But, in the case of matters not within common knowledge of laymen] the standard of conduct becomes one of 'good medical practice', which is to say, what is customary and usual in the profession. It has been pointed out often enough that this gives the medical profession, and the others, the privilege, which is usually emphatically denied to other groups, of setting their own legal standards of conduct, merely by adopting their own practices." (Prosser, 1971, footnotes omitted.)

because *compliance officers*–i.e. personnel responsible for monitoring the conformance of professionals within an organization to regulatory guidelines and certain minimum standards of service–need to examine the records and interpret them without necessarily having access to the original author.

The solution set
This would record the types of financial products recommended, the spread of money between the products and any temporally determined steps, i.e. further investment or disinvestment at a certain time in the future.

Possibly, a brief rationalization of the choice of the design
Such a statement is often presented to the client to justify the advice. It may be noted that even if this is not recorded, an expert can recreate and provide this at the time of knowledge acquisition.

It is for expert financial planners and knowledge engineers to abstract the requisite information from records such as the above, so that it may be represented in the computer. The system decribed below is being developed on the basis that experts can provide this information to the developers by way of a relatively limited vocabulary. They are expected to use labels representing concepts such as marital status, the client's tax band, number of dependants and the anticipated capital gains tax liability in the currents year. Many of these labels correspond to those found in actual records of the information, such as the attitudes and *objectives* that may be attributed to the client.

Of course, there is some interpretation to be done, particularly if the information in the case records is sketchy or incomplete. However, it must be remembered that most expert financial planners are already able to fit the details of cases into pro forma work sheets with a predefined vocabulary; qualification by means of annotation is essential only in a small proportion of cases. It is an essential part of the expertise to distil key personal criteria and stereotypical objectives from the client's situation, and thenceforth primarily rely on those objectives during planning. This restriction is prevalent in practice for a variety of reasons, including the need for unambiguous communication of information between the different professionals who may be involved in the planning, either at the same stage or at different times over a period of years.

It is not being suggested that one should mirror in a knowledge system the entire vocabulary of personal criteria and objectives that is used by professionals in practice. Rather, that the process of reducing the open-textured features of a problem situation to a limited collection of labels is not alien to financial planners. Moreover, the planning process is so open-ended that *a part of the excercise of skill is in reducing the open-textured features to formal criteria and objectives*. For example, the martial status of a client can take a number of values in the real world. Where divorce is involved, the client may fall in any one of these categories:

- planning to divorce
- divorcing, court proceedings launched but no decree obtained
- divorcing, decree *nisi* obtained but not converted to a decree absolute
- divorced and paying alimony
- divorced and receiving alimony
- ...

It is the job of the contributing expert to translate this information about martial status into the limited vocabulary defined for this purpose while the knowledge-base (or that part of it which is called the *case-base*) is being constructed. The user-professional has to do a similar exercise (albeit one requiring less skill) of assigning situational facts to the limited vocabulary when the case-base is being used.

13 The architecture of FINESSE

The architecture of a financial planning system called FINESSE, which is being developed at the Centre for Computers in Law and Finance at Brunel University, accords with the case-based design

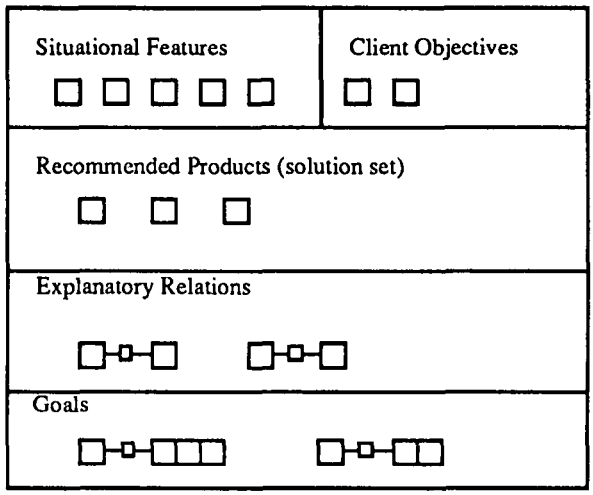

Figure 2 A case as an aggregate object

approach identified above[15]. The conceptual representation of a case, as an object, is shown in Figure 2. Each case consists of:

- The relevant situational features.
- The objectives of the client.
- The solutions or recommendations that are said to have achieved the objectives.
- The explanatory relations which map one of the above elements of the case to another such consituent or to an object outside the case.
- The goals, which are relations mapping one of the objectives in the case to one or more of the situational features of the case.

