Preference Factoring for Stochastic Trees

Gordon Hazen

Department of Industrial Engineering and Management Sciences, Northwestern University, Technological Institute,
Evanston, Illinois 60208-3119
hazen@iems.nwu.edu

S tochastic trees are extensions of decision trees that facilitate the modeling of temporal uncertainties. Their primary application has been to medical treatment decisions. It is often convenient to present stochastic trees in *factored* form, allowing loosely coupled pieces of the model to be formulated and presented separately. In this paper, we show how the notion of factoring can be extended as well to preference components of the stochastic model. We examine *updateable-state* utility, a flexible class of expected utility models that permit stochastic trees to be rolled back much in the manner of decision trees. We show that preference summaries for updateable-state utility can be factored out of the stochastic tree. In addition, we examine utility decompositions which can arise when factors in a stochastic tree are treated as attributes in a multiattribute utility function.
(*Expected Utility; Medical Decision Making; Stochastic Trees; Multiattribute Utility; Time Preference; Quality-Adjusted Lifetime*)

Stochastic trees are graphical modeling tools which extend decision trees by allowing the explicit depiction of temporal uncertainty. They are especially well suited to modeling temporal issues in medical treatment decision analyses (Hazen 1992). My colleagues and I have constructed stochastic tree models for total hip replacement and knee replacement decision analyses (Chang et al. 1996, Pellissier et al. 1996, Gottlob et al. 1996) and for breast cancer decision analysis (Hazen et al. 1999). We have investigated risk-sensitive utility rollback for stochastic trees (Hazen and Pellissier 1996), and have explored preference assessment techniques that are tailored to address risky temporal tradeoffs (Hazen et al. 1991, Pellissier and Hazen 1994). We have found the notion of *factoring* a stochastic tree (Hazen 1993) to be very useful for model formulation and presentation, and have developed software (available at ⟨www.iems.nwu.edu/~hazen⟩) with a graphical interface for formulating and rolling back factored stochastic trees. All stochastic tree illustrations in this article are screen captures from that software.

Formulating a stochastic tree in a factored fashion can greatly simplify the model construction process. For example, it is usually beneficial to factor out background mortality and consider the rest of the model separately (Hazen 1993, Hazen et al. 1998). Further factoring is usually helpful as well, for example, to separate essentially independent processes such as drug side effects and disease progression. We give examples of stochastic tree factoring below.

In this paper, we show how the notion of factoring can be extended as well to preference components of the stochastic model. When features involving patient preference are factored out, the remaining stochastic model is often significantly simpler. We take up preference factoring in §2, where we examine *updateable-state* utility models. These arise from the *state-trajectory* preference summaries recently introduced by Hazen and Sounderpandian (1999), and extend the *Markovian* utility models introduced by Hazen and Pellissier (1996). The major result presented in §2 is that every updateable-state utility model is equivalent

Figure 1   A Stochastic Tree Model

λ = 0.25 /yr
μ = 0.04 /yr
p = 15%
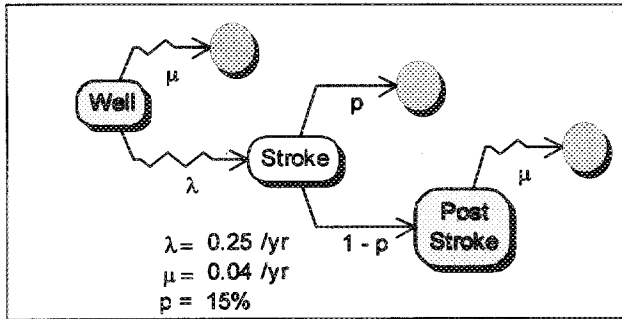
Note: Unlabeled shaded nodes represent death. A patient begins in the Well state with mortality risk due to unrelated causes occurring at average rate μ, and competing risk of stroke occurring at average rate λ. Should death occur first, the process terminates. Should stroke occur first, the patient moves to the state *Stroke*, at which there is a chance p of immediate mortality. Should the patient survive, he moves to the state *Post Stroke*, from which he is subject only to mortality risk at rate μ.

Figure 2   Using Rollback to Calculate Mean Survival Duration in a Stochastic Tree

λ = 0.25 /yr
μ = 0.04 /yr
p = 15%

Note: Numerical values at nodes are mean survival durations in years. For example, mean survival time at Post Stroke is $1/\mu = 25$ years, mean survival time at Stroke is $25 \times (1 - p) = 21.2$ years. Overall mean survival time is 21.8 years using the rollback formula (2).

to a Markovian utility model in a stochastic tree augmented by factors which record preference summary states. The latter are in this sense *factored* from the model.
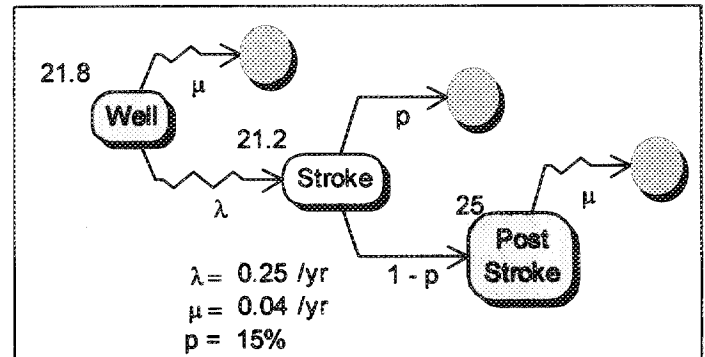
Of course, the factoring process establishes a natural multiattribute structure to any stochastic tree model, and this suggests that preference decomposition issues be investigated. In §3 we examine utility independence and related conditions, and show how such assumptions can be used to decompose updateable-state utility models. We begin in the next section with a review of factoring stochastic trees.

1. Factoring Stochastic Trees

Background
In its simplest form, a stochastic tree is merely a continuous-time Markov chain unfolded into a tree structure, with possible chance or decision nodes added. Straight-line arcs emanate from chance or decision nodes, and arcs from chance nodes have attached probabilities, just as in a decision tree. Wavy arcs emanate from *stochastic* nodes, and denote transitions which take time to occur. Wavy arcs have transition rates attached. Figure 1 displays a simplified stochastic tree model of transient isch-

emic attack adapted from Matchar and Pauker (1986).

Stochastic trees can be rolled back much like decision trees. Figure 2 illustrates how rollback can be used to calculate mean survival times. One may also calculate mean *quality-adjusted* survival times, in which time spent in a state is weighted by its relative worth compared to an equal time in the well state. This is illustrated in Figure 3.

The generic rollback formula is given by Hazen

Figure 3   Using Rollback to Calculate Mean Quality-Adjusted Lifetime

λ = 0.25 /yr
μ = 0.04 /yr
p = 15%
vPostStroke = 0.3
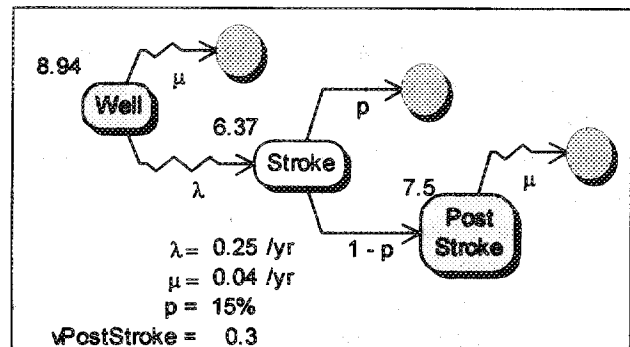
Note: For example, mean quality-adjusted lifetime at Post Stroke is $(1/\mu) \times vPostStroke = 7.5$ years, and mean survival time at Stroke is $7.5 \times (1 - p) = 6.37$ years. Overall mean quality-adjusted lifetime is 8.94 years using the rollback formula (2).

