Stochastic Trees:

A New Technique for Temporal Medical Decision Modeling

GORDON B. HAZEN, PhD

This paper introduces stochastic trees, a new modeling approach for the class of medical decision problems in which risks of mortality and morbidity may extend over time. A stochastic tree may be regarded as a continuous-time version of a Markov-cycle tree, or alternately, as a multi-state DEALE model. Optimal decisions in stochastic trees can be determined by rollback, much in the same fashion as decision trees. The author discusses how age-dependent mortality rates and declining incidence rates may be modeled using stochastic trees. Concepts are illustrated using examples from the medical literature. It is argued that stochastic trees possess important advantages over Markov-cycle trees for medical decision modeling. *Key words*: stochastic trees; DEALE models; decision analysis; Markov cycle trees. **(Med Decis Making 1992;12:163–178)**

Medical treatment decisions often have uncertain consequences that stretch months or years into a patient's future. From a decision-analytic perspective, it is therefore desirable to construct stochastic models (models in which uncertainty unfolds over time) of patient prognosis. The simplest and most commonly used stochastic model is the discrete-time Markov chain. There have been a significant number of published stochastic analyses of medical decisions, some of which are summarized in table 1. Many of these have appeared in Pauker's continuing series, "Clinical Decision Making Rounds at the New England Medical Center," published in this journal.

Nearly all these analyses employ as stochastic model the discrete-time Markov chain.* The descriptor "dis-

Received August 29, 1991, from the Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois 60208-3119. Revision accepted for publication October 17, 1991.

Address correspondence and reprint requests to Dr. Hazen.

*This terminology conforms to that used in the stochastic-processes literature. There, Markov chains in which transition probabilities are time-dependent are called *nonstationary* Markov chains, and a chain is assumed to be stationary unless nonstationarity is explicitly indicated. The term *Markov process* refers to the general class of Markov chains that may be stationary or not, and have either discrete or continuous time domain.[20] However, certain authors refer to the smaller class of continuous-time, stationary Markov chains as Markov processes (e.g., Cinlar[21]). In unfortunate contrast, the medical decision making literature appears to follow Beck and Pauker,[3] who refer to discrete-time nonstationary Markov chains as Markov processes.

Table 1 • Published Stochastic Analyses of Medical Management Decisions

Reference	Medical Scenario
Pauker[1]	Coronary artery surgery
Beck and Pauker[2]	Anticoagulation in the brachycardia–tachycardia syndrome
Beck and Pauker[3]	Anticoagulation for a patient with an artificial heart valve and recent cerebral hemorrhage
Ransohoff et al.[4]	Prophylactic cholecystectomy for silent gallstones
Sonnenberg[5]	Endoscopic screening for gastric stump cancer
Sonnenberg[6]	Treatment of duodenal ulcer
Eckman et al.[7]	Eaton–Lambert syndrome and small-cell lung cancer
Hillner et al.[8]	Treating osteoporosis using postmenopausal estrogens
Matchar and Pauker[9]	Endarterectomy for transient ischemic attacks
Simon[10]	Use of cyclosporine in cadaveric kidney transplantation
Eckman et al.[11]	Redical cystectomy after recent myocardial infarction
Plante et al.[12]	Pyriform sinus carcinoma
Wong et al.[13]	Valvular thrombosis and cerebral hemorrhage
Ellwein and Farrow[14]	Urinary cytology screening
Fleming et al.[15]	Management of subarachnoid hemorrhage
Roach et al.[16]	Prostatic cancer and HIV infection
Tsevat et al.[17]	Anticoagulation therapy for dilated cardiomyopathy
Hagen et al.[18]	Management of an aortic aneurysm
Mooney et al.[19]	Magnetic resonance imaging in suspected multiple sclerosis

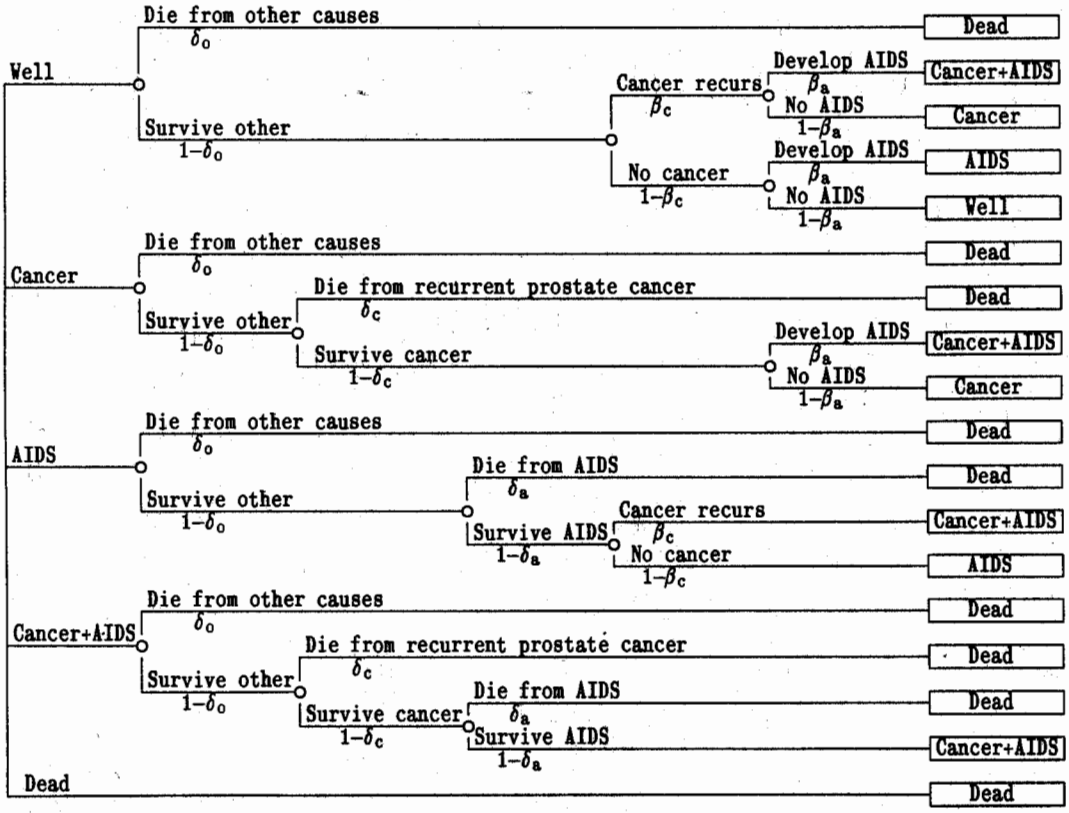

FIGURE 1. Markov-cycle tree for competing cancer and AIDS risks.

crete-time" indicates that the model treats time as a sequence of successive epochs (e.g., months, days, years), during which the state of the system remains fixed. This state may change only at the end of an epoch. By contrast, a continuous-time Markov chain treats time as a continuum, and jumps from one state to another may occur at any instant (e.g., Ross[22]). Continuous-time Markov chains have not been used extensively in the modeling of medical treatment decisions.

A convenient method for the formulation and display of discrete-time Markov chains is the Markov-cycle tree, originally devised by Hollenberg.[23] (For an illustration, see figure 1.) The purpose of this paper is to introduce a continuous-time analog of the Markov-cycle tree, which we call the *stochastic tree*. A stochastic tree is a graphic display of a continuous-time Markov chain. Because such chains have constant (not time-dependent) incidence and mortality rates, a stochastic tree may be regarded as a multi-state DEALE model.[24] The stochastic tree diagram:

1. incorporates familiar features from decision trees and Markov-cycle trees;

2. suggests a natural computational approach for solution; and

3. facilitates presentation and comprehension of the stochastic model.

Because of property 1, the stochastic tree may (like the Markov-cycle tree) be combined with the decision tree into a single diagram. Because of property 2, mean quality-adjusted lifetime and other mean durations of interest may be calculated by simply rolling back the stochastic tree (or the combined diagram) much as one would a decision tree. Markov-cycle trees also have this property, but the change from discrete to continuous time has the advantage that the rollback may directly use transition rates, thereby avoiding the nuisance translation from rates to probabilities required for discrete-time Markov modeling. Property 3 results because compared with Markov-cycle trees, stochastic trees are simultaneously more compact and less cluttered representations.

The stochastic tree concept is original to this author, and has not been previously published. It was motivated by (but does not appear in) the author's work in the Markov modeling of gout treatment.[25,26] Knowledgeable readers will recognize that a stochastic tree is merely a transition diagram for a continuous-time Markov chain that has been "unfolded" into a tree structure. Moreover, stochastic tree rollback is identical to value iteration (the method of successive approximations) in Markov decision processes.[22,27,28]

The following presentation is expository, the major concepts being illustrated by examples from the medical literature. We begin by giving a short review of Markov-cycle trees and pointing out the connection

between tree rollback and cohort analysis. Next we introduce stochastic trees, and stochastic tree rollback. Following this we examine the modeling of age-dependent mortality rates using stochastic trees. Finally, we show how declining incidence rates can be accounted for using stochastic trees.

