Representing Trees Using Microsoft Doughnut Charts

Robert F. BORDLEY

Trees are widely used to represent and solve problems in Bayesian statistics, risk analysis, marketing statistics, reliability theory, and Markov chains. But they have three limitations: (1) for moderately sized problems, the tree becomes unwieldy; (2) while the tree is very good at visually representing qualitative relationships, it does not allow visual representation of quantitative information; and (3) the tree is not easily represented on a PC without special software. In the spirit of Tufte, this article presents an alternative, the circular decision tree, which visually represents all the information—problem structure, strengths of relationship, and values of nodes—using colors and arc lengths. Mathematical computations in the tree are accomplished by "mixing colors" from adjacent segments. Furthermore, it is easily programmable in Excel. Hence it is especially suitable for presenting problems to students and nontechnical audiences.

KEY WORDS: Decision trees; Excel modeling; Fault trees; Markov chains; Value trees.

1. INTRODUCTION

A tree starts with a single "parent" node which then gives rise to several "children" nodes at the second layer of the tree. Each of these children nodes then gives rise to their own children at the third layer of the tree, and so on, as indicated in Figure 1.

A wide variety of statistical phenomenon are represented using trees. For example:

1. Decision trees are used in data mining to create segments which are homogeneous as possible with respect to some dependent variable. Each successive layer in the tree represents a partitioning of the population into subgroups based on the different values of some independent variable. Thus, the population 21, at level 2, is split into subpopulations 121 and 221 based on how these two subpopulations differ on their response to some specific question.

2. Decision trees are used in risk analysis (Bernstein 1998) to make optimal decisions in the presence of uncertainty (Baron and Brown 1991; Berger 1985; Bordley 2001; Chernoff 1987; DeGroot 1970; Lindley 1990; Skinner 1999) and lattice trees are used in financial analysis to evaluate how a stock's price changes over time. In both these applications, each successive layer in the tree represents different possible states of the world at successive points in time. The nodes 121 and 221 in Figure 1 represent the two possible states of the world at the third point in time given you were at node 21 at the second point in time.

3. Fault trees are used in reliability theory (Barlow and Proschan 1996; Gertsbakh and Gertsbakh 2000; Hoyland and Rausand 1994; Tobias and Trindate 1995) to identify how various component failures contribute to the overall failure of a system. The nodes 121 and 221 represent the two possible failures at layer three which either individually, or in combination, could lead to the failure of node 21 at the second layer.

4. Value trees are used in marketing to specify how customer satisfaction at an overall level relates to customer satisfaction at more specific levels. They specify, for example, how satisfaction with attributes 121 and 221 at level 3 relates to satisfaction with the more general attribute 21 at level 2.

These are only a few of the many applications of trees.

As the number of layers expands, the tree often grows exponentially into a gigantic bush (Raiffa 1967). Since such a tree is uninterpretable, some practitioners never present trees to their clients; other practitioners keep their trees artificially small by preceding construction of the tree with sensitivity analysis aimed at identifying and eliminating all but the most critical factors. To solve this problem of uninterpretability in data mining, Johnson (1993) proposed replacing the traditional node and arrow representation of the tree by a circular tree. SAS implemented this solution in some versions of its enterprise miner.

The author found that Johnson's solution likewise provides a compact way of representing trees in finance, marketing, risk analysis, and reliability. But as this article shows, the circular solution offers some additional benefits in these applications:

1. It allows all the quantitative information associated with the tree to be completely represented visually. Tufte (1992, 1997) established the importance of representing quantitative information visually.

2. The recursive computations commonly associated with "solving" trees can be done visually, on the circular tree, without the explicit use of any mathematics. (Hence even mathematically unsophisticated audiences can create and solve trees.)

3. The circular tree can easily be programmed in Microsoft Excel (Microsoft Excel 97, 1997) using Excel's Doughnut chart. Since many statistical courses presume access to Excel (e.g. Levine, Berenson, and Stephan 1998), this allows the user to do tree analyses without specialized software.

These features are important because many potential clients are:

1. *Unfamiliar with statistics*. Hence highly quantitative formalisms are not helpful.

2. *Skeptical of black boxes*. While clients in the past may have accepted a statistician's conclusions out of blind faith, today's

Robert F. Bordley is Technical Director, General Motors Corporate Strategy; Adjunct Professor at University of Michigan, Dearborn; and Adjunct Professor at Oakland University, MC482-D08-B24, P.O. Box 100, Renaissance Center, Detroit, Michigan 48265-1000 (E-mail: Robert.bordley@gm.com). I am indebted to Joseph Zaccagni, two anonymous referees, and the editor and associate editor for very helpful comments. I thank Gene LaForest for technical help with the Visual Basic macro.