Semantically, situational features denote that information about the client or the environment which provides the context for the financial-legal design. Examples of situational features are: age, sex, marital status, income tax band, accrued capital gains tax, inflation rate, etc. In the case object, each situational feature is an instance of class objects which are arranged in a hierarchy, a portion of the hierarchy is shown in Figure 3[16].

There are a number of ways in which the concepts in a domain may be organized in a hierarchy. The precise shape of the hierarchy and decisions such as which are primitive concepts and which are to be indirectly inferred depend on analyzing the way in which the task is to be performed. Consequently, no universally true organization is implied in the particular hierarchy in Figure 3.

Objectives are financial-legal cliches or stereotypical, high-level aims that the client (or the client's adviser) wishes to be achieved. Examples are:

- achieve security of investments
- defer capital gains tax liability
- provide for family when breadwinner dies
- provide for payment of inheritance tax liability.

[15]The architecture underlying FINESSE is being developed by a consortium funded under the ESPRIT initiative (project no. 6763) for conceptual retrieval of documents in the litigation support application. Though it may not immediately be apparent, there are several common elements in the way case-based financial planning and document retrieval in litigation support are approached. Handling documents for litigation support and financial planning by relying on records of problem solving activities share the function of retrieving information based on the conceptual significance, rather than mere word patterns or attached labels. Also, the significance is judged on the basis of the ability of the retrieved information to serve some purpose in a particular manner or context.

[16]Henceforth, we will generally ignore the distinction between class and instance objects. Clearly, a mention of components or attributes of a case object refers to instances of class objects.