(1992) or Hazen and Pellissier (1996). At an arbitrary stochastic fork

$$H = \quad y \overset{\lambda_1}{\rightsquigarrow} K_1 \quad \overset{\lambda_2}{\rightsquigarrow} K_2 \quad \overset{\lambda_3}{\rightsquigarrow} K_3 \quad (1)$$
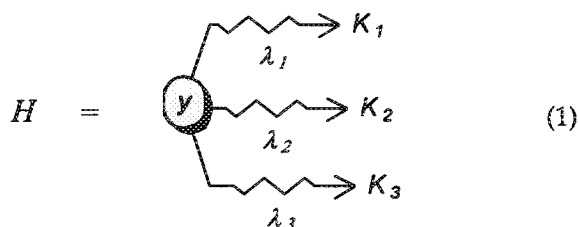
in which subtrees K_i are reached from state y at competing rates λ_i, the rollback equation for mean discounted quality-adjusted lifetime $L(\cdot)$ takes the simple form

$$L(H) = \frac{v(y) + \sum_i \lambda_i L(K_i)}{a(y) + \sum_i \lambda_i}. \quad (2)$$

Here $v(y)$ is the quality weight for state y, and $a(y)$ is the state-specific discount rate at state y. When $a(y) = 0$ for all y, expected utility is mean quality-adjusted duration, and when $a(y) = \alpha$ for all y, expected utility is quality-adjusted duration discounted at rate α. At chance forks, expected utility computation is identical to the usual probability-weighted averaging done in decision trees. For reasons which will become clear, we call (2) the *Markovian* rollback formula.

Factoring

In the model of Figure 1, a background mortality rate μ is present for both the Well and Post Stroke states.

Figure 4   Factoring Background Mortality Out of the Stroke Model of Figure 1
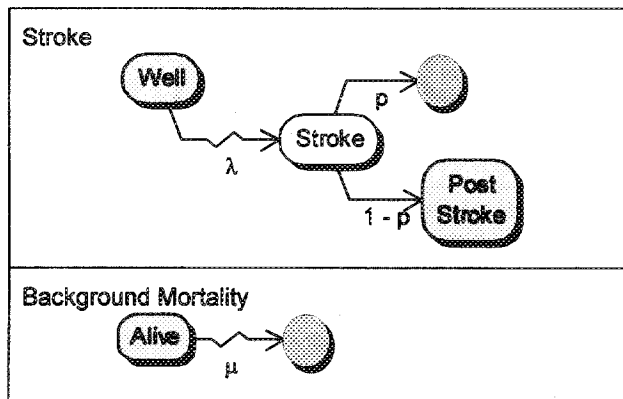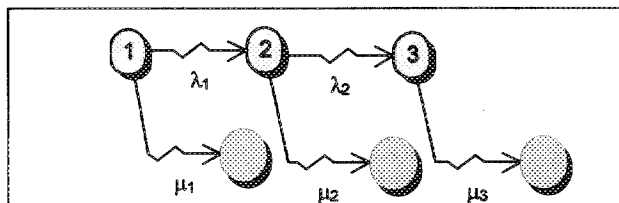

Figure 5   A Coxian Mortality Factor, in Which Mortality Rate depends on the Stage of the Process

Background mortality is really a process operating in parallel with the stroke process. It can therefore be *factored* from the stroke model. The result appears in Figure 4. We say that the transient ischemic attack model of Figure 1 is the (*Cartesian*) *product* of the Stroke and Background Mortality factors in Figure 4. One advantage of this factoring is that it simplifies the model construction process. The modeler may first deal with issues involving stroke without having to worry about background mortality.

Similarly, background mortality can be dealt with independently. For example, the background mortality model in Figure 4 is too simple to be realistic because actual human mortality rates are age dependent, not constant. A more accurate model of human mortality can be constructed by using the Coxian model (Cox 1955, Hazen et al. 1998) depicted in Figure 5. The parameters and the number of stages in a Coxian model can be arbitrarily chosen to well approximate human survival durations. The Coxian factor can be substituted for simple background mortality in Figure 4, resulting in an improved model. The Cartesian product tree need never be explicitly constructed (but see Hazen et al. 1998 for a picture). Hazen (1993) discusses a rollback procedure for factored stochastic trees.

Hazen (1993) constructs a six-factor stochastic tree model of the use of an anticoagulant drug warfarin to treat dilated cardiomyopathy. The model is based on Tsevat et al. (1989). The *Systemic Embolism* factor for this model appears in Figure 6, and the *Anticoagulant Status* factor in Figure 7. Three useful modeling features are illustrated: First, stochastic trees can have *transition cycles*, which allow the possibility of repeated visits to a state. Second, rates or probabilities in one factor can depend on the states of other factors.

391

Figure 6   The Systemic Embolism Factor from a Six-Factor Model of the Use of Warfarin in Dilated Cardiomyopathy
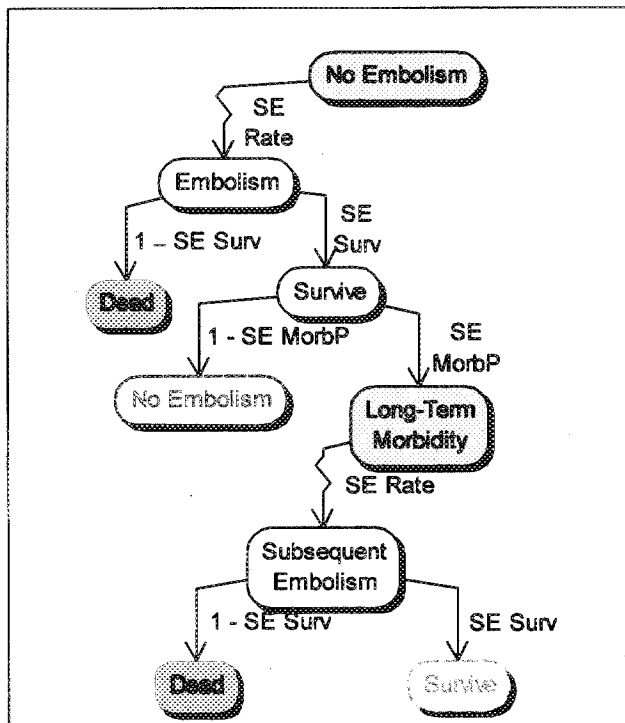
Note: Nodes with dashed-line borders refer to previous nodes in the tree. For example, from the node *Subsequent Embolism* transition is possible (with probability *SE Surv*) to the prior node *Survive*. So repeated visits to *Survive*, *Long-Term Morbidity* and *No Embolism* are possible. Moreover, the rate of systemic embolism (*SE Rate*) depends on the state of the *Anticoagulant Status* factor (Figure 7).

Finally, transitions in one factor can trigger transitions in other factors.

2. Factoring Preference Components from a Stochastic Tree

### Updateable-State Preference Summaries

The Markovian rollback formula (2) for mean quality-adjusted duration presupposes that every state y can be assigned a quality weight $v(y)$ as well as a discount rate $a(y)$. However, sometimes the quality weight at a given state y should depend not only on y but also on the process history prior to entering y. For example, in Figure 6 the process begins in the state *No Embolism*, and may return there after visiting the state *Long-Term*

Morbidity. Because the effects of long-term morbidity are permanent, the quality weight assigned to *No Embolism* should depend on whether *Long-Term Morbidity* has been previously visited. This cannot be accounted for in the Markovian rollback procedure (2).

The dependence of quality weights (and discount rates as well) on the prior history of the process can be incorporated into an efficient rollback procedure by introducing the notion of an *updateable-state preference summary* (Hazen and Sounderpandian 1999, Hazen and Pellissier 1996). Suppose that the quality weight at state y can depend on the entire sequence $\mathbf{y} = (x_0, x_1, \ldots, x_k, y)$ of states visited up to and including y. Suppose the same is true for the discount rate at y. A *state preference summary* is a function $e(\cdot)$ which assigns a *preference summary* $r = e(\mathbf{y})$ to every sequence of states $\mathbf{y}$ so that the quality weight and the discount rate at $\mathbf{y}$ depends only on $r = e(\mathbf{y})$. We say that $e(\cdot)$ *has updateable states* if there is a function $\theta(q, y)$ such that if $\mathbf{y} = (\mathbf{x}, y)$ and $q = e(\mathbf{x})$, then $e(\mathbf{y}) = \theta(q, y)$. In other words, the preference summary at $\mathbf{y} = (\mathbf{x}, y)$ depends only on the last state y and the preference summary at $\mathbf{x}$. When an updateable-state preference summary exists, then rollback operations analogous to (2) may be performed. The notion of preference summary was introduced for discrete-time models by Meyer (1976), who called it *state descriptor*.

