

SUPPORT FOR POLICY MAKERS: FORMULATING LEGISLATION WITH THE AID OF LOGICAL MODELS

T. J. M. Bench-Capon
Department of Computing
Imperial College
180 Queen's Gate
London SW7
England

Introduction

The work described in this paper is being undertaken as part of a large five year project, currently at its half way stage, involving collaboration between academics, industry and Government, which is funded under the United Kingdom Alvey Programme. The project as a whole is addressing the provision of knowledge based decision support to large legislation based organisations; the work described here, the provision of support to policy makers, is one strand among several. The particular organisation providing the focus of the project is the UK's Department of Health and Social Security, which is responsible, amongst other things, for making and administering law relating to Social Security benefits. The law in this area has a number of features attractive to the project; in particular it is complex and its parts are highly inter-related, and much of it is definitional in character. Moreover, it is also an area of law which is subject to constant change. As will be seen, these features make it an ideal domain for the exploration of the issues to be discussed in this paper.

In this paper I shall first discuss the nature of the application. I shall then describe the work done within the project to date, giving details of the approach taken, together with a worked example. I shall conclude by indicating some of the technical AI issues raised, and the expected future directions of the project.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0-89791-230-6/87/0500/0181 \$0.75

The nature of the problem

Most attempts to use AI in relation to the law have taken the legislation as something fixed that needs to be reasoned with and applied. Of course, legislation is not fixed. Legislation is passed so as to achieve certain aims; it often fails to achieve these aims and so it needs to be amended. Equally the aims themselves change over time, and the legislation must be amended to reflect these changes. Some legislation is relatively stable, but other parts are in a state of constant flux. Where the fact that legislation continually changes has been noticed by those working in AI they have usually made some claims as to the ease of amendment of AI programs over conventional programs, or that AI programs can be used to assist the act of drafting. This paper is about something else. It is about providing support to the people who are charged with deciding what the legislation should be if they are to effect their policies - the people who are responsible for telling the draftsman what to draft. The practical potential of such a system is enormous. The cost of mistakes at this level, in terms both of actual cost of subsequently correcting them, and in the social confusion, hardship and inequity that can result from bad decisions, is so enormous that even a slight improvement would be highly cost effective. Moreover, and perhaps more importantly from a research perspective, consideration of systems that could play a role here throws up an abundance of problems for AI, some of a character quite different from those found in other legal applications.

One word of caution: the task we are looking at is highly complicated, skilled, and one for which people must be held accountable. We have no intention of producing a system which could perform the task on its own; on the contrary we see that as both unrealistic and undesirable. What we have in mind is rather the provision of a set of tools that will provide some support for sub-tasks within the overall

process. The system may be able to say what the consequences of a change are, and detect some that the user would have missed, but it must remain the user's responsibility to say whether those consequences are desirable or undesirable, acceptable or unacceptable.

Who are policy makers and what do they do?

Social Security policy makers (in the UK at least) do not have formal legal training; their expertise is supposed to lie in an appreciation of the (not necessarily coherent) policies of their masters (ministers of the Government), and their ability to translate these policies into a clear, coherent and consistent policy statement. This statement can then form the basis for instructions to legal draftsmen, who produce the legislation required to fulfil the specification. Policy makers are not engaged full-time in extending and amending legislation; they are also required to explain and justify policies and the current state of the legislation, and to monitor the legislation to ensure that it is working as they desired it should. Their job requires that they have a wide ranging knowledge of their area of the law and the world more generally. As well as the legislation itself they need to know the aims which the legislation is intended to achieve, facts about the world in which the legislation is intended to operate, theories about the world, such as how people will respond to legislation, and how the legislation will be applied.

Inevitably there will be problems somewhere in this process, causing the legislation to fail to achieve the desired results; this may be because their theories are inadequate, or because of some deficiency in the legislation, or because society has changed, or even because the aims of the Government have changed. These problems may be termed policy problems. When confronted by a policy problem a policy maker will have to do a number of things.