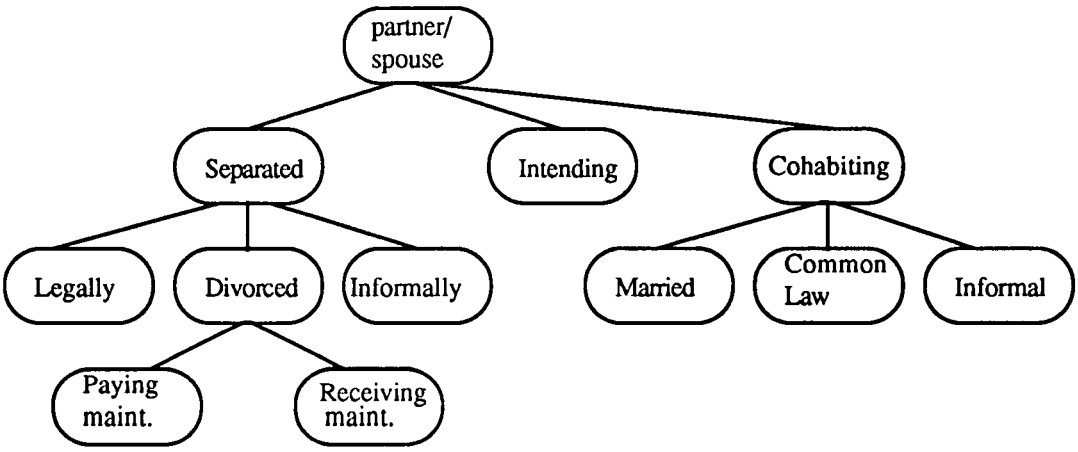

Figure 3 Situational features

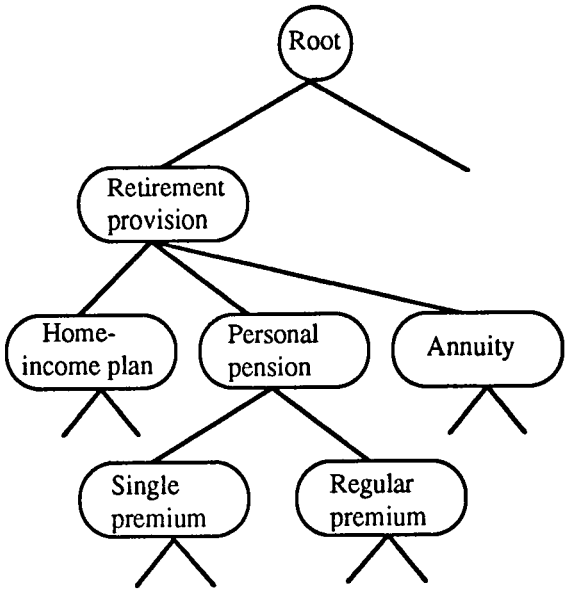

Figure 4 Solutional elements

Objectives too are arranged in a hierarchy. An example of one objective being a generalization of another can be given:

Specialization: 'provide for payment of IHT from gross estate'
Generalization: 'provide for payment of inheritance tax liability'

Solutional elements refer to those financial products or financial-legal devices which are recommended in order to achieve the objectives of the client. These are the standard components of which the financial plan is constructed. The precise mix of the solutional elements will be determined by the scope of the service and the market stance of the particular organization using the system. An example of how these concepts may be hierarchically arranged is given in Figure 4.

An explanatory link is a two-place relation which may be implemented as an extended association link object. The link connects a situational feature, objective or situational element in the case to either:

- another situational feature, objective or solutional element in the same case; or
- any instance of a class object of type situational features, objectives or solutional elements.

It may be noted that explanatory links are *not* universal truths. Instead, they are held to be true only within the context of the case in which they are specified. The following links are used:

ALTERNATIVE link: denotes that a case feature is an arbitrary choice made from several equally applicable alternatives. For instance, though it is not a universal relation, in a case the following may be true:

single premium insurance bond *is alternative of* unit trusts

CONCOMITANT link: this connects either two solutional elements in a case or two objectives, signifying that one is there as the result of the other and not in its own right. For example, the following may be true of a particular case:

security of investment *is concomitant with* moderate growth

OPTIONAL link: which shows that at least one of the pair so related must be present. For example:

whole life policy *is optional with* term assurance

INCOMPATIBLE link: connects a solutional element in a case with a primitive concept (situational feature, solutional element or objective) that is not otherwise found in the case. For instance:

unit trust *is incompatible with* insurance bond

There is no attempt, incidentally, to try and define all those links needed to define a case completely in the computer representation. This is because the role of the explanatory links is not to represent all the information in a case, but only to increase somewhat the chances of a case being retrieved by matching.

Finally, goals[17] are two-place relations that link an objective in a case with one or more situational features in the same place. Semantically, the situational features are said directly to influence the achievement of the related objective by means of one or more of the solutional elements in the case. However, the relevant solutional elements are not linked or identified with an individual goal; it is the overall solution set that achieves the objectives.

Goals too are not universal truths. They have relevance only within the context of the case(s) in which they are specified. If goals were universally valid knowledge structures, then it would be possible to obtain heuristic rules which would allow the synthesis of a solution by combining solutional elements step-by-step, for each objective. Instead, the stance taken here is that different situational features may be relevant to an objective in different contexts, without it being possible to specify rules which can predict this. An example of a goal is given in Figure 5.

Indexing cases

While goals are valid only within the context of a particular case, it is inevitable, indeed an essential predicate of the approach, that a number of cases will contain goals which are identical in structure or vary in only one or two attributes. Advantage is taken of the point, well recognized in CBD, that there are far fewer types of goals than there are situations dealt with in cases. As such, goals can be an efficient means of indexing cases. Further, as shall be discussed below, that two cases are related by a goal has special significance in denoting the usefulness of a case in the context of finding a solution to the problem presented by the user.