Updateable-state preference summaries are plentiful. If quality weights and discount rates depend only on the current state, as we have been assuming, then $e((\mathbf{x}, y)) = y$ is an updateable-state preference sum-

Figure 7   The Anticoagulant Status Factor in the Warfarin Model
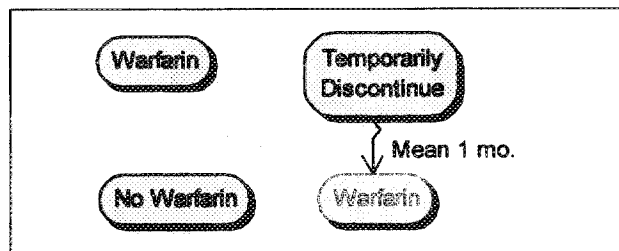
Note: An embolism occurrence in the *Systemic Embolism* factor triggers a transition from *No Warfarin* to *Warfarin* in this factor. A hemorrhage in the *Systemic Hemorrhage* factor (not shown here) triggers a transition here from *Warfarin* to *Temporarily Discontinue*.

mary, which might aptly be called a *Markovian* preference summary (hence the terminology *Markovian* rollback for (2)). The update function is $\theta(q, y) = y$. Other examples of updateable-state summaries are the *count of visits* to a particular state, *whether* (Yes or No) a state has ever been visited, and *which* of a subset of states has been most recently visited. If states are ordered from best to worst in some fashion, then the *best* state visited so far is an updateable-state summary, but the *second-best* state so far is not (if there are more than two states).

Hazen and Sounderpandian (1999) derive the following rollback equation for mean discounted quality-adjusted duration $L(H|q)$ at the generic stochastic fork H in (1) when there is an updateable-state preference summary q prior to H:

$$L(H|q) = \frac{v(r) + \sum_i \lambda_i L(K_i|r)}{a(r) + \sum_i \lambda_i}, \quad r = \theta(q, y). \quad (3)$$

At the analogous chance fork in which branches are labeled by probabilities p_i, the rollback formula is

$$L(H|q) = \sum_i p_i L(K_i|r), \quad r = \theta(q, y). \quad (4)$$

Note that these equations require the computation of mean discounted quality-adjusted duration $L(H|q)$ for every possible preference summary q which could occur prior to the initial node y of H.

Figure 8 shows an example rollback for the *Systemic Embolism* stochastic tree introduced above. The state preference summary is whether the state *Long-Term Morbidity* has yet been visited, and consists of the two possibilities {*Well, LT_Morb*}. There are therefore *two* rollback values at each node in Figure 8, except at *Long-Term Morbidity* and *Subsequent Embolism*, where the only possible value of the preference summary is *LT_Morb*.

Factoring to Recover Markovian Rollback

One of the main results of this paper is that by suitably factoring a stochastic tree, one may always transform the updateable-state rollback (3) to Markovian rollback (2) on the factored tree. We formally derive this result in the next subsection. Here we illustrate how it may be done for the systemic embolism example of Figure 8.

Figure 8 Rollback Using the Updateable-State Preference Summary of Whether the State Long-Term Morbidity Has Yet Been Visited

SE Rate = 0.5 /yr
SE Surv = 75%
SE MorbP = 50%

q	v(q)	a(q)
Well	1	0
LT_Morb	0.32	0
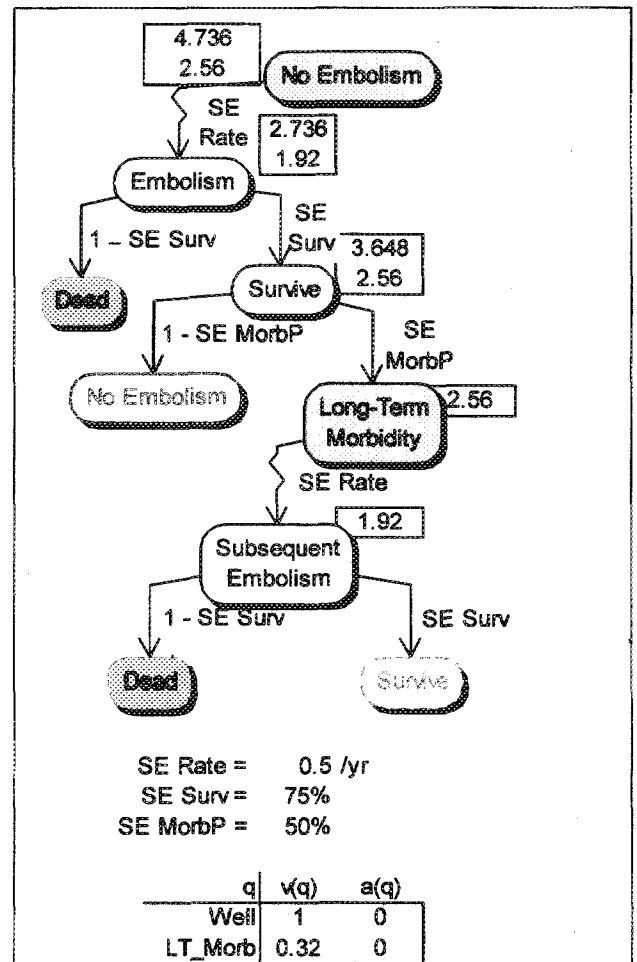
Note. There are two rollback values at every state except *Long-Term Morbidity* and *Subsequent Embolism*, corresponding to the two possible values (*Well* or *LT_Morb*) of the preference summary.

Generally speaking, the idea is to factor all preference summary information from the stochastic tree. In the systemic embolism example, the preference summary is whether the state *Long-Term Morbidity* has yet been visited. Therefore, we create a new factor *Morbidity Status* with possible states *Well* and *Long-Term Morbidity*, and we remove *Long-Term Morbidity* from the original systemic embolism tree. The result is displayed as a two-factor stochastic tree in Figure 9. We have in effect *factored out* the essential preference information. Note that one of the transitions *Survive* →

Figure 9 Factoring Morbidity Status from the Systemic Embolism Model of Figure 8
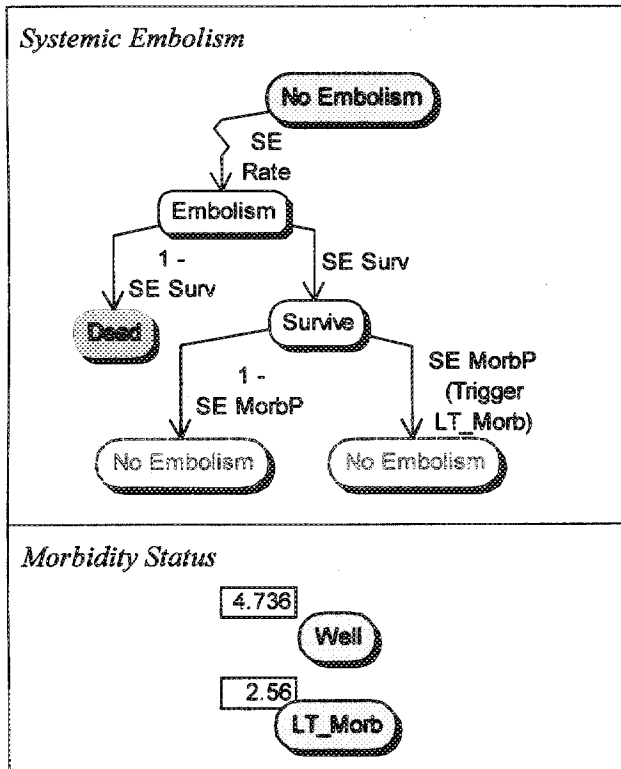
Note: Transition from *Well* to *LT_Morb* in the *Morbidity Status* factor can occur only when triggered by the transition *Survive → No Embolism* in the *Systemic Embolism* factor. Mean quality-adjusted life years for the product of these two factors are shown next to the appropriate states. These quantities were calculated using the Markovian rollback formulas and are identical to the results in Figure 8 obtained using updateable-state rollback formulas.

No Embolism in the new *Systemic Embolism* factor now triggers transition from *Well* to *LT_Morb* in the *Morbidity Status* factor.

