

PREFERENCE FACTORING FOR STOCHASTIC TREES

GORDON HAZEN

Northwestern University

May 1998

Stochastic trees are extensions of decision trees which 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 utility models which permit stochastic tree rollback. 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.

S*tochastic trees* are graphical modeling tools which extend decision tree 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, Pellissier and Hazen 1996;

Pellissier, Hazen and Chang 1996; Gottlob *et al.* 1996), and for breast cancer decision analysis (Hazen, Morrow and Venta 1998). We have investigated risk-sensitive utility rollback for stochastic trees (Hazen and Pellissier 1996), and have explored preference assessment techniques which are tailored to address risky temporal tradeoffs (Hazen, Hopp and Pellissier 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 with a graphical interface for formulating and rolling back factored stochastic trees. (For further information on software, contact the author.)

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, Pellissier and Sounderpandian 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 Section 2, where we examine *updateable-state* utility models. These arise from the *state-trajectory* preference summaries recently introduced by Hazen and Sounderpandian (1997), and extend the *Markovian* utility models introduced by Hazen and Pellissier (1996). The major result presented in Section 2 is that every updateable-state utility model is equivalent 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 Section 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

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. Stochastic trees can be rolled back much like decision trees. Figure 1 displays a simplified stochastic tree model of transient ischemic attack, adapted from Matchar and Pauker (1986), along with the rollback values at each node. For further introductory material on

stochastic trees, see Hazen (1992) or Hazen and Pellissier (1996).




Figure 1: A stochastic tree model of transient ischemic attack, having two stochastic nodes (shaded) and one chance node. Shaded terminal nodes represent death. The rollback values shown at nodes are discounted quality-adjusted life years, calculated at discount rate 5%/yr. when $m = 0.04/\text{yr.}$, $l = 0.25/\text{yr.}$, $p = 15\%$, and the quality rate for Post Stroke equal to 30% of Well.

In the model of Figure 1, a background mortality rate μ is present for both the Well and Post Stroke states. 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 2. We say that the transient ischemic attack model of Figure 1 is the (*Cartesian*) *product* of the Stroke and Background Mortality factors in Figure 2. 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.

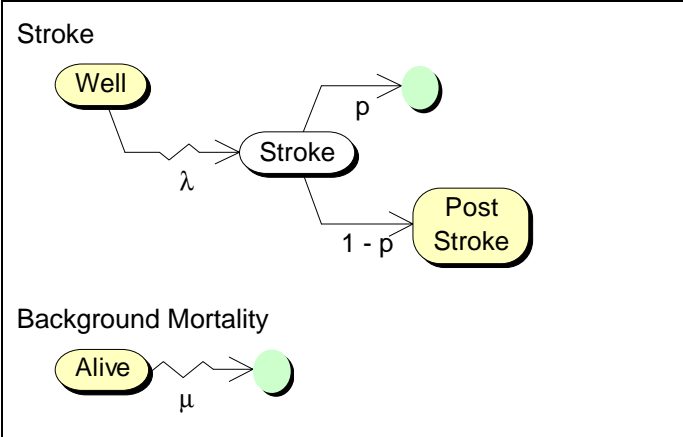


Figure 2: Factoring background mortality out of the transient ischemic attack model of Figure 1.

Similarly, background mortality can be dealt with independently. For example, the background mortality model in Figure 2 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, Pellissier and Sounderpandian 1998) depicted in Figure 3. The parameters and the number of stages in a Coxian model can be chosen to arbitrarily well approximate human survival durations. The Coxian factor can be substituted for simple background mortality in Figure 2, resulting in an improved model. The Cartesian product tree need never be explicitly constructed (but see Hazen, Pellissier and Sounderpandian 1998 for a picture). Hazen (1993) discusses a rollback procedure for factored stochastic trees. Software for constructing and rolling back multi-factor stochastic trees can be obtained from the

author. All stochastic tree illustrations in this article are screen captures from that software.




Figure 3: A Coxian mortality factor, in which mortality rate μ depends on the stage of the process.

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 4, and the *anticoagulant-status* factor in Figure 5. Three useful modeling features are illustrated: First, stochastic trees can have *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. Finally, transitions in one factor can trigger transitions in other factors.

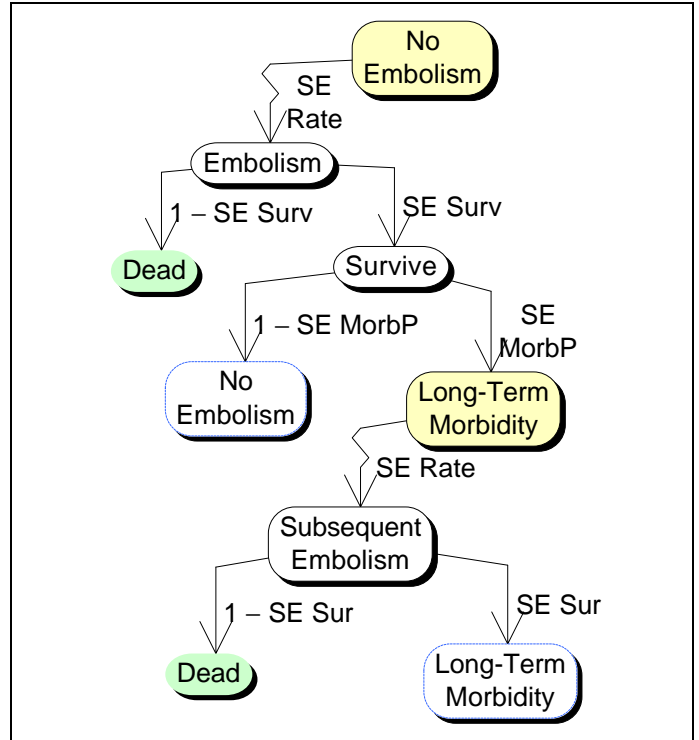


Figure 4: The systemic-embolism factor from a six-factor model of the use of warfarin in dilated cardiomyopathy. The tree has cycles - repeated visits to *No_Embolism* and to *Long-Term_Morbidity* are possible. The rate of systemic embolism (*SE_Rate*) depends on the state of the anticoagulant-status factor (Figure 5).