A secondary index to cases is provided by solutional elements: because there are relatively few of them, they too form a good index, though without the same semantic significance as for goals. The indexing through solutional elements is used for a very limited purpose only: to prune the

[244]It is necessary to warn here that there is no uniform understanding of what a 'goal' is in a design system. Various knowledge constructs or structures have been so termed.

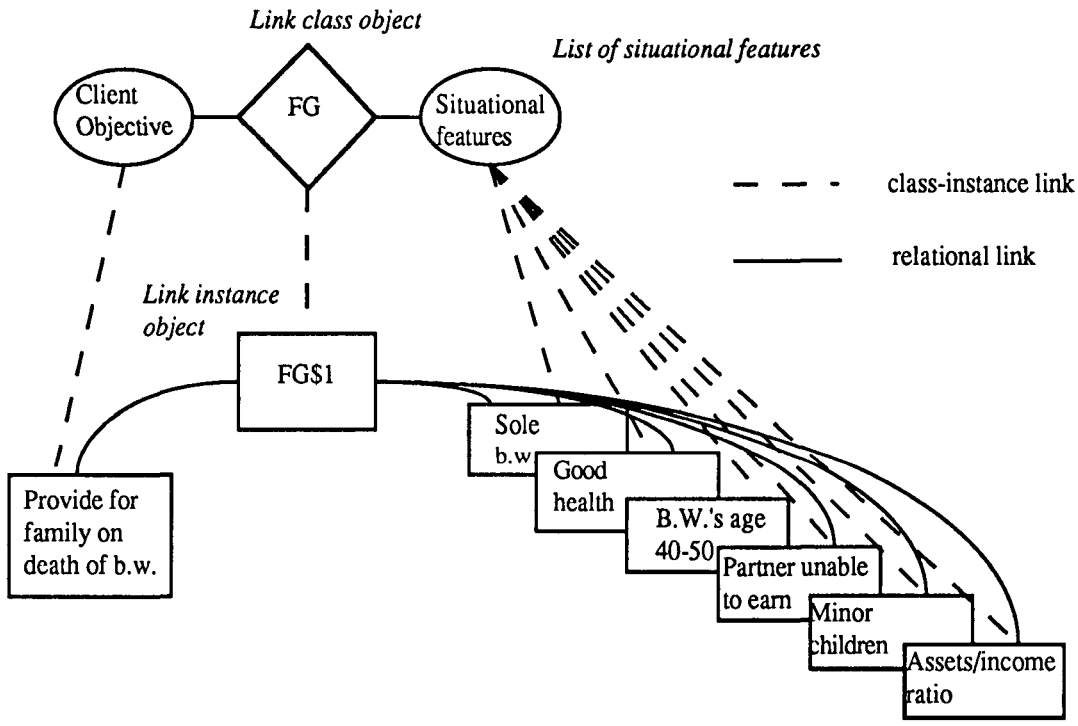

Figure 5 A goal [b.w. = bread winner of a family)

search space in a preliminary manner, without being determinative of which cases are eventually presented as part of the solution set.

The first step in searching for the cases which will eventually be matched through goals is to prune the search space in a manner which does not diminish the stance that generalized heuristics should not be used to synthesize the solution.

In addition to the case-base, the system has to contain the kinds of heuristic rules which were described above in the context of investment advisory systems. These rules are of limited power—they can identify the types of products which *may* be relevant to the client's situation, but cannot take into account all the complex interactions and constraint satisfactions that a comprehensive planning process must. The rules are initially applied to the situational features and client objectives in the problem case—in a manner identical to that used in the simplest investment advisory systems—and a list of all products (solutional elements) that may *possibly* be found in the target cases is created. It may be noted that this list can contain solutional elements which are mutually exclusive or which violate constraints when considered together. No account is taken during the creation of the list of any interactions between the heuristics used to select suitable products. Nor is it sought to exclude the possibility that some of the invoked rules should properly have been defeased or made invalid because there are present exceptional facts of which the rules themselves do not take account. In other words, all that is produced is a cumulative list of possible solutional elements.

