

Preference Summaries for Stochastic Tree Rollback

Gordon B. Hazen

Northwestern University

and

Jayavel Sounderpandian

University of Wisconsin-Parkside

Key words: Medical decision making, QALY, Stochastic tree, Rollback, Preference summary

Abstract: A stochastic tree is a convenient structure to represent the future health process of a patient, and it can be used to make complex medical decisions or to carry out cost effectiveness analyses of expensive medical treatments. Several types of von Neumann utility functions defined on stochastic trees are recursive in that they allow rollback of the stochastic tree, as in the case of decision trees. Most recursive utility functions also admit preference summaries that can be used to decompose a stochastic tree into preference elements and probabilistic elements. Through an example, we describe this decomposition and subsequent utility computations.

1. Introduction

The projection of the health status of a patient into the future can be modeled as a stochastic process so that medical decisions can be made based on the patient's preferences among possible realizations of that process. A *stochastic tree*, introduced by Hazen (1992), is a convenient structure to represent such a stochastic health process. Essentially, a stochastic tree is a continuous-time Markov chain with chance nodes added. Figure 1, taken from Hazen and Pellissier (1996), depicts a stochastic tree based on the analysis of transient ischemic attack conducted by Matchar and Pauker (1986). Wavy arrows out of a node denote competing exponentially distributed transition times with rate given on the arrow. Nodes from which wavy arrows emanate are called *stochastic*. Straight arrows emanate from *chance* nodes, are labeled with probabilities. Because they contain chance nodes and, if necessary, decision nodes, stochastic trees are generalizations of decision trees.

Suppose a patient's preferences over possible realizations of future health process have been assessed using the techniques discussed in Hazen, Hopp and Pellissier (1991). If these preferences are assumed or verified to satisfy von Neumann axioms, then a von Neumann utility function could be defined and used to select the best medical decision (E.g., Hazen and Sounderpandian 1996, Hazen, Pellissier and Sounderpandian 1997).

Using brute force, the expected utility at the start node of a stochastic tree can be computed for any type of von Neumann utility. But the familiar rollback method, as in the case of decision trees, will be considerably simpler and faster. The type of utility that permits rollback is termed *recursive*. Recursive utility functions for stochastic trees are analyzed in Hazen and Pellissier (1996).

A feature used in computing recursive utility on stochastic trees is the *preference summary* at a node, which summarizes just the essentials of past history up to that node. The notion of preference summary was introduced by Meyer who called it *state descriptor*. For efficient computation, a preference summary must be easy to update so that the summaries at all the nodes

of a tree can be computed by a forward pass over the tree. Also, the information content of the preference summary at any node should be sufficient to compute the expected utility of the subtree following that node. Then, through a backward pass, or rollback, the expected utilities of subtrees can be properly aggregated to yield the expected utility at the start node.

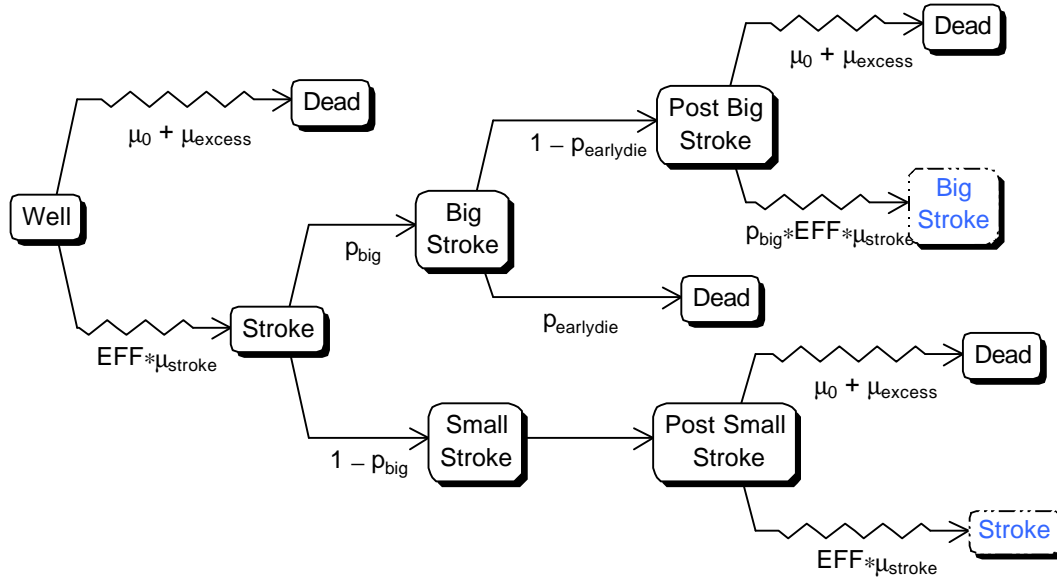


Figure 1. A stochastic tree depicting treatment results for transient ischemic attack (Matchar and Pauker 1986).

