Circular Decision Trees

Robert F. Bordley

Abstract: While essential to many branches of operations research, the classic decision tree becomes unwieldy and cluttered for moderately sized problems. It also requires us to introduce probabilities and payoffs into the tree by writing numbers below or at the end of various branches of the tree. This paper presents an alternative way of representing decision trees which visually represents all the information in a tree without getting unwieldy or explicitly writing down numbers.

1. DECISION ANALYSIS AND THE BUSHY MESS

In its original form, decision analysis recommended that we represent an individual's decisions and uncertainties by:

A- Preparing the Decision Basis

(A.1) Listing each decision and each uncertainty

(A.2) Constructing a decision-risk timeline which specified the order in which decisions had to be made and in which information on each uncertainty would arrive

(A.3) Identifying possible alternative choices for each decision and possible outcomes for each uncertainty

B- Structuring the Information Using a Decision Tree

We begin with the first decision/uncertainty in the decision-risk timeline, write it down and then write down branches corresponding to each possible outcome of that decision or uncertainty. We then go to the end of each branch and repeat this procedure with whatever decision/uncertainty comes next. When the analyst is finished, each complete branch of the tree represents a possible choice of decisions and a possible set of outcomes for each uncertainty. The probability associated with each possible outcome of an uncertainty is generally written under the part of the tree corresponding to that outcome.

C- Quantitative Evaluation of the Tree

The payoffs associated with each branch on the decision tree are written at the end of each branch. We then used the familiar `folding back' operation of decision analysis to compute the value associated with each of the preceding nodes on the tree. Decision analysts generally report the overall value assigned to each decision in the tree. They also often report the `risk profile', i.e., the distribution of possible payoffs associated with each decision.

Thus suppose our decision is whether or not to pursue an R&D project. Given we decide to pursue the R&D project, we have some uncertainty about whether the project will be successful. Given the project is successful, there is some uncertainty about the commerical value associated with implementing the project. If we assume there's a:

(1) 20% chance of the project being successful

(2) A 50% chance of high commercial value given a successful project(and a 50% chance of low commercial value)

then the corresponding decision tree would have the form:

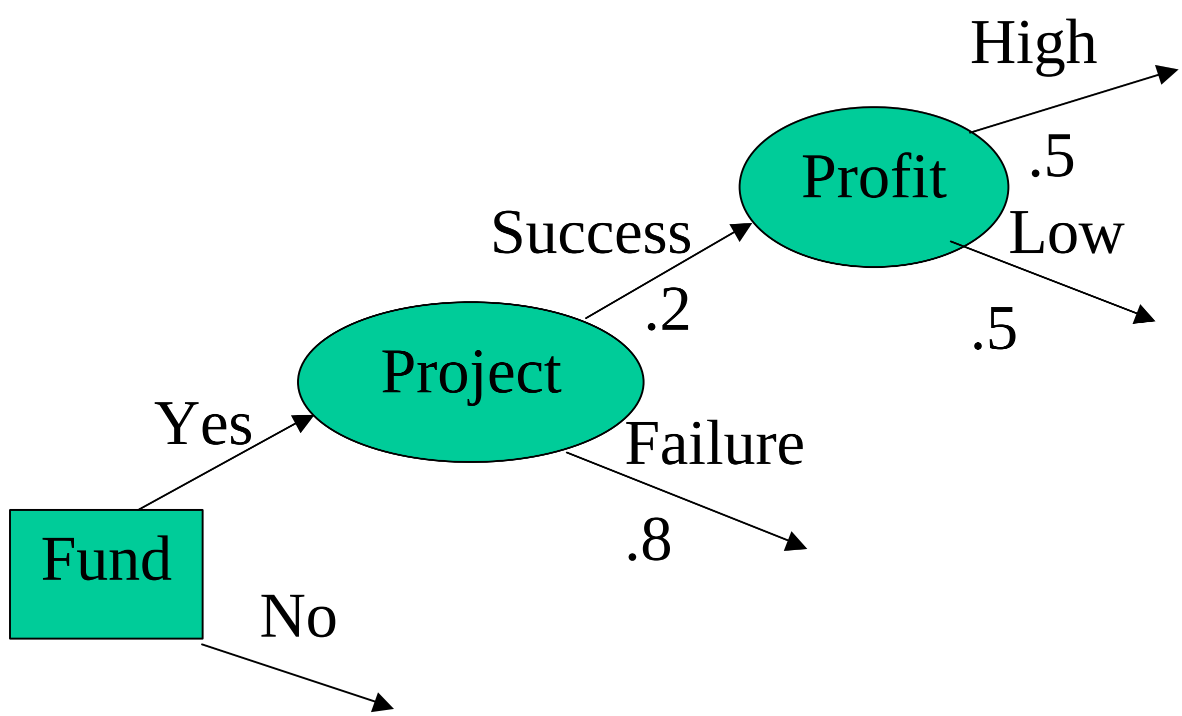

We can think of this tree as having three layers---the first layer corresponding to the decision, the second layer corresponding to the project uncertainty and the third layer corresponding to the commercial uncertainty. In this tree, there are six branches, i.e., six possible distinct consequences. Raiffa noted that one of the problems with this representation is that the number of endpoints increases rapidly with the number of layers. As a result the righthandside of the tree can become quite cluttered. Raiffa described this as the tree becoming a bushy mess.

