[25] Tony Leal

Research supported in part by
the National Science Foundation
under Grant GJ-42732

UCLA-ENG-7665
JUNE 1976

A COMPUTER SYSTEM FOR CONVERSATIONAL ELICITATION OF PROBLEM STRUCTURES

A. LEAL
J. PEARL

UCLA • SCHOOL OF ENGINEERING AND APPLIED SCIENCE

UCLA-ENG-REP-7665
JUNE 1976

A COMPUTER SYSTEM FOR CONVERSATIONAL ELICITATION OF PROBLEM STRUCTURES

By

Antonio Leal*

and

Judea Pearl

School of Engineering and Applied Science

University of California, Los Angeles

* * *

* *

*

*Portions of this work are contained in Antonio Leal's Ph.D. dissertation in the School of Engineering and Applied Science at UCLA, June, 1976.

Work reported herein was supported in part by NSF Grant GJ-42732.

ABSTRACT

An interactive computer program has been designed and implemented that elicits a decision tree from a decision maker in an English-like conversational mode. It emulates a decision analyst who guides the decision maker in structuring and organizing his knowledge about a particular problem domain.

The objectives of the research were: (1) to provide the decision analysis industry with a practical automated tool for eliciting decision structures where manual elicitation techniques are either infeasible or uneconomical, (2) to cast the decision analyst's behavior into a formal framework in order to examine the principles governing the elicitation procedure and gain a deeper understanding of the analysis process itself, and (3) to provide experimental psychologists with an automated research tool for coding subjects' perception of problem situations into a standard and formal representation.

The approach centers on the realization that the process of conducting an elicitation dialogue is structurally identical to conducting a heuristic search on game trees, as is commonly practiced in Artificial Intelligence programs. Heuristic search techniques, when applied to tree elicitation, permit real-time rollback and sensitivity analysis as the tree is being formulated. Thus, it is possible to concentrate effort on expanding those parts of the tree which are crucial for the resolution of the solution plan. The program requires the decision maker to provide provisional values at each intermediate stage in the tree construction, that estimate the promise of future opportunities open

to him from that stage. These provisional values then serve a role identical to a heuristic evaluation function in selecting the next node (scenario) to be explored in more detail.

The program is domain independent, as it assumes no prior knowledge specific to any problem environment, and can therefore be used as a universal decision-aiding tool. Although the program makes almost no effort toward Language-Understanding, the conversation seems to follow a natural discourse. This is due to the simplicity of the structure underlying decision trees. The user's verbal responses are mapped directly into one of the following data types: events, actions, likelihood relations, value estimates, and experiment descriptions. The final result of the computer interview is a "solution plan" that recommends an action for all anticipated contingencies.

The paper describes the design philosophy and procedural characteristics of the program and demonstrates its operation using a sample interview session.

I. INTRODUCTION

Decision analysts are often called upon to assist in the solution of planning problems ranging over a large variety of domains. In most such cases the decision analysts possess less specific knowledge about the problem-domain than their customers, and their contributions are confined primarily to the phases of formalization, and optimization. While optimization is usually performed on electronic computers, the formalization phase invariably has been accomplished manually, using lengthy interviews with persons intimately familiar with the problem domain.

This paper describes an initial attempt to automate the formalization phase using an interactive computer system which guides the decision maker through a structured English-like dialogue and constructs a decision tree from his responses. The objectives of this work are three-fold: 1) to provide the decision analysis industry with a practical automated tool for eliciting decision-trees in cases where manual elicitation techniques are either infeasible or non-economical, 2) to cast the decision analysts' behavior into a formal framework in order to examine the principles governing the elicitation procedure and gain a deeper understanding of the analysis process itself, and 3) to provide experimental psychologists with an automated research tool for coding subjects' perception of problem situations into a standard and formal representation.

The absence of an established formal procedure of eliciting decision trees is not difficult to understand. Tree construction is the first step in the decision making endeavor: the formal representation of informal concepts. Since the informal concepts reside in the decision maker's perception of the problem environment, the translation process consists of discussions and interviews as well as attempts to educate the decision maker as to the type of information he is to supply. It often requires a special insight and ingenuity on the part of the analyst to direct the conversation and phrase the queries in a way that would yield both informative and reliable responses.

From a practical viewpoint, though, the major drawback of manual interviews is their length and cost. Since real-time analysis of decision trees is beyond the limitation of human computational capability, it invariably happens that many hours of interviews are spent on eliciting portions of the decision tree which do not have decisive bearing on the problem(s) at hand. This fact can only be discovered at a later stage once the problem structure is formalized, and a sensitivity analysis has been conducted on an electronic computer. During the interview itself, however, it is impossible for the analyst to process the entire information obtained by him up to that point, and to select the optimum course for conducting his future inquiries. For example, in eliciting utilities it is a common practice to extract indifference curves among several value attributes. Often the decision analyst is forced to elicit these curves over wide ranges of attributes, only to find out later that eliciting preferences among a few selected points would have been sufficient to determine the entire problem[1]. Similar situations occur in the process of eliciting conditional probabilities and in the expansion of very complex trees. Thus, our inability to perform real-time analysis of the information at hand forces us to waste precious time on inconsequential detailed queries.

A direct man-machine interface could provide three distinct advantages. First would be the capability of real-time sensitivity analysis, which in turn could be used to guide the growth of the decision tree in only the more promising directions. Detailed queries could be reserved, then, for those branches of the tree which, on the basis of the information obtained thus far, seem most crucial to the main decision-related goals.

The second advantage of direct man-machine elicitation would be the ease of updating the system with new knowledge. It is expensive (if not impractical) in many situations to solicit the assistance of a decision analyst each time the user gains a new piece of knowledge. A conversational system, on the other hand, would provide an intimate media that could be updated quickly even by the non-technical decision maker.

The third advantage would lie in the feasibility of incorporating real-time Delphi methods for aggregating the opinions of several experts. Decision structures elicited from other members of the team could be interrogated at will and displayed during the elicitation process to help an expert enrich or revise his previous structure. Disagreement could be detected, isolated, and brought up for further team discussion.

This paper describes the design philosophy operating characteristics of an initial version of a tree-eliciting program implemented at UCLA. Section II describes the method used to direct the conversation toward an efficient tree-expansion. Section III provides a detailed description of the system components and Section IV demonstrates a sample interview session.

II. TREE-EXPANSION PROCEDURE

a. Heuristic judgment, and judgment refinement

The art of directing a tree-eliciting dialogue is governed by two conflicting goals. On one hand the analyst attempts to bring to bear the most complete knowledge the decision maker (DM) may possess relevant to his current problem. On the other hand, he desires to do so with the least number of queries. Cutting down on the number of queries could only be accomplished at the expense of reducing the reliability of the information obtained. The analyst could, for example, limit the queries to holistic, global judgments, avoiding painstaking detailed queries, or he may terminate the tree prematurely, ask for value judgments on the terminal nodes, and begin the optimization phase. Such schemes, though, would defy the very purpose of decision analysis.