Preference summaries can be useful not only to model relatively sophisticated preferences but also to “factor out” preference-based elements from a stochastic tree, as we shall illustrate below. By separating preference-based elements from non-preference-based elements, the model can be formulated and presented in a simpler fashion. In this paper, we shall review some general results about preference summaries and corresponding update functions. We will then discuss in detail a type of preference summary, which we call the *state trajectory*, and demonstrate its use through an example.

2. Preference summaries and statewise exponential utility

Define a *history* as a finite sequence $h = x_1^{t_1} x_2^{t_2} \dots x_n^{t_n}$ of health states x_1, x_2, \dots, x_n and durations t_1, t_2, \dots, t_n , where it is understood that state x_1 is first occupied for duration t_1 , followed by state x_2 for duration t_2 , and so on. Hazen and Pellissier (1996) recursively define a *statewise exponential utility function* $u(h)$ over histories $h = x^t k$ which represents an initial sojourn t in state x followed by history k as

$$u(h) = \int_0^t v(x) e^{-a(x)s} ds + e^{-a(x)t} u(k) \quad (1)$$

where $u(k) = 0$ whenever k is null (i.e., when h is simply x^t). The quantity $v(x)$ is a *quality rate* which assigns quality $v(x)ds$ to each small interval ds spent in state x . The quantity $a(x)$ may be

thought of as a state-dependent discount rate, or alternatively, as a coefficient of risk aversion for an exponential utility function over duration in x . Using the statewise exponential utility function, the overall utility of a stochastic tree can be evaluated by the following simple recursive rollback procedure (Hazen and Pellissier 1996). The expected utility of a stochastic subtree H




Figure 2. A Stochastic Subtree H

which begins at state x and makes transitions to one of the succeeding subtrees K_1, \dots, K_n satisfies the following recursive equation:

$$(2)$$

provided (otherwise $E[u(H)]$ equals $+\infty$). Equation (2) can be used recursively

to calculate expected utility at any stochastic node in the tree from utilities at subsequent nodes. (At chance nodes, the usual averaging formula applies.) In the risk-neutral case where $a(x) = 0$, expected utility is equivalent to mean QALY with quality rate $v(x)$. If $a(x)$ is positive, then the decision maker is risk averse for uncertain lifetimes spent in state x , and if $a(x)$ is negative, the decision maker is risk seeking. Note that risk attitude may depend on the health state x

Given a stochastic tree, we use the notation $\langle x \rangle$ to represent a path in the tree which successively visits states x_1, x_2, \dots, x_n with intervening rates or probabilities $\lambda_1, \lambda_2, \dots, \lambda_n$. A preference summary is a function $e(\langle x \rangle)$ which assigns to each path a summary value $q = e(\langle x \rangle)$. The preference summary $e(\langle x \rangle)$ is *updatable* if there is a function q which assigns a new summary value $q \oplus = q(q, x^1)$ to summary values q and single-step paths x^1 , such that

$$e(\langle x \rangle) = q(e(\langle x \rangle), x^1).$$

The results we present below yield a generalization of the statewise exponential utility function (1). If the initial state x of stochastic subtree H has preference summary state q , then the rollback equation (2) for stochastic trees becomes

$$(3)$$

In case the start node of H is a chance node where branch i leads to K_i with probability p_i , then the recursion formula would be

These recursion formulas can be easily implemented on a spreadsheet.

3. The state-trajectory preference summary

We first reproduce a theorem from Hazen and Pellissier (1996) which proves the existence of recursive utility functions and corresponding preference summaries. The nontriviality mentioned in the theorem means that for every history h , there exists a history g such that either gh is preferred to h or h is preferred to gh . The symbol \emptyset denotes the empty history. The utility function U is conditioned on histories and u is conditioned on preference summaries.