To quite Raiffa,

Trees that exhibit the structure of real problems have a nasty habit of getting rapidly out of control---branches seem to profilerate everywhere and the tree never seems to strop growing…the further we look ahead and the more refined our analysis becomes, the more complex the tree becomes, and if we carry matters to an extreme, the tree begins to resemble a gigantic bush/"(Raiffa,pg 239).

The common solution to this problem is to develop a model using influence diagrams and then use sensitivity analysis to identify a small number of critical decisions and uncertainties. The decision tree is then structured using only those critical decisions and uncertainties. In other words, we attempt to approximate the `dream tree' by a tree which we can understand and interpret. Like any approximation, this has its limits and, in some cases, can lead us to serious distortions of the problem.

For example, much of the real options revolution arises because such an approximation can lead us to neglect downstream decisions. The real options revolution emphasizes how much neglecting downstream decisions can distort the decision problem . So what do we do if adequately representing the decision problem requires a decision tree with more than four or five layers?

In this paper, we provide a novel solution to this problem: the circular decision tree.

2. STRUCTURING THE INFORMATION

The first step in constructing a decision tree is to draw a tree representing the temporal relationships among the various factors.   To represent this information with a circular decision tree, we begin by drawing a small circle

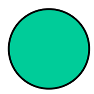

In our example, there were four layers in our decision tree. We now draw four increasingly larger concentric circles about this circle. This creates four rings. The first inner ring corresponds to the first layer, the second inner ring corresponds to the second layer, the third inner ring corresponds to the third layer and the fourth ring corresponds to the fourth layer.

We now focus on the first inner ring. This corresponds to our first layer which is a decision with two possible choices.   We now draw a line cutting this ring, and all larger rings in half.   This divides our first ring into two. We label one side of the ring `Fund' and the other side `Don't Fund' corresponding to the two possible choices for this decision.

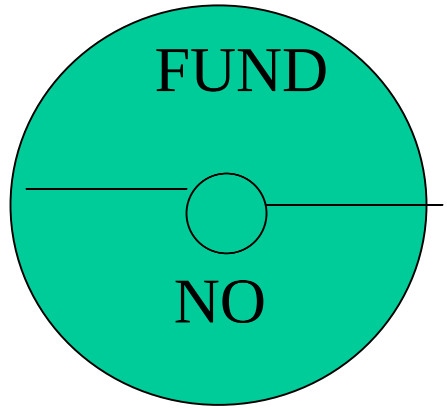
FUND
NO

We now move to the second inner ring. This second inner ring has already been divided into two parts because of the first ring. We now focus on the part which is adjacent to the `Fund' portion of the first ring. Our decision tree indicates that the fund decision is following by an uncertainty about whether or not the project will be successful or not. The two possible outcomes are `project success' and `project failure'. The probability of `project success' was 20% and the probability of project failure was 80%. We now cut this segment into two pieces---one corresponding to project success and one corresponding to project failure. As before, we again this cut so that it cuts all rings containing this second ring. The length of the segment corresponding to `project success' is one quarter of the length of the segment corresponding to `project failure'.

Similarly we move to the portion of the second ring adjacent to `Don't Fund'. In this case, there is no uncertainty. Hence there is no need to cut the ring adjacent to `don't fund'.