The list is then used to select from the case-base all those cases which include recommendations of one or more of these solutional elements. These selected cases form the initial search space.

Finding a solution to the problem case

The user can present the problem to the system in terms of the situational features, the objectives, and the goals. The point of finding similarity between the problem case and cases in the search space is to see what solutions are fit to achieve the objectives of the problem case in light of its situational features, i.e. which cases have goals that match the goals of the problem case. The best answer is to find one case that is fully similar to the problem case; when this is not possible, a set of partially similar cases that together cover all the goals of the problem case is to be found.

The process of search starts by looking for a case which has the maximum number of matches with the goals in the presented problem. If all the goals in the problem cannot be matched with one case, the remaining goals are represented to the system as a new problem so as to find another case with the maximum possible match, and so on. It is acknowledged that this is not aimed at producing a set that is most optimal. It is sufficient that the user is provided with the basis on which to construct a reasonable solution. More sophisticated algorithms could be implemented in the FINESSE architecture so that different kinds of sets of useful cases are presented to the user.

To see what constitutes a match, take two cases, A and B. Case A is said to be fully similar to case B if case A has such goals that some of them match all the goals of case B. This is a undirectional similarity because case A may have additional goals which are not utilized for matching. Partial similarity is when case A contains at least one goal that matches a goal in case B; there are degrees of partial similarity.

Goal x trivially matches goal y if x and y are either identical or is the attributes of one are specializations or generalizations of some of the attributes of the other. In other words, the matching relies on the generalization-specialization or class-subclass relationships between the primitive concepts (situational features and objectives) found in the goals to be matched. Siblings in hierarchies of primitive concepts are considered to match.

Matching of goals can be expanded by looking at the explanatory links contained in the cases being matched. For instance, consider the situation where goal G1 in case A and G2 in case B are being considered for matching, and that goals G1 and G2 have the same objective but are not siblings, the difference being that G1 has situational features SF1.1 and SF1.2 while G2 has SF2.1 and SF2.2 (other situational features match). Now, there may be an explanatory link attached to case B saying that SF2.1 *is optional*; also that SF2.2 *is concomitant with* another feature which has already been matched. In light of these explanatory links the matching engine will consider goals G1 and G2 to be matched. Other kinds of explanatory links may mean that two goals can be matched though they do not share even an objective.

13.1 The central theme: functional or purposive similarity

FINESSE seeks to establish what may be called 'functional' or 'purposive' similarity between the description of the problem and one or more cases in the case-base. This task is central to what the system is meant to do and FINESSE may properly be called a simple *similarity engine*. It is necessary to explain what is meant by functional or purposive similarity, both more generally and in FINESSE in particular. This would allow us to appreciate just what it is that the system searches for.

The process of making analogies between two states of affairs allows us to infer from the fact that there are some similarities between the states, that there must be other similarities. In this sense, a similarity between two situations is a commonality at some level of abstraction. For instance, a student learning about heat transfer can map the knowledge that water falls from a high elevation to a lower one into the heat transfer domain and from there derive an understanding as to the direction in which heat flows between two bodies at different temperature levels (Gentner & Toupin, 1986). To make this analogy the student has to conclude at least that:

- The two situations involved relational chains that were structurally similar.
- That the relation 'is at a higher elevation than' is mappable across to the relation 'is at a higher temperature than'.

The above is an example of what is often called 'without-domain' analogy. There are complex issues involved here such as systematicity (Gentner & Toupin, *op cit*) and the correspondence of the process of analogizing with human learning and discovery. Fortunately, we are dealing with 'within-domain' analogy only. Finding a similarity between two situations is the crucial step in making a within-domain analogy. Where all concepts to be considered belong to the same domain

and it can be taken for granted that identical predicate structures have the same semantic significance throughout, then:

> "Object similarity can potentially be reduced to predicate similarity: two objects are similar to the extent they serve as arguments of similar predicates." (Holyoak & Thagard, 1989).