Types of Knowledge Used by Policy Makers

Firstly, and most importantly, a policy maker must have some understanding of the legislation with which he deals. The depth of his understanding of a particular piece of legislation will, however, vary according to the closeness of the piece of legislation to his area of responsibility. Thus a policy maker responsible for a particular contributory benefit, say Sickness Benefit, will have a detailed knowledge of

the legislation relating to whether or not a person is to be treated as incapable of work. He will also know that to get Sickness Benefit a person must satisfy certain contribution conditions. The precise nature of these conditions, however, is the responsibility of a different policy maker, and the person responsible for Sickness Benefit need have no knowledge of the details. It is enough for him to know that they exist and must be satisfied, and are broadly supposed to indicate that a person satisfying them is normally part of the labour force. Other parts of the legislation will be of no concern to the Sickness Benefit person at all; he may know nothing of War Pensions, for example, beyond that such a benefit exists.

Thus the policy makers' knowledge of the legislation is something of an abstracted and idealised model of the legislation; he will have a general understanding of the whole scope, but a detailed understanding of only the part which is his concern.

Secondly a policy maker must have an understanding of the intention underlying the various benefits. He will know that Sickness Benefit is meant partially to replace income that a person cannot earn whilst he is sick, rather than, for example, to meet any extra expenses attendant upon sickness. This understanding will be of extreme importance in determining his response to complaints about the operation of the benefit and proposals for change.

As well as the specific intentions underlying particular benefits, the policy maker must also be aware of the more general policy aims of the Government. Thus there may be an aim to reduce reliance on means-tested benefits, or to provide incentives for people to work rather than claim benefits. Some of these aims will be sufficiently specific and pressing as to necessitate legislative change to give them effect, others will operate as constraints on the desirability of proposed changes.

Next the policy maker must know a number of facts about the world. He must know such mundane things as that a person cannot be both male and female, and that men cannot get pregnant. Without knowledge such as this he could not begin to frame legislation to fulfil his aims.

Moreover he will also need a number of theories about the way the world works. Some of these will be behavioural theories about what people may do in

response to certain situations. Others will be social theories. The UK Social Security scheme can be traced back to the Beveridge Report (1942), and reading that report gives a clear sense of the social theories operating at that time. These theories are helpful in explaining some of the features of the Social Security scheme introduced in 1948. For example it is assumed that a married woman would always look to her husband for financial support. In consequence, married women could opt out of paying contributions and receiving benefits, and husbands could claim additional benefit for their wives, whereas wives could not claim additional benefit for their husbands. To take another example, unemployment was considered to be a short term catastrophe, and the working week was considered to be six days, Monday to Saturday. For this reason Unemployment Benefit is calculated on a daily basis and is payable for one year only. Further no Unemployment Benefit is payable in respect of Sundays. Now these things change over time, and these theories which motivated some of the structure of the Social Security system may not be true today. Many policy problems arise because benefits were designed to operate in the context of certain theories being applicable, whereas this is no longer the case.

Lastly the policy maker must know a good deal of statistical information. He must, for example, be able to estimate how many people are affected by a given circumstance so that he can gauge the seriousness and extent of the problem, and estimate the cost of solving it.

Life Cycle of a Policy Problem

The policy maker may be alerted to the existence of a potential problem in a number of ways. Most often this will be by complaints from people about the operation of the system. These complaints may arise from letters to the Government department, by Parliamentary Questions, or by articles in the press. The first thing that the policy maker must do is verify that the complaints are well founded and that there is actually a problem. To take an example, a person may complain that they do not receive a pension even though they are over pensionable age. This is not in itself a problem, since age is not the only condition relevant to receive a pension. In addition a person must be retired and must have paid the relevant National Insurance contributions. The underlying aim of Retirement Pension is to pay a pension to those

of a certain age who have retired and who have paid the relevant contributions. There is no problem involved in the fact that some people of a certain age do not receive such a pension. There is a problem only if some people of the relevant age who are retired and who have paid the contributions fail to receive the pension. Many complaints do not raise problematical issues, simply because the people who made them fail to recognise the existence of certain other conditions that were always present within the aims of the legislation. Such people can be answered by an explanation of the rules applying to their case and the justifying policy aims.