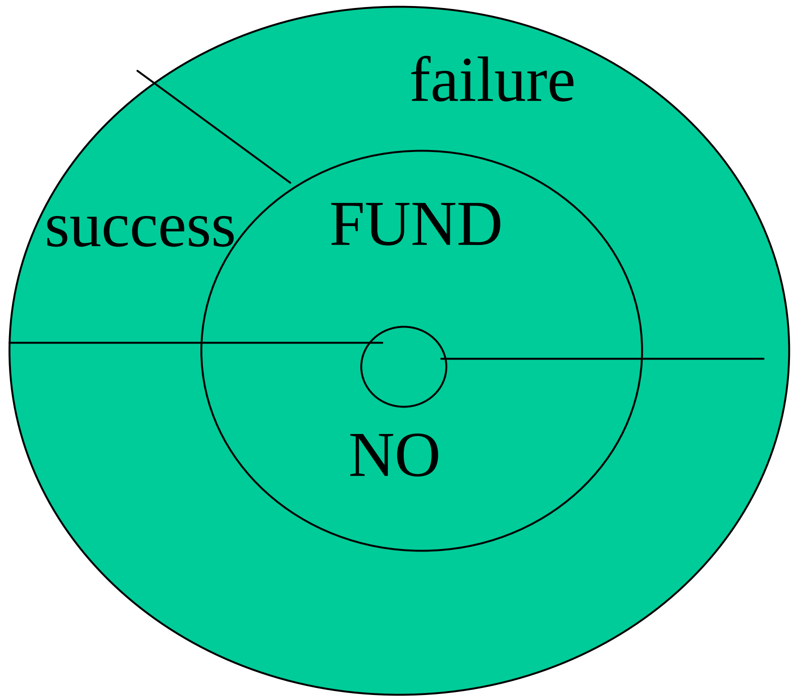

Continuing in this fashion, we proceed to the third ring. This ring has already been cut into three pieces because of the cuts arising from the first ring and second ring. We focus first on that portion of the third ring which is adjacent to `project success'.  The third layer of the decision tree indicates that there is one uncertainty following project success---which has two possible outcomes`High' or `Medium'.  Both are equally likely. As before, we now cut this portion of the third ring into two equal parts, with one part labeled `High' and the other part labeled `Medium.'  We then move to the portion of the ring adjacent to `project failure.' In this case, there is no uncertainty following project failure so that, again, there is no need to split this segment further. We then move to that portion adjacent to the ring which is adjacent to the decision `don't fund.' As before, there are no uncertainties. Hence there is no reason to cut this ring.

This gives us the final circular decision tree given below.

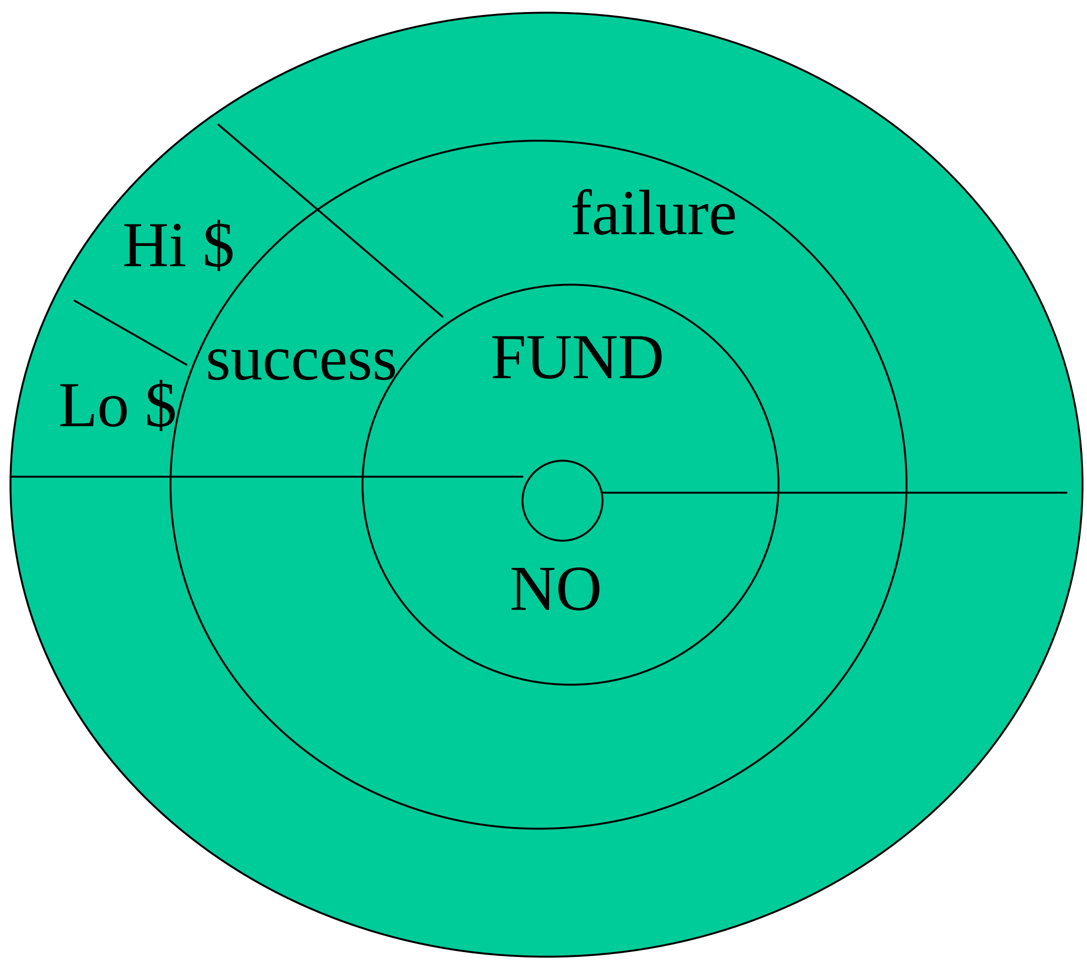

2. QUANTITATIVE EVALUATION

(2.1) Representing Values Using Colors

Note that our circular tree has already represented the probabilistic information in the decision tree using the lengths associated with segments. Hence it provides a visual representation of an important piece of quantitative information commonly used on a decision tree.

The next quantitative information used in a decision tree is the payoffs associated with various branches. To represent this information on our tree rings, we---following standard decision analysis practice---note the highest payoff and the lowest payoff . We now map these outcomes into various shades of colors. For example, we might associate the lowest payoff with dark shading and the highest payoff with no shading.