Then the treatment of each case as a separate object with an explicit structure means that:

> "[T]he problem is already analyzed into a problem structure . . . Thus, it is straightforward to identify the same functional components (e.g. goals and constraints)." (Shinn, 1988).

However, whilst finding within-domain similarity when a object-oriented knowledge representation is used may be much simpler than finding without-domain similarity, it is still considerably more complex than merely looking for siblings or close relatives in an object class hierarchy. FINESSE's cases cannot themselves be said to be organized in a *simple* hierarchy, it is only goals which are so arranged and a case may have several goals. While the goals can be arranged in a hierarchy, as each case may have several goals, the cases themselves are in a richer organization. The question then is: What is the significance of attempting to determine similarity through goals?

A goal in FINESSE maps an objective which the case is considered to be achieving to (a) the solutional elements which achieve the objective, and (b) the situational facts considered to be relevant to the achieving of the objective by means of those solutional elements. The situational facts are those which, in an expert's opinion, directly influenced the outcome or shape of that (undefined) part of the solution which achieves the particular objective. Of course, there is room for doubt, and we have not been rigorous about what is meant by directly influenced, but that does not entirely vitiate the mechanics used.

FINESSE's way of representing goals is somewhat different from that used in other CBD and CBP systems. The difference arises from the nature of the domain, rather than from any new thesis of design. For instance, in the planning system CHEF (Hammond, 1988) typical goals are:

- Make a stir fry dish
- Include broccoli in the dish
- The broccoli should be crisp.

The case (plan) consists of the steps required to cook a dish. The goals that the case satisfies can be derived from its knowledge of stir frying. CHEF also knows what goals *should* be satisfied if the plan is to be considered as having ended successfully. As such, the goals are not directly stated within the case; they are derived using the domain knowledge. In fact, in many other CBD (Navinchandra, 1988; Barletta & Mark, 1988) and CBP (Hammond, 1988) systems, a causal explanation of the goals and subgoals of the case is stored with the case. Establishing similarity then requires finding a common causal relationship among attributes in the test situation and the target case, rather than a commonality of attributes alone (Gentner & Toupin, 1986). In essence, what such CBD (or CBP) systems do is to *assume that explanations attached to goals are independent and free of interactions and can be combined in another context.*

This cannot be assumed in the financial planning domain where a case has a normative status: the case goes to establish reasonableness; it may not be possible to attribute the same status to individual elements or steps in the case. In FINESSE the goals that the case is supposed to have satisfied are *directly specified* by the contributing expert during the acquisition of the case from the expert; the satisfaction of particular goals cannot be inferred merely by using the knowledge possessed by the system. Consequently, FINESSE is not able to abstract explanational knowledge from a particular case and applying it to another case. That is a job for the user. Further, it does not undertake to modify or amalgamate cases to formulate a new solution: the task is to find a superset of cases which individually or collectively are capable of achieving the conjunctive goals set by the problem case.

The explanatory links in FINESSE's cases are not of the same level of detail as found in many other systems, and their function is to broaden the scope of the match when similarity is being judged. In contrast, generally in CBD, explanatory links consist of a chain of known relations connecting a feature (or conjunction of features) to another feature or class (Bareiss, 1989).

It should at this point be mentioned that the architecture that is presented here is just one way of using CBD techniques in financial planning. CBD is a field that is rapidly developing, and researchers are constantly proposing and implementing novel formulations. There is a great likelihood that others building systems for financial planning or related applications will deviate significantly from the proposed route.

14 Planning using 'prototypes'

For the sake of completeness, it is interesting to note that even in the financial planning domain, or to be precise, estate planning, some researchers have taken an approach which does not shirk from the use of generalized knowledge, as opposed to non-generalized, specific instances of problem solving.

Scholobohm and McCarty (1989) have proposed a system called EPS II which also relies on estate planning precedents. The differences with FINESSE are, however, profound. In EPS II the precedents for estate plans are represented by means of the Language for Legal Discourse (LLD), developed by McCarty (1989). The idea is to find universal primitives of legal concepts including relationships like 'owns', 'parent of', 'resident of', various kinds of property interests and the actions that change these interests.