Markov-cycle Trees

Discrete-time Markov chains were introduced to the medical decision making community by Beck and Pauker.[3] Most current analyses depict the Markov chain using the Markov-cycle tree devised by Hollenberg.[23] To illustrate Markov-cycle trees, we use the following example, taken from Roach et al.,[16] who consider the choice between radical prostatectomy and radiation therapy for a 58-year-old man who has asymptomatic HIV infection and stage B1 prostatic cancer. The Roach et al. analysis incorporates the discrete-time Markov chain having the structure given in figure 1, which is its Markov-cycle tree representation. It portrays the patient's possible states of health, in this case one of the five states Well, Cancer, AIDS, Cancer + AIDS, and Dead; as well as the possible monthly transitions between them. The patient begins the month in one of these five states on the left of the cycle tree, and depending on path taken on the sequence of chance nodes, finds himself in one of the states on the right of the cycle tree at the end of the month. The quantities labeling the branches in figure 1 are:

δ_0 = the monthly mortality probability due to other causes
β_c = the monthly probability of cancer recurrence
δ_c = the monthly mortality probability due to cancer
β_a = the monthly probability of developing AIDS
δ_a = the monthly mortality probability due to AIDS

Each transition path from left to right in figure 1 has an associated probability. For example, the probability that a patient in state Well will, in one month, move to state Cancer is given by a product of probabilities corresponding to the branches traversed in the Markov-cycle tree:

$$P(\text{Well} \rightarrow \text{Cancer})$$

$$= P\begin{pmatrix}\text{Survive death}\\ \text{from other causes}\end{pmatrix} P\begin{pmatrix}\text{Cancer}\\ \text{recurs}\end{pmatrix} P\begin{pmatrix}\text{No}\\ \text{AIDS}\end{pmatrix}$$

$$= (1 - \delta_0)\beta_c(1 - \beta_a) = \bar{\delta}_0\beta_c\bar{\beta}_a$$

Here the bars denote the complementary probabilities. Quantities such as $\bar{\delta}_0\beta_c\bar{\beta}_a$ are called *one-step* (or *one-cycle*) transition probabilities. They are all that is really required for computational purposes. In fact, computation is quicker if one works with the Markov-cycle tree containing only one-step transitions. This tree is depicted in figure 2, and is obtained by collapsing the nested chance nodes in the original cycle tree and combining all transition paths from a given state that lead to the same new state. For example, there are in figure 1 three ways to move from the state Cancer + AIDS to the state Dead in one cycle. In figure 2, these are combined into one transition path with probability label equal to the sum of the one-cycle transition probabilities derived from figure 1.

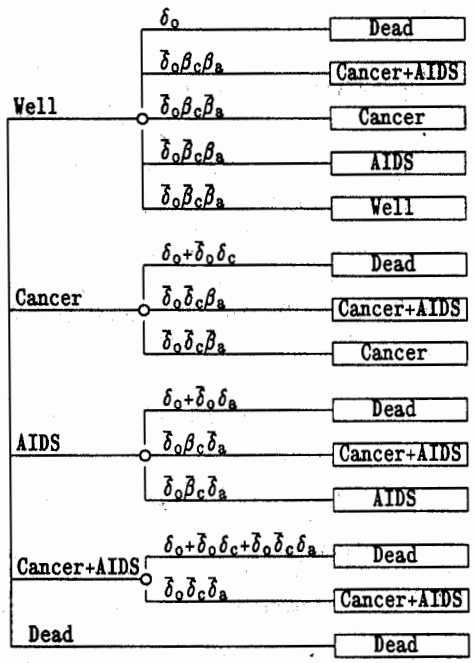

FIGURE 2. Markov-cycle tree depicting only the one-cycle transitions.

Cohort Analysis and Cycle Tree Rollback

Beck and Pauker[3] introduced a computational method called cohort analysis to calculate such mean durations as quality-adjusted lifetime in Markov chains. In cohort analysis one begins with 100% of a hypothetical population (the cohort) in the Well state initially, and calculates, using the one-cycle transition probabilities, the proportion of the population in each of the possible health states at the end of cycle 1, then cycle 2, and so forth. The cumulative proportion of patient cycles spent in each health state may then be derived, and from this, the mean quality-adjusted lifetime (or any other mean duration).

However, if one is interested only in a particular mean duration, say mean quality-adjusted lifetime, then the calculation of the cumulative proportion of patient cycles in each health state may be bypassed in the following way. For each health state y, let

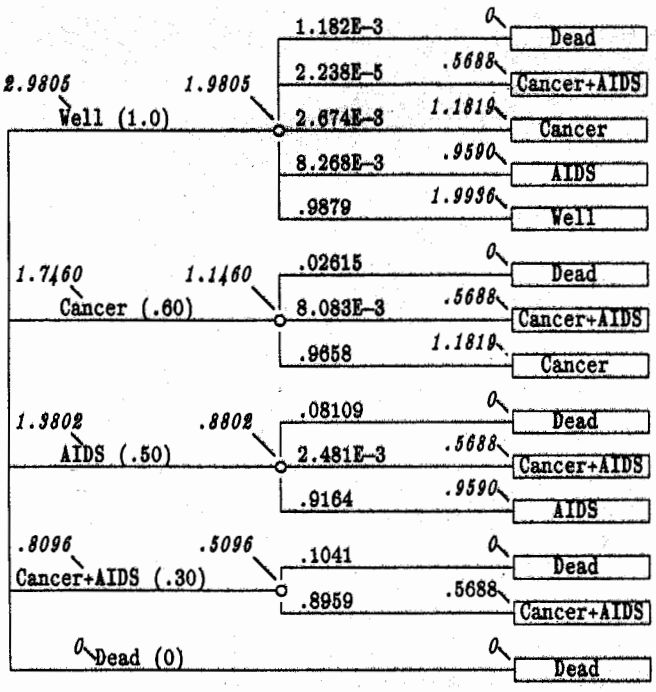

FIGURE 3. Cycle 3 of value iteration in the Markov-cycle tree of figure 2.

$L_n(y)$ = mean quality-adjusted duration accrued through n cycles, if the first cycle is begun in state y

Let $v(y) = L_1(y)$ be the quality factor associated with state y. Then the following recursive equation holds for $n > 1$:

$$L_n(y) = v(y) + \Sigma_z p_{yz} L_{n-1}(z) \qquad (1)$$

Here the sum is taken over all possible states z that may be reached from y, and p_{yz} is the one-cycle transition probability from y to z. (E.g., if y = Well and z = Cancer, then $p_{yz} = P(\text{Well} \rightarrow \text{Cancer}) = \tilde{\delta}_0\beta_c\tilde{\beta}_a$.)

The recursion (equation 1) is called *value iteration* in the theory of Markov decision processes, and has apparently not been used previously in the medical literature. Value iteration may be graphically depicted on the Markov-cycle tree. To give a concrete illustration of this process, let the parameters δ_c, δ_a, β_c, β_a take the values

$$\beta_c = 0.0027 \qquad \beta_a = 0.0083$$

$$\delta_c = 0.025 \qquad \delta_a = 0.080$$

used by Roach et al. under the assumption that radical prostatectomy has been done. Moreover, let δ_0 be constant and equal to the current background monthly mortality probability used by Roach et al.[16] for the patient (a 58-year-old white man), namely

$$\delta_0 = 1.182 \times 10^{-3}$$

(This value corresponds to a mortality rate of 0.01419/year.) Then the one-cycle transition probabilities take on the values shown on the branches in figure 3. In addition, suppose the quality factors v(y) are those used by Roach et al., namely

$$v(\text{Well}) = 1 \qquad v(\text{Cancer}) = 0.60$$

$$v(\text{AIDS}) = 0.50 \qquad v(\text{Cancer + AIDS}) = 0.30$$

We attach these values in parentheses next to the states on the left side of the cycle tree.

The graphic version of value iteration proceeds as follows. Assuming one has already computed $L_{n-1}(z)$ for all states z, begin cycle n of value iteration by labeling the states z on the right side of the cycle tree with these values. For example, figure 3 depicts value iteration cycle 3, and the rightmost labels are the $L_2(z)$ values. There $L_2(\text{Cancer}) = 1.1819$. Next roll back the cycle tree just as one would roll back a decision tree, to obtain average values at the chance nodes. For example, in figure 3, the value 1.1460 at the chance node following Cancer is obtained by averaging the three values 0, 0.5688, 1.1819. Finally, to each average value add the associated quality factor v(y) and label the associated state with the result, which is $L_n(y)$. For example, the label 1.7460 on Cancer is obtained by calculating

$$L_3(\text{Cancer}) = v_3(\text{Cancer}) + 1.1460$$

$$= 0.60 + 1.1460 = 1.7460$$

Of course, the original $L_2(z)$ values in figure 3 were obtained by initially labeling the rightmost states with the quantities $L_1(z)$, and performing the rollback procedure just described. And to start the procedure, each

Table 2 • Value Iteration for the Markov-cycle Tree of Figure 3

State	Cycle						
	1	2	3	10	100	1000	1100
Well	1.0	1.9936	2.9805	9.6809	63.019	90.473	90.473
Cancer	0.6	1.1819	1.7460	5.2311	17.637	18.206	18.206
AIDS	0.5	0.9590	1.3802	3.5035	6.068	6.069	6.069
Cancer + AIDS	0.3	0.5688	0.8096	1.9222	2.883	2.883	2.883
Dead	0	0	0	0	0	0	0

$L_1(z)$ is by definition set equal to the given quality factor $v(z)$. Value iteration may continue for arbitrarily many cycles. The values $L_n(y)$ will converge as n becomes large to the true mean quality-adjusted duration $L(y)$ beginning in state y. The convergence process is illustrated in table 2. On the order of 1,000 cycles are required before the $L_n(y)$ values stabilize. Mean quality-adjusted lifetime beginning in the Well state is 90.47 months.