Recall that the overall stochastic tree is now the *Cartesian product* of the *Systemic Embolism* factor and the *Morbidity Status* factor in Figure 9, so has states of the form (y, q), where q is morbidity status. Therefore, what was previously a preference summary q in the original model of Figure 8 has become part of the state information in the factored model of Figure 9. For this reason, the preference summary for the factored model is Markovian, and the Markovian rollback formulas can be employed. The numerical results of Markovian rollback are displayed in Figure 9, and are

of course identical to the values in Figure 8. Computationally, the Markovian rollback is identical to the previous updateable-state rollback. However, the factored form of the model in Figure 9 gives much greater insight into model structure: The *Systemic Embolism* factor in Figure 9 succinctly summarizes probabilistic structure, and the *Morbidity Status* factor summarizes preference structure.

Compared to the simple Markovian preference summary, updateable-state preference summaries appear to allow much greater flexibility in preference representation. However, this appearance is illusory because the same flexibility can be acquired by including the preference summary q in the state of the stochastic tree and using a Markovian summary.

At first sight, this result appears to diminish the significance of updateable-state utility. From a computational standpoint one need never implement the updateable-state rollback Equations (3) and (4) because Markovian rollback (2) over augmented-state stochastic trees will do. However, from the viewpoint of model construction and presentation, it is still advantageous to think in terms of updateable-state preference summaries because these may be *factored* from the model in a way that eases model construction and simplifies model presentation.

As promised, the preference factoring methodology discussed here by example is presented in a general setting in the following section. Readers not interested in the mathematical details can skip the next section without loss of continuity.

General Results on Preference Factoring

The Markovian rollback formulas (2) can be generalized to allow for *tolls* and *discount multipliers*, as I now explain. Let $EU(H|x)$ be the expected Markovian utility of the stochastic tree H in (1) given that y is reached from preceding state x. From Hazen and Pellissier (1996), the rollback Formulas (2) generalize to

$$EU(H|x) = w(y|x) + \Delta w(y|x)EU_0(H), \quad (5a)$$

$$EU_0(H) = \frac{v(y) + \sum_i \lambda_i EU(K_i)}{a(y) + \sum_i \lambda_i}. \quad (5b)$$

Here $w(y|x)$ is the negative of the toll from x to y, and $\Delta w(y|x)$ is the discount multiplier from x to y.

Updateable-state utility can be generalized to allow for tolls and discount multipliers as well. From Hazen and Sounderpandian (1999), one may derive the following generalization of the updateable-state rollback formulas (3):

$$Eu(H|q) = w(y|q) + \Delta w(y|q)Eu_0(H|q), \quad (6a)$$

$$Eu_0(H|q) = \frac{v(r) + \sum_i \lambda_i Eu(K_i|r)}{a(r) + \sum_i \lambda_i}, \quad r = \theta(q, y). \quad (6b)$$

Again, $w(y|q)$ and $\Delta w(y|q)$ are the negative toll and the discount factor respectively at y given preference summary q. Both the Markovian rollback formulas (5) and the updateable-state rollback formulas (6) are true expected utility formulas, that is, each results from taking the expectation of a utility function defined over sample paths.

Consider the situation described in the *systemic embolism* example of the previous section, in which one or more states are present not for any stochastic reason but only for their impact on preference. We formalize this notion as follows. Given a stochastic tree H with states in Y, we say that an equivalence relation $\equiv$ over Y is a *stochastic equivalence relation* if whenever $x \equiv x'$, then there is a one-to-one correspondence between the arcs immediately succeeding x and the arcs immediately succeeding x' such that

(a) corresponding arcs have the same attached rate (or probability),

(b) if y_i is a node successor of x and y'_i is the corresponding node successor of x' then $y_i \equiv y'_i$.
Note that if $x \equiv x'$, then starting the process in state x is stochastically identical to starting the process in state x'.

Let $x^{\equiv}$ be the set of all nodes x' stochastically equivalent to x. By merging all stochastically equivalent states, and identifying the corresponding successor arcs, one can form the stochastic tree $H^{\equiv}$ (sometimes denoted $H/\equiv$), whose states and arcs are equivalence classes modulo $\equiv$ of states and arcs in H.

For example, in the *systemic embolism* tree H of Figure 6, notice that

$$No\ Embolism \equiv Long\text{-}Term\ Morbidity, \quad (7a)$$

$$Subsequent\ Embolism \equiv Embolism. \quad (7b)$$

When these equivalent states and their corresponding arcs are merged, the result is the *Systemic Embolism* factor $H^{\equiv}$ of Figure 9.

If one's only goal is to display probabilistic structure or to compute probabilities, then it is better to use $H^{\equiv}$ because it has fewer nodes and arcs. However, if one wishes to compute expected utilities under an updateable-state preference structure, then one must revert to H, because stochastically equivalent states need not have the same preference summaries. For instance, in the *Systemic Embolism* model H, the state *Embolism* has possible preference summaries {*Well, LT_Morb*} but the stochastically equivalent state *Subsequent Embolism* has only *LT_Morb* as a possible preference summary.

Our main result of this section follows. We show that the method of the previous section works in general; that is, one may reduce an updateable-state expected utility computation to a Markovian expected utility computation by factoring out preference summaries and identifying stochastically equivalent states. Here are the details.

Construct a product tree $H^{\equiv} \otimes Q$ as follows: Let $H^{\equiv} \otimes Q$ have states $(x^{\equiv}, q)$, where q is a possible preference summary *up to and including* some $x' \in x^{\equiv}$. For every arc $x^{\equiv} \to y^{\equiv}$ in $H^{\equiv}$, every $x' \in x^{\equiv}$, and every preference summary q which can occur at x', add in $H^{\equiv} \otimes Q$ the arc $(x^{\equiv}, q) \to (y^{\equiv}, r)$ where $r = \theta(q, x')$. This product tree can be visualized in factored form: The transition $x^{\equiv} \to y^{\equiv}$ in $H^{\equiv}$ triggers the transition $q \to r$ in Q.

Suppose $w(y|q)$, $\Delta w(y|q)$, $v(q)$, and $a(q)$ are the component functions of the original updateable-state utility function $u(\cdot)$. We make the assumption that *tolls and discount factors depend only on preference summaries*, that is, $w(y|q)$ and $\Delta w(y|q)$ are functions only of q and $r = \theta(q, y)$. For convenience we write

$$w(y|q) = w(r|q), \quad \Delta w(y|q) = \Delta w(r|q).$$

Define the components of a Markovian utility function $U(\cdot)$ over product trees $H^{\equiv} \otimes Q$ as follows:

$$w((y^{\equiv}, r)|(x^{\equiv}, q)) = w(r|q), \quad (8a)$$

$$\Delta w((y^{\equiv}, r)|(x^{\equiv}, q)) = \Delta w(r|q), \quad (8b)$$

$$v(y^=, r) = v(r), \tag{8c}$$

$$a(y^=, r) = a(r). \tag{8d}$$

We have the following result.

THEOREM 1. *Suppose preference over stochastic trees having possible states $y \in Y$ is represented by an update-able-state utility function $u(\cdot)$ whose tolls and discount factors depend only on preference summaries. Let H be an arbitrary stochastic tree with states in Y, and let $=$ be a stochastic equivalence relation over H. Suppose the stochastic tree $H^= \otimes Q$ and Markovian utility $U(\cdot)$ are constructed as described above. For any $(y^=, r)$ in $H^= \otimes Q$, let $H(y^=, r)$ be the subtree with initial state $(y^=, r)$, and let $H(y)$ be the subtree of H with initial state y. Then*

$$EU(H(y^=, r)|(x^=, q)) = Eu(H(y)|q).$$

Arthritis Modeling with Updateable-State Utility

My colleagues and I (Chang et al. 1996) examined the cost-effectiveness of total hip replacement surgery (total hip arthroplasty, or THA). An estimated 120,000 hip replacements are performed per year in North America (Harris and Sledge 1990). THA is an elective, high cost procedure which reduces disability but does not extend life, and which has less expensive short-term alternatives. It is therefore particularly vulnerable to questions of cost-effectiveness. Nevertheless, we concluded that when *quality of life* is taken into account, THA is one of the most cost-effective of medical procedures, comparable or superior to well-accepted procedures such as cardiac bypass or renal dialysis.