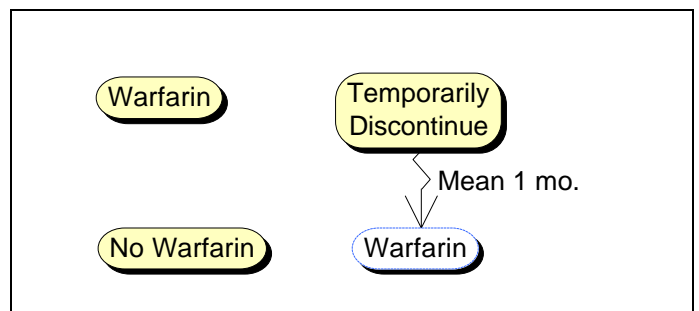


Figure 5: The anticoagulant-status factor in the warfarin model. An embolism occurrence in the systemic embolism factor triggers a transition from *No_Warfarin* to *Warfarin* in this factor. A hemorrhage in the systemic hemorrhage factor (not shown) triggers a transition here from *Warfarin* to *Temporarily_Discontinue*.


2. FACTORING PREFERENCE COMPONENTS FROM A STOCHASTIC TREE

Updateable-State Utility

The rollback operations performed in Figure 1 are the recursive evaluation of the *memoryless* utility function, defined recursively as

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

(Hazen and Pellissier 1996). Here y^t denotes a duration- t sojourn in state y , and h is a (possibly null) sequence of states and sojourns. The quantity $v(y)$ is the quality rate in state y , and $a(y)$ is the discount rate in state y , or equivalently, the risk attitude for time spent in state y . When $a(y) = 0$ for all y , expected utility is mean quality-adjusted duration, and when $a(y) = a$ for all y , expected utility is quality-adjusted duration discounted at rate a . At a stochastic fork



in which subtrees K_i are reached from state y at competing rates l_i , the rollback equation can be shown to take the simple form

$$E[U(H)] = \frac{v(y) + \sum_i l_i E[U(K_i)]}{a(y) + \sum_i l_i}. \quad (2)$$

At chance forks, expected utility computation is identical to the usual probability-weighted averaging done in decision trees.

Hazen and Pellissier introduce the notion of *preference summary* as follows. Let $\mathbf{x}^\lambda = x_1^{l_1} \cdots x_n^{l_n}$ denote a path through a stochastic tree, consisting of a sequence of states x_i and intermediate rates l_i . We allow $n = 0$, in which case \mathbf{x}^λ is the empty sequence \emptyset . Given a utility function U over stochastic trees, and a path \mathbf{x}^λ through a stochastic tree, the *conditional* utility function $U(\mathbf{x}^\lambda)$ is defined by

$$U(h|\mathbf{x}^\lambda) \sim_h E[U(\mathbf{x}^\lambda h)] - E[U(\mathbf{x}^\lambda)].$$

That is, $U(\mathbf{x}^\lambda)$ is a strategically equivalent version of U over the set of stochastic trees beginning with path \mathbf{x}^λ , rescaled so that $U(\emptyset|\mathbf{x}^\lambda) = 0$. If $U(\mathbf{x}^\lambda)$ depends on \mathbf{x}^λ only through some function $e(\mathbf{x}^\lambda)$ of \mathbf{x}^λ , then we say that $e(\cdot)$ is a *preference summary* for U . We say that $e(\cdot)$ is *updateable* if there is a function q such that $e(\mathbf{x}^\lambda y^m) = q(q, y^m)$ whenever $e(\mathbf{x}^\lambda) = q$.

A preference summary captures aspects of past history which suffice to determine preference over subsequent stochastic trees. The notion of preference summary was introduced for discrete-time models by Meyer (1976), who called it *state*

descriptor. For memoryless utility, conditional utility given \mathbf{x}^λ can be obtained by invoking (2):

$$\begin{aligned} U(h|\mathbf{x}^\lambda) &\sim_h E[U(\mathbf{x}^\lambda h)] \\ &\sim_h \frac{v(x_1) + l_1 E[U(\mathbf{x}_{(2)}^{(2)} h)]}{a(x_1) + l_1} \\ &\sim_h E[U(\mathbf{x}_{(2)}^{(2)} h)]. \end{aligned}$$

where $\mathbf{x}_{(2)} = (x_2, \dots, x_n)$ and $\lambda_{(2)} = (l_2, \dots, l_n)$. Repeating this reduction $n-1$ times produces $U(h|\mathbf{x}^\lambda) \sim_h U(h)$. Therefore the empty summary $e(\mathbf{x}^\lambda) = \emptyset$ is a preference summary (hence the name *memoryless*). Of course, more complex preference summaries are possible. Consider, for example, the *Markovian* preference summary $e(\mathbf{x}^\lambda)$ which specifies the *last* state visited in \mathbf{x}^λ . Markovian preference summaries are updateable. Hazen and Pellissier identify all utility functions having Markovian preference summaries. They show that when $x = e(\mathbf{x}^\lambda)$, the corresponding conditional utility function must be of the form $U(y^t h | \mathbf{x}^\lambda) \sim u(y^t h | x)$, where

$$u(y^t h | x) = w(y/x) + Dw(y/x)u_0(y^t h)$$

and u_0 is the memoryless utility function

$$u_0(y^t h) = \int_0^t v(y)e^{-a(y)s} ds + e^{-a(y)t} u(h).$$

Markovian utility extends memoryless utility by allowing *tolls* (or bonuses) $w(y/x)$ and *discount multipliers* $Dw(y/x)$ for transitions from x to y .