Theorem 1. Assume nontriviality holds and $U(\emptyset | h) = 0$ for all h . If U has an updatable preference summary $\langle e, q \rangle$, then there exists a summary utility function $u(h | q)$ such that $u(\emptyset | q) = 0$ and $u(\cdot | q) \sim U(\cdot | h)$ whenever $e(h) = q$. Moreover, there are summary discount factors $\Delta u(y^\mu | q) > 0$ satisfying the affine restriction

$$\dots\dots\dots (5)$$

the concatenation restriction

$$\Delta u(y^{\mu^*v} | q) = \Delta u(y^\mu | q) \Delta u(y^v | q(q, y^\mu)) \dots\dots\dots (6)$$

and the recursive equation

$$u(y^\mu h | q) = u(y^\mu | q) + \Delta u(y^\mu | q) u(h | q(q, y^\mu)) \dots\dots\dots (7)$$

for all h .

The implications of the above theorem are explained in Hazen and Pellissier (1996) by way of some special types of preference summaries, namely, memoryless, Markovian and Semi-Markovian. In this paper, we shall concentrate on another special type of preference summary, which we call *state trajectory*. A state trajectory simply remembers the sequence of health states visited so far. For this summary, the function e is defined over histories of sequentially distinct states $x_1 \ \# \ x_2 \ \# \ \dots \ \# \ x_n$ by

$$= (x_1, x_2, \dots, x_n).$$

To apply the function e to histories of sequentially *nondistinct* states, concatenate identical adjacent states and then apply the definition above. For any state trajectory $q = (x_1, x_2, \dots, x_n)$ and state x , let $qx = (x_1, x_2, \dots, x_n, x)$ denote the concatenation of q and x . Also define

$$q \wedge y =$$

The state trajectory update function q can then be defined as

$$q(q, x^\lambda) = q \wedge y$$

The following theorem establishes the type of utility function that has state trajectories as preference summaries.

Theorem 2. Assume U is continuous in the time argument, nontrivial, and $U(\emptyset | h) = 0$ for all h . If U admits a state trajectory preference summary then there exist real valued functions $w(y | q)$, $\Delta w(y | q) > 0$, $v(q)$ and $a(q)$ such that the summary utility function u for U satisfies

$$u(y^t | q) = w(y | q) + \Delta w(y | q) \int_0^t v(q \wedge y) e^{-a(q \wedge y)s} ds \quad (8)$$

$$\Delta u(y^t | q) = \Delta w(y | q) e^{-a(q \wedge y)t} \quad (9)$$

where if the last state in q is y , then $w(y | q) = 0$ and $\Delta w(y | q) = 1$.

Moreover, for a distribution μ on $[0, \infty)$,

$$\Delta u(y^\mu | q) = \Delta w(y | q) \int e^{-a(q \wedge y)s} d\mu(s). \quad (10)$$

The function $w(y | q)$ can be thought of as a bonus gained, or a toll paid when entering state y with state trajectory q . The function $\Delta w(y | q)$ is a discount factor. The proof of Theorem 2 may be found in the appendix.

3.1 Coarsening State-Trajectory Summaries

Whereas we have presented Theorem 2 with state trajectories as preference summaries, a summary could be more coarse than (that is, a deterministic function of) a complete trajectory, and the theorem will still apply to it as long as it is updatable. Because *any* preference summary that involves only previously visited states (but not previous sojourn rates) is coarser than the state trajectory summary, it follows that Theorem 2 applies to it as well. For such preference summaries (e, q) , the update function $q(q, y^1)$ cannot depend on l , so we write $q(q, y^1) = q(q, y)$. We have the following result.

Corollary: Given an updatable preference summary (e, q) that involves no preceding sojourn rates but only previously visited states, Theorem 2 applies with the state-trajectory update function $q \wedge y$ replaced by $q(q, y)$

We make a note of this because preferences among most medical conditions seen in practice depend only on relatively simple summaries, such as whether or how many times a state has been previously visited.. For example, if a patient has encountered a crippling medical condition, his or her preferences regarding future states could be different than if the condition had not been encountered. The summary in this case will just be yes or no depending upon whether or not the condition has been encountered. This summary is updatable and Theorem 2 will apply.

In some cases, we need to know how many times a particular state has been encountered. For example, following total hip arthroplasty (THA), revision surgeries may be required due to loosening or infection (see Chang, Pellissier and Hazen 1996). The patient's state of health can be modeled as a function of the number of loosening and the number of infections, so it is natural to use these two numbers as a preference summary. Again, this summary is updatable, and Theorem 2 will apply.