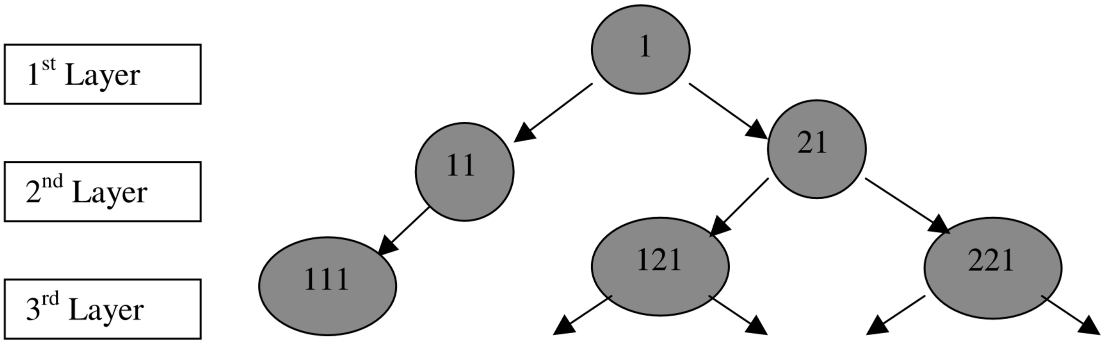

Figure 1. A general tree.

clients often insist not only on understanding the results but also the reasoning, data, and assumptions that led to the results.

3. *Self-reliant*. Some, though not all, clients prefer to conduct analyses internally, even if specialized in-house technical resources are lacking.

For these reasons, being able to construct and solve trees visually using Microsoft Excel could greatly facilitate the use of trees in marketing, risk analysis, finance, and reliability.

This article illustrates our approach to decision risk analysis trees (Barabba 1991; Howard 1988; Kusnic and Owen 1992; Skinner 1999) applied to technology projects (Bordley 1998; Sharpe and Keelin 1998). Technology projects typically involve several phases:

1. A research phase in which concept feasibility is established;
2. a concept development phase aimed at establishing proof of concept;
3. a preliminary design phase where a high-level concept is developed and tested;
4. a detailed design phase;
5. the early commercialization phase; and
6. the later commercialization phase.

At the end of each phase, a project is evaluated and is either cancelled or allowed to proceed to the next phase.

The author was commissioned to evaluate 100 R&D projects in a major corporation. This involved collecting data on the probability of each project passing from one stage to another and on the potential value of the project if it successfully completed every phase. (These probabilities were estimated either from data on related projects, from stock market prices—using options pricing techniques—or from expert judgment.) Decision trees were constructed for every project.

Although the effort involved hundreds of projects, this article considers a single project involving eight scientists working over four years. For illustrative purposes, we initially focus on a highly oversimplified description of the project in which:

1. The cost of the project, including the salaries of the scientists and their project expenses, was $10 million.
2. The four developmental phases of the project are treated as a single phase with two major technical hurdles, one with a 50% chance of being overcome, the other with a 40% chance of being overcome. Treating these two uncertainties as independent gave a 20% chance of being technically successful.
3. The two commercialization phases are treated as a single phase. If the economy is strong, then demand for the product resulting from the project will be good and the firm could expect to earn an overall profit of $100 million. If the economy is weak, then the firm would expect to earn only $10 million. Analysis suggests that there is a 50% chance of the economy being strong; otherwise it will be weak.

This representation is oversimplified but it allows us to describe the circular decision tree methodology. At the end of the article, we illustrate how a more realistic version of this problem would be represented using circular decision trees.

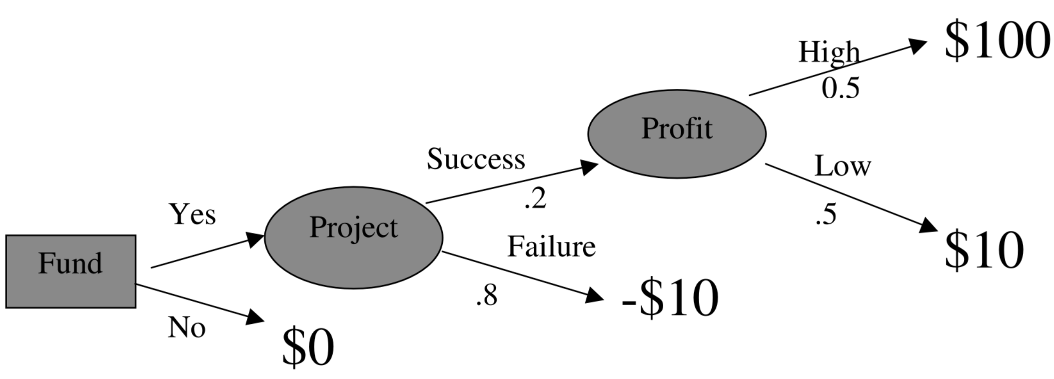

Figure 2. A decision tree for a simple R&D problem.

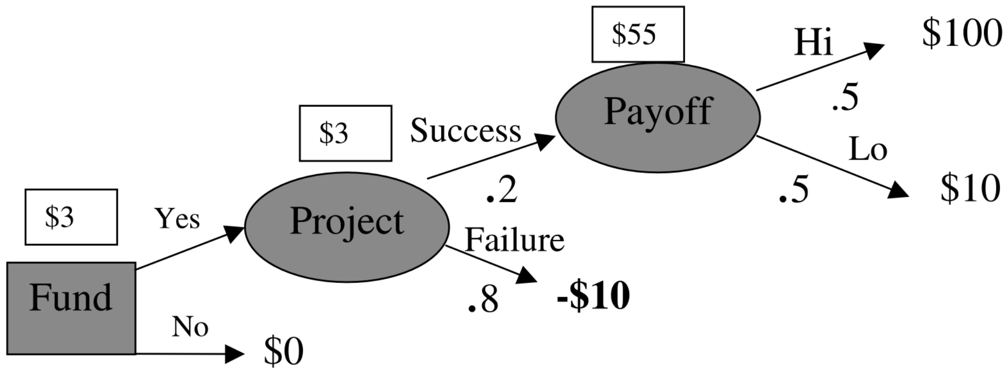

Figure 3. Solving the simple R&D decision tree.