The cases or 'prototypes'—they are more like exemplars than records of situations which professionals happened to face and the solutions they provided—are thought to be applications to specific facts of generalized legal knowledge, such as the Internal Revenue Code. Relevant parts of the IRC are also formalized using the LLD representation. Once a prototype close to that required for the client is found, the generalized knowledge is applied to it to deform the prototype step-by-step until the desired result is achieved.

The idea is, doubtless, most powerful and has profound implications. However, the mechanisms that might be used have not yet been stated by the researchers. Consequently, it is not possible to say what compromises are going to be made to prevent the computational intractability which soon follows in a complex domain when concepts are sought to be represented in terms of universal primitives.

15 Conclusions

One way of studying related systems is through seeing whether the developers have avoided or minimised the so-called 'representational mismatch'.

The 'representational mismatch' is the incongruity between the knowledge representational formalisms and the formulation of the task by domain experts (Gruber & Cohen, 1987). The best-known cause of the representational mismatch, the use of a relatively non-discriminatory formalism like production-rules, has been discussed at length elsewhere (Keravnou & Johnson, 1986). The knowledge expressed in rules is often a collection of heuristics relevant to a particular task without any explicit representation of the underlying structure or principle of behaviour of the problem domain (Hudlicka, 1988). However, the mismatch may exist even where a more sophisticated representation like frames is used, because there exist issues such as which links to pursue for inheritance, which are the artifacts of the representation without any natural connection with the control issues natural to the task (Chandrasekaran, 1985).

As such, we have focused on the architectural and representational principles behind the systems, while at the same time pointing out the precise real world tasks sought to be carried out.

Another view point that might be taken is that obtained by characterizing systems as 'deep' or otherwise. Deepness is often equated with causal models or reasoning from first-principles. However, reasoning at a lower level of granularity does not necessarily make a model more powerful (Bylander & Chandrasekaran, 1987). In any case, it is rarely possible in domains without an agreed underlying theory to say what are indeed the first principles: one person's first principles

may be another's empirical associations. Moreover, as will have been seen for the case-based reasoning approach outlined above, there may be reasons peculiar to the domain that require that little weight be given to reasoning which is deep in the present sense.

A better way of thinking about deepness may be to treat it as a relative term with a number of measures, principally, the explicitness of factual knowledge and reasoning knowledge and the understandability of the solution progression at intermediate stages, (Keravnou and Washbrook, 1989). This type of deepness is, essentially, the consequence of recognizing the relevant task structures and modelling them in the best manner permitted in the current state of the technology. It may be seen from the details of systems presented here that, as the market has demanded increased sophistication, developers have sought to build progressively deeper systems in the latter sense.

References

Bareiss, R, 1989. *Exemplar-Based Knowledge Acquisition* Academic Press, San Diego.

Barletta, R, and Mark, W, 1988. "Explanation-based indexing of cases" In: *DARPA Case-Based Reasoning Workshop* pp 50–60.

Brown, DC and Chandrasekaran, B, 1989. *Design Problem Solving: Knowledge Structures and Control Strategies* Morgan Kaufman, San Mateo, CA.

Butler, A and Chamberlin, G, 1988. "The Aries Clus experience of expert systems in insurance and investment" In: DS Moralee (ed.), *Research and Development in Expert Systems IV, Proc. Expert Systems '87* Cambridge, Cambridge University Press, pp 246–257.

Bylander, T and Chandrasekaran, B, 1987. "Generic tasks for knowledge-based reasoning: the 'right' level of abstraction for knowledge acquisition" *International Journal of Man-Machine Studies* **26** pp 131–144.

Chandrasekaran, B, 1985. "Generic tasks in knowledge-based reasoning: characterizing and designing expert systems at the 'right' level of abstraction" *Second Conference of Artificial Intelligence Applications* pp 294–300.

Chierici, A, Filippini, MG and Minati, M, 1989. "PORTAFOGLIO: A portfolio advisor application" In: N. Shadbolt (ed.), *Research and Development in Expet Systems VI, Proc. Expert Systems '89* Cambridge, Cambridge University Press, pp 140–152.