This value is not in agreement with Roach et al. because we have assumed a constant background mortality probability δ_0. If we allow δ_0 to depend on age, then it must take on a new value during each cycle of value iteration (or of cohort analysis, if that is performed). When this is done with the correct values of δ_0, mean quality-adjusted lifetime is reduced to only 83.13 months.

How does one obtain the event probabilities (the δs and βs in figure 1) required in the cycle tree? Typically data are available to estimate the mean number of occurrences per unit time of the event of interest. If events occur randomly over time at some rate μ (events/unit time), then the probability δ of at least one occurrence in duration t is given by

$$\delta = 1 - e^{-\mu t} \qquad (2)$$

For example, if the mortality rate is $\mu = 0.014191$ per year (appropriate for a 58-year-old white male), then the probability of death in one month ($t = 1/12$ year) is

$$\delta = 1 - e^{-0.014191/12} = 0.0011818$$

which is the value of δ_0 used above. The typical strategy is to convert available data to an annual rate, manipulate the rate if necessary, and then use equation 2 to convert the annual rate to a monthly probability of occurrence. A key property is that rates from disjoint causes may be added (whereas the corresponding occurrence probabilities may not). For example, Roach et al. use this fact to remove background mortality μ_{70} from the overall rate $\mu_c + \mu_{70}$ driving mortality among cancer patients. (The average age of prostatic cancer patients is 70 years.) Although probabilities rather than rates are used in the Markov model computations, conversion to rates is required to successfully manipulate the relevant data.

Stochastic Trees

A stochastic tree may be regarded as the continuous-time version of a Markov-cycle tree. For example, the stochastic tree model corresponding to the Markov-cycle tree of figure 1 is given in figure 4. In figure 4, the wavy arrows denote uncertain *durations*, whose magnitudes are determined by the rate parameters

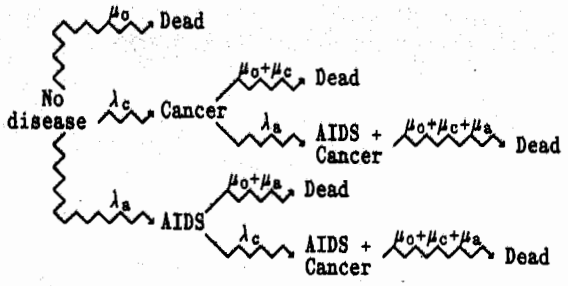

FIGURE 4. A stochastic tree for competing cancer and AIDS risks.

that label the individual arrows. For instance, from the "No disease" state, there is a mortality rate μ_0 equal to the mean number of deaths in a one-year period for individuals of this type. Similarly, there is a cancer recurrence rate λ_c equal to the mean number of cancer recurrences per year, and an annual AIDS onset rate μ_a as well. The diagram represents a situation in which the "No disease" state will endure until the first of the three events death, cancer, or AIDS occurs. The rates μ_0, λ_c, λ_a are constants (not time-dependent), which implies that the corresponding durations are exponentially distributed. A stochastic tree is in this sense a multi-state DEALE model.[24]

Stochastic trees improve on Markov-cycle trees in several ways. Most apparently, they usually give a more compact representation of system dynamics. The stochastic tree of figure 4 contains the same information as the Markov-cycle tree of figure 1, but in dramatically less space. It contains more information than the collapsed cycle tree of figure 2, but still is more compact. The inefficient use of space by the Markov-cycle tree is due to time discretization: The cycle tree must squeeze all possible event sequences into a cycle (e.g., not just cancer occurrence or AIDS onset, but also both in one cycle). It must also depict the transitions from a state to all successors that can be reached in one cycle, whereas a stochastic tree depicts transitions only to immediate successors.

Like Markov-cycle trees, stochastic trees may be rolled back to give mean quality-adjusted lifetime and other appropriate mean durations. The remainder of this section is devoted to a discussion of this procedure. As we shall see, rollback is more efficient in stochastic trees than in Markov-cycle trees, and in simple cases may even be done with pencil, paper, and calculator.

The requirement of constant event rates, and the corresponding exponentially distributed durations, is less of a restriction than it seems. A stochastic tree may in fact depict any *sum* of exponentially distributed durations, and the distributions of these sums are flexible enough to match nearly any desired distribution. A more detailed discussion is presented in a subsequent section of the paper.

MANIPULATING STOCHASTIC TREES

We describe first a fundamental manipulation that

may be used to introduce or delete chance nodes from stochastic trees. Consider a stochastic tree

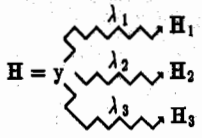

(3)

consisting of an initial state y from which there are *n* possible transitions (here $n = 3$), the ith transition occurring with rate λ_i and leading to stochastic subtree $\mathbf{H}_i$. The fundamental manipulation is

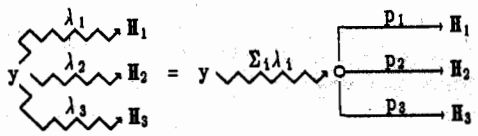

(4)

where $p_i = \lambda_i/\Sigma_i\lambda_i$. In other words, waiting in state y until the first of three competing transitions occurs is equivalent to waiting in y until an anonymous transition occurs and then determining its identity (type i with probability p_i). Stated another way, the first transition occurs at a rate equal to the sum of the rates of the competing transitions, and its identity is determined by probabilities proportional to the competing rates. In the stochastic-processes literature, the transformation in equation 4 from left to right is known as *superposition* of Poisson processes, and the reverse transformation (right to left) is known as *decomposition* of Poisson processes.[21,22]

There is a second manipulation that is sometimes useful. Suppose the stochastic tree **H** of equation 3 has two identical subtrees $\mathbf{H}_1 = \mathbf{H}_2 = \mathbf{K}$. Then

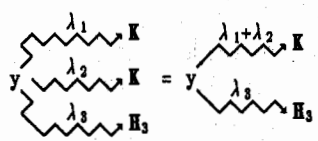

(5)

It is not hard to show that the manipulation of equation 5 is actually a consequence of equation 4.

ROLLING BACK STOCHASTIC TREES

To see how stochastic tree rollback works, consider a fragment

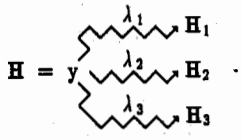

of a stochastic tree in which state y is occupied until the first of three events with competing rates λ_1, λ_2, λ_3 occurs, after which occur subtrees $\mathbf{H}_i$. Suppose subtree $\mathbf{H}_i$ has associated mean quality-adjusted duration $L(\mathbf{H}_i)$ (which may already have been determined by rolling back subsequent portions of the tree). Invoking the superposition rule of equation 4, we may write

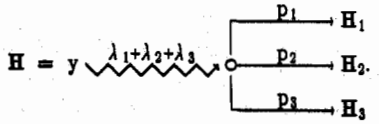

The mean quality-adjusted duration at the chance node is $\Sigma_i p_i L(\mathbf{H}_i)$, which equals

$$\frac{\Sigma_i \lambda_i L(\mathbf{H}_i)}{\Sigma_i \lambda_i}$$

Because the mean time in state y is $1/\Sigma_i\lambda_i$, the mean quality-adjusted duration spent in state y is

$$\frac{v(y)}{\Sigma_i \lambda_i}$$

where v(y) is the quality factor associated with state y. Therefore, the total quality-adjusted duration beginning at y is

$$L(\mathbf{H}) = \frac{v(y) + \Sigma_i \lambda_i L(\mathbf{H}_i)}{\Sigma_i \lambda_i} \qquad (6)$$

Interestingly, $L(\mathbf{H})$ is equal to the mean duration $1/\Sigma_i\lambda_i$ in state y multiplied by a quality factor

$$v(y) + \Sigma_i \lambda_i L(\mathbf{H}_i)$$

which is in turn obtained from the quality factor of y and the natural combination of the quality-adjusted durations $L(\mathbf{H}_i)$. The similarity to Markov-cycle tree rollback (equation 1) becomes evident if we write equation 6 as

$$L(\mathbf{H}) = v(y) \cdot \frac{1}{\Sigma_i \lambda_i} + \Sigma_i p_i L(\mathbf{H}_i)$$

To illustrate, let us roll back the stochastic tree of figure 4, using the rate values