Hazen and Sounderpandian (1997?) extend this result from Markovian preference summaries to *updateable-state* preference summaries. We say that $e(\cdot)$ is an updateable-state preference summary if $e(\cdot)$ is updateable and $e(\mathbf{x}^\lambda)$ depends only on the states \mathbf{x} and not on the distribution information λ . If $e(\cdot)$ is updateable-state, then its update function $q(q, y^m)$ does not depend on m so we write it as $q(q, y)$.

Updateable-state preference summaries include the Markovian summaries and much more. Other examples of updateable state summaries are the *count of previous 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).

We may ask which utility functions over stochastic trees have updateable-state preference summaries. Hazen and Sounderpandian answer that question as follows.

Theorem 1: *Suppose the utility function U over stochastic trees has updateable-state preference summary $e(\mathbf{x})$ with update function $q(q, y)$. Then*

there exists a summary utility function $u(h/q)$ such that $U(\mathbf{x} | \mathbf{x}^1) \sim u(\mathbf{x} | q)$ whenever $e(\mathbf{x}^1) = q$, and

$$u(y^1 h/q) = w(y/q) + Dw(y/q)u_0(y^1 h/q)$$

where

$$u_0(y^1 h/q) = \int_0^t v(q(q, y)) \cdot e^{-a(q(q, y)) \cdot s} ds + e^{-a(q(q, y)) \cdot t} u(h/q(q, y)).$$

◆

In what follows, we shall for simplicity not explicitly distinguish between the original utility function U and its strategically equivalent summary utility function u . The goal of an expected utility computation is to derive the value of $E[u(H_0/q_0)]$, where H_0 is a particular stochastic tree model and $q_0 = e(\emptyset)$ is the initial summary state. For updateable-state utility functions, methods from Hazen and Pellissier (1996) may be adapted to derive recursive formulae for this computation, which are given in the next theorem.

Theorem 2: *Let u be an updateable-state utility function u with update function $q(q, y)$. Then expected utility given summary state q at a stochastic fork (1) in a stochastic tree can be computed recursively according to the formulae*

$$Eu(H | q) = w(y | q) + \Delta w(y | q)Eu_0(H | q)$$

$$Eu_0(H | q) = \frac{v(q(q, y)) + \sum_i p_i Eu(K_i | q(q, y))}{a(q(q, y)) + \sum_i p_i}$$

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

$$Eu(H | q) = \sum_i p_i Eu(K_i | q(q, y)).$$

◆

Factoring Out Updateable-State Summaries

The preliminary identification of an updateable preference summary can simplify the subsequent formulation of a stochastic tree by allowing the exclusive focus on probabilistic aspects of the model rather than preference aspects. For example, consider the stochastic factor of Figure 4. Here, the state Long-Term Morbidity plays no probabilistic role – it is present only to account for degradations in quality of life due to systemic embolism. In fact, whether Long-Term Morbidity has occurred is an updateable-state preference summary which can be accounted for at rollback time, so why not formulate the model without it? Figure 6 shows how this might be done. Without having to keep track of whether long-term morbidity has occurred, the model simplifies nicely.

The point of using an updateable-state preference summary is to simplify model formulation and presentation. There is no associated computational advantage or disadvantage. In fact, an updateable-state utility function over stochastic trees with states y can be equivalently modeled by a Markovian utility function over stochastic trees with augmented states (y, q) , as we now show.

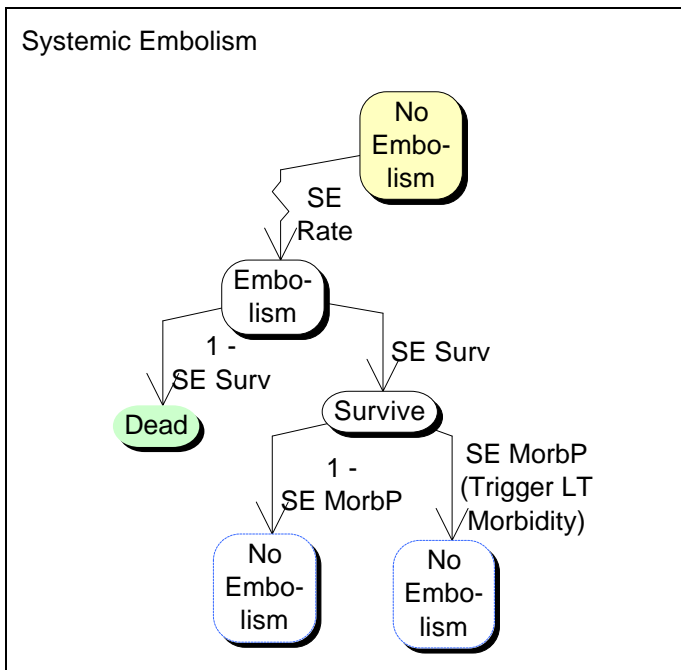


Figure 6: The Systemic Embolism factor of Figure 4 formulated without attempting to account for whether Long-Term Morbidity has previously occurred. The latter can be included as an updateable-state preference summary at rollback.

Given an updateable-state preference summary $e(\cdot)$ over sequences of states $y \in Y$, let Q be the set of possible summary states. Given a stochastic tree H with states in Y , construct a product tree $H \otimes Q$ with states (x, q) , where q is a possible

preference summary at x , and which, for every transition $x \text{ fl } y$ in H , contains the transition $(x, q) \rightarrow (y, r)$ in $H \otimes Q$, where $r = \mathfrak{Q}(q, y)$. The factor Q has no explicit arcs, and only undergoes transition $q \rightarrow r = \mathfrak{Q}(q, y)$ when triggered by transition $x \rightarrow y$ in H . Define a Markovian utility function over product trees $H \otimes Q$ as follows. Suppose $w(y|q)$, $Dw(y/q)$, $v(q)$ and $a(q)$ are the component functions of the original updateable-state utility u (Theorem 1). Then let

$$w_M(y, r/x, q) = w(y/q)$$

$$Dw_M(y, r/x, q) = Dw(y/q)$$

$$v_M(y, r) = v(r)$$

$$a_M(y, r) = a(r)$$

be the components of a Markovian utility u_M over stochastic trees $H \otimes Q$. We have the following result.