The conventional decision tree would represent this problem as shown in Figure 2.

This decision tree represents three points in time:

- At time 1, we make a decision about whether to fund.
- At time 2, we learn about the immediate implications of what happened at time 1; that is, whether the project was funded and successful, whether it was funded and failed, or whether it was not funded.
- At time 3, we learn the immediate implications of what happened at time 2; that is, whether the project was successful and had high payoff, was successful and had low payoff, was not successful, or was not funded.

Decision analysis then assigns an intermediary payoff to each node of the tree based on the possible payoffs arising immediately after that node. If the node represents an uncertainty, the intermediary payoff is the expectation of the possible payoffs occurring after that node. If the node represents a decision, the intermediary payoff is the maximum of the possible payoffs occurring after that node. Thus, the intermediary payoff attached to the "uncertainty" node "Profit" is 50% of $100 million plus 50% of $10 million or $55 million. The intermediary payoff attached to the uncertainty node, "Project," is 20% of the value assigned to "Profit" (or $55 million) plus 80% of the value assigned to project failure(or −$10 million) which equals $3 million. Finally the payoff attached to the decision node, "Fund," is the maximum of the payoff attached to the project (or $3 million) and the payoff attached to no funding (or 0), which equals $3 million. This gives us Figure 3.

Representing this simplistic problem with a decision tree involves the explicit use of 11 numbers (representing the more realistic problem in Section 4 involves 75 numbers). As a result, the decision tree representation will be unattractive to executive audiences.

2. THE CIRCULAR TREE: STRUCTURING

2.1 Description of the Circular Tree

The circular decision tree, like the classical decision tree, models the problem beginning with the first layer (corresponding to the first point in time), the second layer (corresponding to the second point in time), and so on. It starts by drawing a small circle.

To represent the first layer of the decision tree (corresponding to the first point in time), we draw a ring around this circle. In our example, there is a decision at this layer with two possible choices. We assign 50% of the ring to the first choice and 50% of the ring to the second choice. This involves splitting the ring into two segments, one segment corresponding to the first choice (funding the R&D), and one segment corresponding to the second choice (not funding the R&D.)

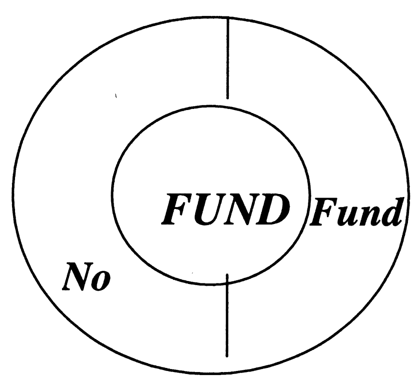

This represents a one-layer decision tree.

To represent the second layer (corresponding to the second point in time), we draw a second ring around this first ring. We then extend whatever cuts were drawn in the first inner ring into this second ring. In our example, this causes the second inner ring to be divided into two segments. We now focus on the segment 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. 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 will later extend this cut to all rings containing this second ring. Of the total area of the original segment, 20% is assigned to the subsegment corresponding to "project success" and 80% to the subsegment corresponding to "project failure."

Now consider the portion of the second ring adjacent to "*no funding*." Since there is no uncertainty, we do not cut the ring adjacent to "no funding." This gives us

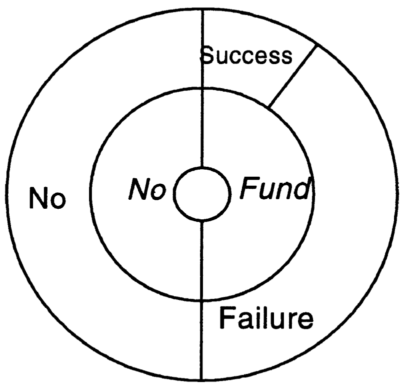

(Outcomes of decisions will be labeled with bold while outcomes of an uncertainty will be labeled in standard font.)

We now turn to the third and last layer in our example. As before, we represent this layer by drawing a third ring about the second ring. Because we extended the cuts in the second ring into the third ring, this third ring will have already been cut into three pieces. 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 "low." 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 "low." We then focus on 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 focus on that portion

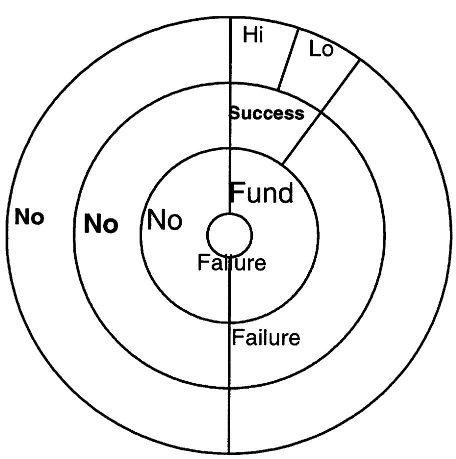

adjacent to the ring which is adjacent to the decision not to fund. Since there are no uncertainties, we do not cut the ring.

This leads to the following representation.