$\lambda_c = 0.03250$/year $\quad \mu_c = 0.3081$/year

$\lambda_a = 0.10$/year $\quad \mu_a = 0.9970$/year

taken from the Roach et al.[16] analysis for the surgical option. We use the same age-independent mortality rate $\mu_0 = 0.014191$/year as in our initial Markov-cycle tree analysis above. We substitute these values into

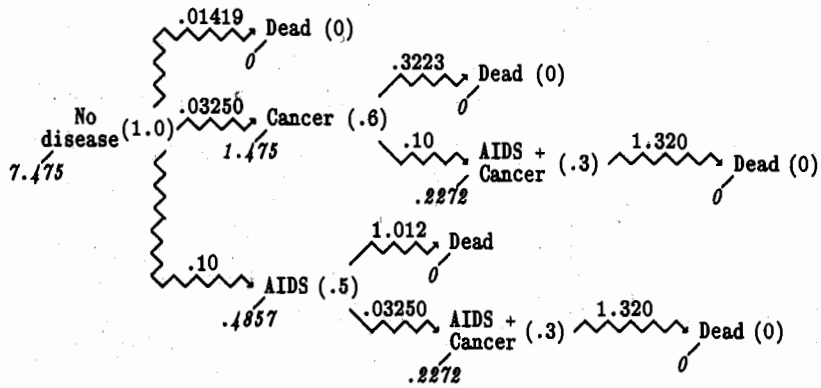

FIGURE 5. Stochastic tree rollback (surgical option).

figure 4, and also attach quality factors (in parentheses) to each state to obtain figure 5. We roll back using equation 6, with calculated mean quality-adjusted durations (in years) attached in italics to each state. For example, for the state y = AIDS + Cancer, equation 6 becomes

$$L(\text{AIDS + Cancer}) = \frac{0.3 + 1.320 \cdot 0}{1.320} = 0.2272 \text{ year}$$

For state y = Cancer, equation 6 gives

$$L(\text{Cancer}) = \frac{0.6 + (0.3323) \cdot 0 + (0.10)(0.2272)}{0.3323 + 0.10}$$

$$= 1.475 \text{ years}$$

For the initial state y = No disease, we obtain

$$L(\text{No disease}) = \frac{1.0 + (0.01419) \cdot 0 + (0.03250)(1.475) + (0.10)(0.4857)}{0.01419 + 0.03250 + 0.10}$$

$$= 7.475 \text{ years} = 89.70 \text{ months}$$

This figure differs slightly from the mean quality-adjusted duration 90.47 months obtained by rolling back the Markov-cycle tree because of time discretization. In the cycle tree model, all morbidities and mortalities are assumed to occur at the end of a cycle; a patient therefore appears to spend more time in healthy states, increasing his or her quality-adjusted life expectancy.

Notice that in this case, stochastic tree rollback is a finitely terminating process, which could have been (and was) done with only pencil, paper, and calculator. In contrast, Markov-cycle tree rollback (or cohort analysis) converges but does not terminate, and requires a computer.

Stochastic tree rollback could also be used to calculate life expectancy: Merely set all the quality factors v(y) equal to 1. Similarly, if one wishes the mean duration in a particular state y_0, set $v(y_0) = 1$ and $v(y) = 0$ for all other health states y.

This stochastic tree model uses a constant rather than an age-dependent mortality rate μ_0. We postpone the discussion of age dependencies until the next section.

ROLLING BACK CYCLIC STOCHASTIC TREES

Many stochastic models of patient status contain *cycles*, that is, health states that can be visited for some duration, vacated, and later revisited. An example is the stochastic tree model of stroke occurrence depicted in figure 6, which is based on a corresponding discrete-time Markov model formulated by Matchar and Pauker.[9] In figure 6, parenthesized state names indicate one is re-entering a previously depicted state. For example, when one reaches the terminal node "(Big stroke)," one is really revisiting the previous state "Big stroke." The same comment applies to the states "(Stroke)" and "Stroke." One can therefore cycle through the tree repeatedly.

Rollback of cyclic stochastic trees is more complicated, but still fairly intuitive: Initially assign the states "(Stroke)" and "(Big stroke)" cumulative quality-adjusted duration of zero. Roll back the stochastic tree to obtain cumulative mean quality-adjusted durations for all other states, including "Stroke" and "Big stroke." Assign these new values to "(Stroke)" and "(Big stroke)," respectively, and repeat the rollback, obtaining new cumulative values for "Stroke" and "Big stroke." Continue this cyclic process until the values assigned converge. This process is a continuous-time version of the Markov-cycle tree rollback discussed previously.

We illustrate this cyclic rollback process for the stochastic tree of figure 6, using the following data from Matchar and Pauker:

$$\mu_{excess} = 0.065/\text{year}$$
$$\mu_{stroke} = 0.05/\text{year}$$
$$p_{earlydie} = 0.38$$

(The quantity μ_{excess} represents additional mortality due to the patient's severe coronary disease.) Matchar and Pauker use an age-dependent background mortality rate μ_0. For this illustration we use a constant

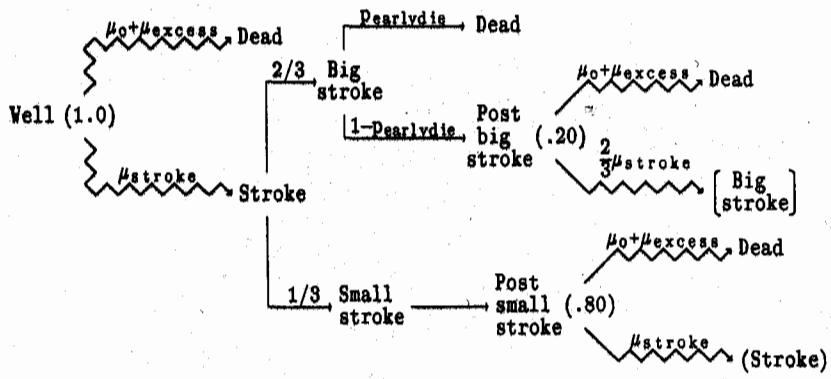

FIGURE 6. Stochastic tree model of recurrent stroke.

mortality rate $\mu_0 = 0.01106$/year, appropriate for a 55-year-old white male. (We treat age-dependent mortality below.) The cycles of value iteration are displayed in table 3.

Unlike the previous example, the stochastic tree of figure 6 contains what might be termed instantaneous states, namely Stroke, Big stroke, and Small stroke. The purpose of such states is to indicate events whose durations are so short as to be immaterial. Branches emanating from instantaneous states represent chance events, and are labeled with probabilities, not rates. Rollback at such states consists of merely averaging the quality-adjusted durations following each branch, using the specified probabilities as weights. For example,

$$L(\text{Big stroke}) =$$

$$p_{\text{earlydie}} \cdot L(\text{Dead}) + (1 - p_{\text{earlydie}}) \cdot L(\text{Post big stroke})$$

and

$$L(\text{Stroke}) = \frac{2}{3} \cdot L(\text{Big stroke}) + \frac{1}{3} \cdot L(\text{Small stroke})$$

In this respect, instantaneous states are identical to chance nodes in a decision tree.

As can be seen from table 3, convergence to three decimal places has occurred by cycle 6. Mean quality-adjusted duration beginning in the Well state is 9.325 years. In dramatic contrast, the Markov-cycle tree calculations in table 2 for the Roach et al.[16] model required on the order of 1,000 cycles to converge. Although we have not depicted Markov-cycle tree rollback for this example, it also requires roughly 1,000 cycles to converge.

Modeling Age-dependent Mortality with Stochastic Trees

One apparent shortcoming of stochastic tree modeling of patient prognosis is that while human mortality rates are age-dependent, the rate parameters in a stochastic tree are constant over time. For example, consider the simple stochastic tree model for human life length

This is identical to the DEALE approach[24] to modeling life length: Using the rollback formula of equation 6 gives

$$L(\text{Alive}) = \frac{v(\text{Alive}) + \mu_0 L(\text{Dead})}{\mu_0} = \frac{v(\text{Alive})}{\mu_0}$$

which is the DEALE formula for quality-adjusted lifetime. However, because μ_0 is constant rather than time-dependent, the mean duration in the Alive state of the stochastic tree will only poorly approximate the true additional mean human lifetime. For example, the mortality rate $\mu_0 = 0.01419$/year appropriate for a 58-year-old white male gives mean additional lifetime $1/\mu_0$ of over 70 years. This overestimate occurs because the increase of μ_0 with age is ignored. Although one may obtain the correct life *expectancy* by choosing some "average" value of μ_0, this value will overestimate

Table 3 • Value Iteration for the Stochastic Tree of Figure 6

State	Cycle						
	0	1	2	3	4	5	6
Well	0	9.072	9.279	9.317	9.324	9.325	9.325
Stroke	0	2.871	3.393	3.489	3.507	3.511	3.511
Big stroke	0	1.134	1.348	1.388	1.396	1.397	1.397
Post–big stroke	0	1.828	2.174	2.239	2.251	2.254	2.254
Small stroke	0	6.346	7.485	7.692	7.730	7.737	7.739
Post–small stroke	0	6.346	7.485	7.692	7.730	7.737	7.739

true mortality for early years and underestimate it for later years.