We now assign intermediate colors to intermediate payoffs. We consider two approaches toward assigning colors to payoffs:

Direct: The psychophysicist, Stevens, suggested that an individual can consistently estimate the relative value of some quantity by comparing it to the relative magnitude of some other quantity. This would suggest that an individual could directly assign some level of shading to the intermediate outcome which reflects how it compares to the best and worst outcomes. For example, suppose our best outcome were `happy family, grateful and successful kids, loving parents.' Suppose our worst outcome were `bitterly divided dysfunctional family, self-destructive children and parents who hate one another.'   It might be easier to assign a color to the intermediate outcome `divorced parents who are civil to one another'   than it would be to assign a number.

Indirect: The indirect approach recommends assigning a value of 1 to the best outcome, 0 to the worst and an intermediate number, p, to the intermediate outcome reflecting how it compares to the best and worst.   If   the density of shading of the worst outcome is some value, d, per square millimeter, then the intermediate outcome would get a shading with density (1-p)d.

(3.2) Solving the Circular Decision Tree

The classical decision tree assigns payoffs to the endpoints of the decision tree. In the circular decision tree, we color the segments in the outer ring according to their payoffs. Thus our example would give us

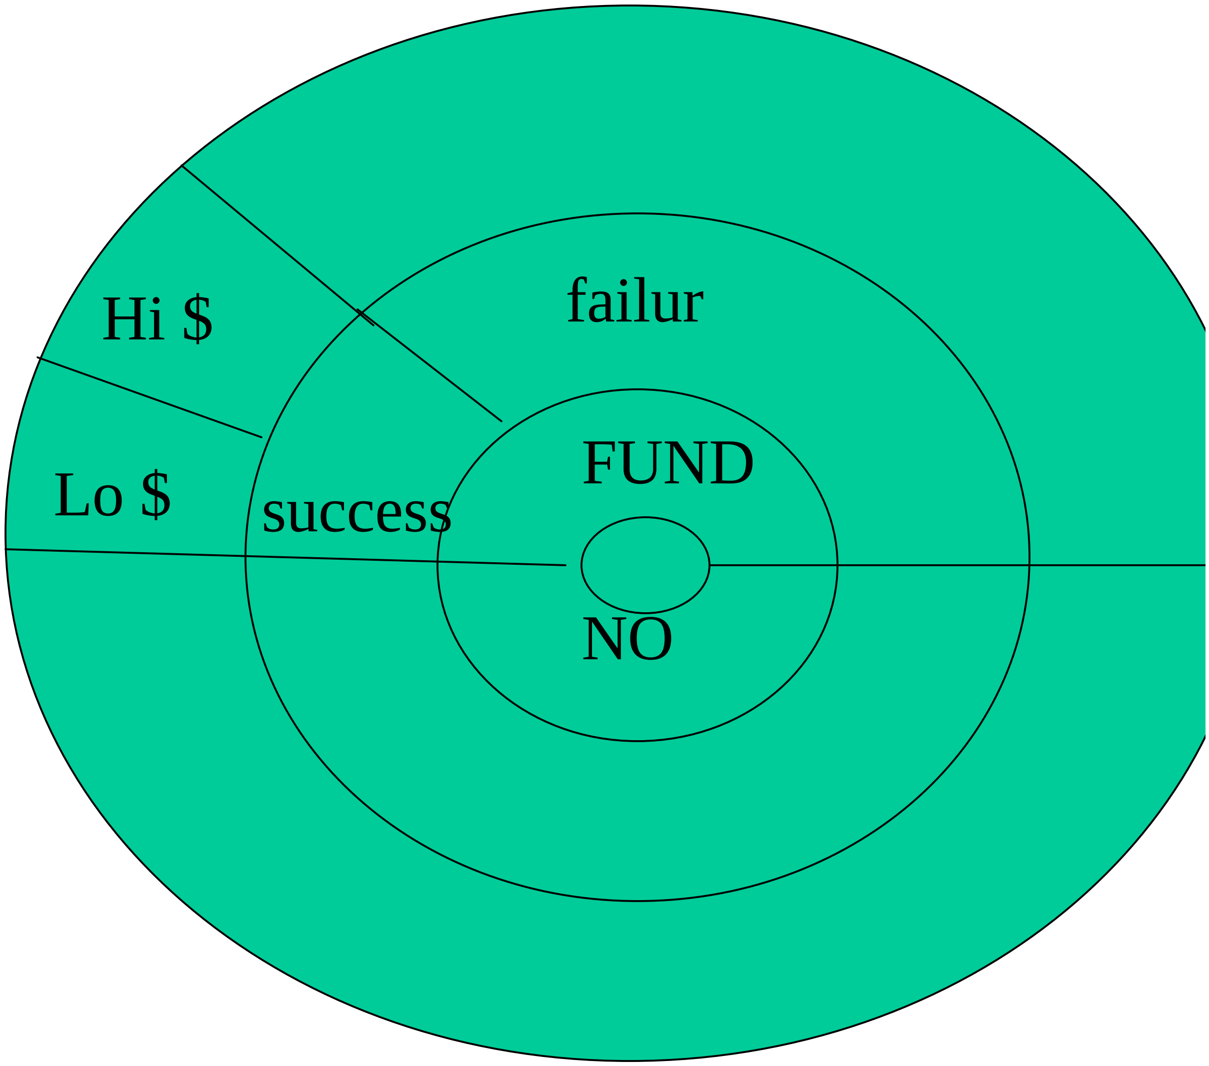
Hi $
failur
Lo $
FUND
success
NO

After assigning payoffs to the endpoints on the tree(i.e. to the last layer), classical decision theory then solves the decision tree by proceeding to the next to the last layer and, for each node, in that layer

(1) Writes a number on that node which---if the node is followed by an uncertainty---is the expected value of the endpoints arising from that node

(2) Writes a number on that node which---if the node is followed by a decision---is the maximum value associated with any of the endpoints arising from that node

This is known as `folding back' the decision tree.

Since there are no explicit numbers on the tree ring, we cannot use this procedure. Instead, after coloring the outer ring, we now go the next outer ring. We will now color the various segments in this next outer ring. To color a segment, we look at the color of the segments that are immediately above this ring. If this immediately higher ring represents an uncertainty, then we color the segment by `mixing' the colors in proportion to the quantity of each color. Hence if a segment lies below two equally sized segments, one of which is `bright yellow' and the other of which is `orange', then we color the segment `yellowish orange.' If a segment lies below a segment which is completely `dark red', then we color that segment `dark red'.