Sometimes, however, it can be established that a genuine problem does exist; in other words that there is a group of people who were intended to benefit from a provision who fail to benefit from it. It is then necessary to say exactly who it is that constitutes this group and to explain why the legislation fails to meet their case. The individual complainants may differ widely in their circumstances, but there will be some common factor which explains why they fall into the problem category. It is the identification of this common factor which enables the policy maker to see precisely who it is that is suffering from the defects in the legislation, and to explain why this is so.

One reason for the problem may be that some of the underlying theories that were used in the original design of the legislation are not correct. Suppose that there was no retirement condition on retirement pension, perhaps because there had been a theory which said that no one would work if they could receive a pension, or perhaps because there had been a theory that no one over pensionable age would be able to retain a job. This would mean that some people in full time work over pensionable age would additionally be able to claim a retirement pension. This is in conflict with the aim that such a pension should be payable only to those people without a job, and the legislation would be defective. Given a conflict between the theories and the fact that there are people over pensionable age who are in full time employment, the theory must be modified. Once the theory is modified or abandoned, it becomes clear that some extra condition analogous to the retirement condition is required to enable the legislation to achieve its desired aim.

The next stage in the process is to decide what can be done to cure the problem. Usually there will be

several different ways of amending the legislation to reach a state where the problem no longer occurs. The retirement condition in the example above, for example, could be expressed in a variety of ways; any work at all might be made a disqualification, or else an hourly limit on work done could be imposed, or else the pension could be reduced in accordance with earnings in a variety of ways, pound for pound, with a cut off at a certain earnings threshold, or by a certain proportion of earnings. The policy maker should at this point generate as many candidate solutions as his ingenuity will allow.

Having generated the solutions the policy maker must next arrive at some kind of evaluation of them. He does this by exploring the consequences of the various solutions, so as to arrive at the good and bad points of each. Some will have to be rejected because they have consequences that conflict with other aims. As an example if there was an aim to remove the need for means testing for benefits in the case of the elderly, then the retirement condition could not be made to depend on earnings. Other solutions may generate other problems because of an interaction with other pieces of legislation. A specific example of this will be treated in detail later. Some solutions may be better than others on grounds of cost or administrative simplicity. All these consequences need to be weighed by the policy maker, but, of course, the consequences need to be identified before they can be weighed.

The policy maker will now submit a proposal to the Government ministers, setting out several options for changes that would eradicate the problem, together with a recommendation based on the advantages and disadvantages of the various options. The minister will select one option, which the policy maker will be expected to progress.

The final stage comes with the turning of the selected option into legislation. This is done by specialist legal draftsman. The policy maker's task is to specify the solution in such a way that the draftsman is able to turn it into legislation. Clarity and lack of ambiguity are paramount here.

To summarise we have identified the following six stages in the activity:

- 1) verify that a problem exists
- 2) arrive at a precise characterisation of the problem

- 3) [optionally] modify theories
- 4) generate candidate solutions
- 5) explore the consequences of the solutions
- 6) specify solution for legal draftsman

Our analysis thus enabled us to identify the stages in the activity and to see what tasks we wished to support and the different kinds of knowledge used in tackling them. An approach which seemed promising and natural was to form a logical model of this information and support the tasks by proving consequences from and about our model. Clearly the next step was to build such a logical model and to provide a means of proving consequences from and about this model. This would enable us to explore the usefulness of these techniques, and to discover extensions to them, by using a prototype system.

Approach Taken

At the heart of such a system lies a formalisation of the various elements that make up a policy maker's knowledge. This representation of this knowledge is itself a demanding task, both in terms of the size of the knowledge base that is required in a realistic system, and in terms of the demands that will be made on the representation. We cannot predict at this point what all these requirements will be, but we know that more general theorem proving techniques than those provided by, for example, PROLOG will be required. A simple example will suffice to show this. In the UK a person is over pensionable age if and only if he is aged over 65, or he is a woman aged over 60. A rule might be something like

```
OfPensionableAge(X) <=>
  (Age(X,A) & ((A >= 65) v
    (Sex(X,Female) & A >= 60)))
```

If we are using the Horn Clause subset of logic for execution with a PROLOG like theorem prover, we will need to commit ourselves to one of the elements being the head, perhaps

```
OfPensionableAge(X) <- Age(X,A) & A >= 65
OfPensionableAge(X) <-
  Age(X,A) & Sex(X,Female) & A >= 60.
```

This formalisation of the rule would allow us to show that anyone aged over 65 is of pensionable age. But we might equally well need to show that anyone over pensionable age is aged over 60. This we cannot do

from the Horn clauses above, although it is an indisputable consequence of the original biconditional.