In summary, we structure a problem using a circular decision tree as follows:

1. If the problem involves n layers(or points in time), draw $n + 1$ concentric circles.
2. If the first layer of the problem involves a decision with m possible choices, split the first ring—and all rings containing it—into m equal parts. If this layer involves an uncertainty with m possible outcomes, then split the first ring—and all rings containing it—into m parts with the area of each part proportional to the probability of the associated outcome occurring.
3. The segmentation of this layer induces a segmentation of the next higher layer. We now focus on each segment induced in this next higher layer. If the segment corresponds to the occurrence of a decision with m^* possible outcomes, then split the segment—and all segments immediately above it—into m^* equal portions. If the segment corresponds to the occurrence of an event with m^* possible outcomes, then split the segment—and all segments immediately above it—into m^* parts with the area proportional to the probability of each outcome occurring.
4. Repeat this procedure with each successive layer.

2.2 Representation of the Tree in Microsoft Excel

To create this circular decision tree in Microsoft Excel, we create one column in Excel for each of the layers in the decision tree. Our example involved three layers. In the first layer, there were two outcomes: "*funding*" or "*no funding*." Since this was a decision layer, each outcome was assigned an equal proportion of the space. Hence in the first column, we enter the proportions, 0.5 and 0.5, for these two outcomes. In the second column, there are three outcomes, "no funding," "funding and success," "funding and failure." The proportions corresponding to those three outcomes are $0.5, 0.5 \times 0.2 = 0.1$, and $0.5 \times 0.8 = 0.4$. We enter these three numbers in the second column. In the third layer, there were four outcomes: "no funding," "funding and failure," "funding and technical success and low sales," "funding and technical success and high sales." The proportions corresponding to those four outcomes are 0.5, 0.5×0.8, $0.5 \times 0.2 \times 0.5$, and $0.5 \times 0.2 \times 0.5$. We enter those four numbers into the third columns. Hence our spreadsheet will look like

	A	B	C	D
1	0.5	0.1	0.05	
2			0.05	
3		0.4	0.4	
4	0.5	0.5	0.5	
5				

These three columns represent three series in Excel. We click *Insert/Chart/Doughnut* on the Excel toolbar to get Excel to create the corresponding doughnut chart. Since Excel creates an unlabeled doughnut chart with arbitrary colors and a large doughnut hole, getting the specific doughnut chart described in this article involves three further steps:

1. Click on each sector in the chart to change the colors Excel assigned to "white."

2. Click on the chart and then on options to reduce the size of the doughnut center to the smallest possible value. We can also rotate the chart.

3. Click Insert/Picture/Wordart, create the words, "No," "Yes," and "Fund" and attach these labels to the appropriate spaces in the chart. This gives the doughnut representation described previously.

The next section focuses on quantification of the tree—which corresponds to choosing how rings will be colored.

3. THE CIRCULAR TREE: QUANTIFICATION

Just as there are two ways to structure the tree, so there are two corresponding ways of quantifying the tree. We first discuss a manual approach to quantifying the tree. We then discuss an approach using Excel and a small Visual Basic macro.

3.1 Manual Quantification of the Circular Tree

Our circular tree represents the probabilistic information in the decision tree using the width associated with segments. Hence probabilities do not need to be written down on the circular decision tree. The other piece of quantitative information used in a decision tree is the payoffs associated with various branches. To represent this information on our circular tree, we follow standard decision analysis practice in ranking the payoffs from best to worst. We deviate from standard practice in associating colors (or shadings) with each payoff. The best payoff is colored "white" and the worst payoff is colored "black." (Obviously the reader can alter these conventions freely.) Intermediate payoffs are assigned intermediate shades of color in proportion to their relative value. When payoffs are financial or easily quantified, we might translate the numerical value into a degree of shading. In this example, the four possible payoffs of \$100 million, \$10 million, zero, and −\$10 million are translated into *white*, *gray*, *dark gray*, *very dark gray*, and *black*, respectively. When payoffs are difficult to quantify, we might choose to bypass explicit quantification and have the individual directly assign colors using principles of cross-modality matching (Stevens and Galanter 1957).

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. In the case of our technology project example, this gives us:

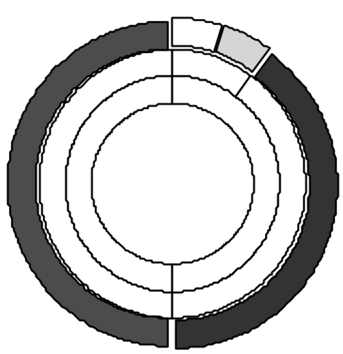

After assigning payoffs to the endpoints on the tree (i.e. to the last stage), the classical tree proceeds to the next to the last layer and, for each node, in that layer

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.

Continuing recursively in this way eventually leads to an assignment of a number to every node. This is known as "folding back" the decision tree. As we now show, the circular decision tree provides an alternate way of "folding back" the decision tree which makes no explicit use of calculations.

Suppose we have colored the segments in the outer ring. Consider the ring immediately within this outer ring. To color a segment in this ring, we look at the color of the segments that are immediately above this ring. Consider the second-layer segment which is immediately below a third-layer segment that has been colored "black." We color that second-level segment "black." Suppose that there are several segments above the segment in question and that these segments represent the different outcomes of an uncertainty. Then we color the segment in question by proportionately "mixing" the colors of these segments. For example, consider the second-layer segment below two equally sized third-layer segments, one of which is "white" and the other of which is "gray." We color this second-level segment "light gray."