Coats, PK, 1988. "Why expert systems fail" *Financial Management* pp 77–86, Autumn.

Denna, EL, Hansen, JV and Meservy, RD, 1991. "Development and application of expert systems in audit services" *IEEE Transactions on Knowledge and Data Engineering* **3**(2) pp 172–184.

Duda, RO, Hart, PE, Reboh, R, Reiter, J and Risch, T, 1987. "SYNTEL: Using a functional language for financial risk assessment" *IEEE Expert* **2**(3) pp 18–31.

Freedman, RS, 1991. "AI on Wall Street" *IEEE Expert* pp. 3–9, April.

Gardner, AvdL, 1987. *An Artificial Intelligence Approach to Legal Reasoning* MIT Press, Cambridge, MA.

Gentner, D and Toupin, C, 1986. "Systematicity and surface similarity in the development of analogy" *Cognitive Science* **10** pp 277–300.

Gruber, TR and Cohen, PR, 1987. "Design for acquisition: principles of knowledge-system design to facilitate knowledge acquisition" *International Journal of Man-Machine Studies* **26** pp 143–159.

Hammond, K, 1988. *Case-Based Planning* Morgan Kaufman, San Mateo, CA.

Heuer, S, Kock, U and Cryer, C, 1988. "INVEST: An expert system for financial investments" *IEEE Expert* pp 60–68, Summer.

Holyoak, KJ and Thagard, P, 1989. "Analogical mapping by constraint satisfaction" *Cognitive Science* **35** pp 81–125.

Hudlicka, E, 1988. "Construction and use of a causal model for diagnosis" *International Journal of Intelligent Systems* **3**(3) pp 315–349.

Hunter, T, 1991. "Father sues over school fees investment plan" *The Guardian*, February 9, p. 15.

Keravnou, ET and Johnson, L, 1986. *Competent Expert Systems* Kogan Page, London.

Keravnou, ET and Washbrook, J, 1989. An analysis of the architectural requirements of second-generation expert systems *The Knowledge Engineering Review* **3**(4) pp 205–233.

McCarty, LT, 1989. "A language for legal discourse, I, Basic features" *Second International Conference on Artificial Intelligence and Law.*

Michaelson, RH, 1984. "An expert system for Federal tax planning" In: *Expert Systems* **1** pp 149–167.

Mital, V and Johnson, L, 1990. "Professional negligence and the reasonableness defence in financial-legal expert systems: a developer's perspective" *Second National Conference on Law, Computers and Artificial Intelligence.*

Mital, V and Johnson, L, 1992. *Advanced Information Systems for Lawyers* Chapman & Hall, London.

Navinchandra, D, 1988. "Case-based reasoning in CYCLOPS, a design problem solver" In: *DARPA Case-Based Reasoning Workshop* pp 286–301.

Prosser, WL, 1971. *Handbook of The Law of Torts* West Publishing Co., St. Paul, MN.

Schank, RC, 1982. *Dynamic Memory—A theory of Reminding and Learning in Computers and People* Cambridge University Press, Cambridge.

Scholobohm, DA and McCarty, LT, 1989. "EPS II: Estate planning with prototypes" In: *Second International Conference on Artificial Intelligence and Law* pp 1–10.

Shinn, HS, 1988. "Abstractional analogy: a model of analogical reasoning" *Proc. DARPA Case-Based Reasoning Workshop* pp 370–387.

Shortliffe, EH, 1976. *Computer-Based Medical Consultations: MYCIN* Elsevier, New York.

Stansfield, JL and Greenfield, NR, 1987. "PlanPower: A comprehensive financial planner" *IEEE Expert* pp 51–60, Fall.

Stevenson, H, 1989. "Expert systems in the UK financial services sector: a symbolic analysis of the hype" In: GI Doukidis, F Land and G Miller (eds.), *Knowledge-based Management Support Systems* Ellis Horwood, Hemel Hempstead.