There was available to us no ready made system that would allow us to try out our approach. Since our aim was only to explore the potential of our approach we needed to construct a simple theorem prover that we could use to derive consequences about our model. In order to make the construction of a theorem prover as simple as possible at this stage we made the important simplifying restriction that we would consider only single individuals with no dependants. This enabled us to adopt an approach in which the sentences of the model could consist of propositions linked by the standard logical connectives. The propositions were not, however, entirely unstructured. Propositions either comprised an attribute and a value or range of values, for that attribute, or an arithmetic expression comprising an equation of some attributes using arithmetic operators. Thus we were allowed propositions such as *Income is £63*, *Income is between £60 and £65* and *Income = Benefit + Earnings*.

Because the simplification of the domain of consideration allowed us to disregard relationships between people, it was unnecessary to include reference to the owner of the attributes (all would be owned by the same person). With each attribute was associated either a list of the possible values for the attribute, or, in the case of attributes with a numeric value, a range of possible numeric values. This enabled the negation of a proposition to be treated as an assertion of the complement of the value of the proposition with respect to the possible values of the attribute concerned. For example *not age over 65* is equivalent to *age less than or equal to 65*. These restrictions enabled the construction of a relatively simple theorem prover which would take in the full syntactic range of logical expressions and derive all valid inferences from them, and which allowed the integration of declarations of arithmetic relationships between attributes. Despite its limitations, the theorem prover thus had the right properties to allow the exploration of the types of support that a theorem prover could provide for the policy making task. A fuller discussion of the theorem prover is given in section 3.1 of [1] and [2].

The basic operation of the theorem prover was to receive a series of assertions of values, or ranges of

values, for attributes, and to deduce all the consequences of those assertions. If a contradiction was detected that would, of course, be reported. The underlying intuition was that the series of assertions corresponded to a description of a class of people, and the consequences represented things that were true of all persons who satisfied the description. Thus the initial range for age is 0 to 137; it is the case that a person's age lies in that range for any person whatsoever. If we now assert that pensionable age is true, age is deduced to lie in the range 60 to 137; that is true of all persons of pensionable age. If we further assert that sex is male, then age is deduced to lie in the range 65 to 137; that is true of all males of pensionable age. Thus the logical model of the relationship between the attributes together with the description of a class of people constrains the range of possible values of other attributes.

Given this theorem prover we produced a logical model of the legislative provisions of the social security system, together with a certain amount of common sense information such as that men can not get pregnant. The coverage of our model was broad in that we covered all the major social security benefits, but shallow in that we used many attributes that had the potential for more detailed definition within an expanded system. This gave us a sufficient degree of interaction between the rules to pursue our exploration.

The following are two typical rules from the prototype system; the first describes the top level entitlement condition for the short term scale rate of Supplementary Benefit (STSR) and the second describes the amount of supplementary benefit that a claimant would receive.

```
R1 (EntitledToSTSR is True) if-and-only-if
  not ((Sex is Female)
        and (LivingWithPartner is True))
  and (Age is-greater-than-or-equal-to 16)
  and (not (Job is FullTime) or
        (RequiredToRegisterForWork is False))
  and (FullTimeEducation is False)
  and (Capital is-less-than 2500)
  and (NormallyResidentInGB is True)
  and (ResourcesForSB is-less-than 2570).
```

```
R2 if (EntitledToSB is True)
    then (AmountOfSB =
          RequirementsForSB - ResourcesForSB)
```

Note that R1 gives necessary and sufficient conditions, whereas R2 gives only sufficient conditions. The consequent of R2 might be fortuitously true while the antecedent is false, since, someone might not be entitled to Supplementary benefit at all, and yet happen to have requirements equal to resources.