Moreover, as Hagen, Eckman, and Pauker[18] point out, using a constant mortality rate μ_0 can cause modeling errors when there are other competing sources of mortality. Their model of mortality due to abdominal aortic aneurysmorrhaphy (AAA) may be depicted as the stochastic tree

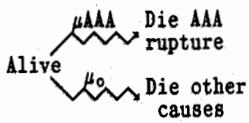

Invoke the superposition rule of equation 4 to produce the equivalent stochastic tree

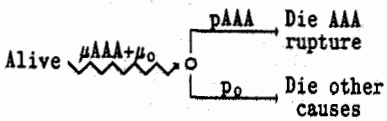

As an unexpected and undesirable consequence of this diagram, those dying of AAA rupture have the same life expectancy, namely $1/(\mu AAA + \mu_0)$, as those dying of other causes. Moreover, this is true regardless of the relative values of μAAA and μ_0. As Hagen et al.[18] verified using cohort analysis, this phenomenon is due to the constancy of μ_0: When the proper age-dependent Gompertz rates are used, there is a difference in life expectancy between the two cohorts (shorter for those dying of AAA rupture), and the difference does change when μAAA changes. We return to this example below.

THE ERLANG MORTALITY MODEL

Although modeling human mortality using these simple stochastic trees is inappropriate, more sophisticated stochastic tree models can approximate human mortality quite well. Essentially this is due to the fact that a stochastic tree can depict any *sum* of independent exponentially distributed durations, and properly chosen sums of this type have distributions that quite closely approximate human lifetime distributions. One such sum of exponentials which is particularly simple is the Erlang, with parameters n and γ. By definition, an Erlang (n,γ) duration is the sum of n independent exponential (γ) durations, and is therefore the time until death in the stochastic tree

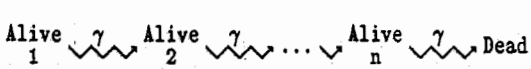

The successive "Alive" states in this tree may be viewed as life stages, or perhaps as consecutive declining states of health.

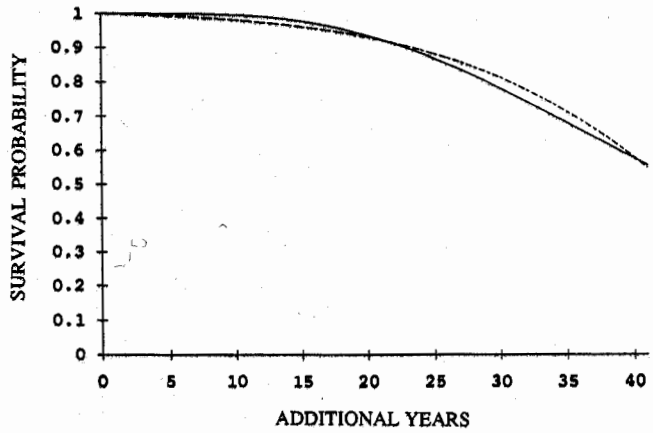

FIGURE 7. Erlang approximation (*solid line*) to 40-year-old white female survival probabilities (*broken line*).

The Erlang distribution can be quite accurate in approximating human lifetime distributions. For example, figure 7 compares the Erlang survival function, given by

$$S(t) = P(\text{Survive at least t years}) = \sum_{k=0}^{n-1} \frac{(\gamma t)^k}{k!} e^{-\gamma t}$$

using $n = 5$ stages and $\gamma = 0.1068$, with the survival function for a 40-year-old white female.[29] The maximum difference between the curves is 0.0251. Similar accurate approximations may be achieved for other gender–race–age combinations, as summarized in table 4.†

The maximum difference in survival probabilities is a useful indicator of accuracy, for the following reason. Suppose we wish to approximate a background mortality process with true survival function $S_0(t)$ using an approximate survival function S(t) such as the Erlang. Let

$L_0(t)$ = true mean quality-adjusted life years accrued in [0,t], calculated using $S_0(t)$;

$L(t)$ = approximate mean quality-adjusted life years accrued in [0,t] calculated using S(t).

We would hope that L(t) would be close to $L_0(t)$. Let

†For each number n of stages, bisecting-line search over values of λ was performed to minimize the deviation of the Erlang from the true survival function. The calculations were programmed in Pascal by Declan O'Riordan, an undergraduate industrial engineering major at Northwestern.

Table 4 • Erlang Approximations to Human Survival Probabilities

	Initial Age (Years)	Male			Female		
		n	γ	Max Error	*n*	γ	Max Error
White	20	11	0.1979	0.03839	11	0.1688	0.02320
	25	10	0.1976	0.03782	10	0.1672	0.02356
	30	8	0.1733	0.03573	8	0.1440	0.02237
	35	7	0.1698	0.03563	7	0.1395	0.02503
	40	6	0.1656	0.03833	5	0.1068	0.02512
	45	4	0.1228	0.03698	4	0.09430	0.02502
	50	3	0.1055	0.03845	3	0.07748	0.02728
	55	3	0.1300	0.03695	3	0.09434	0.02492
	60	2	0.09771	0.02773	2	0.06607	0.02138
	65	2	0.1274	0.02824	2	0.08544	0.01898
	70	2	0.1710	0.03766	2	0.1146	0.02446
	75	1	0.07685	0.01650	1	0.04390	0.01126
Black	20	6	0.1220	0.05118	8	0.1358	0.03014
	25	5	0.1118	0.05223	7	0.1297	0.03016
	30	5	0.1257	0.05041	6	0.1221	0.03082
	35	4	0.1112	0.04374	5	0.1121	0.03039
	40	3	0.09357	0.04879	4	0.09900	0.02876
	45	3	0.1084	0.03979	3	0.08177	0.03108
	50	2	0.08123	0.04725	3	0.09701	0.03084
	55	2	0.09614	0.03109	2	0.06924	0.02803
	60	2	0.1203	0.03715	2	0.08537	0.02628
	65	2	0.1499	0.04034	2	0.1064	0.02801
	70	1	0.07134	0.03506	1	0.04479	0.02940
	75	1	0.08790	0.01525	1	0.05575	0.01380

$L^*(t)$ = mean quality-adjusted life years accrued in [0,t] in the absence of the background mortality process.

Then the following result may be derived.

Theorem: Assume each health state has non-negative quality factor. Then for all $t \geq 0$

$$|L(t) - L_0(t)| \leq L^*(t) \cdot \max_{0 \leq u \leq t} |S(u) - S_0(u)|$$

Stated another way, the proportionate error

$$\frac{|L(t) - L_0(t)|}{L^*(t)}$$

in calculating quality-adjusted lifetime over [0,t] is bounded above by the maximum difference between the exact and the approximate survival functions over this interval.

AN EXAMPLE: COMPETING MORTALITY SOURCES

To illustrate a nonnumeric use of the Erlang mortality model, we return to Hagen, Eckman, and Pauker's[18] discussion of competing mortality rates in abdominal aortic aneurysmorrhaphy (AAA). Assuming a two-stage Erlang model for mortality, the proper stochastic tree is

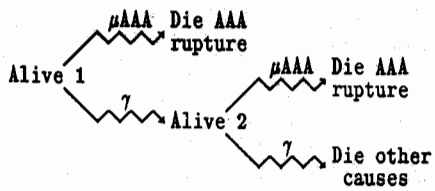

Notice that the two-stage Erlang model is incorporated as the bottommost path in the tree. We wish to compare life expectancies of those who die of AAA rupture with life expectancies of those who die of other causes. We shall show that the conclusions obtained numerically by Hagen et al. using cohort analysis may also be reached in a completely symbolic fashion by manipulating the stochastic tree.

Proceed in the following way. First use the superposition rule of equation 4 to obtain the equivalent stochastic tree:

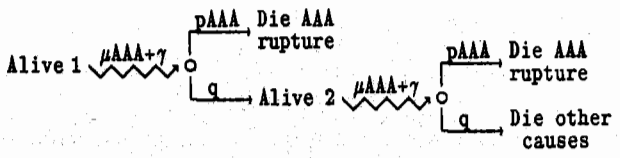

where

$$pAAA = \frac{\mu AAA}{\mu AAA + \gamma} \qquad q = \frac{\lambda}{\mu AAA + \gamma}$$

Next, shift the chance nodes to the beginning of the tree to obtain

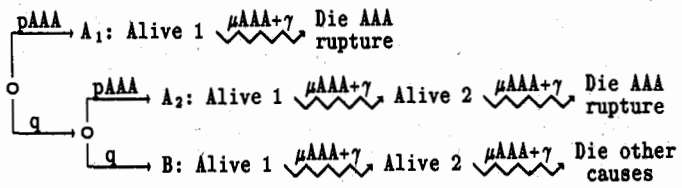

Here we have inserted the labels A_1, A_2, B and will use them from here on as abbreviations for the parts of the tree they precede. The goal of these manipulations is to segregate A_1, A_2 from B in the tree, so that the life expectancies under different causes of death may be compared. Do this with the chance node manipulations

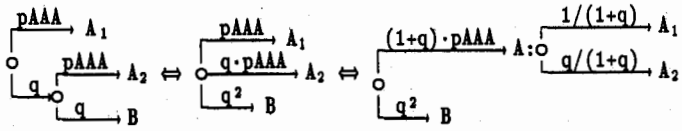

where the additional label A represents the part of the tree in which the cause of death is AAA rupture.