In the case we will consider in the next section, the summary remembers the worst condition encountered so far. Because it is updatable, we can again use the results from Theorem 2.

4. Application to modeling transient ischemic attack

In this section, we consider the nonsurgical treatment of transient ischemic attack described in Matchar and Pauker (1986). The stochastic tree model and computations for this model can be found in Hazen and Pellissier (1996). We will apply preference summaries to this case to show how the number of states and therefore the computational effort can be significantly reduced.

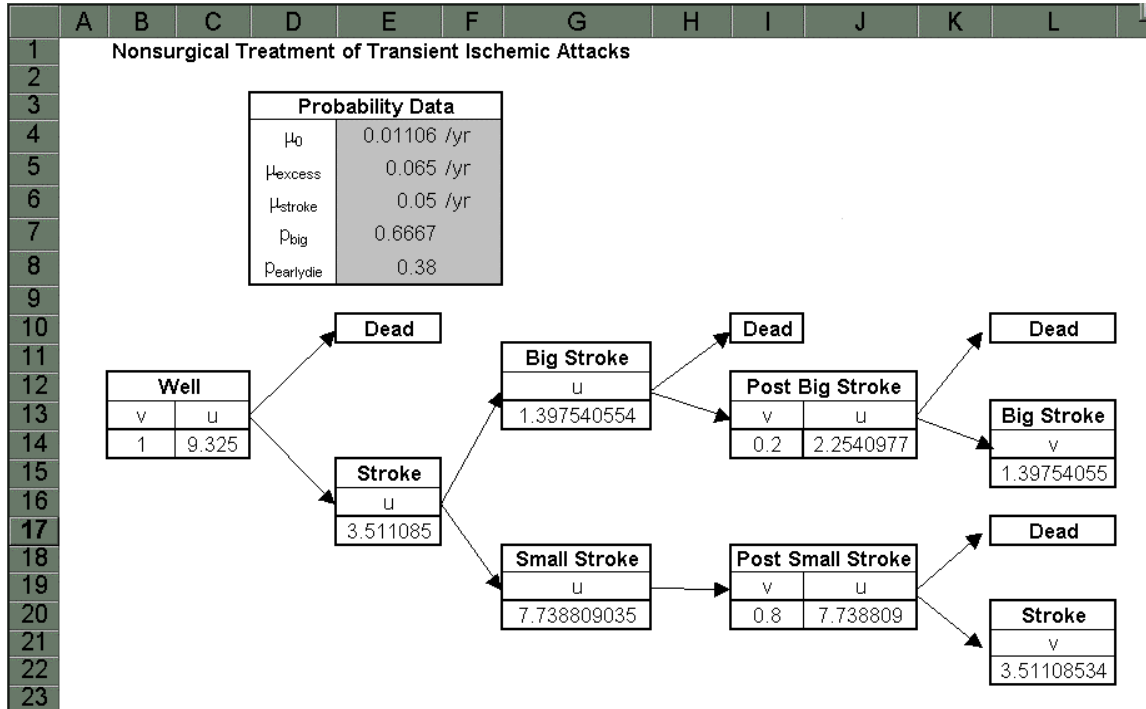


Figure 3. A Spreadsheet Computation

The original stochastic tree model taken from Hazen and Pellissier is reproduced in Figure 1. Note that there are seven states, namely, **Well**, **Stroke**, **Small Stroke**, **Big Stroke**, **Post Small Stroke**, **Post Big Stroke** and **Dead**. Utility rollback for this model is accomplished via equation (2) used iteratively for all the seven states. An implementation of this computation on a Microsoft Excel spreadsheet is shown in Figure 3. For simplicity, we have assumed that all coefficients of risk aversion $a(x)$ are zero.

The rollback formulas are circular and therefore require iterative calculation. Using Tools|Options|Calculation command, the spreadsheet must be set to manual calculation with, say, maximum iterations 100 and maximum change 0.000001. These iterative calculations are called *value iteration* in the dynamic programming literature. Convergence guarantees can be found, for example, in Denardo (1982).