Theorem 3: Suppose preference over stochastic trees having possible states $y \in Y$ is represented by a utility function with updateable-state preference summary $e(\cdot)$ and summary utility function u given in Theorem 1. Let H_0 be an arbitrary stochastic tree with states in Y , and suppose the augmented-state stochastic tree $H_0 \sim Q$ and Markovian utility u_M are constructed as described above. For any transition $x \text{ fl } y$ in H_0 , let H be the subtree of H_0 with initial state y , and

let q be a summary state which can occur at x .
Then

$$E[u(H|q)] = E[u_M(H \otimes Q | x, q)].$$

◆

For example, suppose H is the stochastic tree of Figure 6, and $e(\cdot)$ is the updateable-state summary which indicates whether the state `Begin_LT_Morbidity` has been encountered. Then the product $H \otimes Q$ of H with the factor Q consisting of two states `{No_LT_Morbidity, LT_Morbidity}` is equivalent to the stochastic tree of Figure 4. Initially the state `No_LT_Morbidity` is occupied in Q , but a transition to `LT_Morbidity` is triggered when the right branch of the `Survive` state is traversed in H .

Theorem 3 is important for several reasons. Compared to the simple Markovian preference summary, updateable-state preference summaries appear to allow much greater flexibility in preference representation. In fact this appearance is partly illusory because according to Theorem 3, the same flexibility can be acquired by employing a Markovian preference summary over augmented-state stochastic trees. The nature of the state augmentation is crucial: One should form the product $H \otimes Q$ of the stochastic tree H of interest with the set Q of possible summary states. This product is best done implicitly, that is, it is best to

simply consider H and Q as factors with the appropriate transitions in Q triggered by transitions in H . In this sense we can say that preference structure has been *factored out* of the model.

At first sight, Theorem 3 appears to diminish the significance of updateable-state utility. From a computational standpoint one need never implement the updateable-state rollback equations of Theorem 2 because Markovian rollback 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. We illustrate this process in the next section.

Arthritis Modeling with Updateable-State Utility

My colleagues and I (Chang, Pellissier and Hazen 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, Pellissier and Sounderpandian (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 functional status depends on a bivariate preference summary

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

Summary states are listed in Figure 7, which depicts the two stochastic factors arising from this updateable-state preference summary. The exact dependence of functional status on Failure History and Septic History is given in Table 1.

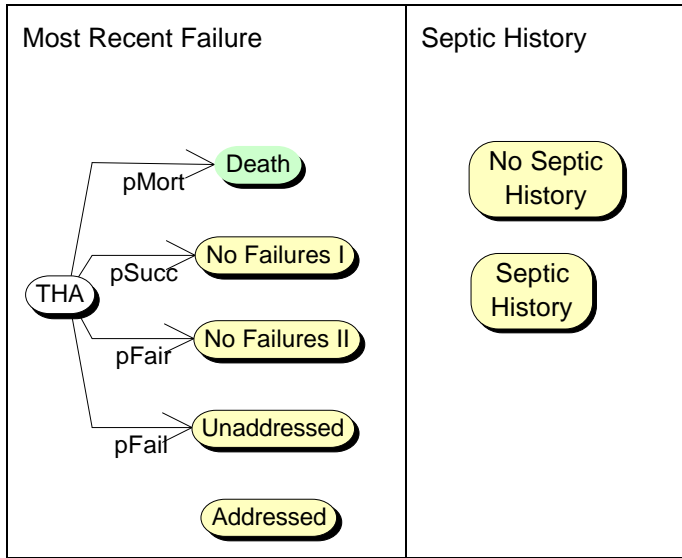


Figure 7: Factors of the THA model corresponding to the bivariate preference summary (Failure History, Septic History). The initial state of the Most Recent Failure factor is determined by the outcome of initial THA, as indicated, whereas the initial state of the Septic History factor is No_Septic_History. Overall utility in the THA model is a direct function of time spent in the combined states of these two factors.

Failure History	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 1: Functional status depends on Failure History and Septic History.

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

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

Table 2: The quality factor for Markovian utility depends on functional status. These are approximate values based on empirical work by ourselves and others.

The primary stochastic factor for the THA model is shown in Figure 8. Transitions in this factor affect utility by triggering transitions in the Failure History 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 clarify preference modeling assumptions and open the model to inspection and critique by others.