We wish to compare the life expectancies L(A) and L(B). Taking v(Alive 1) = v(Alive 2) = 1 and using the rollback of formula of equation 6 gives

$$L(B) = \frac{2}{\mu AAA + \gamma}$$

which is the same value as $L(A_2)$. Moreover, $L(A_1) = 1/(\mu AAA + \gamma)$. Therefore,

$$L(A) = \frac{1}{1+q} \cdot L(A_1) + \frac{q}{1+q} L(A_2)$$

$$= \frac{1}{1+q} \cdot \frac{1}{\mu AAA + \gamma} + \frac{q}{1+q} \cdot \frac{2}{\mu AAA + \gamma}$$

Evidently L(A) is the average of two quantities, one equal to L(B) and the other half that amount. Therefore, L(A) is smaller than L(B), that is, life expectancy under AAA rupture is shorter than life expectancies under other causes of mortality, and this is true regardless of the value of μAAA. These were precisely the numerical conclusions of Hagen et al.

ROLLBACK WITH IMPLICITLY MODELED ERLANG MORTALITY

Although the Erlang mortality model can accurately approximate human mortality rates, its use multiplies the number of states in the stochastic tree. The reason is that for every health state y in the "naive" stochastic tree, an *n*-stage Erlang mortality model requires states (y,1), (y,2), . . . , (y,*n*), where state (y,k) denotes the occupation of health state y while in Erlang stage k.

Fortunately, an appropriately augmented rollback process that incorporates Erlang mortality can be carried out on the naive stochastic tree. The rollback formula may be derived as follows. Suppose y is a state in the naive tree having successors y_1, y_2, y_3 that can be reached with corresponding transition rates λ_1, λ_2, λ_3. Then including Erlang mortality gives an *augmented* stochastic tree that at state (y,k) looks like

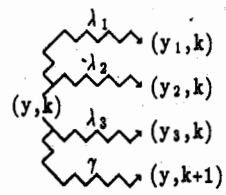

Here the extra bottom transition represents change from stage k to k + 1 in the Erlang model (assuming $k < n$; if $k = n$ the bottom transition would be absent). Let

$L_k(y)$ = mean quality-adjusted duration beginning in state (y,k)

Then according to rollback formula of equation 6, we have

$$L_k(y) = \frac{v(y) + \gamma L_{k+1}(y) + \Sigma_i \lambda_i L_k(y_i)}{\gamma + \Sigma_i \lambda_i} \qquad (7)$$

(if $k = n$ then the term $\gamma L_{k+1}(y)$ would be missing). One need not draw the augmented Erlang tree to use this formula. Instead, one can simply record the values $L_1(y)$, . . . , $L_n(y)$ near state y in the naive tree.

An example should be helpful. Figure 8 presents the naive stochastic tree for the Roach et al.[16] model of competing cancer and AIDS risks (figs. 4 and 5). Background mortality (the rates μ_0 in fig. 4) is to be treated implicitly using the Erlang model, so does not appear in the figure. The patient in this case is a 58-year-old white man. The nearest Erlang approximations in table 4 are to 60- and 55-year-old white males. We use the latter to more clearly illustrate the rollback calculations, that is, we use an Erlang (e,γ = 0.1300) approximation to the patient's remaining lifetime.

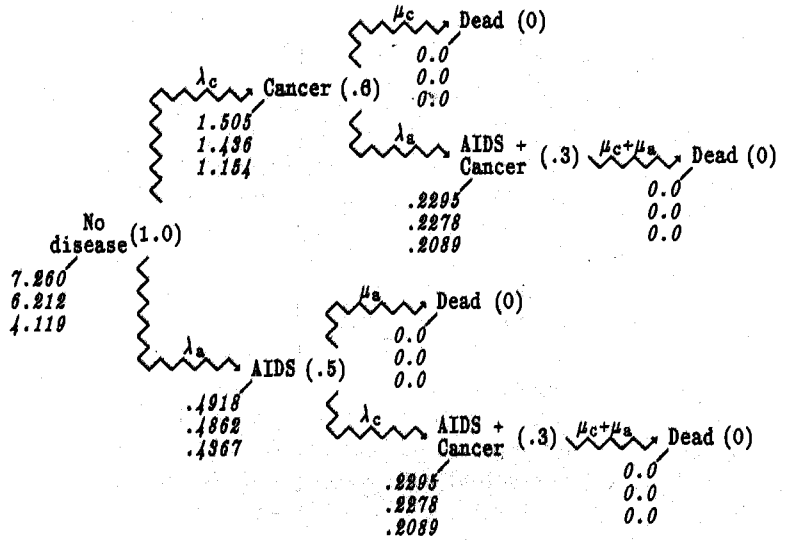

FIGURE 8. Stochastic tree rollback with implicit three-stage Erlang mortality.

The results of the rollback procedure are indicated in the figure by attaching the quantities $L_1(y)$, $L_2(y)$, $L_3(y)$ to each state y. The calculations at state y = AIDS + Cancer, for example, are

$$L_3\begin{pmatrix}\text{AIDS +}\\ \text{Cancer}\end{pmatrix} = \frac{0.3 + \gamma\cdot 0 + (\mu_c + \mu_a)\cdot 0}{\gamma + \mu_c + \mu_a}$$

$$= 0.2089$$

$$L_2\begin{pmatrix}\text{AIDS +}\\ \text{Cancer}\end{pmatrix} = \frac{0.3 + \gamma\cdot(0.2089) + (\mu_c + \mu_a)\cdot 0}{\gamma + \mu_c + \mu_a}$$

$$= 0.2278$$

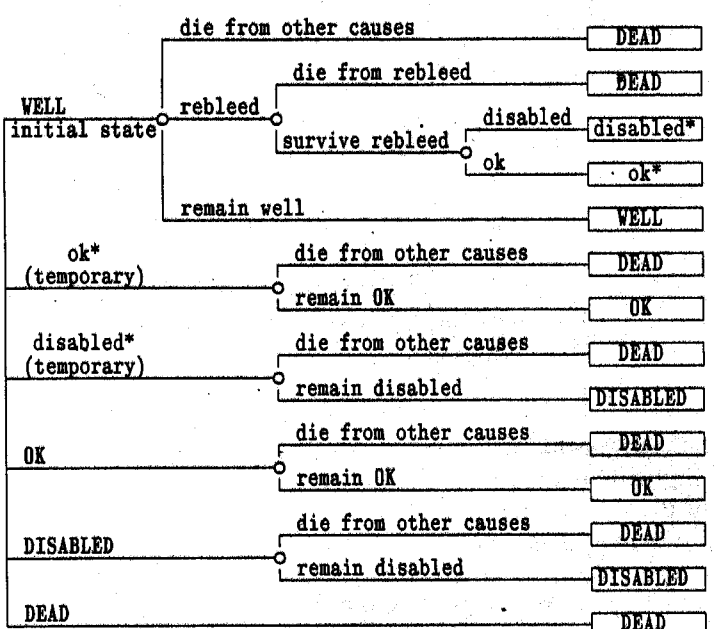

FIGURE 9. Markov-cycle tree depicting risk of aneurysm rupture.

$$L_1\begin{pmatrix}\text{AIDS +}\\ \text{Cancer}\end{pmatrix} = \frac{0.3 + \gamma\cdot(0.2278) + (\mu_c + \mu_a)\cdot 0}{\gamma + \mu_c + \mu_a}$$

$$= 0.2295.$$

The calculations at state y = Cancer are

$$L_3(\text{Cancer}) = \frac{0.6 + \gamma\cdot 0 + \mu_c\cdot 0 + \lambda_a\cdot(0.2089)}{\gamma + \mu_c + \lambda_a}$$

$$= 1.154$$

$$L_2(\text{Cancer}) = \frac{0.6 + \gamma\cdot(1.154) + \mu_c\cdot 0 + \lambda_a\cdot(0.2278)}{\gamma + \mu_c + \lambda_a}$$

$$= 1.436$$

$$L_1(\text{Cancer}) = \frac{0.6 + \gamma\cdot(1.436) + \mu_c\cdot 0 + \lambda_a\cdot(0.2295)}{\gamma + \mu_c + \lambda_a}$$

$$= 1.505.$$

Once calculations are complete, the key statistic is L_1(No disease) = 7.260 years = 87.12 months, the mean quality-adjusted duration beginning with no disease in Erlang stage 1. This compares to the overestimate of 89.70 months calculated previously assuming constant mortality rate. When the age-60 Erlang model ($n = 2$, $\gamma = 0.09771$) is used, the result is a mean quality-adjusted duration of 81.17 months.