Next we shall explore the possibility of simplification using preference summaries. An examination of the model reveals that the mortality and stroke *rates* remains the same regardless of whether or how many strokes have previously occurred. Therefore, the only purpose of the states **Post Big Stroke** and **Post Small Stroke** are to capture the effects of disability due to stroke. At

any point in the tree, if we know *the worst stroke experience so far*, then the quality factor $v(x)$ is determined. This suggests using a preference summary equal to the worst stroke experience so far, with states **None**, **Small**, **Big**. This preference summary is updatable, and Theorem 2 applies.

The update function $q(\cdot)$ is:

		Summary state q		
		Big	Small	None
State x	$q(q,x)$			
	Well	Big	Small	None
	Stroke	Big	Small	None
	Big Stroke	Big	Big	Big
	Small Stroke	Big	Small	Small
Dead	Big	Small	None	

The states **Post Big Stroke** and **Post Small Stroke** are no longer needed in the stochastic tree, which simplifies to the tree depicted in Figure 4. Also included in Figure 4 is an influence diagram depicting the calculation of the preference summary $q =$ worst stroke so far and the quality factor $v(q)$. The use of this preference summary has allowed us to factor the preference elements out of the model.




Figure 4. Stochastic tree with preference elements factored out

Figure 5. Spreadsheet Computation with Preference Summary

Theorem 2 in conjunction with the fact that the transitions are exponential enable us to use equation (3) recursively and iteratively. In Figure 5 we present a spreadsheet for the rollback of the stochastic tree in Figure 4. The values of $v(q)$ for all summary states q , as well as all probabilities and Poisson frequencies of all transitions have been entered as data into the spreadsheet.

The results agree with the previous results. Although the number of states in the new tree is less, the computational effort is the same because multiple utility values are calculated at each node.

5. Conclusion

We have presented a new recursive utility decomposition based on the state-trajectory preference summary, and have illustrated its use in factoring preference elements out of stochastic trees. This factoring idea is similar in spirit to our earlier work (Hazen 1992,1993) in factoring stochastic trees into Cartesian products of smaller stochastic trees. As with this earlier work, the advantages of factoring preference information out of a stochastic tree are in formulation and presentation: A modeler need only capture probabilistic issues in formulating the stochastic tree, and can treat preference issues separately. This usually reduces the size of the stochastic tree. Moreover, a portrayal of the preference summary using an influence diagram, as illustrated above, makes preference issues graphically explicit where they before were not. This simplifies the presentation of the model and conveys greater insight into its structure.

6. Appendix: Proof of Theorem 2

In order to prove the theorem, we shall use the following Lemma proved in Hazen and Pellissier (1996).

Lemma: The only continuous functions $g(t)$ and $f(t)$ satisfying

$$g(s + t) = g(s) + e^{-as} f(t)$$

for real $s, t > 0$ and a are

$$f(t) = k$$

$$g(t) = c + k$$

where c and k are real constants.

Returning to the proof of Theorem 2, we consider the concatenation restriction implied by Theorem 1, which in the present case becomes

$$\Delta u(y^{s+t} / q) = \Delta u(y^s / q) \Delta u(y^t / q \wedge y) \quad (11)$$

In case the last state in q is y we obtain

$$\Delta u(y^{s+t} / q) = \Delta u(y^s / q) \Delta u(y^t / q)$$

from which we conclude (e.g., from Section 1 of Aczel (1987)),

$$\Delta u(y^t / q) = e^{-a(q)t} \quad \text{if } q_{final} = y \quad (12)$$

for some discount factor $a(q)$. Taking natural logarithms in (11) and applying the Lemma with $a = 0$, we get

$$\Delta u(y^s / q) = \exp(k_{q,y} s + c_{q,y}) \quad (13a)$$

$$\Delta u(y^t / q \wedge y) = \exp(k_{q,y} t) \quad (13b)$$

Combine (12) and (13b) to conclude $k_{q,y} = -a(q \wedge y)$. Substitute this to get

$$\Delta u(y^s / q) = \exp(-a(q \wedge y) s + c_{q,y}) \quad (14a)$$

$$\Delta u(y^t / q \wedge y) = \exp(-a(q \wedge y) t) \quad (14b)$$

6.1.1.1 Substitute $q \wedge y$ for q in (14a) and equate the result to (14b) to conclude $c_{q,y} = 0$ if $q_{final} = y$. Now defining

$$\Delta w(y / q) = \begin{cases} e^{c_{q,y}} & \text{if } q_{final} \neq y \\ 1 & \text{if } q_{final} = y \end{cases}$$

allows (14) to be reexpressed as

$$\Delta u(y^t / q) = \Delta w(y / q) e^{-a(q \wedge y)t} \quad (15)$$

as desired.