The prototype model contained some 130 such rules; because the flexibility of format means that different rules may contain significantly different amounts of complexity it might be a fairer measure of size to think of them as equivalent to sets of if-then clauses with single propositions as heads. Thus a rule of the form $P \Leftrightarrow Q$ is equivalent to the four clauses if P then Q, if Q then P, if -P then -Q and if -Q then -P. The 130 rules were equivalent to some 1500 such clauses.

It should be stressed here that no special claims are being made for the theorem prover described above; it was merely used as a tool to carry forward the analysis. Any theorem prover with the desired properties could be used to replace it in a finished application. Indeed one of our aims was to identify what kind of theorem prover would be needed in a practical implementation.

The theorem prover described above supports the basic operation of assertion and examination of consequences. We provided in addition a number of other facilities for the examination of the formalisation itself. By combining these primitive facilities to ask more complex queries, the user was provided with the ability to explore the formalisation so as to assist in the policy makers' tasks identified above.

The basic facilities provided are as follows:

assert: this enables the user to assert that a proposition had a certain value or range of values and so cause the values of other attributes to be constrained.

examine: this allows the user to examine the current range of possible values for an attribute.

why: this enables the user to get a trace of the proof which has led to the constraining of a possible range of values for an attribute, so that he can see which

rules and assertions had been used to produce the effect.

what: this enables the user to see what assertions would cause the possible values of a particular attribute to be constrained, and to which value it would be constrained if the assertions were made.

which: this enables the user to see which sets of assertions would, if asserted, constrain the value of a specified attribute to a specified value.

As an example of how these facilities can be combined, consider the question as to whether a goal of the form "all people satisfying description D should have value V for attribute A" is satisfied. We use **assert** to assert the components of D, and then **examine A** to see if its value has been constrained to V. If it has, the goal is satisfied; if not, we may ask a **which** question and receive a set of further descriptions D1, ..., Dn, the addition of any of which to D would cause the the goal to be satisfied. In effect this tells us for which people satisfying D the goal is satisfied. Alternatively we may ask a **which** question with respect to A having some value other than V, and receive a set of descriptions, any of which would, if true of a person, explain why they failed to have V as the value of A.

Another package of questions allowed the graphical display of the relationship between two attributes. Suppose we wish to see how benefit varies with age for a man incapable of work. We first assert the description (Sex is Male) and (IncapableOfWork is True). We then **assert** a series of values for age and **examine** the maximum and minimum values of benefit. This gives a set of points which can be used to plot a graph. Such a graphical presentation can be very useful in spotting anomalies.

Example Treatment of A Problem Using the Prototype

I shall now give a detailed description of the use of the prototype described above in the examination and solution of a particular policy problem. This example will show that even with the relatively simple facilities outlined above, a substantial exploration of the problem is possible.

A specific problem, namely "the invalidity trap" formed a target problem for the prototype to tackle.

This problem involved the interaction between two benefits known as Supplementary Benefit and Invalidity Benefit. Supplementary Benefit is the so-called "safety net" of the UK Social Security System, and is meant to represent a minimum level of benefit, payable to those who do not qualify for other benefits. Supplementary Benefit is payable at two rates; the short term scale rate (STSR), payable for the first twelve months on supplementary benefit and the long term scale rate (LTSR) payable thereafter (except in the case of the unemployed). LTSR was not part of the original scheme; it was introduced because it was recognised that people who had been on Supplementary Benefit for a substantial period of time were suffering considerable hardship. This was explained by a theory to the effect the needs of those who have been on a low level of income for a considerable time were greater because things such as clothes and furniture would begin to wear out and need to be replaced. Whilst such needs could be ignored in the short term, they needed to be satisfied in the long term. Invalidity benefit (IVB) was a contributory benefit introduced in response to a desire to increase financial help for the long term sick. But some people in receipt of IVB were entitled to an amount greater than STSR but less than LTSR. This meant that they could not receive STSR, and so could never qualify for LTSR. This in turn meant that whilst they were better off for the first twelve months, after that period they were receiving less than LTSR, which was the accepted minimum for someone not unemployed who had been on benefit for twelve months or more. Thus the idea which underlay the introduction IVB, which was to give more money to the long term sick who had contributed to the National Insurance scheme was frustrated. The introduction of IVB the perverse effect of making some of its recipients worse off if their sickness lasted more than a year, which was the very opposite of what had been intended. This was an unfortunate and unforeseen anomaly in the legislation, and was eventually dealt with by making IVB as well as STSR a passport to LTSR after twelve months.