Modeling Declining Incidence Rates with Stochastic Trees

While mortality rates increase over time, instances in which incidence rates decrease over time are also common. For example, the incidence rate of adverse drug side effects may decrease from the time of initial administration of the drug, because each subsequent day without side effects reduces the likelihood the patient is sensitive to the drug. An example we shall

consider in detail was treated by Fleming et al.,[15] who modeled a declining incidence of aneurysm rupture following an initial subarachnoid hemorrhage. They constructed a discrete-time Markov model in which the rupture occurs with probability 0.10 in week 1, 0.08 in week 2, 0.05 in week 3, and thereafter exponentially declines to a value of 0.0044 in week 26 (see their table 3).

We reproduce their Markov-cycle tree in figure 9. The tree contains four long-term states: WELL, OK, DISABLED, DEAD. The state OK depicts a patient surviving aneurysm rupture without permanent disability. There are two temporary states, ok* and disabled*, which represent morbidity of aneurysm rupture in the period immediately following hemorrhage. It is assumed that further bleeding episodes will not occur following a rupture.

Intuitively, the incidence rate for aneurysm rupture should decrease over time because each passing day without rupture increases the likelihood that the patient was in fact never at risk for a rupture. This intuitive viewpoint can be explicitly modeled. Figure 10 depicts a stochastic tree in which the patient may, with probability p_H, be at high risk for aneurysm rupture, or may, with probability $p_L = 1 - p_H$, be at low risk for rupture. High-risk patients have a high rate α_H of rupture, and low-risk patients have a low rate α_L. Although the tree may suggest otherwise, it is intended that the patient's risk status (high or low) is never actually revealed. Further details in the tree are discussed below.

The survival function for such patients is given by

$$S(t) = P(\text{duration without rupture exceeds } t)$$
$$= p_H e^{-\alpha_H t} + p_L e^{-\alpha_L t} .$$

from which the time-dependent incidence rate $\alpha(t)$ may be calculated (e.g., see Kalbfleish and Prentice[30]):

$$\alpha(t) = -\frac{S'(t)}{S(t)} = \frac{p_H \alpha_H e^{-\alpha_H t} + p_L \alpha_L e^{-\alpha_L t}}{p_H e^{-\alpha_H t} + p_L e^{-\alpha_L t}}$$

This function is graphed in figure 11 when $\alpha_H = 1.6$/week, $\alpha_L = 0.5$/week, $p_H = 0.9091$. It declines from an initial value of

$$\alpha_0 = \alpha(0) = p_H \alpha_H + p_L \alpha_L$$

to a limiting value of

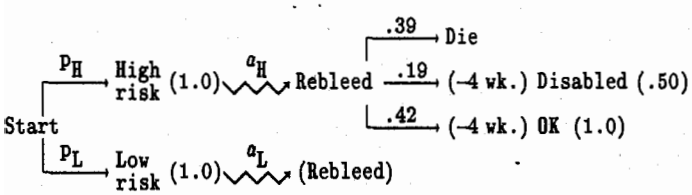

FIGURE 10. Stochastic tree depicting declining rebleed rate.

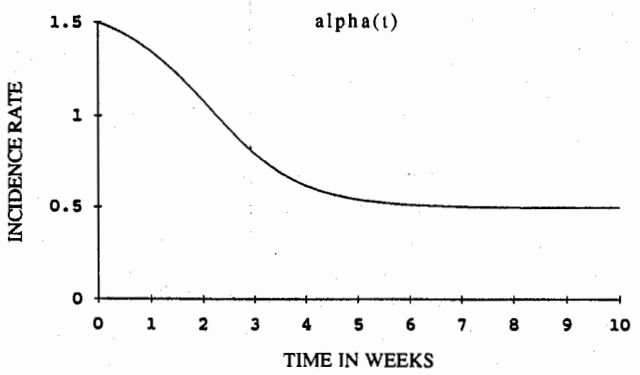

FIGURE 11. Declining incidence rate for a dichotomous population of two constant risks.

$$\alpha_\infty = \lim_{t \to \infty} \alpha(t) = \alpha_L$$

Its derivative at zero is

$$\alpha_0' = \alpha'(0) = (p_H \alpha_H + p_L \alpha_L)^2 - (p_H \alpha_H^2 + p_L \alpha_L^2)$$

The three parameters α_H, α_L, p_H may be adjusted so that $\alpha(t)$ approximates a desired declining incidence rate. For this purpose, it helps if $\alpha(t)$ is expressed equivalently as a function of the parameters α_0, α_∞, α_0' that are directly meaningful with respect to the graph of $\alpha(t)$. This may be accomplished by solving the three previous equations for α_H, α_L, p_H in terms of α_0, α_∞, α_0'. The solution is

$$\alpha_H = \alpha_0 - \frac{\alpha_0'}{\alpha_0 - \alpha_\infty}$$

$$\alpha_L = \alpha_\infty$$

$$p_H = \frac{(\alpha_0 - \alpha_\infty)^2}{(\alpha_0 - \alpha_\infty)^2 - \alpha_0'}$$

Fleming et al.[15] state that the risk of recurrent bleeding declines eventually to a baseline rate of 2–3% per year, or 0.044% (their exact figure) per week. We therefore set $\alpha_\infty = 0.00044$/week, and by manipulating α_0 and α_0', were able to produce weekly probabilities of recurrent bleeding roughly equal to those used by Fleming et al., when $\alpha_0 = 0.125$ and $\alpha_0' = -0.03$. For the first five weeks, these probabilities were:

Week	Probability of rebleed
1	10.4%
2	7.3%
3	5.1%
4	3.5%
5	2.4%

(Notice that the probability of rebleed in week 1 is 1 − S(1); in week 2, S(1) − S(2); in week 3, S(2) − S(3),

and so on.) The corresponding values of α_H and p_H were 0.3658/week and 0.3409, respectively (and of course $\alpha_L = \alpha_\infty = 0.00044$/week).

Using these values, we may roll back the stochastic tree of figure 10. Notice that the tree contains *tolls*, one-time quality-adjusted-duration penalties of 4 weeks associated with surviving a rebleed. During rollback, these tolls are subtracted from the mean quality-adjusted durations that immediately follow in the tree, as indicated below. It is interesting that the states disabled* and OK*, which are necessary in the Markov-cycle tree to account for this one-time penalty, are unnecessary in the stochastic tree. As mentioned previously, rollback at an instantaneous state consists of averaging the quality-adjusted durations following each branch, just as is done with chance nodes in a decision tree.

This stochastic tree does not include background mortality. We will conduct the rollback using implicit Erlang mortality, as described in the previous section. Because the patient in this case was a 34-year-old white woman, we use a seven-stage Erlang model with $\gamma = 0.1395$/year $= 2.674 \cdot 10^{-3}$/week (see table 4). Because of the large number of Erlang stages, pencil–paper–calculator rollback becomes tedious. Calculations were done in command mode of the language APL, which permits arithmetic operations on the $\mathbf{L}(y)$ arrays that arise during implicit Erlang rollback.

The numeric rollback process went as follows. Beginning with the state Disabled, the Erlang rollback formula (equation 7) was used in seven successive stages to obtain

$$\mathbf{L}(\text{Disabled}) = \begin{bmatrix} L_1(\text{Disabled}) \\ L_2(\text{Disabled}) \\ L_3(\text{Disabled}) \\ L_4(\text{Disabled}) \\ L_5(\text{Disabled}) \\ L_6(\text{Disabled}) \\ L_7(\text{Disabled}) \end{bmatrix} = \begin{bmatrix} 1309.14 \\ 1122.12 \\ 935.10 \\ 748.08 \\ 561.06 \\ 374.04 \\ 187.02 \end{bmatrix}$$

Because the state OK has twice the quality factor of the state Disabled, $\mathbf{L}(\text{OK})$ is simply twice $\mathbf{L}(\text{Disabled})$. Here and subsequently, all durations are in weeks. Next, the four-week tolls are subtracted from $\mathbf{L}(\text{Disabled})$ and $\mathbf{L}(\text{OK})$, and the results are averaged to obtain $\mathbf{L}(\text{Rebleed})$:

$$\mathbf{L}(\text{Rebleed}) = (0.39) \cdot 0 + (0.19)[-4 + \mathbf{L}(\text{Disabled})] + (0.42)[-4 + \mathbf{L}(\text{OK})] = \begin{bmatrix} 1345.97 \\ 1153.34 \\ 960.71 \\ 768.08 \\ 575.45 \\ 382.82 \\ 190.19 \end{bmatrix}$$

The Erlang rollback formula is now successively applied to give

$$\mathbf{L}\begin{pmatrix}\text{High}\\ \text{risk}\end{pmatrix} = \begin{bmatrix} 1347.30 \\ 1154.67 \\ 962.04 \\ 769.41 \\ 576.78 \\ 384.15 \\ 191.52 \end{bmatrix} \quad \mathbf{L}\begin{pmatrix}\text{Low}\\ \text{risk}\end{pmatrix} = \begin{bmatrix} 2069.68 \\ 1814.74 \\ 1549.55 \\ 1272.42 \\ 981.39 \\ 674.15 \\ 348.06 \end{bmatrix}$$

Finally, these values are averaged to give

$$\mathbf{L}(\text{Start}) = p_H\mathbf{L}\begin{pmatrix}\text{High}\\ \text{risk}\end{pmatrix} + p_L\mathbf{L}\begin{pmatrix}\text{Low}\\ \text{risk}\end{pmatrix} = \begin{bmatrix} 1823.43 \\ 1589.74 \\ 1349.28 \\ 1100.96 \\ 843.46 \\ 575.30 \\ 294.70 \end{bmatrix}$$

Therefore, the mean quality-adjusted duration beginning at Start is 1823.43 weeks = 34.95 years.