Decision Analysis if founded on the paradigm that people possess reliable procedures for detecting, storing, and retrieving fragments of knowledge, but possess much less reliable procedures foraggregating these fragments into a global inference[2]. If it were not for this deficiency there would be no purpose served by analysis, it would then be sufficient to ask the decision maker, "Which alternative action seems most valuable to you?" and when he responds, advise him "Do it!". The fact that the analyst refrains from asking direct value judgments on actions but prefers to derive them mechanically from judgments on events, consequences, and scenarios, reflects the analyst's hope that mechanically constructed judgments, using Bayes' rollback procedures, are more faithful to the decision maker's experience than direct "holistic" judgments internally processed by him. Thus, by asking more detailed questions and expanding the depth of the decision tree, the analyst expects to obtain a refinement of judgments that could otherwise be obtained at an earlier stage. It is this tradeoff between the cost of each query and its contribution to the overall judgment accuracy that underlies the style analysts' select in conducting their dialogues.

This tradeoff bears a striking similarity to that which governs tree expansions in Artificial Intelligence programs[3]. Game-playing, robot planning, and theorem proving programs all seek to obtain close to optimal solutions without exhaustively searching through the underlying problem trees. In these applications tree pruning is achieved via the use of a heuristic function: a simple rule, externally provided by the programmer, which computes a crude estimate of the value (strength or promise) of any tree node. In the game of chess, for instance, the rule may prescribe the way in which the various aspects of any board configuration (e.g. material advantage, mobility, number of pieces threatened, etc.) should be combined to give a rough estimate of the overall strength of that configuration. The true value of each configuration (i.e. "WIN", "TIE" or "LOSE") cannot, of course, be determined exactly unless the game tree is expanded exhaustively down to its terminal positions, and a minimax rollback procedure is applied. Since such a search is utterly impractical, the strength of each position is determined by its rollback value after the (static) heuristic evaluation function is assigned to the terminal nodes of a truncated (look-ahead) tree. The strength could, of course, be determined by direct application of the static evaluation function to the positions in question, however, such evaluation would be far less reliable than the one obtained by rollback over several levels of look-ahead. Thus, the purpose of tree expansion is, as in the case of decision trees, to obtain a more reliable estimate of a node's value. The static (or provisional) heuristic evaluation function may be considered to consist of the "true" value plus an error factor or noise which is reduced by tree expansion.

The availability of a heuristic evaluation function is one of the advantages that heuristic search has compared to decision tree expansion. That function is available <u>at each node during expansion</u> and estimates its relative potential for achieving the search objective. Thus, not only is this measure used to select among alternative action-plans, but it also determines the order of node expansion. Heuristic search procedures select for expansion that node which, on the basis of the provisional evaluation function, seems most likely to separate among the top contending plans. This method yields substantial savings in node expansion, even with very crude static evaluation functions.[4]

In contrast, in manual decision tree elicitation the analyst usually structures the entire tree (to some comfortable depth) before values are placed on the tip nodes. Drawing an analogy with heuristic search techniques it seems that decision tree analysis could be enhanced considerably if provisional values were available on nodes as the tree is being expanded. What would such values represent? They would be no different than values placed at the tip nodes of the completed tree: estimates of the relative worth of the opportunity provided by each node. Since all utilities are merely mental estimates of the rollback value of a complete theoretical tree emanating from that particular point on[5], the difference between provisional values and terminal utilities would only be their degree of accuracy. The former could be regarded as consisting of the latter plus a noise term due to deficiencies in mental aggregation procedures. Nevertheless, if the decision maker is requested to provide values for nodes as they are being expanded, the information can be used to determine a node expansion order. As the tree expands, these provisional values become (by rollback) more "refined", that is, closer to the 'true' mechanically processed value.

Figure 1 depicts the analogy between decision tree elicitations and heuristic search methods for game trees. The provisional utility estimates are used in the same way as the values of the heuristic function. The difference is only that the heuristic function is supplied initially by the programmer and is defined over the entire state space, while the utility values are supplied by the decision maker interactively with the elicitation program.

b. Expansion Order and Sensitivity Analysis

The actual order of node expansion can be determined by a simple sensitivity analysis. Since the only effect of expansion is to refine the provisional value assigned to the expanded node, it is reasonable to focus the expansion effort on those nodes which, if they suffer a change of value, would be most likely to influence the plan selection, namely those most likely to impact the top contending alternative courses of action. This is accomplished by ranking all the available tip nodes in order of <u>sensitivity</u>.

The sensitivity measure consists of first estimating the amount of change (differential) in a given provisional node value necessary to cause a change in the currently best initial decision. For example, in Figure 2, a partial tree is shown with initial decision branches b_1, b_2, and b_3. Branch b_2 is shown with an expanded event node that has two outcomes A and B. Assume that from a previous rollback calculation, the values of the three decision branches are 5, 3, and 2 respectively. Thus, b_1 represents the currently most promising decision. To calculate the sensitivity differential of node A, the following question is posed, "How much should the value of node A (currently 5) be raised so that the value of the initial branch leading to node A (i.e. b_2) will equal the currently highest branch (i.e. b_1)?" Branch b_2 must obviously be incremented by at least 2, but node A is only contributing 20% of its value to it. Node A must be raised to a total higher than 15 in order to cause b_2 to exceed b_1. Thus, the sensitivity differential of node A is 10. Similarly, B must be raised to 5 (assuming no other changes) in order to cause b_2 to become preferred to b_1. Since A must be raised more than B, A may be said to be more "robust" than B. The general procedure for finding the sensitivity differential $\Delta(n)$ for any node n is given by the following recursion relation:

$$(1) \qquad \Delta(\Gamma(n)) = \begin{cases} \dfrac{\Delta(n)}{P(n)} & \text{for an event node } n \\ \\ \Delta(n)+V(n)-V(\Gamma(n)) & \text{for a decision node } n \end{cases}$$

where $\Gamma(n)$ is a successor of node n and $P(n)$ is the probability along the branch from n to $\Gamma(n)$.

At first glance, the node with the lowest sensitivity differential may appear to be crucial and should be chosen as the next to expand. It may be argued, however, that the factor which determines the node to be selected for expansion is not the absolute sensitivity differential $\Delta(n)$, but the *relative* sensitivity $S_r(n)$:

$$(2) \qquad S_r(n) = \frac{\sigma_v(n)}{\Delta(n)}$$

where $\sigma_v(n)$ is the anticipated variation in the provisional value of node n which is likely to take place by further refinements. $\sigma_v(n)$ represents, therefore, the magnitude of the error present in the provisional value estimate $V(n)$ and $S_r(n)$ represents the likelihood that this error would result in a change of plan. The error $\sigma_v(n)$ may depend on the magnitude of the value, for example, a node with a (provisional) value of 10 is less likely to be raised to 15 by refinement than a node with a value of 110 is to be raised to 115.