In our example, this specifies how to color the second ring

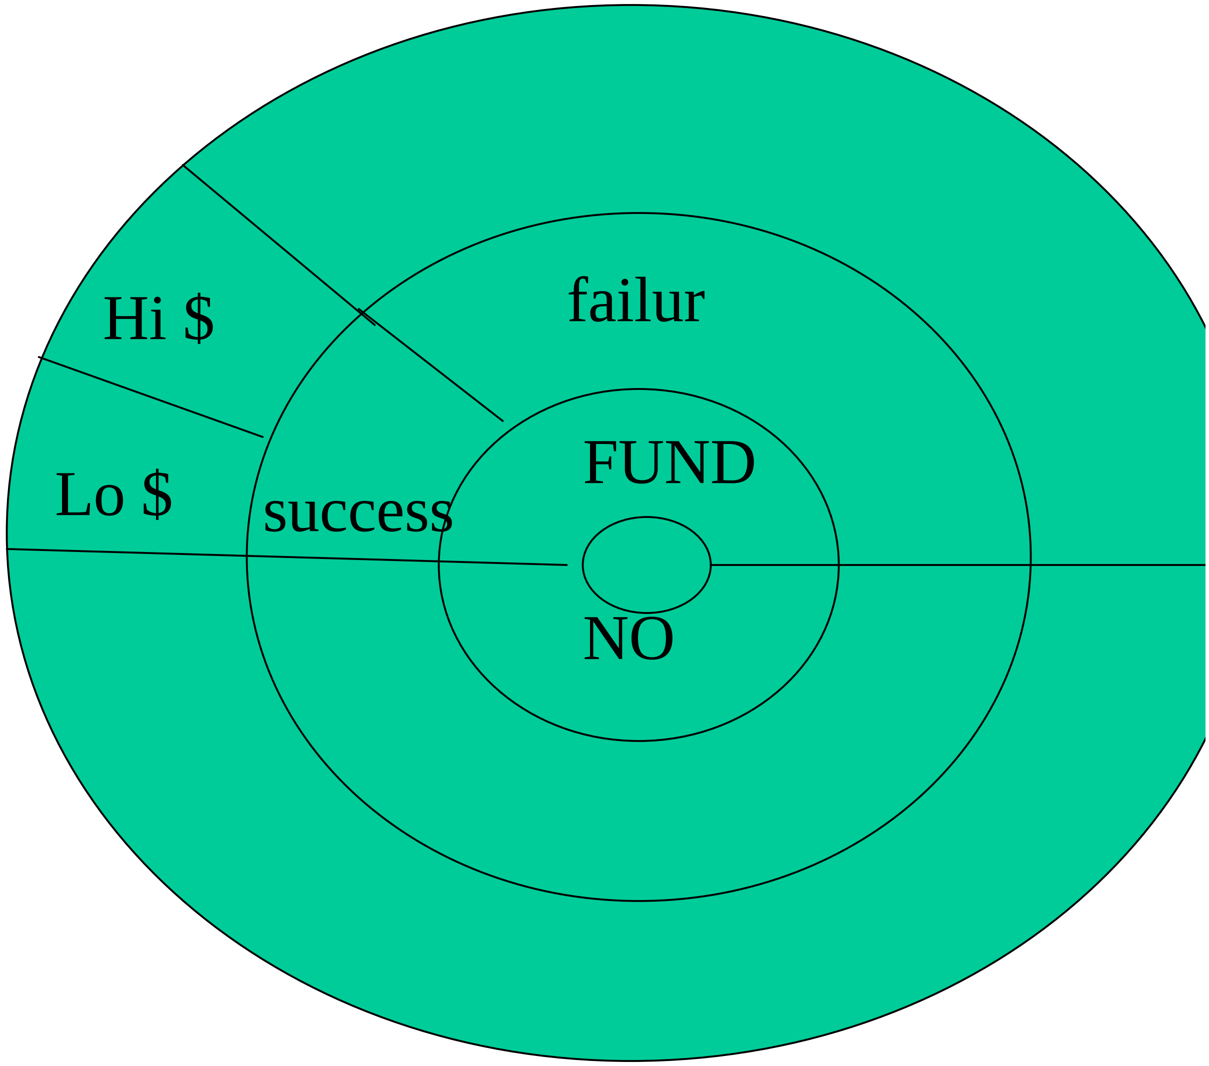
Hi $
failur
FUND
Lo $
success
NO

We repeat the procedure in order to color the second ring. Note, for example, that one segment lies below two segments., one of which is `yellowish orange' and the other of which is `dark red.' The length of the yellowish orange segment is one quarter of the length of the `dark red' segment. When we mix the colors, we get `medium red'. Hence we color the segment `medium red' . As before, if a segment lies below a segment which is completely `dark red', then we color that segment `dark red'.