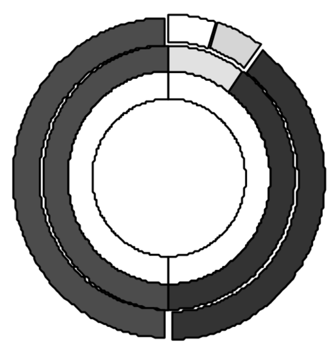

We repeat this procedure to color the first-layer ring. There is one first-layer segment lying below two second-layer segments, one of which is "light gray" and the other of which is "black." The length of the "light gray" segment is one quarter of the length of the "black" segment. Mixing the colors gives "dark gray." Hence we color the first-layer segment "gray." Now consider the first-layer segment that lies completely below a second-layer segment which is colored "very dark gray." That first-layer segment is colored "very dark gray."

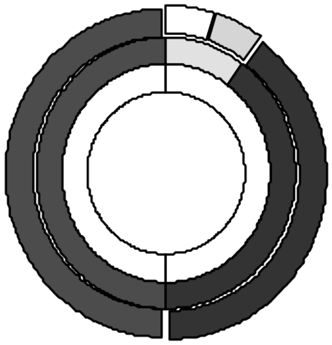

When we get to the center, we look at the colors of the segments that surround the center. Since the segment is a decision segment, the color of the center will be the lighter of the two colors present in the decision segment. Hence we color the center "dark gray." This gives

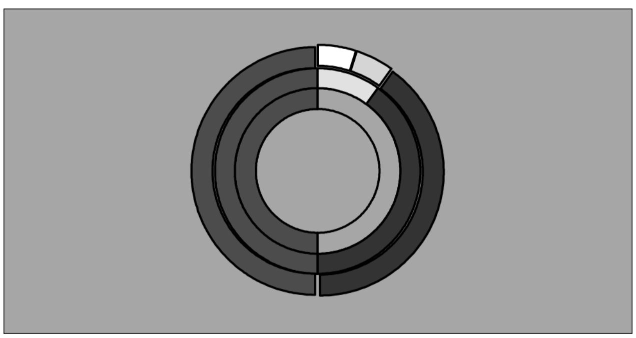

Thus, the value of this decision tree is dark gray—which corresponds to a payoff of 3.

We have "solved" the decision tree without explicitly using any numbers! Note that in examining this circle, the human eye can immediately discern a small sliver of "white" against a background of dark colors. This indicates that the R&D represents a small chance of a very good outcome and a large chance of a poor outcome.

3.2 Quantification of the Tree using Excel and a Visual Basic Macro

As Section (2.2) noted, the circular decision tree can be structured in Excel. We can also color the tree in Excel by clicking on each region and then changing the color of that region to the desired color. But since the colors in the inner layers are determined by the colors on the outer layers, a small Visual Basic macro was created that colors the inner layers automatically.

Table 1. State Matrix

Layer	Fund?	Success?	Payoff?
No funding	0.5	0.5	0.5
Funding & succcess & high payoff	0.5	0.1	0.05
Funding & success & low payoff			0.05
Funding & failure		0.4	0.4

Table 2. Payoff Matrix

Layer	Fund?	Success?	Payoff?
No funding	0	0	0
Funding & succcess & high payoff	3	55	100
Funding & success & low payoff			10
Funding & failure		−10	−10

This macro involved the creation of three matrices. The first "state" matrix (Table 1) corresponded to the matrix of three columns created in Section 2.2. Every cell in this matrix with a positive number corresponds to a region in the doughnut chart.

We then created a second "payoff" matrix (Table 2) with an entry for every entry in the uncertainty matrix. The last column of the payoff matrix, which corresponded to the outer layers of the Doughnut chart, is filled in with the appropriate payoffs. The next to last column of this matrix is filled in by computing the appropriate probability weighted average of entries from the last column. (If the next to the last column represents a decision, the appropriate number corresponds to the maximum of the appropriate entries from the last column.) The column preceding this next to last column is constructed in a similar fashion. We continue until the matrix was complete.

We also created a third "utility" matrix (Table 3) that corresponded to the outcome matrix with all elements rescaled to lie between zero and one.

Finally we created a legend that mapped numbers from zero to one into various shades of color. The numbers 0 to 0.05 were mapped into white. The numbers 0.95 to 1 were mapped into black. Intermediate numbers were assigned intermediate shades of gray.

Our Visual Basic code used the state matrix to structure the doughnut chart and used the utility matrix and the legend to color the doughnut chart.

4. A MORE REALISTIC EXAMPLE

4.1 Problem Description

Actual R&D problems are often considerably more complicated than the problem considered here. To understand how the circular decision tree would handle more realistically complicated problems, consider a technology evaluation problem in which:

1. There are three competing ways of designing a project for eight scientists:

- one possible project design (Design A) involves low technical risk;

Table 3. Utility Matrix

Layer	Fund?	Success?	Payoff?
No funding	0	0	0
Funding & succcess & high payoff	0.11	0.59	1
Funding & success & low payoff			0.18
Funding & failure		0	0