The role of expression (2) in determining the order of expansion is based on the following noise model. Let v be the provisional value of some tip node as reported during the elicitation interview. Let v^* be the "true" value that would result if this node could be theoretically expanded completely. The difference $\Delta v = |v^*-v|$ is the error due to "noise". Treating v^* as a random variable, we wish to find the probability that expanding a node with value v will cause a change in the initial decision. Let v_0 be the required value for a decision change, as calculated by sensitivity analysis, we wish to find $P(v^*>v_0|v)$. Assuming that $P(v^*>v_0|v)$ is some monotonic increasing function of σ_v/v_0-v where σ_v^2 is the variance of v^*, leads to equation (2) as a proper basis for node expansion.

The value $\sigma_v(n)$ cannot, of course, be computed exactly, It can, however, be estimated either by asking the decision maker directly to assess the reliability of his value judgment $V(n)$ (e.g. in the form of a utility interval), or by assuming a reasonable reliability model in the form of a functional relationship $\sigma_v(n) = f(V(n))$ connecting $\sigma_v(n)$ to the magnitude of the value estimate $V(n)$. In the current version of the program a linear noise model is assumed: $\sigma_v(n) = a+bV(n)$ reflecting the fact that greater inaccuracies are anticipated in assessing scenarios involving higher stakes. By comparing the provisional value $V(n)$ with its rollback value over many nodes, it is possible to collect statistics on the factors which determine the reliability of human assessments. These could be incorporated to construct more refined models of value reliability for use in subsequent runs.

Once $\sigma_v(n)$ is determined, the value of the relative sensitivity $S_r(n)$ can be computed for all the tip nodes of the partial tree under analysis. The one with the lowest value would be selected for expansion. Needless to say, such analysis cannot be performed during a manual interview, as it involves real-time manipulation of the entire tree at hand.

III. SYSTEM CHARACTERISTICS

a. Elicitation Procedure Cycle

The currently implemented elicitation program constructs and expands a decision tree by gaining information from a decision maker interactively at a computer terminal. Figure 3 shows the main parts of the procedure for expanding each node. The first step (1) is to choose a node for expansion on the basis of sensitivity analysis (see Section II). The decision maker, after being alerted to shift his attention to a specific area of the tree, is asked to make a determination of the node type (2). The actual questions are phrased so that the decision maker is unaware that a tree is being constructed. Some example questions that ask for node type are:

IS THERE A DECISION TO BE MADE AT THIS POINT?

ARE EVENTS ABOUT TO HAPPEN OVER WHICH YOU HAVE NO CONTROL?

DO YOU WISH TO STOP EXPLORING FURTHER IN THIS DIRECTION?

The node type is classed as a decision or an event and elicitation of alternatives or event outcomes, respectively, can begin (3):

PLEASE LIST THE ALTERNATIVES THAT YOU HAVE.

or: EXACTLY WHAT EVENTS COULD OCCUR?

Using the decision maker's response ("X" below) for feedback:

LET'S CONSIDER "X" FOR A MOMENT,

the provisional values (4), probabilities (5), and costs(6) are requested (if necessary):

TRY TO PLACE A NUMERIC VALUE ON THIS SITUATION.

WHAT ARE THE CHANCES OF THIS EVENT HAPPENING?

IS THERE ANY IMMEDIATE COST ASSOCIATED WITH THIS ALTERNATIVE?

After a request (see below) for experimentation (7), the chosen node is said to be "expanded". Enough information has been obtained to permit rollback and the elicitation procedure is repeated when a new node is chosen.

b. Informational Nodes

To give the decision maker an opportunity to improve his probability estimates, a mechanism is provided in the elicitation program to represent the option of gathering information about uncertain events. This information gathering usually takes the form of an experiment with the following properties:

(1) The possible observations (experimental results) are known and are mutually exclusive.

(2) The relations between the observations and the uncertain event are expressible in probabilistic terms.

(3) If the experiment has a cost, it must be known.

The "experiment" may represent either an actual physical act (e.g. call your broker, consult the literature) or an internal mental act of recalling pertinent information (e.g. analyze the clues that led to a certain belief). The interactive program elicits the above information from the decision maker and incorporates it into the current tree by inserting an "experiment node" at the proper place. Then, the decision maker is requested to supply conditional probabilities relating each observation to each event outcome. This allows calculation (by Bayes rule) of the value of the experiment (if it involves a physical act) and of the refined (*a posteriori*) probability once the observation becomes known.

Information gathering acts usually result in excessive repetition , since the trees emanating from each experimental outcome are usually identical in structure. To eliminate difficulties connected with subtree duplication we chose to represent the entire experiment structure by a single (diamond shaped) node with a single branch. Since the probability and value labels on the duplicated subtrees are not identical, but vary with the experimental outcome, we represent these labels in a form of vectors with one entry dedicated to each experimental outcome. This representation results in a 'cleaner' tree structure which more closely matches the user's perception. Its use is demonstrated in the example of Section IV (see figures 4, 5, 6).

IV. SAMPLE INTERVIEW SESSION

The following hypothetical situation was chosen as an example for demonstrating the program operation. We imagine a scientist who is facing the dilemma of sending his proposal either to governmental agency X or to agency Y. (The example is not altogether unrealistic from the environment surrounding the funding of *this* research.) It is assumed, because of inter-agency code of conduct, that he cannot send the proposal to both agencies simultaneously. Agency Y has already indicated interest in his work and a willingness to support it at a level of $50,000. Agency X, on the other hand, has not yet had an opportunity to appraise it and further, is not committing funds until the end of the year (i.e. two months hence). The potentially available funds from agency X are much higher: full funding at a level of $100,000 or partial funding at $70,000 (see Figure 4).

The basic decision that the scientist must make at the moment is whether to send the proposal to agency Y immediately or to hold it for two months for the purpose of improving it to fit some of the specific needs of agency X. After the two months has elapsed, the opportunity to send it to agency Y is still open. However, some erosion in certainty would result due to the delay in submission.

During the elicitation of the probability estimates concerning the funding from agency X, the scientist feels uncertain about the values he reports. He realizes that the delay in submission will also offer him the opportunity to gain information by soliciting the opinion of an official from agency X as to the technical compatibility of the proposal with respect to the overall agency objectives. The official cannot be expected to give more than a "favorable" or "adverse" response. Through past experience, the scientist is aware of the official's tendency to respond overoptimistically even regarding proposals which are not about to be funded. This tendency can be formulated in terms of the event-conditional probabilities of the experiment associated with evaluating the proposal.

The following interview session demonstration is reproduced directly from the computer output. An asterisk is the signal from the computer that it is ready to receive input from the user. Thus, text immediately following an asterisk is considered to have been typed by the scientist. Following the interview are two figures showing the resulting decision structure. Figure 5 shows all of the duplicated subtrees due to the experiment and Figure shows the condensed structure using the (diamond shaped) experiment nodes.

WHAT'S YOUR PROBLEM?

*SENDING THE PROPOSAL TO AGENCY Y OR WAITING.

PLEASE LIST THE ALTERNATIVES THAT YOU HAVE, ONE AT A TIME.

1. *SUBMIT TO AGENCY Y IMMEDIATELY

2. *WAIT FOR TWO MONTHS

3. *NO MORE

ARE THE ALTERNATIVES MUTUALLY EXCLUSIVE?