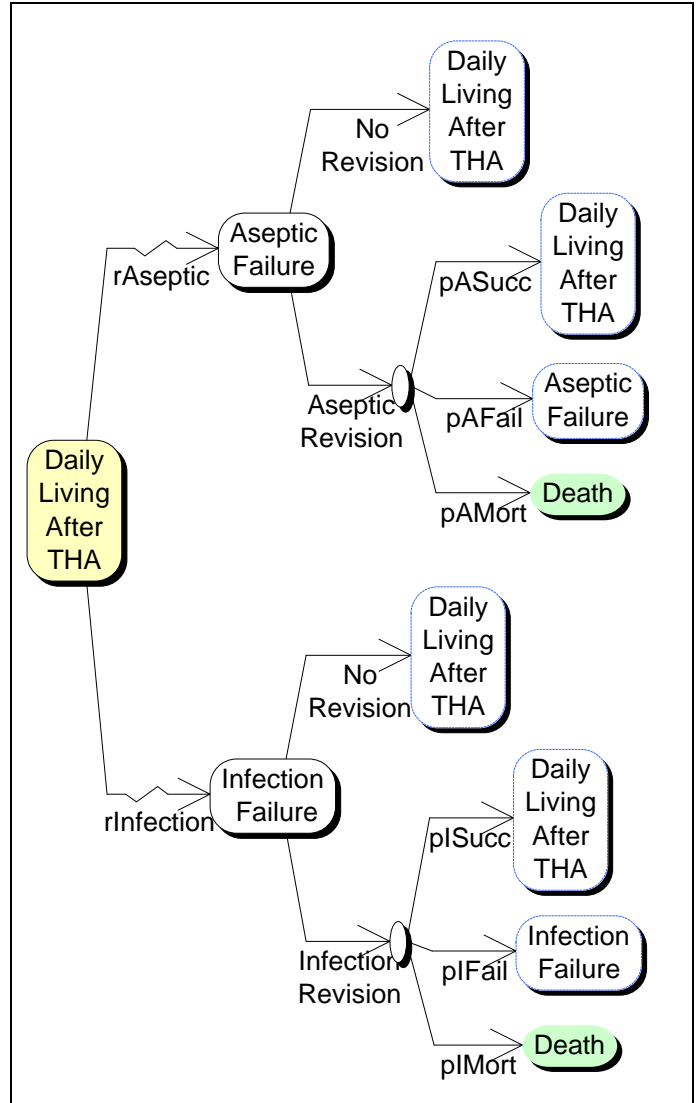


Figure 8: The primary stochastic factor for the THA model. Transition to *Infection_Failure* triggers a transition to *Unaddressed in the Most Recent Failure* factor, and a transition to *Septic_History* in the *Septic History* factor. Infection revision is performed if the total revision count (factor 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).

	Discounted quality-adjusted years	
	THA	No THA
White female age 60	13.70	6.82
White male age 85	4.16	2.16


Table 3: Comparison of THA versus No THA for two different individuals. An annual discount rate of 5% was used.

3. MULTIATTRIBUTE DECOMPOSITION OF MARKOVIAN UTILITY


In this section we examine utility decompositions which can arise when factors in a stochastic tree are treated as attributes in a multiattribute utility function. We consider only Markovian utility functions, noting that the results have implications for state-updateable utility as well because of the connection established by Theorem 3.

We consider the case of two-factor stochastic trees, since the multi-factor case follows naturally. We denote states in the first factor by y, y_0, y_1 and so on; and states in the second factor by z, z_0, z_1 and so on. Let y^* and z^* be distinguished states (e.g., Well) and for Markovian utility. We assume $v(y^*, z^*) > 0$.


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 to 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*

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

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

◆

Denote by x^1 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

$$u(x^1) = E\left[v(x) \int_0^T e^{-a(x)t} dt\right] = v(x)m(\lambda + a(x))$$


For constant hazard rate λ , the result is

$$u(x^1) = \frac{v(x)}{\lambda + a(x)}.$$

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, Hopp and Pellissier 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 $m = 1 + c$ for some nonnegative constant c).

The failure of full utility independence

In any multiattribute context, it is natural to attempt to invoke utility independence. 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. However, it is still not a useful definition because 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

$$\frac{\text{y} \xrightarrow{\lambda + \mu}}{\text{z} \xrightarrow{\mu}} = \frac{\text{y} \xrightarrow{\lambda}}{\text{z} \xrightarrow{\mu}} = \frac{\text{y}}{\text{z} \xrightarrow{\lambda + \mu}}$$

holding because all three are equal to the product tree

$$\text{y,z} \xrightarrow{\lambda + \mu}$$

Utility independence of factor 1 from factor 2 would imply, for example, that the indifference

$$\frac{\text{y} \xrightarrow{\lambda}}{\text{z} \xrightarrow{\mu}} \sim \frac{\begin{array}{c} \text{p} \rightarrow \text{y} \xrightarrow{\lambda^*} \\ 1 - \text{p} \rightarrow \end{array}}{\text{z} \xrightarrow{\mu}}$$

if holding for one z, m combination, must hold for all. In product tree terms, this indifference translates to

$$\text{y,z} \xrightarrow{\lambda + \mu} \sim \begin{array}{c} \text{p} \rightarrow \text{y,z} \xrightarrow{\lambda^* + \mu} \\ 1 - \text{p} \rightarrow \end{array}$$

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

$$\frac{1}{a(y,z) + l + m} = \frac{p}{a(y,z) + l^* + m}$$

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

Non-interfering 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 *non-interfering utility independence*. It would imply, for example, that the indifference

$$\frac{\text{y} \xrightarrow{\lambda}}{\text{z}} \sim \frac{\begin{array}{c} \text{p} \rightarrow \text{y} \xrightarrow{\lambda^*} \\ 1 - \text{p} \rightarrow \end{array}}{\text{z}}$$

if holding for one nonfatal z , would hold for all nonfatal z . In product form, this indifference is

$$\text{y,z} \xrightarrow{\lambda} \sim \begin{array}{c} \text{p} \rightarrow \text{y,z} \xrightarrow{\lambda^*} \\ 1 - \text{p} \rightarrow \end{array}$$

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 non-interfering 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 4: *Under Markovian utility each factor is non-interfering utility independent of the other if and only if there exist constants $v^* > 0$ and a^* , marginal quality-rate functions $v_1(y)$, $v_2(z)$, marginal tolls $w_1(y/y_0)$, $w_2(z/z_0)$, and marginal discount multipliers $Dw_1(y/y_0) > 0$, $Dw_2(z/z_0) > 0$ such that*

$$(i) a(y,z) = a^* \text{ independent of } y,z;$$

$$(ii) \frac{v(y,z)}{v^*} = \frac{v_1(y)}{v^*} \frac{v_2(z)}{v^*}$$

$$\text{where } v^* = v(y^*,z^*) = v_1(y^*) = v_2(z^*);$$

$$(iii) w(y,z/y_0,z) = w_1(y/y_0) \cdot v_2(z) \cdot v^*$$

$$w(y,z/y,z_0) = v_1(y) \cdot w_2(z/z_0) \cdot v^*;$$

$$(iv) Dw(y,z/y_0,z) = Dw_1(y/y_0)$$

$$Dw(y,z/y,z_0) = Dw_2(z/z_0).$$

◆

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

The fact that non-interfering 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 non-constant $a(y,z)$.