Our cost-effectiveness analysis was performed using a stochastic tree model of hip surgery and its consequences. Our model at that time did not exploit updateable-state preference summaries. However, it is easy to recast our efforts in those terms, and the resulting formulation and presentation given below, while equivalent, gives considerably greater insight than the original. A summary of the model as originally conceived may be found in Hazen et al. (1998).

We chose to characterize the effectiveness of THA in terms of functional outcome measured by the four-state American College of Rheumatology (ACR) functional status classification. The four classes on the ACR scale are as follows:

I. Complete ability to carry on all usual duties without handicaps.

II. Adequate for normal activities despite handicap of discomfort or limited motion in one or more joints.

III. Limited only to little or none of duties of usual occupation or self-care.

IV. Incapacitated, largely or wholly bedridden or confined to wheelchair; little or no self-care.

A candidate for THA is typically in functional class III. The result of the THA procedure is usually an immediate transition to functional class I or II. However, short- or long-term complications may arise. The presence of infection might cause the prosthesis to fail over time (*septic* failure). The prosthesis might fail for a variety of other reasons such as mechanical loosening, prosthesis breakage, or dislocation (*aseptic* failure). Should any of these failures occur, revision surgery is required which is often less successful than the initial THA. A septic failure can cause permanent degradation in functional status. Small mortality rates accompany both initial THA and revision surgeries.

How can we construct an updateable-state preference summary for this problem? Clearly the ACR functional status measure should somehow be involved. We made several relevant assumptions early in the modeling effort. First, we assumed that all revision surgeries which the patient survives would either *succeed* or *fail*. Therefore, surviving patients would have either *successfully addressed* or *unsuccessfully addressed* prosthesis failures. We secondly assumed that in the absence of septic history, a successfully addressed aseptic failure would place the patient into functional class II, whereas an unsuccessfully addressed aseptic failure would result in functional class III. In the presence of septic history, the respective outcomes would be classes III and IV. Finally, a successfully addressed septic failure would place the patient into functional class III, and an unsuccessfully addressed septic failure would result in functional class IV. Patients for whom revision fails may undergo further revisions, but due to limitations in available bone stock, we decided it was not realistic to allow in the model more than three revision surgeries.

Based on these comments, it is apparent that func-

Figure 10 Factors of the THA Model Corresponding to the Bivariant Preference Summary (Failure History, Septic History)
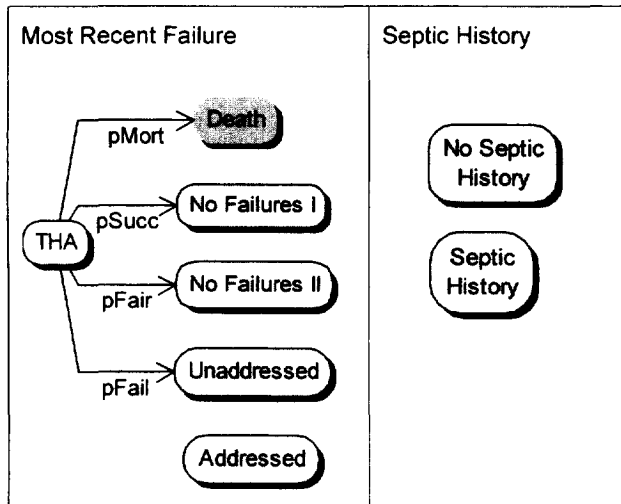
Note: The initial state of the *Most Recent Failure* factor is determined by the outcome of initial THA, as indicated, whereas the initial state of the *Septic History* factor is *No_Septic_History*. Overall utility in the THA model is a direct function of time spent in the combined states of these two factors.

tional status depends on a bivariate preference summary

$$q = (Most\ Recent\ Failure,\ Septic\ History).$$

Summary states are listed in Figure 10, which depicts the two stochastic factors arising from this updateable-state preference summary. The exact dependence of functional status on Most Recent Failure and Septic History is given in Table 1.

We can now assign quality factors $v(q)$ for Markovian utility by assessing patients' quality factors for

Table 1 Functional Status Depends on Failure History and Septic History

Most Recent Failure	Septic History	Functional Status
No Failures I	–	I
No Failures II	–	II
All Failures Addressed	No Septic History	II
	Septic History	III
Unaddressed Failures	No Septic History	III
	Septic History	IV

Table 2 The Quality Factor for Markovian Utility Depends on Functional Status

Functional Status	I	II	III	IV
Quality factor	1.0	0.8	0.5	0.3

Note: These are approximate values based on empirical work by ourselves and others.

functional classes I, II, III, IV. Based on empirical work by ourselves and others, we assigned the approximate values indicated in Table 2. Of course these values will differ across individuals.

The primary stochastic factor for the THA model is shown in Figure 11. Transitions in this factor affect utility by triggering transitions in the Most Recent Failure and Septic History factors.

Other factors were required in the model as well. A revision count factor was necessary to limit the number of revisions to three. Factors which count aseptic revisions and septic revisions separately and a factor which counts the most recent revision type (aseptic or septic) were needed because subsequent failure rates and revision success probabilities depend on these. Finally, the model includes an age- and gender-specific Coxian mortality factor.

The results of rollback for this stochastic tree model are shown in Table 3 for two individuals of differing gender and age. For both, THA roughly doubles discounted expected remaining quality-adjusted lifetime.

The purpose of preference factoring is to explicitly present features of the model on which overall utility most directly depends. As we hope this example illustrates, an analyst who employs this technique can

Table 3 Comparison of THA Versus No THA for Two Different Individuals

|	Discounted Quality-Adjusted Years	
	THA	No THA
White female age 60	13.70	6.82
White male age 85	4.16	2.16

Note: An annual discount rate of 5% was used.

clarify preference modeling assumptions and open the model to inspection and critique by others.

3. Multiattribute Decomposition of Markovian Utility

In this section we examine utility decompositions which can arise when factors in a stochastic tree are

Figure 11   The Primary Stochastic Factor for the THA Model
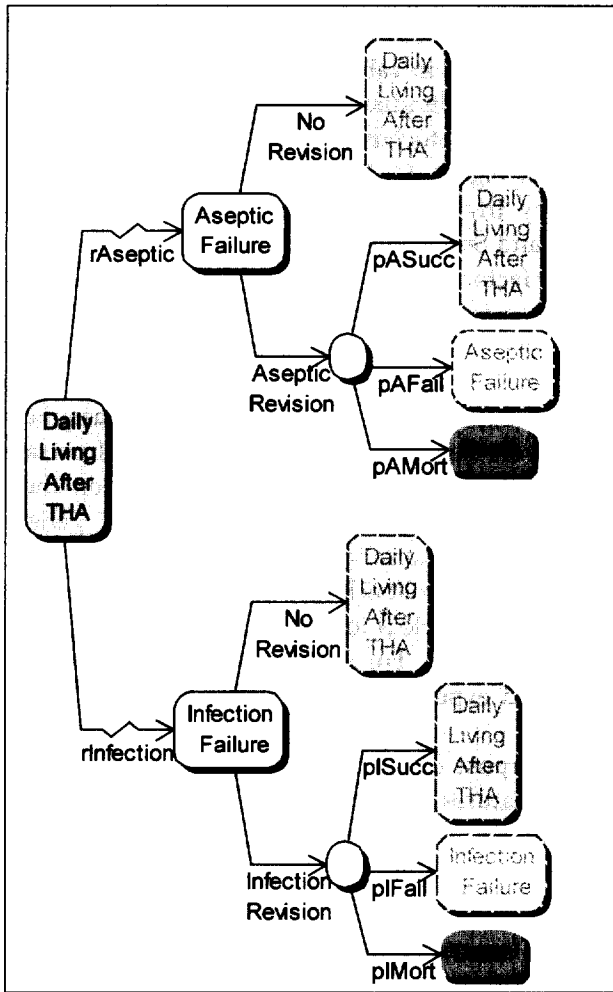
Note: Transition to *Infection_Failure* triggers a transition to *Unaddressed* in the *Most Recent Failure* factor, and a transition to *Septic_History* in the *Septic History* factor. Infection revision is performed if the total revision count (factor not shown) is less than three. Subsequent successful revision triggers a transition to *Addressed* in the *Most Recent Failure* factor. Similar triggers are invoked upon transition to *Aseptic_Failure*. All rates and probabilities in this factor depend on other factors not shown (*Aseptic Revision Count, Infection Revision Count, Last Revision Type*).

treated as attributes in a multi-attribute utility function. We consider only Markovian utility functions, noting that the results have implications for updateable-state utility as well because of the connection established by Theorem 1.