*YES

NOW CONSIDER "SUBMIT TO AGENCY Y IMMEDIATELY".

ASSUMING THAT THE CURRENT SITUATION WERE TRUE, TRY TO ESTIMATE HOW GOOD IT IS INDEPENDENTLY OF OTHER POSSIBILITIES. YOU MAY USE MONEY AS A SCALE OR ONE OF YOUR OWN CHOOSING, HOWEVER, YOU MUST BE CONSISTENT AND USE THE SAME SCALE THROUGHOUT THE INTERVIEW.

DO YOU WISH TO USE MONEY AS A SCALE?

*YES

ENTER YOUR ESTIMATE.

*50000

ARE THERE ANY HIDDEN COSTS NOT TAKEN INTO ACCOUNT IN THE PREVIOUS VALUE ESTIMATION?

*NO

LET'S LOOK AT "WAIT FOR TWO MONTHS"

WHAT VALUE WOULD YOU GIVE TO THIS SITUATION?

*70000

WHAT IMMEDIATE COST IS EXPECTED?

*0

SUPPOSE THAT YOU HAD CHOSEN TO WAIT FOR TWO MONTHS

CAN YOU THINK OF THINGS THAT MAY HAPPEN AS A RESULT?

*NO

IS THERE A DECISION TO BE MADE AT THIS POINT?

*YES

STATE THE CHOICES THAT YOU HAVE.

1. *SEND TO X

2. *SEND TO Y

3. *DONE

DOES THE CHOICE OF ONE ALTERNATIVE EXCLUDE THE OTHERS?

*YES

NOW CONSIDER "SEND TO X"

TRY TO PLACE A NUMERIC VALUE ON THIS SITUATION.

*70000

WHAT WOULD BE THE IMMEDIATE COST, ASSUMING THIS SITUATION?

*0

LET'S LOOK AT "SEND TO Y"

HOW WOULD YOU EVALUATE THIS SITUATION?

*50000

HOW MUCH EXPENSE IS ANTICIPATED?

*0

ASSUMING THAT "SUBMIT TO AGENCY Y IMMEDIATELY" WAS PICKED,

ARE EVENTS ABOUT TO HAPPEN OVER WHICH YOU HAVE NO CONTROL?

*NO

DO YOU HAVE A CHOICE OF ALTERNATIVES?

*NO

DO YOU WISH TO STOP EXPLORING FURTHER IN THIS DIRECTION?

*OK

SUPPOSE THAT "SEND TO Y" IS YOUR CHOICE,

ARE THERE SOME EVENTS THAT MAY HAPPEN?

*YES

PLEASE LIST THE OUTCOMES.

1. *ACQUIRE FUNDING FROM Y

2. *Y REJECTS PROPOSAL

3. *DONE

ARE THESE EVENTS MUTUALLY EXCLUSIVE?

*YES

NOW CONSIDER "ACQUIRE FUNDING FROM Y"

ESTIMATE THE VALUE IF YOU WERE IN THIS POSITION.

*50000

TRY TO ESTIMATE THE PROBABILITY THAT THIS EVENT WILL HAPPEN.

*0.8

IF THERE IS AN ASSOCIATED COST, WHAT IS IT?

*0

LET'S LOOK AT "Y REJECTS PROPOSAL"

WHAT IS THE SITUATION WORTH TO YOU?

*0

STATE THE IMMEDIATE COST IF THERE IS ONE.

*0

CAN YOU IMPROVE THESE PROBABILITY ASSIGNMENTS

BY PERFORMING AN EXPERIMENT?

*NO

SUPPOSE THAT "ACQUIRE FUNDING FROM Y" HAPPENED.

ARE THERE SOME OPTIONS AVAILABLE TO YOU?

*NO

COULD UNCONTROLLABLE OUTCOMES OCCUR?

*NO

HAS THERE BEEN ENOUGH DETAIL EXPRESSED SO FAR?

*YES

WHAT IF "Y REJECTS PROPOSAL" OCCURS"

WOULD THERE BE OPPORTUNITIES OPEN TO YOU NOW?

*NO

CAN YOU THINK OF THINGS THAT MAY HAPPEN AS A RESULT?

*NO

SHALL WE TERMINATE THE INTERVIEW?

*NO

WHAT IF YOU CHOOSE TO SEND TO X?

ARE EVENTS ABOUT TO HAPPEN OVER WHICH YOU HAVE NO CONTROL?

*YES

EXACTLY WHAT EVENTS COULD OCCUR?

1. *FULL FUNDING FROM X

2. *PARTIAL FUNDING FROM X

3. *X REJECTS PROPOSAL

4. *NO MORE

DOES THE OCCURRENCE OF ONE EVENT EXCLUDE THE OTHERS FROM HAPPENING?

*YES

NOW CONSIDER "FULL FUNDING FROM X"

WHAT VALUE WOULD YOU GIVE TO THIS SITUATION?

*100000

WHAT ARE THE CHANCES OF THIS OUTCOME HAPPENING?

*0.4

WHAT IMMEDIATE COST IS EXPECTED?

*0

LET'S LOOK AT "PARTIAL FUNDING FROM X"

TRY TO PLACE A NUMERIC VALUE ON THIS SITUATION.

*70000

WHAT ARE THE CHANCES THAT THIS EVENT WILL OCCUR?

*0.4

WHAT WOULD BE THE IMMEDIATE COST, ASSUMING THIS SITUATION?

*0

WHAT ABOUT "X REJECTS PROPOSAL"?

HOW WOULD YOU EVALUATE THIS SITUATION?

*0

HOW MUCH EXPENSE IS ANTICIPATED?

*0

ARE YOU UNHAPPY WITH THE ACCURACY OF THESE PROBABILITY ESTIMATES?

*YES

WHAT EXPERIMENT COULD BE PERFORMED?

*EVALUATE THE PROPOSAL

WHEN WOULD IT BE PERFORMED?

*AFTER WAITING TWO MONTHS

PLEASE LIST THE POSSIBLE OBSERVATIONS FROM THE EXPERIMENT.

1. *FAVORABLE RESPONSE

2. *ADVERSE RESPONSE

3. *DONE

CAN JUST ONE OUTCOME HAPPEN AT A TIME?

*YES

SUPPOSE THAT "X REJECTS PROPOSAL" IS ABOUT TO HAPPEN...

WHAT IS THE PROBABILITY OF "FAVORABLE RESPONSE" HAPPENING?

*0.3

SUPPOSE THAT "PARTIAL FUNDING FROM X" IS ABOUT TO HAPPEN...

WHAT IS THE PROBABILITY OF "FAVORABLE RESPONSE" HAPPENING?

*0.9

SUPPOSE THAT "FULL FUNDING FROM X" IS ABOUT TO HAPPEN...

WHAT IS THE PROBABILITY OF "FAVORABLE RESPONSE" HAPPENING?

*1

ASSUME THAT "FULL FUNDING FROM X" HAS HAPPENED.

IS THERE A DECISION TO BE MADE AT THIS POINT?

*NO

ARE THERE SOME EVENTS THAT MAY HAPPEN?