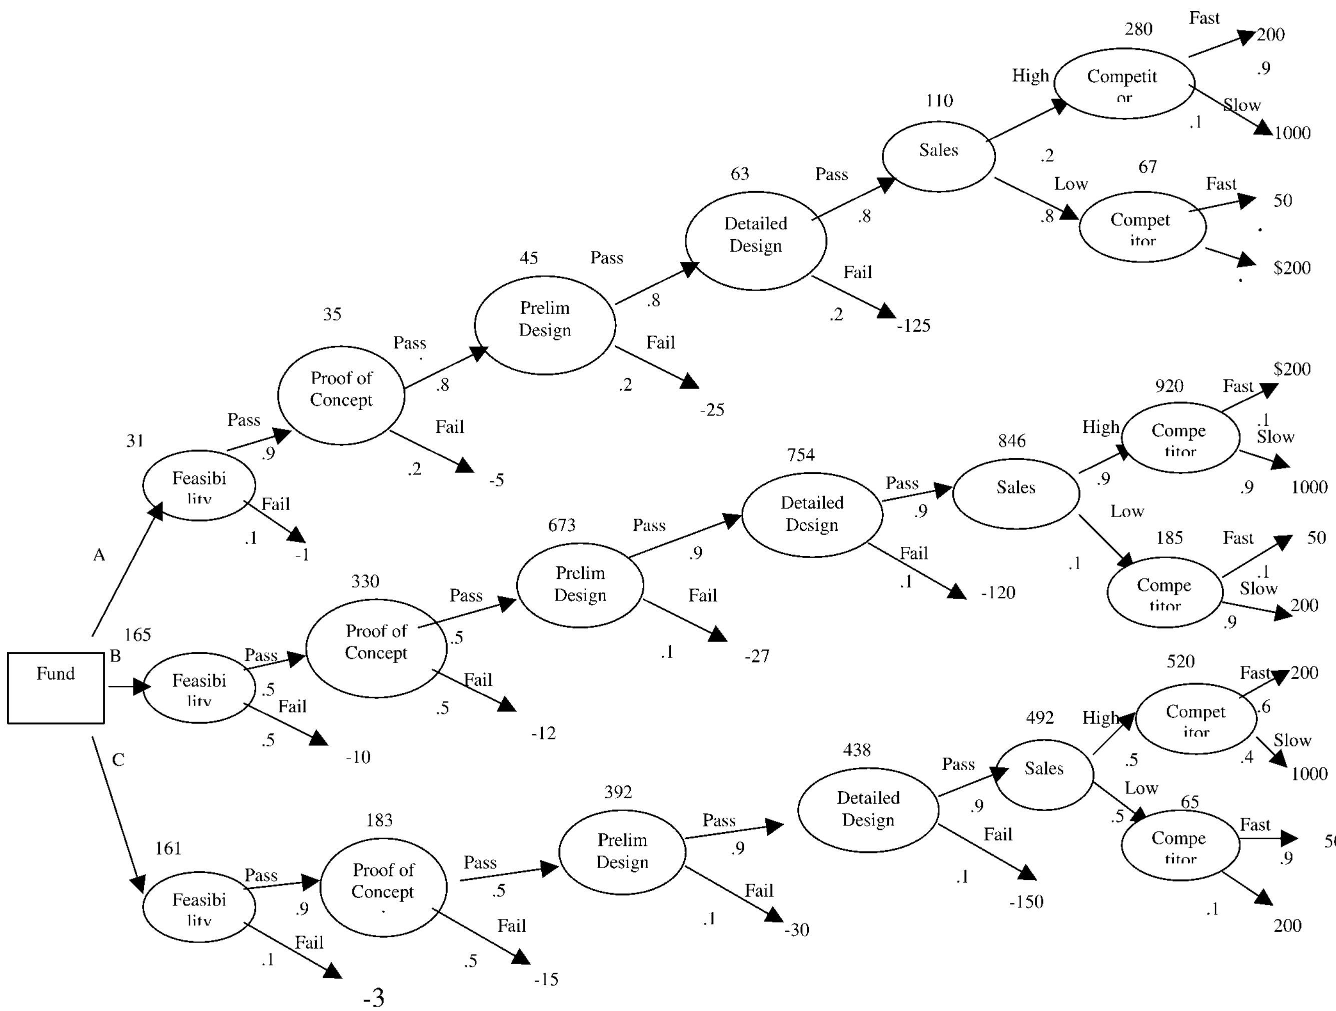

Figure 4. A decision tree for a complex R&D problem.

- a second possible project design (Design B) involves a new technology; and
- a third possible project design (Design C) involves a technology which has been used on a different product.

2. The project must transition through four development phases:

- research to establish technical feasibility;
- concept development to establish proof of concept;
- preliminary design to develop a high-level prototype; and
- detailed design to develop a product.

3. There are two phases in commercialization:

- initial commercialization; and
- late commercialization in which competitors have reacted to the new product.

4.2 Conventional Decision Tree Solution

The decision tree for this more realistic, but still simplified problem, is shown in Figure 4. For the first design, it presumes there is a 90% chance of feasibility being established and a 10% chance of failure to establish feasibility. The cost of failure at this branch is $1 million. Given the project is determined to be feasible, there is an 80% chance of successfully developing a proof of concept. The cost of failure at this stage is $5 million. Given a successful proof of concept, there is an 90% chance of a successful preliminary design. The cost of failure is now $25 million. Given a successful preliminary design, there is an 80% chance of a successful detailed design. The cost of failure at this point is $125 million. (Naturally the costs of failure increase as we go to later and later stages.) Given a successful detailed design, the probability of a good initial sales response is 20%. Regardless of whether there is a good sales response or not, there is a 90% chance of the competition reacting quickly with a product or promotion which steals some of these initial sales. There is a 10% chance that the competition will be slow to develop an effective response.