The problem is particularly interesting because it shows how two pieces of legislation, both of which would achieve their aims in isolation, interact so as to frustrate the aim of one of them.

Using the Prototype

I shall now describe one way in which the prototype could be used by a policy maker confronted with the invalidity trap problem, in terms of the six sub-activities mentioned above. There is, of course, not a single route that could be taken through the problem.

1) verify that a problem exists

To support this task it is necessary only to show that being in receipt of IVB for more than 12 months does not imply that the recipient has income greater than the rate of LTSR. We therefore assert that the person is entitled to IVB and that the duration of his benefit period is greater than 12 months. A number of consequences will be deduced from this. In particular it will be deduced that the income of people satisfying this description will be in the range £31.45 to £156.50. Now we know that the rate of LTSR Supplementary Benefit is £32.70, and so we see that it is possible that to be in receipt of IVB for more than 12 months and get less money than would be received on LTSR. Thus a problem really does exist.

2) Next we must arrive at a precise characterisation of the problem. This involves discovering which people in receipt of IVB for more than 12 months receive less than £32.70. The user therefore asks what affects the value of income for people satisfying the initial description, and he will find that it is affected by the amount of income maintenance benefit (a benefit designed to replace a person's lost earning power, of which IVB is one), other benefits, and earnings. Both other benefits and earnings are irrelevant to the problem under consideration, since if a recipient of IVB was taken over the LTSR rate by these means, this would not result in the satisfaction of the original aim, which was to take his income above this level as a consequence of his unassisted benefit. If we assert that other benefits and earnings are both 0, we will find that income is constrained to the range £31.45 to £38.35.

We may now use a "which" query to ask which people satisfying the amended description have an income less than £32.70. We find that, among other things, if entitlement to Invalidity Allowance (a supplement payable with IVB depending on the age at which someone became incapacitated) is false, then people will fall into this category. Further inquiry as to

what causes entitlement to Invalidity Allowance to be false reveals that it is that incapacity started at age 60 or above in the case of men, or 55 or above in the case of women.

This process of progressively refining the description of the affected class in response to information obtained from the model thus gives us both a precise characterisation of the group affected by the problem, namely those who have become incapacitated within five years of pensionable age, and an explanation of why the problem arises, namely because such people are not entitled to Invalidity Allowance.

3) [optionally] modify theories

The theory used in this problem is that a person who has been on a low level of income for a substantial period of time has greater financial needs because durables and clothes wear out and need to be replaced. It is this theory which means that the aim of providing everyone with an income at least equal to their needs is not satisfied. It would be possible to modify the theory so as to allow the aim to be achieved; either by saying that it doesn't apply to invalids, or to people within five years of pensionable age, and then the problem would dissolve and no amendment to the legislation would be needed. These modifications are not, however, acceptable. We therefore must accept the problem, and leave our theories unchanged.

4) generate candidate solutions

At this point a different (and less well developed) sub-system is invoked which uses the formalisation together with some knowledge of what possible actions are available to the policy maker to suggest ways in which the legislation could be amended to remove the problem. The policy maker is, of course, free to contribute his own ideas as well at this stage. If we invoke the module on the problem of the invalidity trap we will get a number of suggestions; these include increasing the rate of Invalidity Pension, disregarding a portion of Invalidity Benefit for the purpose of calculating resources for the purposes of Supplementary Benefit, making incapacity for work a qualifying condition for Supplementary Benefit, or making entitlement to IVB a passport to LTSR after 12 months. It is worth noting that this last solution was the one eventually

adopted by the UK Social Security policy makers.

