

RECURSIVE UTILITY FOR STOCHASTIC TREES

GORDON B. HAZEN

Northwestern University, Evanston, Illinois

JAMES M. PELLISSIER

Loyola University, Chicago, Illinois

(Received August 1993; revised March 1995; accepted October 1995)

Stochastic trees are semi-Markov processes represented using tree diagrams. Such trees have been found useful for prescriptive modeling of temporal medical treatment choice. We consider utility functions over stochastic trees which permit recursive evaluation in a graphically intuitive manner analogous to decision tree rollback. Such rollback is computationally intractable unless a low-dimensional preference summary exists. We present the *most general* classes of utility functions having specific tractable preference summaries. We examine three preference summaries—*memoryless*, *Markovian*, and *semi-Markovian*—which promise both computational feasibility and convenience in assessment. Their use is illustrated by application to a previous medical decision analysis of whether to perform carotid endarterectomy.

A *stochastic tree* is a graphical modeling approach which combines useful features from decision trees and semi-Markov process transition diagrams. This paper concerns itself with the recursive evaluation of utility functions over stochastic trees, that is, the calculation of an expected utility measure using iterative methods akin to the method of successive approximations (value iteration) in the stochastic dynamic programming literature (e.g., Bertsekas 1976, Ross 1983). Predecessors in this area include the works of Mitten (1974), Sobel (1975), and Sniedovich (1981) on preference-order dynamic programming; Howard and Matheson (1972), and Chung and Sobel (1987) on dynamic programming with an exponential utility function; and Kreps (1977a and b, 1978) on dynamic programming using arbitrary utility functions.

One of the most striking properties of the utility functions we present is that they may be evaluated over any particular stochastic tree using a rollback process akin to that used for decision trees. Such rollback is achievable *in principle* for an arbitrary utility function (a nonobvious fact on which we elaborate below), but is in general computationally intractable unless a low-dimensional *preference summary* exists. The essential contribution of this paper is to present the *most general* classes of utility functions having specific tractable preference summaries. We examine three preference summaries—*memoryless*, *Markovian*, and *semi-Markovian*—which promise both computational tractability and convenience in assessment. For each of these summaries, we derive the corresponding class of utility functions having that summary.

Our approach may be viewed as an addition to the decision analytic literature on utility over time streams (Fishburn 1965 and 1978, Meyer 1976, Bell 1977, Barrager 1980, Harvey 1988). Stochastic tree modeling was motivated

by applications to medical treatment decision making (Hazen 1992, 1993). This paper may therefore also be regarded as a contribution to the literature on health status indices (Torrance 1976, Bodily 1980, Pliskin et al. 1980, Torrance et al. 1982). One distinguishing aspect of our work is that in contrast to nearly all the literature mentioned above, we model time as a continuous rather than a discrete quantity. We feel that modeling in continuous rather than discrete time has significant advantages in formulation and presentation (Hazen 1992).

Because stochastic trees have only recently been introduced and are as yet not widely used, we open the paper in Section 1 with a brief introduction to the topic. Our main results are presented in Section 2. We discuss assessment issues and present an application in Section 3.

1. STOCHASTIC TREES

What is a stochastic tree? In its simplest and most useful form, a stochastic tree is merely a transition diagram for a continuous-time Markov chain, unfolded into a tree structure, and augmented by chance nodes. We have used stochastic trees as modeling tools for the analysis of medical treatment decisions (Hazen 1992 and 1993, Chang et al. 1993, Pellissier and Hazen 1994).

As an example, consider the stochastic tree of Figure 1, which is a model of nonsurgical treatment of transient ischemic attacks, based on Matchar and Pauker (1986). The usual conventions for transition diagrams apply: the nodes represent states (Well, Big Stroke, Post Big Stroke, and so on), and the arrows represent transitions between states. There are two types of arrows, corresponding to the two types of possible transitions. Wavy arrows are labeled

Subject classifications: Utility/preference, theory: utility over time streams. Utility preference/applications: medical treatment decisions. Decision analysis, theory: stochastic trees.

Area of review: DECISION ANALYSIS, BARGAINING AND NEGOTIATION.

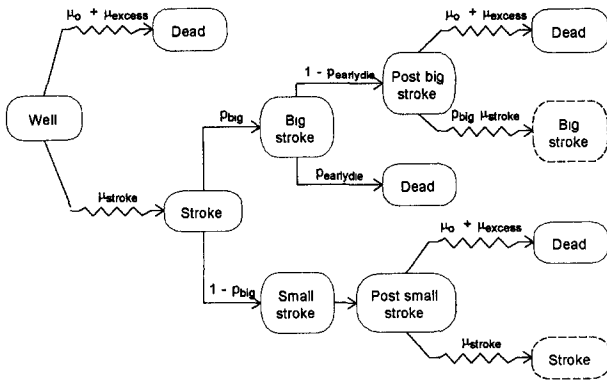



Figure 1. Stochastic tree model for nonsurgical treatment of transient ischemic attacks.

with rates, and signify transitions which take time to accomplish. For instance, from the initial state Well, there is an exponentially distributed duration with rate $\mu_0 + \mu_{excess}$ until transition to Dead, and an exponential (μ_{stroke}) duration until transition to Stroke. The next state occupied is determined by the shorter transition time, just as in any continuous-time Markov chain. Straight arrows are labeled with probabilities, and signify immediate transition to one of the states indicated. For example, from the state Stroke, there is an immediate transition to Big Stroke with probability p_{big} and to Small Stroke with probability $1 - p_{big}$. Wavy arrows emanate from *stochastic nodes*, and straight arrows from *chance nodes*.