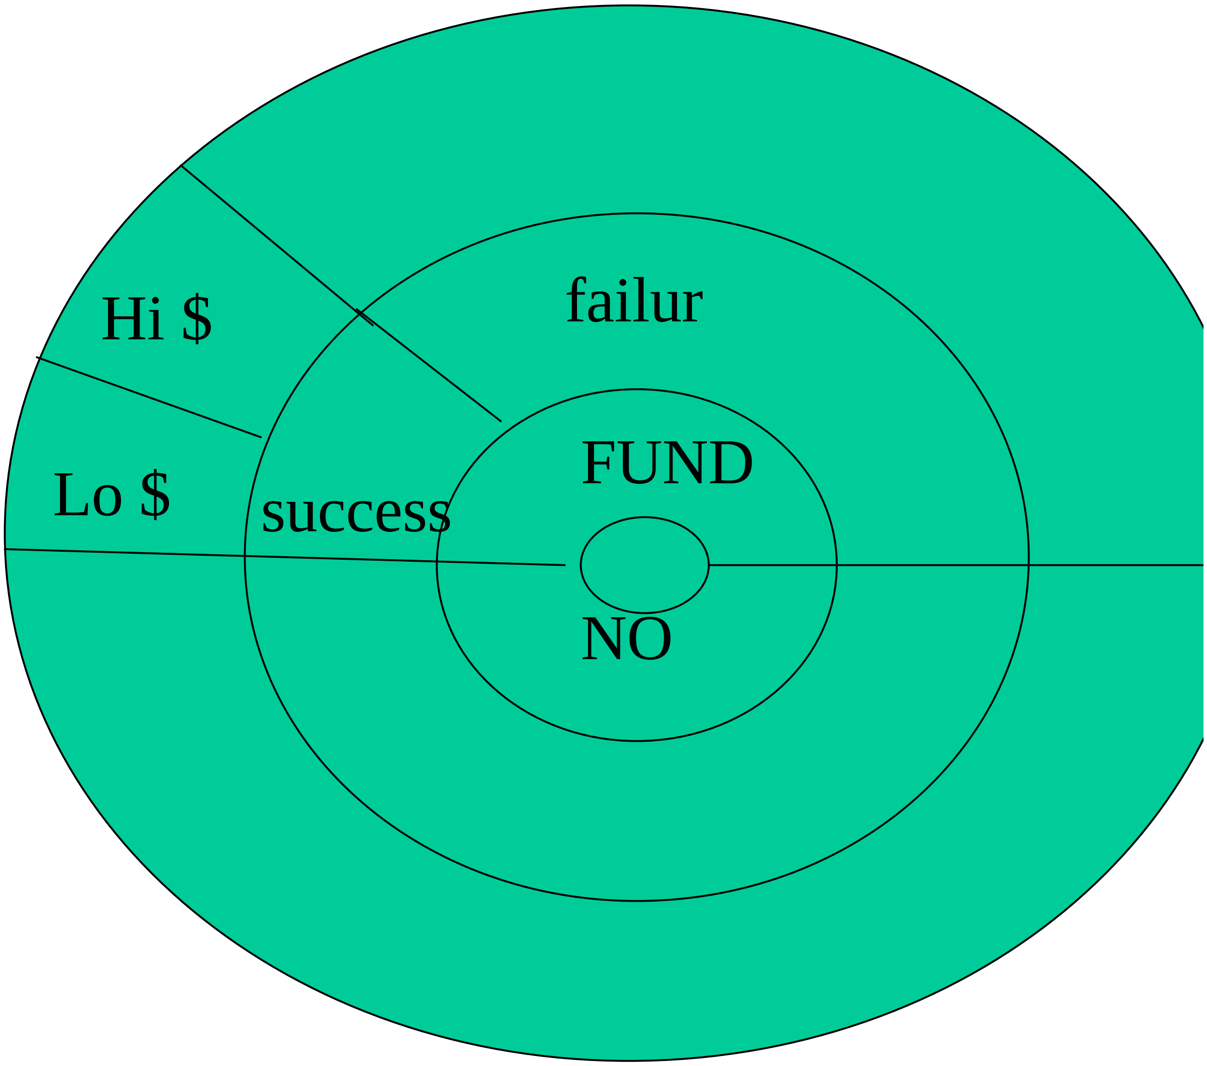
Hi $
failur
Lo $
success
FUND
NO

When we get to the center, we look at the colors of the segments that surround the center. In our example, this segment is a decision segment. Hence the color of the center will be the lighter of the two colors present in the decision segment. Hence we color the center `medium red'.

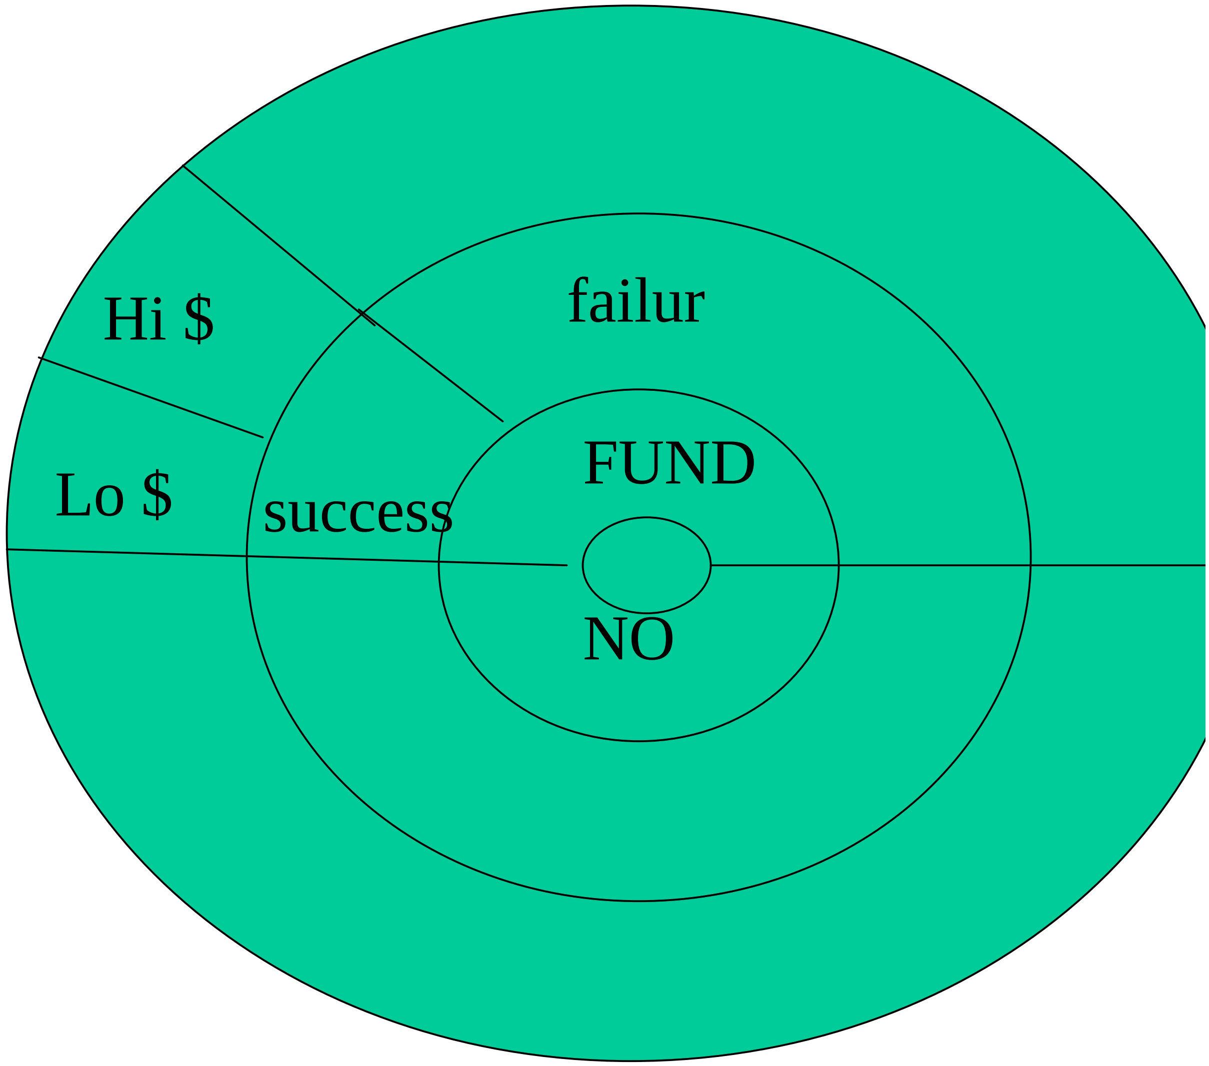

Hi $
Lo $
failur
success
FUND
NO

Note that we have `solved' the decision tree without explicitly using any numbers! Note also that this circle has a small sliver which is `bright yellow', a second sliver which is `orange' with the remainder being `dark red'. Hence it's clear that this ring represents a problem where there's a small chance of a good outcome.