Design B will differ from Design A in having a lower chance of passing the feasibility and proof of concept hurdle and a lower chance of being quickly thwarted by a competitive response. Design C will have a good chance of passing feasibility but a

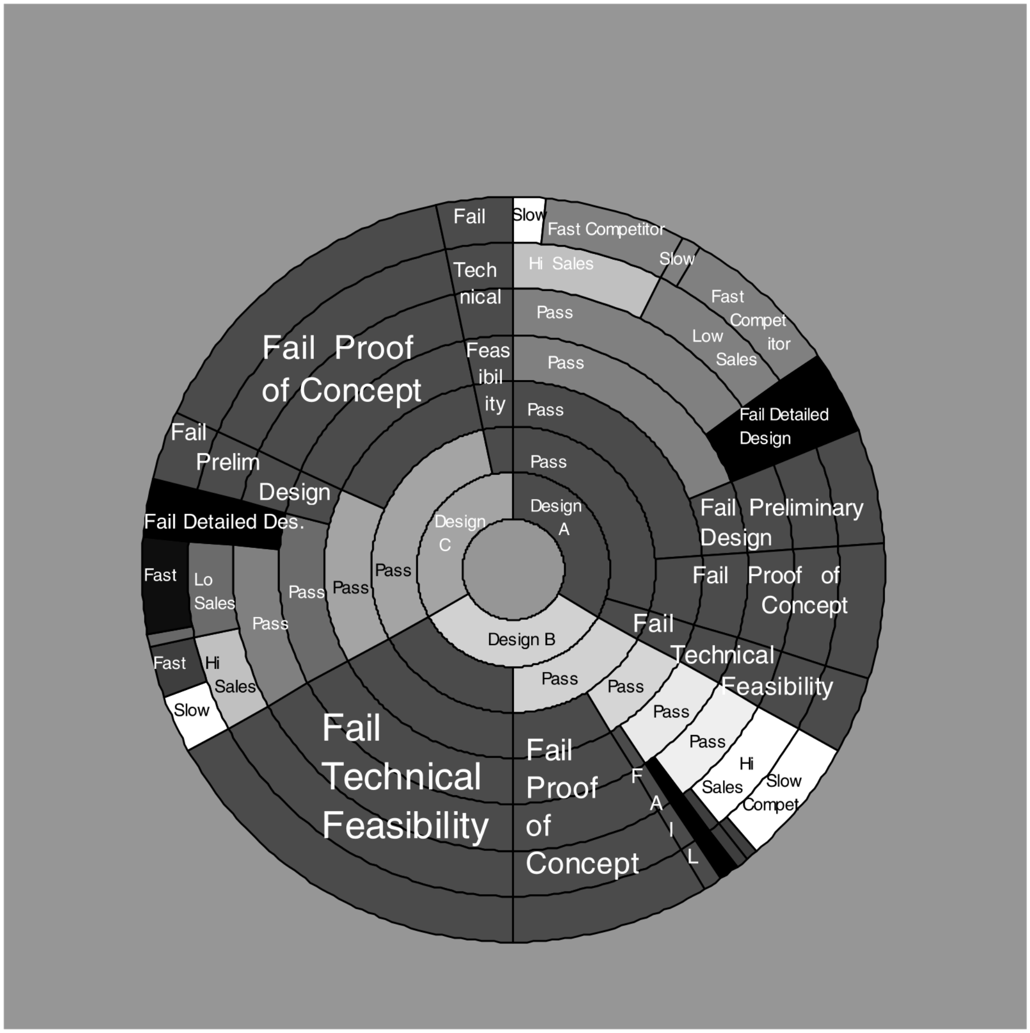

Figure 5. A circular decision tree for a complex R&D problem.

lower chance of reaching proof of concept. It has an intermediate chance of being quickly thwarted by the competition.

We now discuss how a circular decision tree would represent this problem.

4.3 Circular Decision Tree Representation

Since there are three alternatives, the circular decision tree is initially split into three (Figure 5). The lower third of the circle will focus on Design B, the rightmost third on Design A and the leftmost third on Design C.

Consider the region assigned to Design B. If we move down one layer, we either enter a big zone labeled "technical feasibility failure" or a "pass" zone. If we enter the pass zone, then moving further down takes us either into a big zone labeled "proof of concept" failure or a pass zone. If we enter the pass zone, we see that successive layers are mainly light-colored with only tiny failure zones. They terminate in the white zone, "high sales with slow competitive response"—which represents the best possible result. Hence if we avoid the technical feasibility and proof of concept failure zones, we have an excellent chance of moving to the highest payoff "white" region.

To analyze Design A, return to the center and start moving to the right. With Design A, the failure zones are all somewhat small. But the failure zones associated with later layers are still much thicker than in Design B. In particular, we see that one failure zone, "detailed design failure," is colored black and is much larger than the corresponding failure zone in Design B. We also see that the white region is much smaller with Design A than it was in Design B. As a result, the color assigned to the design A layer is darker than the color assigned to the Design B layer; that is, Design A is of lower value.

Hence our analysis appears to favor designs that concentrate most of the project's risk in the early phases of research.

To analyze Design C, return again to the center of the circle and move to the left. In this case, we find that the zone of technical feasibility failure was small. If we get past this layer, the zone of concept failure is fairly large. If we get past this failure zone, the remaining failure zones are small, but not trivial. The white zone associated with the best possible outcome is smaller than with Design B. On balance, Design C is slightly worse than Design B—which is reflected in the color assigned to it.