*NO

PERHAPS WE SHOULD TALK ABOUT SOMETHING ELSE?

*YES

LET US SAY IT WAS "PARTIAL FUNDING FROM X" THAT ACTUALLY TOOK PLACE. DO YOU HAVE A CHOICE OF ALTERNATIVES?

*NO

COULD UNCONTROLLABLE OUTCOMES OCCUR?

*NO

SHALL WE TERMINATE THE INTERVIEW?

*YES

CONCERNING SENDING THE PROPOSAL TO AGENCY Y OR WAITING

YOUR BEST DECISION IS TO WAIT FOR TWO MONTHS

EVALUATE THE PROPOSAL AND...

IF "FAVORABLE RESPONSE" IS OBSERVED, SEND TO X

IF "ADVERSE RESPONSE" IS OBSERVED, SEND TO Y

PROGRAM FINISHED

V. CONCLUSIONS

Although the program has not been in operation for a sufficient length of time to permit exhaustive tests, it was generally found that the elicitation system could provide an adequate and useful tool for decision aiding in problems that naturally lent themselves to decision tree representations. The most desirable features were (1) the ease of operation that resulted once familiarity was gained with the mechanics of interaction; (2) the inherent "direction" of elicitation that leads to the exploration of the most important and appropriate areas for quickly resolving the major decision problem with few questions; and (3) the determination to maintain, as closely as possible, a one-to-one correspondence between the system's internal data structure and the decision maker's perception of the problem situation.

The successful operation of the current program confirms our earlier hopes that due to the structural simplicity of decision trees, only very primitive level of language-understanding would be sufficient to conduct natural, English-like dialogues. The lack of sophisticated language understanding features, aside from accounting for the simplicity of the program, also resulted in several deficiencies. The decision tree demands mutually exclusive decision alternatives and event outcomes, and it is always necessary to confirm this fact because of the ease with which the non-technical user may misinterpret the system's queries. We found that it is very tempting to report *aspects* of the previously expanded node rather than new future situations. For example, in a response to the query: "What may happen if you choose 'X'?, it is not uncommon to find a list such as the following:

1.*I MAY FIND A PLACE TO LIVE

2.*I MAY FIND A JOB

These responses seem to follow just as naturally from the elicitation query and yet they are not mutually exclusive, but rather describe the *aspects* or *attributes* of the scenario emanating from X. The reasons that the decision maker might be tempted to provide a list of attributes in response to event queries is that he focuses on the former as auxilliary tools for mental estimation of value (these aspects usually coincide with what is also called value <u>dimensions</u>). A similar situation also arises when non-mutually exclusive decision alternatives are provided. A trained decision analyst would detect such situations immediately on the basis of his general knowledge of the problem domain. In our program, the adequacy of the responses had to be verified by additional queries (e.g. "Are these mutually exclusive?").

The second deficiency is a subtle problem for any formal decision aiding system and surfaces as a result of representing knowledge in tree form. In many real-world applications, the decision maker may not perceive a problem in the form of a time sequence of decision alternatives and event outcomes, but rather as a static network of influences surrounding *issues*.* Consider, for example, our perception of the environmental pollution problem. The issues of capital investment, energy needs, energy supply, unemployment, public health, etc. all seem to be tightly interwoven in a network of cause-and-effect relationships. The main difficulty in attacking such a problem seems to be that of unraveling the underlying causal network (on the basis of insufficient data) rather than that of planning with action/event scenarios. The major difference in the formal representation required for such problems and the one handled by decision trees is that the atomic entities admitted by the latter representation are restricted to be descriptions of "world states" or decision situations. The decision maker can express relations among these situation units but is unable to express relations between the constituents of these units. For example, he may assess the value of a certain situation consisting of a combination of attributes but he is unable (with the present system) to express the relative value of each individual attribute or cross impacts among them. Likewise, when the decision maker is asked to assess the likelihood of a certain event taking place, he ought to consider the entire state of affairs prior to that occurrence. He cannot, for example, express the belief that pollution is a positive contribu-

*Ward Edwards elucidated this point in personal discussions with him.

tor to cancer independently of other situational factors such as unemployment or some energy embargo.

To facilitate an "issue-oriented" problem elicitation program, the internal machine representation of problem situations should follow a different structure. One such structure is referred to in the Artificial Intelligence literature as "problem reduction" representations[3]. Each node in a tree represents a sub-problem or a sub-goal rather than a state description. The task of resolving each separate "issue", which the decision maker perceives as part of his problem, constitutes a reduction of the global problem into several components. These can be further reduced to their constituencies, and so on. The tree expansion procedure guides the decision maker in selecting the issue to explore next rather than the time-sequenced scenario that he should follow next.

It is believed that an amalgamation of the issue-oriented and the scenario-oriented approaches into a unified problem representation would yield greater matching between human perception and organization of information and would render computers a more effective decision aid to man.

REFERENCES

[1] Geoffrion, A.M., J.S. Dyer, and A. Feinberg, "An Interactive Approach for Multi-Criterion Optimization with an Application to the Operation of Academic Departments", Management Science, Vol. 19, No. 4, December 1972.

[2] Edwards, W., C.D. Phillips, W.L. Hays, and B.G. Goodman, "Probabilistic Information Processing System", IEEE Transactions on System Science and Cybernetics, Vol. SSC-4, No. 3, September 1968. Also Dawes, R.M., "Prediction Models as Guide to Preference", Proceedings of the 1975 Conference on Cybernetics.

[3] Nilsson, N.J., Problem Solving Methods in Artificial Intelligence, McGraw-Hill, New York, 1971, pp. 43-154.

[4] Slagle, J.R., Artificial Intelligence: The Heuristic Programming Approach, McGraw-Hill, New York, 1971.

[5] Pearl, J., "A Framework for Processing Value Judgments," School of Engineering and Applied Science, University of California, Los Angeles, UCLA-REP-7622, March 1976.