We consider the case of two-factor stochastic trees, since the multifactor case follows naturally. We denote states in the first factor by y, y', y_0, y_1 and so on; states in the second factor by z, z', z_0, z_1 and so on; and states in the product tree by $x = (y, z)$, $x' = (y', z')$, etc. Let y^* and z^* be distinguished states (e.g., Well). We assume $v(y^*, z^*) > 0$, where $v(\cdot)$ is the quality weight for Markovian utility.

We use the following graphical notation. The display

denotes the occupation of state y subject to mortality rate λ. For utility assessment purposes, it is often convenient to let λ be time dependent, so this display denotes a duration T sojourn in y followed by process termination (death, in the medical context), where T is a continuous-valued nonnegative random variable with hazard rate function $\lambda(t)$. The display
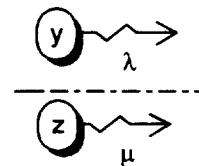
denotes a two-factor scenario in which y is occupied with mortality rate λ in the first factor and z is occupied with mortality rate μ in the second factor. Because termination in either factor forces termination of the entire process, this is equivalent to the product tree
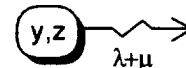
The Risk-Neutral Preference Interpretation for Markovian Utility

Given a hazard rate function $\lambda(t)$, let $m(\lambda)$ denote the mean of the corresponding random variable. The following elementary results are useful.

PROPOSITION 1. *If T has hazard rate function* $\lambda(t)$ *then*

(i) $E\left[\int_0^T e^{-at}dt\right] = m(\lambda + a),$

(ii) $E[e^{-aT}] = 1 - a \cdot m(\lambda + a).$

Markovian utility over sample paths can be expressed recursively. Let $x^t h$ be a sample path beginning with a duration-t sojourn in state y and followed by some other sample path h. If $U(x^t h | x')$ is the Markovian utility of sample path $x^t h$ given preceding state x', then

$$U(x^t h | x') = w(x | x') + \Delta w(x | x') U_0(x^t h),$$

$$U_0(x^t h) = \int_0^t v(x) e^{-a(x)s} ds + e^{-a(x)t} U(h | x).$$

(See Hazen and Pellissier 1996.) As previously mentioned, $v(x)$ is a quality weight at state x, $a(x)$ is a state-dependent discount rate, $w(x | x')$ is the negative of a toll in the transition from x' to x, and $\Delta w(x | x')$ is a discount multiplier for that transition.

Denote by x^λ a duration-T sojourn in state x, where T has hazard rate function λ. Note that for Markovian utility without toll or discount multiplier, we have

$$U(x^\lambda) = E\left[v(x) \int_0^T e^{-a(x)t}dt\right] = v(x)m(\lambda + a(x)).$$

For constant hazard rate λ, the result is

$$U(x^\lambda) = \frac{v(x)}{\lambda + a(x)}.$$

One interpretation of this equation is that under Markovian utility, a risk-sensitive sojourn with hazard rate λ in state x is equivalent to a risk-neutral sojourn with hazard rate $\lambda + a(x)$ in state x. We call this the *risk-neutral interpretation* of Markovian utility.

A strategy we have used (Hazen et al. 1991) to assess the risk-attitude parameter of Markovian utility is to elicit from subjects an immediate mortality probability p such that the subject is indifferent between the current mortality rate and a chance p of surviving immediate death with an improved mortality rate,

that is,
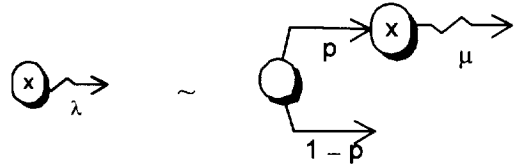
Then equating expected utilities yields

$$v(x)m(\lambda + a(x)) = pv(x)m(\mu + a(x)).$$

This equation can be solved for $a(x)$ in most situations (e.g., when $\mu = \lambda + c$ for some nonnegative constant c).

The Failure of Full Utility Independence

In any multiattribute context, it is natural to attempt to invoke utility independence (Keeney and Raiffa 1976). The naïve extension of utility independence from attributes to stochastic factors would run as follows: Call factor A utility independent of factor B if preference for stochastic trees in factor A do not depend on the particular stochastic tree in factor B. This definition must be amended to include the requirement that there are no parameter dependencies or triggers linking the two factors.[1] However, it is still not a useful definition because despite appearances to the contrary, transition rates do not necessarily attach to particular factors. For example, due to the way the Cartesian product of stochastic trees is defined, we have the equalities
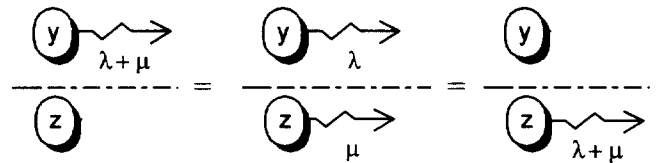

holding because all three are equal to the product tree

---
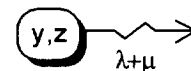
[1] This is an extension of a less familiar but equivalent form of utility independence: Attribute X is utility independent of attribute Y if for any lotteries x, x', y, y' with x, x' stochastically independent of y, y' we have (x, y) preferred to (x', y) if and only if (x, y') preferred to (x', y').

Utility independence of factor 1 from factor 2 would imply, for example, that the indifference
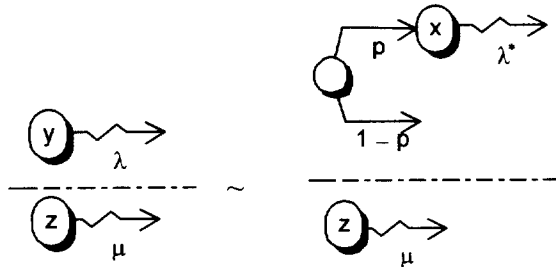
if holding for one z, μ combination, must hold for all. In product tree terms, this indifference translates to
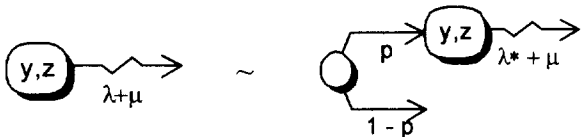

The requirement that this indifference holds for all μ if it holds for one is almost always false. For Markovian utility and constant mortality rates, the resulting expected utility equation is

$$\frac{1}{a(y, z) + \lambda + \mu} = \frac{p}{a(y, z) + \lambda^* + \mu}.$$

If this equation holds for one value of μ, it can hold for no others. Therefore full utility independence cannot hold.

Noninterfering Utility Independence

As a less restrictive assumption, it might be reasonable to require that one factor be utility independent of another factor as long as the other factor never changes state, and therefore never *interferes* with the first factor. Call this *noninterfering utility independence*. It would imply, for example, that the indifference
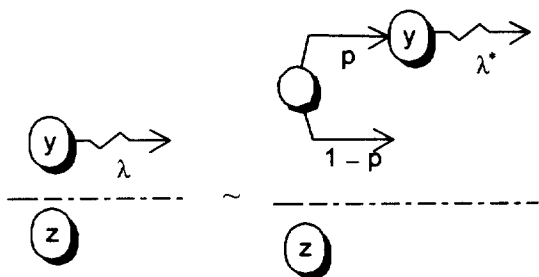
if holding for one nonfatal z, would hold for all nonfatal z. In product form, this indifference is
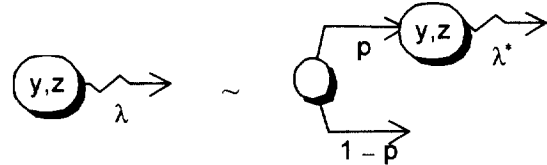
Since this indifference determines the risk attitude parameter $a(y, z)$, it would follow that $a(y, z)$ is independent of z. If each factor is noninterfering utility independent of the other, then $a(y, z)$ can depend on neither y nor z, so is a constant $a(y, z) = a^*$. This establishes the first part of the following result.