4. ESTIMATING THE VALUE OF INFORMATION

Since the circular decision tree is a visual representation of all the information in a decision tree, any operation doable with a decision tree is doable with the circular decision tree. We now focus on computing the value of information,

Suppose we must decide on whether to fund the project. We then learn about whether market response is good or bad. Suppose the payoff for funding the project if market response is good is 4 (and equals 0 otherwise). Suppose the payoff from not funding the project of market response is bad is 0(and equals 1 otherwise.) The first layer of this tree is the decision and the next layer is market response. Hence the corresponding circular decision tree has the form

Suppose we let yellow denote a payoff of 4, light orange represent a payoff of 3, orange represent a payoff of 2, red denote a payoff of 1 and dark red denote a payoff of 0.

Then the tree---once solve for the value of all nodes is

(3.2) Value of Information

Since the circular decision tree is a visual representation of all the information in a decision tree, any operation doable with a decision tree is doable with the circular decision tree. We now focus on computing the value of information,

Suppose we must decide on whether to fund the project. We then learn about whether market response is good or bad. Suppose the payoff for funding the project if market response is good is 4 (and equals 0 otherwise). Suppose the payoff from not funding the project of market response is bad is 0(and equals 1 otherwise.) The first layer of this tree is the decision and the next layer is market response. Hence the corresponding circular decision tree has the form

Suppose we let yellow denote a payoff of 4, light orange represent a payoff of 3, orange represent a payoff of 2, red denote a payoff of 1 and dark red denote a payoff of 0. Then the tree---once solve for the value of all nodes is

Now suppose we can learn the outcome of market response before we have to make the decision. Note that we have draw the circular decision tree so that the segment good market response given we fund and good market response given we don't fund are adjacent. This simplies the process of computing the value of information. We begin by erasing the border between good market response given we fund and good market response given we don't fund. We then flip the rings and get

Now suppose we can learn the outcome of market response before we have to make the decision. Note that we have draw the circular decision tree so that the segment good market response given we fund and good market response given we don't fund are adjacent. This simplies the process of computing the value of information. We begin by erasing the border between good market response given we fund and good market response given we don't fund. We then flip the rings and get

5. Programming in to Microsoft Excel

We can create our circular decision trees as a `donut chart' in Microsoft Excel.   In this section, we briefly describe the required programming

6 . Benefits of the Circular Decision Tree

(4.1) Elimination of   Clutter

The standard decision tree begins with one node on the lefthandside of the tree and ends up with a possibly very large number of nodes on the righthandside.   Spatially, this means that some of the space on the lefthandside of the diagram is wasted while much of the space on the righthandside looks cluttered.

In our representation, there is no waste of space. The lefthandside of the tree is placed in the tiny center of the circle. The more cluttered righthandside of the tree is now spread across the outer ring of the circle.   Hence it provides a more compact way of presenting the relational information presented in the decision tree.

(4.2) Visual Presentation of Probabilities and Payoffs

Probabilities are represented as the lengths of arcs and payoffs as colors. Unlike the decision tree, there is no reason to write down numbers.   This can be a tremendous asset with nontechnical audiences.

(4.3) Visual Presentation of Joint Probabilities

Now consider the area of each segment on the outer ring, e.g., the segment `high market response to a successful project.'   Note that:

(1) This segment takes up one eighth of   area associated with the `fund ' decision. (See diagram.)

(2) If we assumed that the decisionmaker was equally likely to make any one of the possible decisions, then the segment---which makes up one sixteen of the entire circle---has a one sixteenth chance of occurring.

Hence the relative areas on the circular decision tree have direct intuitive implications. This can be quite powerful in drawing a decisionmaker's attention to the most important aspects of a risk distribution.

(4.4) Presenting Expected Value and Risk Profiles

In many applications of decision analysis, decisionmakers often focus too heavily on the overall expected value of the decision and not on the uncertain distribution of the payoffs from that decision.   To counteract this overemphasis, it's also common to report the risk profile which reports the distribution of payoffs associated with each decision.

But in the circular decision tree, the expected value is represented by the color of the inmost circle while the risk profile is described by the distribution of colors along the outmost ring.   Hence both pieces of information are represented simultaneously. Indeed, the diagram represents the risk profiles associated with each of the various decisions in the decision tree.   Because the diagram represents all this information in a single visual format, there's a much better chance that it will be properly factored into the decision.

REFERENCES

(1) Raiffa, H. (1967). Decision Analysis. McGraw-Hill