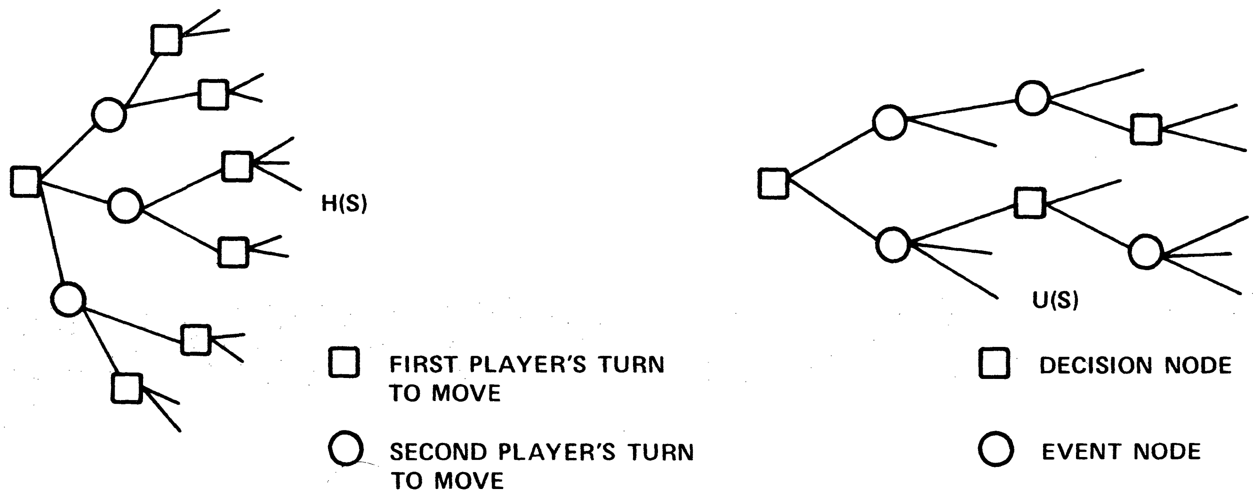

HEURISTIC SEARCH ON GAME TREES

- OBJECT IS TO FIND THE PATH (PLAN) WITH THE HIGHEST HEURISTIC VALUE H(S) WITH THE MINIMUM NUMBER OF NODE EXPANSIONS.
- COMPLETE TREE UNAVAILABLE EXPLICITLY. (IMPLICITLY CONTAINED IN GAME RULES.)
- EXPANSION FOLLOWS STATE TRANSITION RULES (LEGAL MOVES).
- HEURISTIC FUNCTION PROVIDED BY ANALYST.
- HEURISTIC FUNCTION GUIDES SEARCH.
- MINI-MAX ROLLBACK.
- TERMINAL NODES DETERMINED BY RULES. (WIN/LOSS.)

DECISION TREE ELICITATION

- OBJECT IS TO FIND THE PATH (PLAN) WITH THE HIGHEST UTILITY U(S) WITH THE MINIMUM NUMBER OF QUESTIONS.
- COMPLETE TREE UNKNOWN TO THE ANALYST. (RESIDES IN THE DECISION MAKER'S KNOWLEDGE.)
- EXPANSION FOLLOWS THE DECISION MAKER'S PERCEPTION OF EVENT/ACTION RELATIONSHIPS.
- PROVISIONAL VALUES PROVIDED BY DECISION MAKER.
- PROVISIONAL VALUES DETERMINE NEXT QUESTION.
- EXPECTI-MAX ROLLBACK.
- TERMINAL NODES DETERMINED BY DECISION MAKER.

Figure 1. Analog between Heuristic Search on Game Trees and Decision Tree Elicitation.

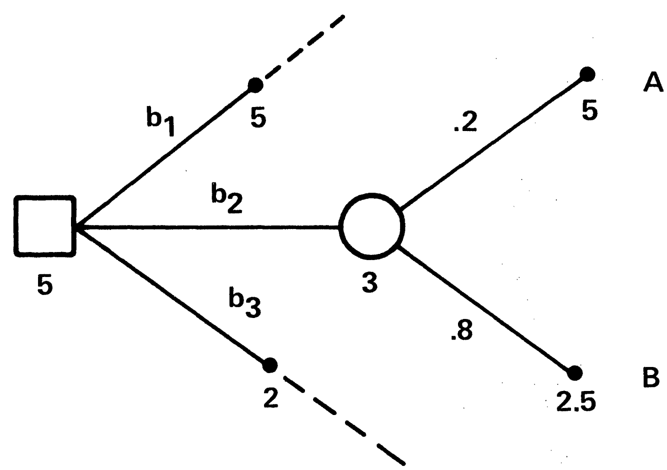

Figure 2. Basic Sensitivity Differential.

1. SELECT NEXT NODE FOR EXPANSION
2. DETERMINE NODE TYPE
3. ELICIT ALTERNATIVES OR OUTCOMES
4. DETERMINE IF MUTUALLY EXCLUSIVE
5. ELICIT PROVISIONAL VALUES
6. ELICIT PROVISIONAL PROBABILITIES
7. ELICIT COSTS
8. REQUEST FOR EXPERIMENTATION

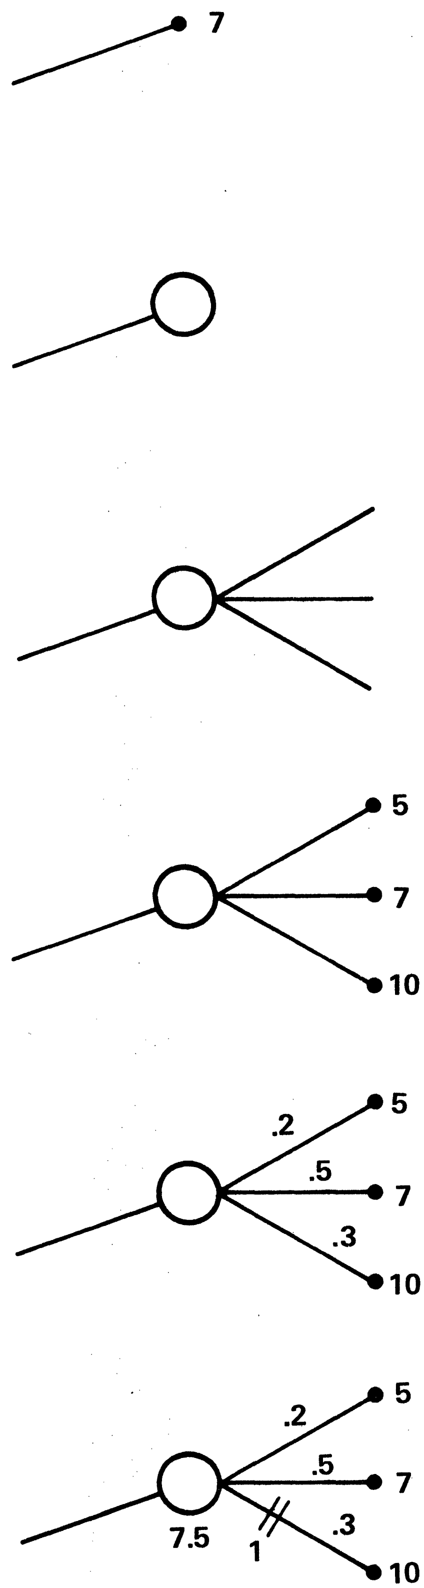

Figure 3. Elicitation Procedure.