This declining-incidence model divides patients into high-risk and low-risk groups, and has three free parameters α_H, α_L, p_H. Should greater flexibility be desired in matching declining rate data, one could divide patients into three categories (high-, medium-, and low-risk), thereby obtaining five free parameters. In general, n risk categories yield $2n - 1$ free parameters that can be varied to match the desired declining-risk pattern.

Discussion

Why should an analyst use a stochastic tree rather than a Markov-cycle tree for the modeling of medical treatment decisions? There are at least three dimensions along which the two tools differ:

Quality of graphic portrayal

Nature of solution algorithm

Breadth of representable problem scenarios

We discuss each of these in turn.

1. Quality of graphic portrayal. By this point it should be evident that the stochastic tree is superior in this regard. To summarize what has already been said, Markov-cycle trees suffer graphically because they must squeeze all possible event combinations into a cycle, and must portray all possible one-cycle transitions from each state, both immediate and non-immediate. In contrast, stochastic trees have no fixed-length cycle restriction and depict only immediate transitions. Because of this, the simple structure present in many medical decision models is clearly portrayed in a sto-

chastic tree, whereas it is often obscured by a Markov-cycle tree diagram. Moreover, the Erlang mortality technique introduced above allows one to omit the graphic portrayal of background mortality, further simplifying the stochastic tree. One-time penalties (tolls) may also be handled in stochastic trees without introduction of the additional states that may be required in a Markov-cycle tree.

2. Nature of the solution algorithm. Both Markov-cycle trees and stochastic trees have natural rollback algorithms suggested by their structures, although this has apparently not been previously noted for the cycle tree. Stochastic tree rollback is, however, faster by several orders of magnitude, because rollback time is proportional to the number of states, whereas cycle tree rollback time is proportional to the number of states and the number of time periods. This speed difference is unlikely to be important except for very large trees. However, cycle tree rollback converges but does not terminate, whereas for acyclic stochastic trees, rollback terminates finitely. In simple acyclic cases (which are not infrequent), stochastic tree rollback may be done with only pencil, paper, and calculator, thereby avoiding computer programming. In addition, quantitative insight may be obtained in a non-computational fashion by symbolic manipulation of the stochastic tree, as illustrated above. Finally, the nuisance translation of rates to probabilities is avoided in stochastic trees, where the rollback uses rate data directly.

3. Breadth of representable problem scenarios. Both modeling approaches are fairly flexible in the types of scenarios that can be treated. The restriction that rates in stochastic trees must be constant rather than age-dependent poses no difficulty for mortality rate modeling if one is willing to use the Erlang mortality model introduced above. Moreover, declining rates can be treated in stochastic trees by postulating categories of constant-high-risk and constant-low-risk patients, as discussed in the previous section. Stochastic trees do, however, have one drawback that has not yet been mentioned: There is no simple way of dealing computationally with known fixed durations. If one is satisfied with treating a fixed duration as uncertain with some small variance, then the Erlang model with n and γ both large may suffice. However, it may be simpler to revert to the cycle tree model, where fixed durations that are multiples of the cycle time are easily handled.

Medical scenarios do arise that are complex enough to tax the representational powers of both models. Methods for displaying large decision trees can be adapted to Markov-cycle trees (e.g., Tsevat et al.[17]). A new technique called *stochastic factoring* for representing large stochastic trees will be introduced in a subsequent report. Utility functions that simultaneously allow both stochastic tree rollback and non-neutral attitudes towards risk will be reported later.

Finally, one comment on the graphic notation we have employed. Although we believe stochastic trees deserve to be widely used for medical decision modeling, they are still new tools, and there are no conventions regarding their graphic portrayal. The only precedent is the "transition-rate diagram" which is standard for continuous-time Markov chains. We have tried to use notation that is simple and suggestive. However, the reader should note that most of this notation is employed for the first time here.

The author thanks Drs. Mark Eckman and Michael Hagen for their helpful comments on an earlier draft of this paper.

References

1. Pauker SG. Coronary artery surgery: the use of decision analysis. Ann Intern Med. 1976;85:8–18.
2. Beck JR, Pauker SG. Anticoagulation and atrial fibrillation in the brachycardia–tachycardia syndrome. Med Decis Making. 1981;1:285–301.
3. Beck JR, Pauker SG. The Markov process in medical prognosis. Med Decis Making. 1983;3:419–58.
4. Ransohoff DF, Gracie WA, Wolfenson LB, Neuhauser D. Prophylactic cholecystectomy or expectant management for silent gallstones. Ann Intern Med. 1983;99:199–204.
5. Sonnenberg A. Endoscopic screening for gastric stump cancer—would it be beneficial? A hypothetical cohort study. Gastroenterology. 1983;87:489–95.
6. Sonnenberg A. Comparison of different strategies for treatment of duodenal ulcer. Br Med J. 1985;290:1185–7.
7. Eckman MH, Robert NJ, Parkinson DR, Latzke G., Pauker SG. Eaton-Lambert syndrome and small-cell lung cancer: side effects and certainty. Med Decis Making. 1986;6:174–86.
8. Hillner BE, Hollenberg JP, Pauker SG. Postmenopausal estrogens in prevention of osteoporosis. Am J Med. 1986;80:1115–7.
9. Matchar DB, Pauker SG. Transient ischemic attacks in a man with coronary artery disease: two strategies neck and neck. Med Decis Making. 1986;6:239–49.
10. Simon DG. A cost–effectiveness analysis of cyclosporine in cadaveric kidney transplanation. Med Decis Making. 1986;6:199–207
11. Eckman MH, McNutt RA, Parkinson DR, Pauker SG. The timing of radical cystectomy after recent myocardial infarction: waiting for Godot. Med Decis Making. 1987;7:52–66.
12. Plante DA, Piccirillo JF, Sofferman RA. Decision analysis of treatment options in pyriform sinus carcinoma. Med Decis Making. 1987;7:74–83.
13. Wong JB, Webb RK, Pauker, SG. Double trouble: a patient with two prosthetic valves and two episodes of intracranial bleeding. Med Decis Making. 1987;7:174–93.
14. Ellwein LB, Farrow GM. Urinary cytology screening: the decision facing the asymptomatic patient. Med Decis Making. 1988;8:110–9.
15. Fleming C, Wong JB, Moscowitz AJ, Pauker SG. A peripartum neurologic event: shooting from the hip. Med Decis Making. 1988;8:55–71.
16. Roach PJ, Fleming C, Hagen MD, Pauker SG. Prostatic cancer in a patient with asymptomatic HIV infection: are some lives more equal than others? Med Decis Making. 1988;8:132–44.
17. Tsevat J, Eckman MH, McNutt RA, Pauker SG. Warfarin for dilated cardiomyopathy: a bloody tough pill to swallow? Med Decis Making. 1989;9:162–9.
18. Hagen MD, Eckman MH, Pauker, SG. Aortic aneurysm in a 74-year-old man with coronary disease and obstructive lung disease. Med Decis Making. 1989;9:285–99.

19. Mooney C, Mushlin AI, Phelps CE. Targeting assessments of magnetic resonance imaging in suspected multiple sclerosis. Med Decis Making. 1990;10:77–94.
20. Feller W. An introduction to probability theory and its applications. Third ed. New York: John Wiley and Sons, 1968.
21. Cinlar E. Introduction to stochastic processes. Englewood Cliffs, NJ: Prentice-Hall, 1975.
22. Ross SM. Introduction to probability models. Fourth ed. San Diego, California: Academic Press, 1989.
23. Hollenberg JP. Markov cycle trees: a new representation for complex Markov processes (abstr). Medical Decision Making. 1984;4:529.
24. Beck JR, Kassirer JP, Pauker SG. A convenient approximation of life expectancy (the "DEALE"): I. Validation of the method. Am J Med. 1982;73:883–8.
25. Schousboe JT, Hazen GB, Manheim LM, Chang RW. The economic implications of the prevention of gouty attacks with allopurinal. Eighth annual meeting of the Society for Medical Decision Making, Chicago, October 1986.
26. Schousboe JT, Hazen GB, Manheim LM, Chang RW. The economic implications of the prevention of gouty attacks with probenecid or allopurinal. American Rheumatism Association regional meeting, Chicago, October 1986.
27. Ross SM. Applied probability models with applications. San Francisco: Holden-Day, 1970.
28. Denardo EV. Dynamic programming: models and applications. Englewood Cliffs, NJ: Prentice-Hall, 1982.
29. U.S. National Center for Health Statistics. Vital statistics for the United States, 1986.
30. Kalbfleisch JD, Prentice RL. The statistical analysis of failure time data. New York: John Wiley and Sons, 1980.