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 multi-factor decomposition of the stochastic tree. Then one may compute expected updateable-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 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 *non-interfering* utility independence leads to a useful utility decomposition. Unfortunately, non-interfering utility independence forces risk attitude to be a constant independent of health state, a severe restriction. We are currently aware of no realistic

decomposition for Markovian utility in which risk attitude is not constant.

REFERENCES

- R.W. Chang, J.M. Pellissier and G.B. Hazen, "A Cost-Effectiveness Analysis of Total Hip Arthroplasty for Osteoarthritis of the Hip", *Journal of the American Medical Association* **275** (1996) 858-865.
- Cox, D.R., "A use of complex probabilities in the theory of stochastic processes," *Proceedings of the Cambridge Philosophical Society*, **51** (1955) 313-319.
- C.A. Gottlob, J.M. Pellissier, R.L. Wixson, S.H. Stern, S.D. Stulberg, R.W. Chang, "The Long-Term Cost-Effectiveness of Total Knee Arthroplasty for Osteoarthritis", Multipurpose Arthritis Center, Northwestern University, Chicago, Illinois.
- W.H. Harris and C.B. Sledge, "Total Hip and Total Knee Replacement," *New England Journal of Medicine* **323** (1990) 725-731.
- G.B. Hazen, "Stochastic Trees: A New Technique for Temporal Medical Decision Modeling," *Medical Decision Making* **12** (1992) 163-178.
- G.B. Hazen, "Factored Stochastic Trees: A Tool for Solving Complex Temporal Medical Decision Models," *Medical Decision Making*, **13** (1993), 227-236.
- G.B. Hazen, W.J. Hopp and J.M. Pellissier, "Continuous-Risk Utility Assessment in Medical Decision Making," *Medical Decision Making* **11** (1991) 294-304.
- G.B. Hazen, M. Morrow and E.R. Venta, "A Decision Analysis for the Treatment of Ductal Carcinoma in Situ", in progress, IEMS Department, Northwestern University, Evanston Illinois.
- G.B. Hazen and J.M. Pellissier, "Recursive Utility for Stochastic Trees", *Operations Research* **44** (1996) 788-809.
- G.B. Hazen and J. Souderpandian, "State-Trajectory Preference Summaries for Stochastic Tree Rollback", July 1997.
- G.B. Hazen, J.M. Pellissier and J. Souderpandian, "Stochastic tree models in medical decision making", forthcoming 1998 in *Interfaces*.
- Keeney, R.L. and H. Raiffa.. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York, 1976.
- D.B. Matchar and S.G. Pauker, "Transient Ischemic Attacks in a Man with Coronary Artery Disease: Two Strategies Neck and Neck", *Medical Decision Making* **6** (1986) 239-249.
- R.F. Meyer, "Preferences over time", in Keeney and Raiffa (1976), Chapter 9.
- J.M. Pellissier, G.B. Hazen and R.W. Chang, "A Continuous-Risk Decision Analysis of Total Hip Replacement", *Journal of the Operations Research Society* **47** (1996) 776-793.
- J.M. Pellissier and G.B. Hazen, "Implementation of Continuous-Risk Utility Assessment: The Total Hip Replacement Decision", *Socio-Economic Planning Sciences* **28** (1994) 251-276.
- J. Tsevat, M.H. Eckman, R.A. McNutt, S.G. Pauker, "Warfarin for Dilated Cardiomyopathy: A Bloody Tough Pill to Swallow?", *Medical Decision Making* **9** (1989) 162-169.

APPENDIX: PROOFS OF THEOREMS

Proof of Theorem 3

From Hazen and Pellissier (1996), or from Theorem 2 using a Markovian preference summary, the generic rollback equation for Markovian utility is

$$Eu(H | x) = w(y | x) + \Delta w(y | x)Eu_0(H)$$

$$Eu_0(H) = \frac{v(y) + \sum_i \lambda_i Eu(K_i | y)}{a(y) + \sum_i \lambda_i}$$

where H is the stochastic fork in (1). For the Markovian utility u_M defined in the theorem, this becomes

$$\begin{aligned} Eu_M(H \otimes Q | x, q) \\ = w_M(y, r/x, q) + Dw_M(y, r/x, q)Eu_{0M}(H \otimes Q) \end{aligned}$$

$$\begin{aligned} Eu_{0M}(H \otimes Q) \\ = \frac{v_M(y, r) + \sum_i \lambda_i Eu_M(K_i \otimes Q | y, r)}{a_M(y, r) + \sum_i \lambda_i} \end{aligned}$$

For the updateable-state utility u of the theorem, we have from Theorem 2,

$$Eu(H | q) = w(y | q) + \Delta w(y | q)Eu_0(H | q)$$

$$\begin{aligned} Eu_0(H | q) \\ = \frac{v(q(q, y)) + \sum_i \lambda_i Eu(K_i | q(q, y))}{a(q(q, y)) + \sum_i \lambda_i} \end{aligned}$$

Let $r = q(q, y)$. By induction on the height of the tree, we can assume $Eu(K_i/r) = Eu_M(K_i \otimes Q/y, r)$.