Coming to the function u , we note that the recursive equation in Theorem 1 yields,

$$u(y^{s+t} / q) = u(y^s / q) + \Delta u(y^s / q) u(y^t / q \wedge y) \quad (16)$$

Substituting the result (15) into (16) we get,

$$u(y^{s+t} / q) = u(y^s / q) + \Delta w(y / q) e^{-a(q \wedge y)s} u(y^t / q \wedge y)$$

Invoking the Lemma again, we get,

$$\Delta w(y / q) u(y^t / q \wedge y) = k_{q,y} \int_0^t e^{-a(q \wedge y)s} ds \quad (17a)$$

$$u(y^t / q) = c_{q,y} + k_{q,y} \int_0^t e^{-a(q \wedge y)s} ds \quad (17b)$$

for some new constants $k_{q,y}$ and $c_{q,y}$. Set $q_{final} = y$ to get

$$u(y^t / q) = k_{q,y} \int_0^t e^{-a(q)s} ds \quad (18a)$$

$$u(y^t / q) = c_{q,y} + k_{q,y} \int_0^t e^{-a(q)s} ds \quad (18b)$$

Conclude that $c(q,y) = 0$ when $q_{final} = y$. Define $v(q) = k_{q,q_{final}}$ and substitute into (18) to get $u(y^t$

$/ q) = v(q) \int_0^t e^{-a(q)s} ds$ when $q_{final} = y$.

Substitute this back into (17a) to get

$$Dw(y/q)v(q \wedge y) \int_0^t e^{-a(q \wedge y)s} ds = k_{q,y} \int_0^t e^{-a(q \wedge y)s} ds .$$

Conclude that $k_{q,y} = Dw(y/q)v(q \wedge y)$.

Substitute this into (17b) along with the definition $w(y/q) = c_{q,y}$ to obtain

$$u(y^t / q) = w(y / q) + \Delta w(y / q) \int_0^t v(q \wedge y) e^{-a(q \wedge y)s} ds$$

as desired.

To establish the last claim of the theorem, substitute the state trajectory preference summary into the affine restriction (5) to get

$$Du(y^m/q)u(h/q \wedge y) = \int_0^\infty \Delta u(y^s | q)u(h | q \wedge y) d\mathfrak{m}(s).$$

and cancel the factor $u(h/q \wedge y)$.

7. References

Aczel, J. 1987. *A Short Course on Functional Equations*. Reidel Publishing, Dordrecht, Holland.

Chang R. W., J.M. Pellissier, and G. B. Hazen. 1996. A cost-effective analysis of total hip arthroplasty for osteoarthritis of the hip. *JAMA*, 275 (11), pp. 858-865.

Denardo E. V. 1982. *Dynamic programming: Models and applications*. Prentice Hall, Englewood Cliffs, NJ.

Hazen G. B. 1992. Stochastic Trees: A new technique for temporal medical decision modeling.

- Med. Dec. Making* 12, pp. 163-178.
- Hazen G. B. 1993. Factored stochastic trees: A tool for solving complex temporal medical decision making problems. *Med. Dec. Making* 13, pp. 227-236.
- Hazen G. B., W. H. Hopp and J. M. Pellissier. 1991. Continuous-risk utility assessment in medical decision making. *Med. Dec. Making* 11, pp. 294-304.
- Hazen G. B., and J. M. Pellissier. 1996. Recursive utility for stochastic trees. *Operations Research*, vol. 44, No. 5, pp. 788-809.
- Hazen, G. B., and J. Souderpandian. 1996. "Stochastic Trees and Medical Decision Making." In *Economic and Environmental Risk and Uncertainty: New Models and Methods*, Robert Nau *et al.* (eds.) Kluwer Academic Press, Dordrecht, Netherlands.
- Hazen, G. B., J. Pellissier and J. Souderpandian. 1997. "Stochastic Tree Models in Medical Decision Making." Submitted to *Interfaces*.
- Matchar D. B., and S. G. Pauker. 1986. Transient ischemic attacks in a man with coronary artery disease: Two strategies neck and neck. *Med. Dec. Making* 6, pp. 239-249.
- Pellissier J. M., G. B. Hazen, and R. W. Chang. 1996. A continuous risk decision analysis of total hip replacement. *Journal of the Operational Research Society* 47, pp. 776-793.