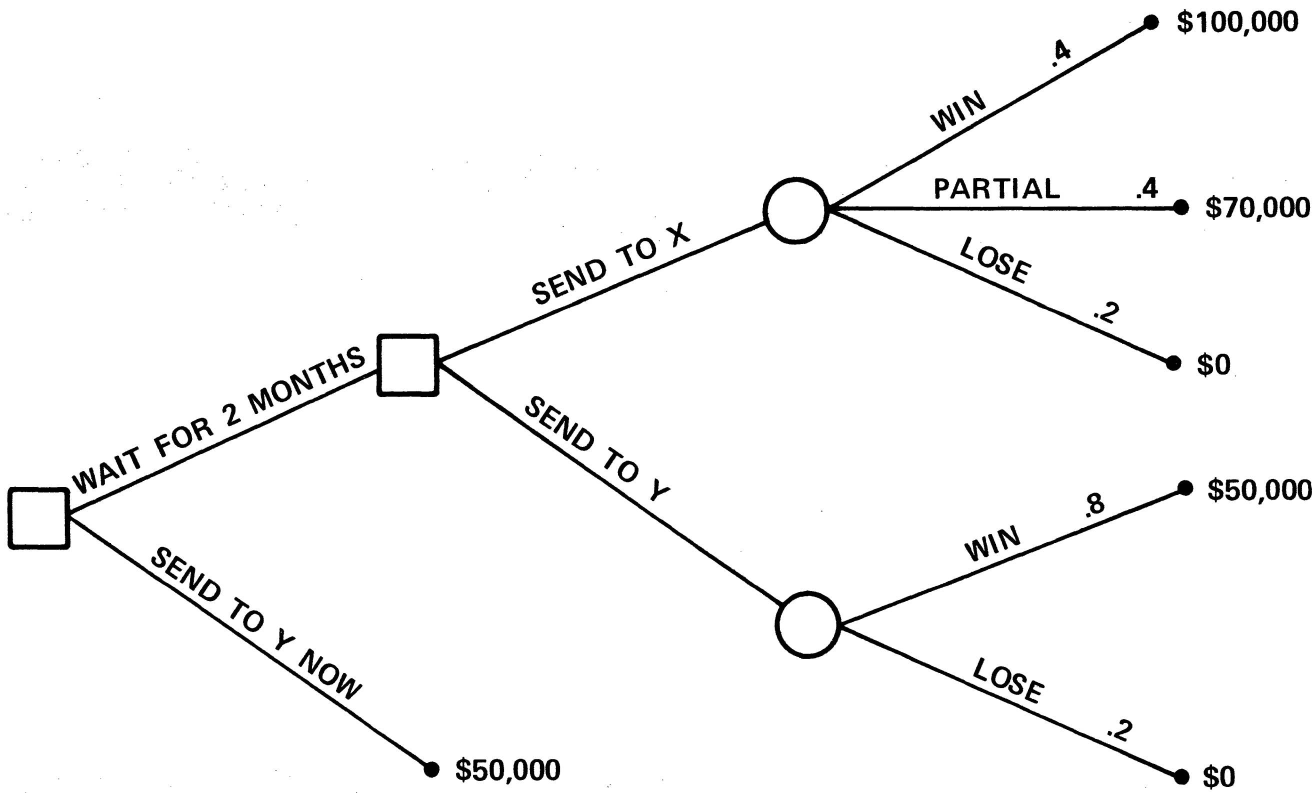

Figure 4. The Basic Decision Structure.