Then the last equation becomes

$$\begin{aligned} Eu_0(H | q) &= \frac{v(r) + \sum_i \lambda_i Eu(K_i | r)}{a(r) + \sum_i \lambda_i} \\ &= \frac{v_M(y, r) + \sum_i \lambda_i Eu_M(K_i \otimes Q | y, r)}{a_M(y, r) + \sum_i \lambda_i} \\ &= Eu_{0M}(H \otimes Q) \end{aligned}$$

Therefore,

$$\begin{aligned} Eu(H | q) &= w(y | q) + \Delta w(y | q)Eu_0(H | q) \\ &= w_M(y, r/x, q) + Dw_M(y, r/x, q)Eu_{0M}(H \otimes Q) \\ &= Eu_M(H \otimes Q/x, q) \end{aligned}$$

as desired. For chance forks H the proof is similar.

◆

Equivalence of Markovian utility functions

Lemma 1: Two Markovian utility functions u and $u\mathfrak{C}$ with $u(\emptyset) = u\mathfrak{C}(\emptyset) = 0$ are strategically equivalent if and only if there is a positive constant c such that

$$\begin{aligned} w\mathfrak{C}(x/x_0) &= cw(x/x_0) & Dw\mathfrak{C}(x/x_0) &= Dw(x/x_0) \\ v'(x) &= cv(x) & a\mathfrak{C}(x) &= a(x) \end{aligned}$$

Proof: Because u and $u\mathfrak{C}$ share a common zero point, strategic equivalence forces $u\mathfrak{C} = cu$ for some positive c . Consider a stochastic tree $x^1 K$. Given Markovian preference state x_0 , expected utilities are

$$E[u(x^1 K|x_0)] = w(x/x_0) + Dw(x/x_0)v(x)m(a(x)+1) + (1 - a(x)m(a(x)+1))E[u(K/x)] \quad (1)$$

$$E[u(\mathbb{C}x^1 K|x_0)] = w(\mathbb{C}x/x_0) + Dw(\mathbb{C}x/x_0)v(\mathbb{C}x)m(a(\mathbb{C}x)+1) + (1 - a(\mathbb{C}x)m(a(\mathbb{C}x)+1))E[u(\mathbb{C}K/x)] \quad (2)$$

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

$$E[u(x^1 |x_0)] = w(x/x_0) + Dw(x/x_0)v(x)m(a(x)+1)$$

$$E[u(\mathbb{C}x^1 |x_0)] = w(\mathbb{C}x/x_0) + Dw(\mathbb{C}x/x_0)v(\mathbb{C}x)m(a(\mathbb{C}x)+1)$$

Then $u\mathbb{C} = cu$ implies

$$w(\mathbb{C}x/x_0) + Dw(\mathbb{C}x/x_0)v(\mathbb{C}x)m(a(\mathbb{C}x)+1) = c(w(x/x_0) + Dw(x/x_0)v(x)m(a(x)+1))$$

Set $l = \infty$ to get

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

Substitute this back into its predecessor and cancel to get

$$Dw(\mathbb{C}x/x_0)v(\mathbb{C}x)m(a(\mathbb{C}x)+1) = cDw(x/x_0)v(x)m(a(x)+1)$$

Let $l(t)$ be a constant l to conclude

$$Dw(\mathbb{C}x/x_0)v(\mathbb{C}x)/(a(\mathbb{C}x)+1) = cDw(x/x_0)v(x)/(a(x)+1)$$

The equality of two functions $k\mathbb{C}(a\mathbb{C}+1)$ and $k/(a+1)$ of $l \in [0, \infty)$ forces $k\mathbb{C} = k$ and $a\mathbb{C} = a$. Therefore we conclude

$$Dw(\mathbb{C}x/x_0)v(\mathbb{C}x) = cDw(x/x_0)v(x) \quad (4)$$

$$a(\mathbb{C}x) = a(x). \quad (5)$$

Substitute (3),(4),(5) back into (2), then use $u\mathbb{C} = cu$ and simplify to get

$$cE[u(x^1 K|x_0)] = cw(x/x_0) + cDw(x/x_0)v(x)m(a(x)+1) + Dw(\mathbb{C}x/x_0)(1 - a(x)m(a(x)+1))cE[u(K/x)]$$

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

$$cDw(x/x_0)(1 - a(x)m(a(x)+1))E[u(K/x)] = Dw(\mathbb{C}x/x_0)(1 - a(x)m(a(x)+1))cE[u(K/x)]$$

Cancel terms to conclude $Dw(x/x_0) = Dw(\mathbb{C}x/x_0)$.

Therefore from (4) we conclude $v(\mathbb{C}x) = cv(x)$. We have thus shown that strategic equivalence of u and $u\mathbb{C}$ imply the four equalities specified in Lemma 1.

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

$$E[u(\mathbb{C}x^1 K|x_0)] = cw(x/x_0) + Dw(x/x_0)(cv(x)m(a(\mathbb{C}x)+1) + (1 - a(\mathbb{C}x)m(a(\mathbb{C}x)+1))E[u(\mathbb{C}K/x)] \quad (6)$$

Inducting on the height of the stochastic tree, if we have $E[u(\mathbb{C}K/x)] = cE[u(K/x)]$, then from (6) we conclude $E[u(\mathbb{C}x^1 K|x_0)] = cE[u(x^1 K|x_0)]$. Hence $u\mathbb{C}$ and u are strategically equivalent. \blacklozenge