THEOREM 2. *Under Markovian utility each factor is noninterfering utility independent of the other if and only if there exist constants $v^* > 0$ and a^*, marginal quality-rate functions $v_1(y)$, $v_2(y)$, marginal tolls $w_1(y|y_0)$, $w_2(z|z_0)$, and marginal discount multipliers $\Delta w_1(y|y_0) > 0$, $\Delta w_2(z|z_0) > 0$ such that*

(i) $a(y, z) = a^*$ *independent of y, z;*

(ii) $v(y, z)/v^* = (v_1(y)/v^*)\,(v_2(z)/v^*)$ *where $v^* = v(y^*, z^*) = v_1(y^*) = v_2(z^*)$;*

(iii) $w(y, z|y_0, z) = w_1(y|y_0) \cdot v_2(z) \cdot v^*$, $w(y, z|y, z_0) = v_1(y) \cdot w_2(z|z_0) \cdot v^*$;

(iv) $\Delta w(y, z|y_0, z) = \Delta w_1(y|y_0)$, $\Delta w(y, z|y, z_0) = \Delta w_2(z|z_0)$.

The proof is given in the appendix. Usually one would rescale v so that $v^* = 1$.

The fact that noninterfering utility independence forces risk attitude $a(y, z)$ to be a constant independent of y, z is disappointing because it reduces the flexibility of the multiattribute Markovian utility model. We have, for example, observed that subjects engaged in immediate mortality scenarios of the type described above have risk attitudes which can depend on health state (Pellissier and Hazen 1994). However, we are aware of no realistic Markovian utility decomposition which allows nonconstant $a(y, z)$.

4. Conclusions

We have shown that the useful notion of *factoring* a stochastic tree model may be extended to the preference domain as well, that is, one may represent

updateable-state preference components as one or more factors in a multifactor decomposition of the stochastic tree. Then one may compute expected up-dateable-state utility by rolling back an equivalent Markovian utility function over the product tree. The Markovian rollback is computationally equivalent to an updateable-state rollback without preference factoring. However, the explicit factoring of preference components from the stochastic tree eases model formulation and presentation, and the resulting transparent form opens the model to inspection and critique by others.

We have also explored multiattribute utility decompositions for Markovian utility. The usual notion of utility independence fails for Markovian utility over stochastic trees. However, the restricted notion of *noninterfering* utility independence leads to a useful utility decomposition.

Appendix. Proofs of Theorems

PROOF OF THEOREM 1. Let $H(y^*, r)$ be a stochastic fork at state (y^*, r) in $H^* \otimes Q$. Suppose (y^*, r) has successors (z_i^*, s_i), where $s_i = \theta(r, z_i)$. For the Markovian utility $U(\cdot)$ defined in the theorem, the rollback equations (5) become

$$EU(H(y^*, r)|(x^*, q)) = w((y^*, r)|(x^*, q))$$
$$+ \Delta w((y^*, r)|(x^*, q))EU_0(H(y^*, r)),$$

$$EU_0(H(y^*, r)) = \frac{v(y^*, r) + \sum_i \lambda_i EU(H(z_i^*, s_i)|(y^*, r))}{a(y^*, r) + \sum_i \lambda_i}.$$

Apply the definitions (8) to obtain

$$EU(H(y^*, r)|(x^*, q))$$

$$= w(r|q) + \Delta w(r|q) \frac{v(r) + \sum_i \lambda_i EU(H(z_i^*, s_i)|(y^*, r))}{a(r) + \sum_i \lambda_i}.$$

By induction on the height of the tree, we can assume

$$EU(H(z_i^*, s_i)|(y^*, r)) = Eu(H(z_i)|r).$$

Then the prior equation becomes

$$EU(H(y^*, r)|(x^*, q))$$

$$= w(r|q) + \Delta w(r|q) \frac{v(r) + \sum_i \lambda_i Eu(H(z_i)|r)}{a(r) + \sum_i \lambda_i}$$

$$= Eu(H_y|q).$$

For chance forks H the proof is similar.

Equivalence of Markovian Utility Functions

LEMMA 1. *Two Markovian utility functions u and u' with $u(\varnothing) = u'(\varnothing) = 0$ are strategically equivalent if and only if there is a positive constant c such that*

$$w'(x|x_0) = cw(x|x_0), \quad \Delta w'(x|x_0) = \Delta w(x|x_0),$$

$$v'(x) = cv(x), \quad a'(x) = a(x).$$

PROOF. Because u and u' share a common zero point, strategic equivalence forces $u' = cu$ for some positive c. Consider a stochastic tree $x^\lambda K$. Given Markovian preference state x_0, expected utilities are

$$E[u(x^\lambda K|x_0)] = w(x|x_0) + \Delta w(x|x_0) \cdot (v(x)m(a(x) + \lambda)$$
$$+ (1 - a(x)m(a(x) + \lambda))E[u(K|x)]), \quad (1)$$

$$E[u'(x^\lambda K|x_0)] = w'(x|x_0) + \Delta w'(x|x_0)(v'(x)m(a'(x) + \lambda)$$
$$+ (1 - a'(x)m(a'(x) + \lambda))E[u'(K|x)]). \quad (2)$$

Set $K = \varnothing$ in the above to get

$$E[u(x^\lambda|x_0)] = w(x|x_0) + \Delta w(x|x_0) \cdot v(x)m(a(x) + \lambda),$$

$$E[u'(x^\lambda|x_0)] = w'(x|x_0) + \Delta w'(x|x_0)v'(x)m(a'(x) + \lambda).$$

Then $u' = cu$ implies

$$w'(x|x_0) + \Delta w'(x|x_0)v'(x)m(a'(x) + \lambda)$$

$$= c(w(x|x_0) + \Delta w(x|x_0) \cdot v(x)m(a(x) + \lambda)).$$

Set $\lambda = \infty$ to get

$$w'(x|x_0) = cw(x|x_0). \quad (3)$$

Substitute this back into its predecessor and cancel to get

$$\Delta w'(x|x_0)v'(x)m(a'(x) + \lambda) = c\Delta w(x|x_0) \cdot v(x)m(a(x) + \lambda).$$

Let $\lambda(t)$ be a constant λ to conclude

$$\Delta w'(x|x_0)v'(x)/(a'(x) + \lambda) = c\Delta w(x|x_0) \cdot v(x)/(a(x) + \lambda).$$

The equality of two functions $k'/(a' + \lambda)$ and $k/(a + \lambda)$ of $\lambda \in [0, \infty)$ forces $k' = k$ and $a' = a$. Therefore we conclude

$$\Delta w'(x|x_0)v'(x) = c\Delta w(x|x_0) \cdot v(x), \quad (4)$$

$$a'(x) = a(x). \quad (5)$$

Substitute (3), (4), and (5) back into (2), then use $u' = cu$ and simplify to get

$$cE[u(x^\lambda K|x_0)] = cw(x|x_0) + c\Delta w(x|x_0)v(x)m(a(x) + \lambda)$$
$$+ \Delta w'(x|x_0)(1 - a(x)m(a(x) + \lambda))cE[u(K|x)].$$

Substitute the right side of (1) for $E[u(x^\lambda K|x_0)]$ in the last equation and simplify to get

$$c\Delta w(x|x_0)(1 - a(x)m(a(x) + \lambda))E[u(K|x)])$$

$$= \Delta w'(x|x_0)(1 - a(x)m(a(x) + \lambda))cE[u(K|x)].$$

Cancel terms to conclude $\Delta w(x|x_0) = \Delta w'(x|x_0)$. Therefore from (4) we conclude $v'(x) = cv(x)$. We have thus shown that strategic equivalence of u and u' imply the four equalities specified in Lemma 1.