5) explore the consequences of the solutions

At this point we will have a number of suggestions for new or amended rules which would allow us to solve the problem of the invalidity trap. Now we can make these changes to the formalisation, and see what the consequences of such a change would be. Some of these will be undesirable and may cause us to reject the solution. In future versions of our system we plan to have additional statistical information on the population as to the numbers satisfying the various conditions, and so on, so that we could go on to cost the various proposed changes, as the cost of a change is, of course, an important factor to be considered in assessing the relative merits of the proposed changes.

6) specify solution for legal draftsman

By using the above process, and selecting the most desirable solution, we will have the amendments required to the formulation expressed in our formal language. This will provide a basis for the specification of the solution to be passed to the legal draftsman for enactment. We therefore see a precise and unambiguous statement of the proposed change fall out as a by-product of the problem solving process.

Conclusion

Although the work is at an early stage, we believe that it already demonstrates that the overall approach is well worth pursuing. This belief has been confirmed by discussion with policy makers to whom we have demonstrated the prototype. Currently we are producing a second prototype, which will be complete by summer 1987. This contains a number of practical improvements, in the presentation of results to the user, in the provision of facilities to maintain and amend different versions of the model, and in the expressiveness of the model, so as to allow for inequalities and the expression of functional relationships within our propositions. This second prototype will be sufficiently usable by policy makers to allow them to experiment with the system, allowing more useful feedback to be given.

It is clear from the above example that one area which requires further support is the exploration of the consequences of proposed solutions. The

invalidity trap is obvious with hindsight, albeit that it was missed when the legislation was passed. Although our current system can be used to identify the cause of the trap, it could not in its present form have predicted the trap without some intuition from the user to guide it. The question we need to address is whether a system of the sort described above could be extended to predict, or help more in the prediction of, problems. To do this there would need to be an explicit representation of aims, such as the aim that anyone on benefit for 12 months or more should receive at least the equivalent of LTSR. We propose to represent aims by representing them, as a first approximation, as sentences that should be provable from the rest of the model.

In addition, theories are often implicit and unarticulated. It would be helpful if our system could assist in bringing them out into the open, so that they can be considered, and used in explanation. One way of doing this would be to take facts that we know to be the case, or aims that we stipulate to be the case, and to try to prove them from our model. The points at which the theorem proving process sticks will be indicative of missing assumptions and theories.

Aims may, and typically will, be in conflict with one another. Indeed a good part of the policy maker's skill lies in the balancing of conflicting priorities. This means that we need to have a means of handling the conflicting aims. Similar problems arise with theories. These may themselves be inconsistent, and we will want to provide a machinery for resolving such conflicts.

To summarise, what this project is about is supporting high level policy decisions by reasoning about a representation of aims, goals, theories about the world, common sense and the properties of legislation. The domain is rich, both in presenting technical problems to tackle, and in the benefits that would accrue from even a partial success. So far we have done little more than map out the problem domain and to identify an approach which would seem to have considerable potential for tackling it.

ACKNOWLEDGEMENTS

I am very grateful to Marek Sergot for many illuminating and enjoyable discussions on logic and law. The work described above was carried out as part of the DHSS Large Demonstrator project,

supported by the Alvey Directorate of the UK Department of Trade and Industry and the UK Science and Engineering Research Council. The project collaborators are ICL, Logica, Imperial College, and the Universities of Lancaster and Surrey. The help of other project members, especially Justin Forder, Alan Bustany, Duncan Lowes and Andrew Taylor is gratefully acknowledged.

The views expressed in this paper are those of the author and may not necessarily be shared by other collaborators.

REFERENCES

- [1] Bench-Capon, T.J.M., Bustany A., Forder, J.M, Taylor, A.D, and Timms, S., (1985). Description of Feasibility Prototype, ALVIN 353, ALVEY-DHSS Demonstrator Policy Application, July 1985.
- [2] Bustany, A. (1985). Numeric Inference. ALVIN 351, ALVEY-DHSS Demonstrator, September 1985.