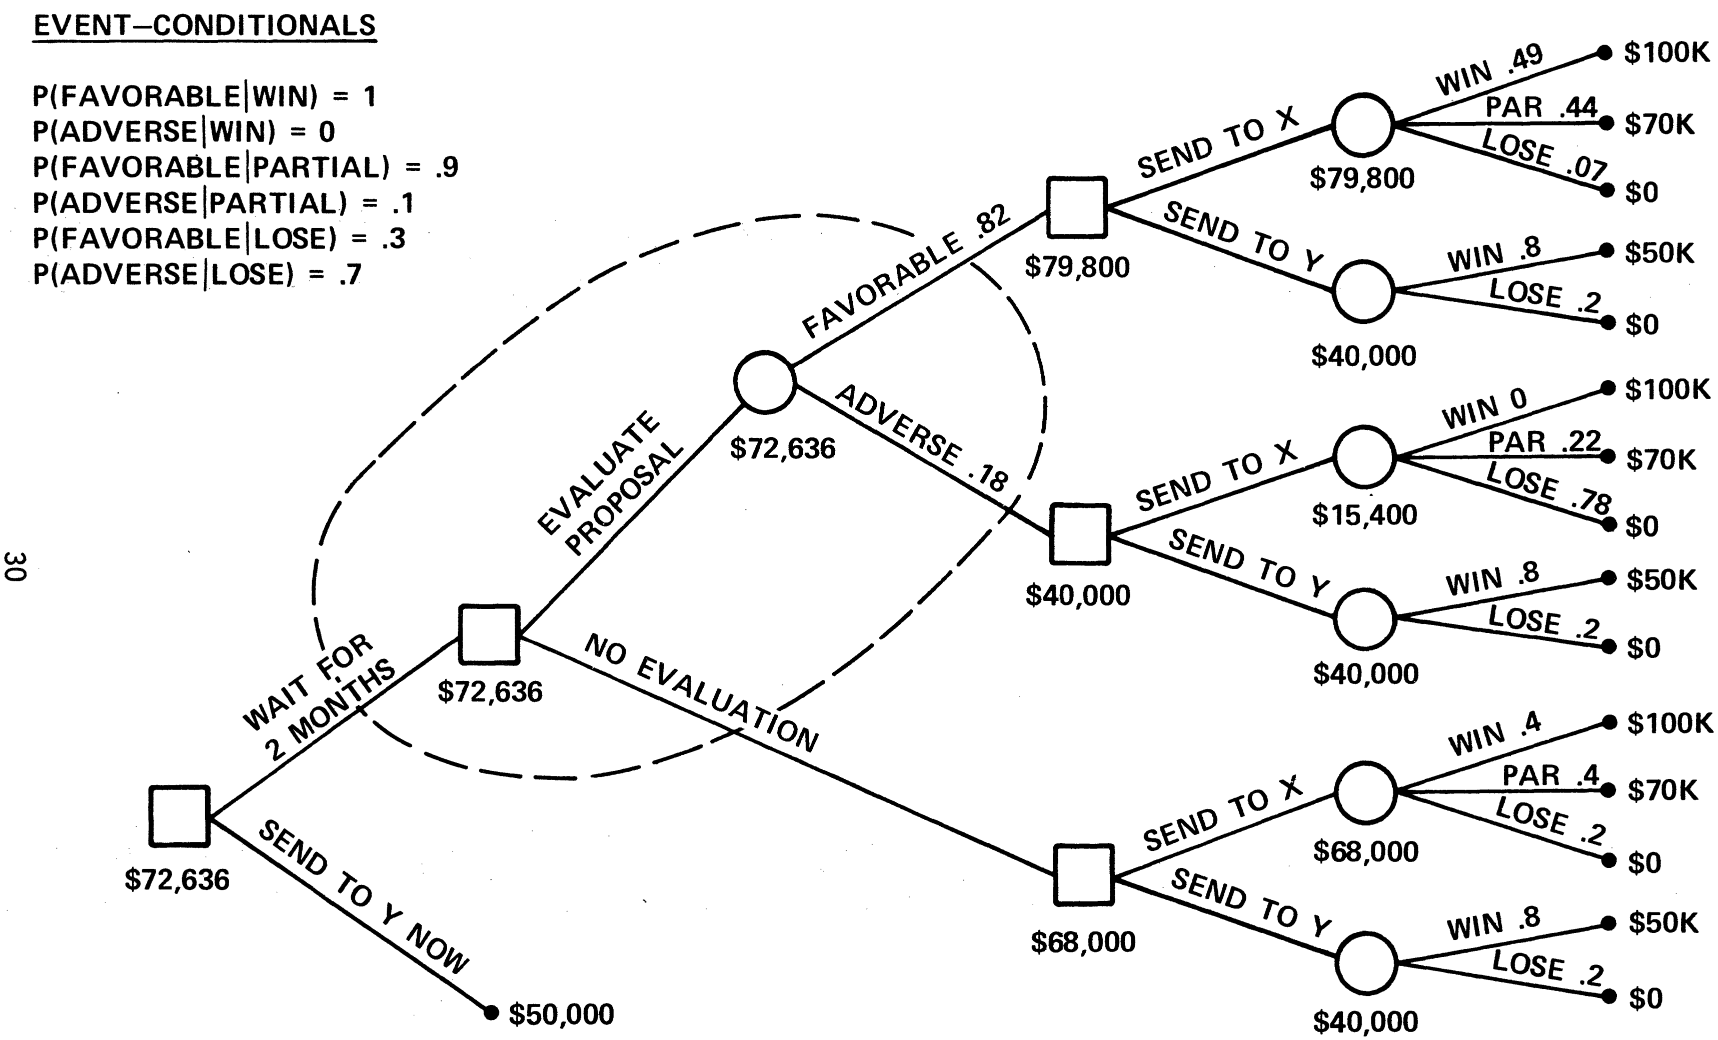

Figure 5. The Full Decision Tree.

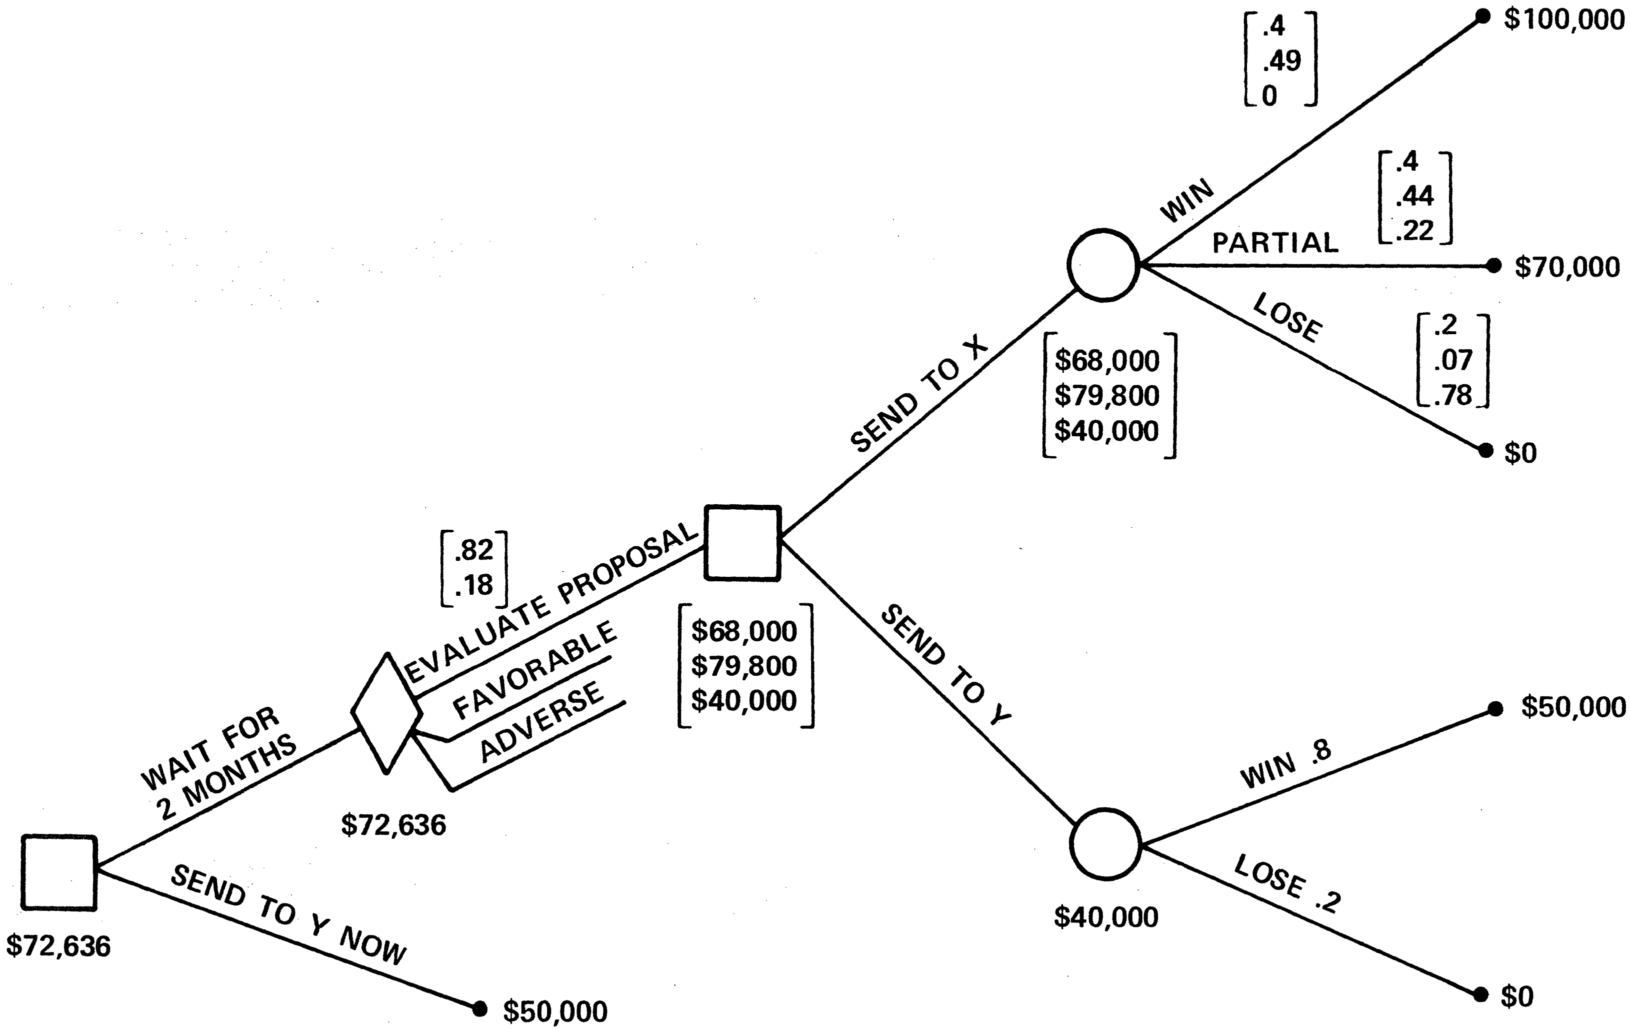

Figure 6. The Condensed Decision Structure.