Conversely, if these four equalities hold, then from (2) we obtain

$$E[u'(x^\wedge K|x_0)] = cw(x|x_0) + \Delta w(x|x_0)(cv(x)m(a'(x) + \lambda)$$

$$+ (1 - a'(x)m(a'(x) + \lambda))E[u'(K|x)]). \quad (6)$$

Inducting on the height of the stochastic tree, if we have $E[u'(K|x)] = cE[u(K|x)]$, then from (6) we conclude $E[u'(x^\wedge K|x_0)] = cE[u(x^\wedge K|x_0)]$. Hence u' and u are strategically equivalent. □

PROOF OF THEOREM 2. Given Markovian utility over stochastic trees with product states (y, z), define the following marginal functions:

$$v_1(y) = v(y, z^*), \quad v_2(z) = v(y^*, z),$$

$$a_1(y) = a(y, z^*), \quad a_2(z) = a(y^*, z),$$

$$w_1(y|y') = w(y, z^*|y', z^*),$$

$$w_2(z|z') = w(y^*, z|y^*, z'),$$

$$\Delta w_1(y|y') = \Delta w(y, z^*|y', z^*),$$

$$\Delta w_2(z|z') = \Delta w(y^*, z|y^*, z'),$$

and let $v^* = v(y^*, z^*)$. Suppose noninterfering utility independence holds. Consider a stochastic tree $(y^\wedge K, z)$. Its Markovian expected utility given preference summary (y_0, z) is

$$E[u(y^\wedge K, z|y_0, z)]$$

$$= w(y, z|y_0, z) + \Delta w(y, z|y_0, z) \cdot (v(y, z)m(a(y, z) + \lambda)$$

$$+ (1 - a(y, z)m(a(y, z) + \lambda))E[u(K, z|y, z)]). \quad (7)$$

Substitute $z = z^*$ to get

$$E[u(y^\wedge K, z^*|y_0, z^*)]$$

$$= w_1(y|y_0) + \Delta w_1(y|y_0) \cdot (v_1(y)m(a_1(y) + \lambda)$$

$$+ (1 - a_1(y)m(a_1(y) + \lambda))E[u(K, z^*|y, z^*)]). \quad (8)$$

Noninterfering utility independence of factor 1 from factor 2 implies the last two expressions are equivalent Markovian utility functions over factor 1. Invoking Lemma 1, we conclude there is a positive $c_2(z)$ such that

$$v(y, z) = v_1(y)c_2(z), \quad a(y, z) = a_1(y),$$

$$w(y, z|y_0, z) = w_1(y|y_0)c_2(z) \quad \Delta w(y, z|y_0, z) = \Delta w_1(y|y_0).$$

Similarly, noninterfering utility independence of factor 2 from factor 1 implies there is a positive $c_1(y)$ such that

$$v(y, z) = c_1(y)v_2(z), \quad a(y, z) = a_2(z),$$

$$w(y, z|y, z_0) = c_1(y)w_2(z|z_0) \quad \Delta w(y, z|y, z_0) = \Delta w_2(z|z_0).$$

Repeatedly invoking these restrictions, we obtain condition (i) of the theorem:

$$a(y, z) = a_1(y) = a(y, z^*) = a_2(z^*) = a^*.$$

In a similar way

$$v(y, z) = v_1(y)c_2(z)$$

$$= v(y, z^*)c_2(z)$$

$$= c_1(y)v_2(z^*)c_2(z)$$

$$= c_1(y)c_2(z)v^*.$$

Set $z = z^*$ in the first equality of this sequence to get $c_2(z^*) = 1$. Similarly, $c_1(y^*) = 1$. Substitute $z = z^*$ into the last equality of the sequence to get $v_1(y) = c_1(y)v^*$. Similarly, $v_2(z) = c_2(z)v^*$. Conditions (ii), (iii), (iv) of the theorem follow immediately.

Conversely, if conditions (i)–(iv) all hold, then (7) becomes

$$E[u(y^\wedge K, z|y_0, z)]$$

$$= w_1(y|y_0)v_2(z)v^* + \Delta w_1(y|y_0) \cdot (v_1(y)v_2(z)v^*m(a^* + \lambda)$$

$$+ (1 - a^*m(a^* + \lambda))E[u(K, z|y, z)])$$

$$= (w_1(y|y_0) + \Delta w_1(y|y_0)(v_1(y)m(a^* + \lambda))v_2(z)v^*$$

$$+ \Delta w_1(y|y_0)(1 - a^*m(a^* + \lambda))E[u(K, z|y, z)],$$

and (8) becomes

$$E[u(y^\wedge K, z^*|y_0, z^*)]$$

$$= w_1(y|y_0) + \Delta w_1(y|y_0) \cdot (v_1(y)m(a^* + \lambda)$$

$$+ (1 - a^*m(a^* + \lambda))E[u(K, z^*|y, z^*)])$$

$$= w_1(y|y_0) + \Delta w_1(y|y_0)v_1(y)m(a^* + \lambda)$$

$$+ \Delta w_1(y|y_0)(1 - a^*m(a^* + \lambda))E[u(K, z^*|y, z^*)].$$

We induct on the height of the stochastic tree. If the induction hypothesis

$$E[u(K, z|y, z)] = E[u(K, z^*|y, z^*)]v_2(z)v^*$$

holds for the subtree K following y then it follows from the last two equations that

$$E[u(y^\wedge K, z|y_0, z)] = E[u(y^\wedge K, z^*|y_0, z^*)]v_2(z)v^*.$$

Therefore, by induction, the last equation holds for all states y and subsequent subtrees K. Therefore factor 1 is noninterfering utility independent of factor 2, as claimed. Similarly, factor 2 is noninterfering utility independent of factor 1. □

References

Chang, R. W., J. M. Pellissier, G. B. Hazen. 1996. A cost-effectiveness analysis of total hip arthroplasty for osteoarthritis of the hip. *J. Amer. Medical Assoc.* **275** 858–865.

Cox, D. R. 1955. A use of complex probabilities in the theory of stochastic processes. *Proc. Cambridge Philos. Soc.* **51** 313–319.

Gottlob, C. A., J. M. Pellissier, R. L. Wixson, S. H. Stern, S. D. Stulberg, R. W. Chang. 1996. The long-term cost-effectiveness of total knee arthroplasty for osteoarthritis. Multipurpose Arthritis Center, Northwestern University, Chicago, IL.

Harris, W. H., C. B. Sledge. 1990. Total hip and total knee replacement. *New England J. Medicine* **323** 725–731.

Hazen, G. B. 1992. Stochastic trees: A new technique for temporal medical decision modeling. *Medical Decision Making* **12** 163–178.

——. 1993. Factored stochastic trees: A tool for solving complex temporal medical decision models. *Medical Decision Making* **13** 227–236.

——, W. J. Hopp, J. M. Pellissier. 1991. Continuous-risk utility assessment *Medical Decision Making* **11** 294–304.

——, M. Morrow, E. R. Venta. 1999. Patient values in the treatment of ductal carcinoma in situ. Society for Medical Decision Making Annual Meeting, Reno, NV.

——, J. M. Pellissier. 1996. Recursive utility for stochastic trees. *Oper. Res.* **44** 788–809.

——, ——, J. Sounderpandian. 1998. Stochastic tree models in medical decision making. *Interfaces* **48** 64–80.

——, J. Sounderpandian. 1999. Preference summaries for stochastic trees rollback. Mark J. Machina, Bertrand Munier, eds. *Beliefs, Interactions, and Preferences in Decision Making.* Kluwer, Dordrecht, The Netherlands 109–120.

Keeney, R. L., H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* Wiley, New York.

Matchar, D. B., S. G. Pauker. 1986. Transient ischemic attacks in a man with coronary artery disease: Two strategies neck and neck. *Medical Decision Making* **6** 239–249.

Meyer, R. F. 1976. Preferences over time. R. L. Keeney and H. Raiffa, eds. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* Wiley, New York.

Pellissier, J. M., G. B. Hazen. 1994. Implementation of continuous-risk utility assessment: The total hip replacement decision. *Socio-Econom. Planning Sci.* **28** 251–276.

——, ——, R. W. Chang. 1996. A continuous-risk decision analysis of total hip replacement. *J. Oper. Res. Soc.* **47** 776–793.

Tsevat J., M. H. Eckman, R. A. McNutt, S. G. Pauker. 1989. Warfarin for dilated cardiomyopathy: A bloody tough pill to swallow? *Medical Decision Making* **9** 162–169.