Informal surveys of several individuals familiar with the company's development process indicated that all found the doughnut representation easier to understand than the conventional decision tree.

5. BENEFITS OF A CIRCULAR TREE

5.1 Visual Presentation of Probabilities and Payoffs

The conventional decision tree is built from one-dimensional lines and nodes. Hence information about probability and value must be represented by writing numbers next to lines and nodes. In contrast, the circular decision tree is constructed from two-dimensional objects. This added dimension allows probabilistic information to be represented by varying the width of these objects and value information by varying their color. No numbers are explicitly needed. In addition, the computational algorithm used to solve decision trees corresponds to mixing colors in an intuitively natural way.

5.2 Visual Solution of Decision Trees

There are many proposed approaches for visually representing the *results* of an analysis. But the circular tree makes the *entire analysis* (*as well as the results*) visual. This may reduce the chances of *rejection* by clients who cannot, or do not have the time to, understand a mathematical analysis. Since the circular tree visually represents the entire analysis, it may also reduce the chances of *oversimplification* by clients tempted to summarize the analysis with a single number like the expected value. It also may address the *not-invented here syndrome* by allowing the clients to do the analysis themselves.

5.3 Identification of Which Uncertainties are Most Critical

In decision analysis, it's common—after the tree is constructed—to do a sensitivity analysis to identify the less critical uncertainties (i.e., those which do not provide good discrimination between good and bad outcomes.) These uncertainties are then removed from the tree. The circular representation suggests an alternate way of treating less critical uncertainties.

First note that trees in data mining always place those uncertainties first, providing the most discrimination between good and bad outcomes. Following this example, we propose that the circular tree, after it is first constructed, be modified by moving the most critical uncertainty layers closer to the center. As a result, adjacent areas of the tree (especially areas in the outer rings) will tend to have similar colors, leading to a visually more compact (and more comprehensible) representation. Less critical uncertainties, instead of being deleted, are assigned to the outer rings. (If we click on the diagram and then on "series order," Excel allows us to interchange the order of the rings.)

5.4 Representation of an Infinite Number of Possible Outcomes

The conventional decision tree requires that the number of outcomes emerging from a given node be limited since an individual cannot be expected to reasonably distinguish between more than five or six branches from any single node. In contrast, the circular tree could allow for a much larger number of outcomes for a given uncertainty. And, in fact, we could represent an infinite number of outcomes by coloring the wedge representing the uncertainty by a continuously varying set of colors.

[Received March 2001. Revised February 2002.]

REFERENCES

Barabba, V. (1991), "Through a Glass Less Darkly," *Journal of the American Statistical Association*, 86, 1–8.

Barlow, R., and Proschan, F. (1996), *Mathematical Theory of Reliability* (Classics in Applied Mathematics), Philadelphia, PA: SIAM.

Baron, J., and Brown, R. (1991), *Teaching Decisionmaking to Adolescents*, Hillsdale, NJ: Erlbaum.

Berger, J. (1985), *Statistical Decision Theory & Bayesian Analysis*, New York: Springer-Verlag.

Bernstein, P. (1998), *Against the Gods: The Remarkable Story of Risk*, New York: Wiley.

Bordley, R. (1998), "R&D Project Generation versus R&D Project Selection," *IEEE Transactions in Engineering Management*, December, 1998.

——— (2001), "Teaching Decision Theory in Applied Statistics Courses," *Journal of Statistical Education*, 9, 2.

Chernoff, H. (1987), *Elementary Decision Theory*, New York: Dover Publications.

DeGroot, M. (1970), *Optimal Statistical Decisions*, New York: McGraw-Hill.

Gertsbakh, I., and Gertsbakh, E. (2000), *Reliability Theory With Applications to Preventive Maintenance*, New York: Springer-Verlag.

Howard, R. A. (1988), "Decision Analysis: Practice and Promise," *Management Science*, 34, 679–695.

Hoyland, A., and Rausand, M. (1994), *System Reliability Theory: Models & Statistical Methods* (Wiley Series in Probability and Mathematical Statistics. Applied Probability and Statistics), New York: Wiley.

Johnson, B. (1993), "Tree-Maps: Visualizing Hierarchical and Categorical Data," PhD Dissertation, Dept. of Computer Science, University of Maryland, College Park.

Kusnic, M., and Owen, D. (1992), "The Unifying Vision Process: Value Beyond Traditional Decision Analysis in Multiple-Decision-Maker Environments," *Interfaces*, 22, 150–166.

Levine, D., Berenson, M., and Stephan, D. (1998), *Statistics for Managers Using Microsoft Excel*, Englwood Cliffs, NJ: Prentice Hall.

Lindley, D. (1991), *Making Decisions*, New York: Wiley.

Microsoft Excel 97 (1997), Redmond, Washington: Microsoft Corporation.

Raiffa, H. (1967), *Decision Analysis*, New York: McGraw-Hill.

Sharpe, P., and Keelin, T. (1998), "How SmithKline Beecham Makes Better Resource-Allocation Decisions," *Harvard Business Review*, March, 45–57.

Skinner, D. (1999), *Introduction to Decision Analysis*, Probabilistic Press.

Tobias, P., and Trindate, D. (1995), *Applied Reliability*, Boca Raton, FL: CRC Press.

Tufte, E. (1992), *The Visual Display of Quantitative Information*, Graphics Press.

——— (1997), *Visual and Statistical Thinking*, Graphics Press.