Proof of Theorem 4

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

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

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

$$w_1(y/y_0) = w(y, z^*/y_0, z^*)$$

$$w_2(z/z_0) = w(y^*, z/y^*, z_0)$$

$$Dw_1(y/y_0) = Dw(y, z^*/y_0, z^*)$$

$$\Delta w_2(z/z_0) = Dw(y^*, z/y^*, z_0)$$

and let $v^* = v(y^*, z^*)$. Suppose non-interfering utility independence holds. Consider a stochastic tree $(y^1 K, z)$. Its Markovian expected utility given preference summary (y_0, z) is

$$\begin{aligned} E[u(y^1 K, z|y_0, z)] &= w(y, z/y_0, z) + Dw(y, z/y_0, z) \\ &\quad (v(y, z)m(a(y, z)+1) + (1 - a(y, z)m(a(y, z)+1))) \\ &\quad E[u(K, z/y, z)] \end{aligned} \quad (7)$$

Substitute $z = z^*$ to get

$$\begin{aligned} E[u(y^1 K, z^*|y_0, z^*)] &= w_1(y/y_0) + Dw_1(y/y_0) \\ &\quad (v_1(y)m(a_1(y)+1) + (1 - a_1(y)m(a_1(y)+1))) \\ &\quad E[u(K, z^*/y, z^*)] \end{aligned} \quad (8)$$

Non-interfering 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

$$\begin{aligned} v(y, z) &= v_1(y)c_2(z) & a(y, z) &= a_1(y) \\ w(y, z/y_0, z) &= w_1(y/y_0)c_2(z) & Dw(y, z/y_0, z) &= Dw_1(y/y_0) \end{aligned}$$

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

$$\begin{aligned} v(y, z) &= c_1(y)v_2(z) & a(y, z) &= a_2(z) \\ w(y, z/y, z_0) &= c_1(y)w_2(z/z_0) & Dw(y, z/y, z_0) &= Dw_2(z/z_0) \end{aligned}$$

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

$$\begin{aligned} v(y, z) &= v_1(y)c_2(z) \\ &= v(y, z^*)c_2(z) \\ &= c_1(y)v_2(z^*)c_2(z) \\ &= c_1(y)c_2(z)v^* \end{aligned}$$

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 follows immediately.

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

$$\begin{aligned} E[u(y^1 K, z|y_0, z)] &= w_1(y/y_0)v_2(z)v^* + Dw_1(y/y_0) \\ &\quad (v_1(y)v_2(z)v^*m(a^*+1) + (1 - a^*m(a^*+1))) \\ &\quad E[u(K, z/y, z)] \\ &= (w_1(y/y_0) + Dw_1(y/y_0)(v_1(y)m(a^*+1))) v_2(z)v^* \\ &\quad + Dw_1(y/y_0) (1 - a^*m(a^*+1)) E[u(K, z/y, z)] \end{aligned}$$

and (8) becomes

$$\begin{aligned} E[u(y^1 K, z^*|y_0, z^*)] &= w_1(y/y_0) + Dw_1(y/y_0) \\ &\quad (v_1(y)m(a^*+1) + (1 - a^*m(a^*+1))) \\ &\quad E[u(K, z^*/y, z^*)] \end{aligned}$$

$$= w_I(y/y_0) + Dw_I(y/y_0)v_I(y)m(a^{*+1}) \\ + Dw_I(y/y_0)(1 - a^*m(a^{*+1}))E[u(K, z^*/y, z^*)]$$

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

$$E[u(y^1 K, z|y_0, z)] = E[u(y^1 K, z^*|y_0, z^*)] v_2(z)v^*$$

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

UNUSABLE MATERIAL

Marginal preference conditions

Even non-interfering utility independence is too strict, in that it forces $a(y,z)$ to be constant in y,z . Next we examine independence conditions under which $a(y,z)$ is not constant, but depends only on $a_1(y), a_2(z)$.

Assuming that risk attitude can be captured separately in each factor, an argument involving equivalent risk neutral preference presented in Figure 9 suggests that the factor risk-aversion coefficients $a_1(y)$ and $a_2(z)$ should combine additively to give $a(y,z)$.

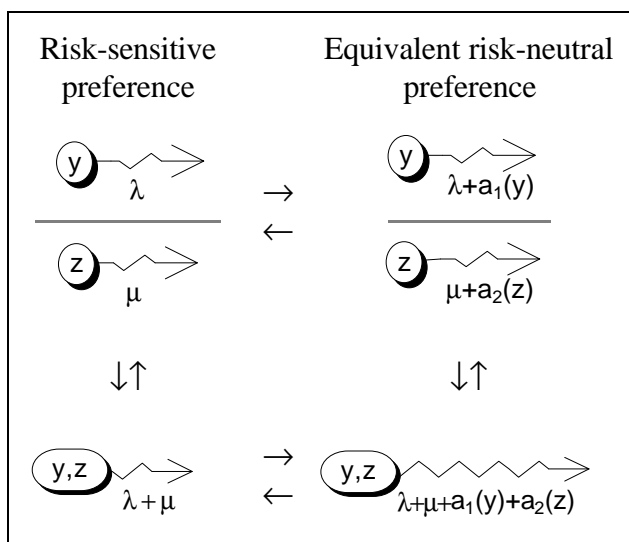


Figure 9: A heuristic argument involving risk-neutral preference suggests that the coefficient of risk aversion $a(y,z)$ should decompose additively.


However, this argument is only heuristic. Preference conditions which yield an additive

decomposition for $a(y,z)$ are given in the following theorem.


Theorem 5: Under Markovian utility, a necessary and sufficient condition for the relation

$$a(y,z) = a_1(y) + a_2(z)$$


is that



and



imply



for all $\lambda, \mu, \lambda_0, \mu_0$.

◆

Theorem 5 is not as satisfying as one would wish because of the special role played by the distinguished states y^* , z^* . If the condition of Theorem 5 holds for y^* , z^* , then why should it not hold for any two states $y \neq z$? But this broader assumption yields

$$a(y,z) = a(y,z^*) + a(y^*,z)$$

for all states y, z, y', z' and from this it follows that $a(y,z) = 0$ for all y, z . So caution is advised in the use of the condition of Theorem 5.