Nodes with dashed-line borders indicate transitions to previously depicted states. We call these *phantom nodes*. For example, from the state Post Small Stroke, the next state visited is either Dead or the previously depicted state Stroke. Note that there are several states to which repeated visits are possible. When this occurs, we say the stochastic tree is *cyclic*. Otherwise it is *acyclic*.

Transformation of Stochastic Trees


There are several ways in which a stochastic tree can be transformed into a different but equivalent representation. One of the most useful is a consequence of the familiar superposition/decomposition rules for Poisson processes (e.g., Cinlar 1975, Ch. 4), and may be depicted as follows:




where

$$\lambda = \sum_y \lambda_y, \quad p_y = \lambda_y / \lambda.$$

The diagram indicates that a sojourn in state x having several competing exponential (λ_y) transitions leading to respective states y is equivalent to an exponential (λ) sojourn in x followed by a chance p_y of transition to y . A second elementary transformation involves the aggregation of several distinct transitions to the same state:



A third useful transformation is the elimination of self-transitions:



Its validity arises from the fact that a sum of geometrically many independent exponential durations is also exponential.

As an example, in the recurrent stroke model of Figure 1, the states Big Stroke and Small Stroke serve only descriptive purposes and can be eliminated. Once this is done, the resulting tree can be transformed using these




Figure 2. Acyclic stochastic tree equivalent to Figure 1.

rules into the acyclic tree of Figure 2. The details are left to the reader. (In Figure 2, the abbreviation μ_s has been used for μ_{stroke} , p_b has been used for p_{big} , and so forth.)


Quality-adjusted Duration and Stochastic Tree Rollback

One popular measure of treatment efficacy used in the medical decision making literature is *mean quality-adjusted duration*. This measure is calculated by weighting each interval of time spent in a particular health state x by a *quality factor* $v(x)$ proportional to the desirability of that state. So if T_x is the total duration for which state x is occupied, then one seeks the mean value of


$$Q = \sum_x v(x) T_x .$$

Typically the Well state is assigned quality factor 1, the Dead state 0, and other states of health are given intermediate values. Note that total lifetime and total time spent in a particular state are instances of quality-adjusted duration with $v(x) = 1$ for the appropriate set of states x and $v(x) = 0$ for the complementary set, so a broad range of conventional measures is subsumed. For models using mean quality-adjusted duration, see for example Weinstein and Stason (1976), Beck and Pauker (1981), Hillner et al. (1986), Plante et al. (1987), Roach et al. (1988), Mooney et al. (1990).

A further advantage of the stochastic tree model is that it allows the recursive computation of mean quality-adjusted duration by “rolling back” the stochastic tree (Hazen 1992), much as one would roll back a decision tree. To see how this works, consider a subtree of a stochastic tree in which an initial state x is occupied until one of several competing transitions with rate λ_y occurs to state y :



Invoking the superposition rule, we can transform this subtree into



where $p_y = \lambda_y/\lambda$ and $\lambda = \sum_y \lambda_y$. It is therefore apparent that beginning in x , a mean time $1/\lambda$ is spent in state x , following which transition to y occurs with probability p_y . Suppose the mean quality-adjusted duration beginning at y is $L(y)$. Mean quality-adjusted duration beginning in x is therefore

$$L(x) = v(x) \cdot \frac{1}{\lambda} + \sum_y p_y L(y) = \frac{v(x) + \sum_y \lambda_y L(y)}{\sum_y \lambda_y} .$$

This formula can be used to recursively evaluate mean quality-adjusted lifetime in any stochastic tree. Rollback for the recurrent stroke example of Figure 2 is depicted in Figure 3. (See Table 2 in Section 3 for the rate and probability values. This example is discussed further there.) Mean quality-adjusted durations are indicated in italics below the corresponding states. Nonsurgical treatment has a mean quality-adjusted lifetime of 9.325 years beginning in the Well state.

Rollback for stochastic trees with cycles may also be performed using the method of successive approximations (value iteration) from dynamic programming. Hazen (1992) gives several examples, and we present a computational example at the end of Section 2.

Risk-averse Preferences in Stochastic Trees

One drawback of the mean quality-adjusted duration measure is that it is risk neutral with regard to longevity in any particular health state. For example, consider a hypothetical choice between a vaccination which will reduce your mortality rate by $c\%$, but unfortunately may result in adverse reaction and immediate death with probability p .

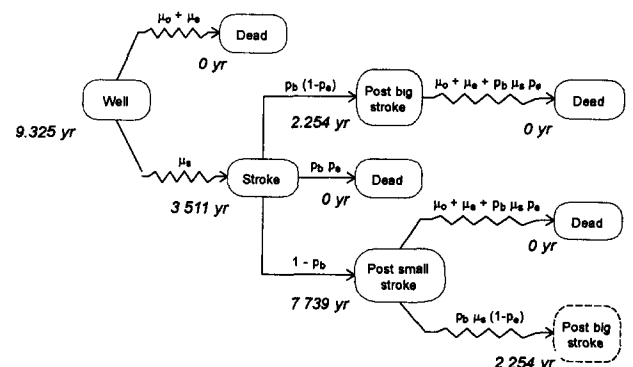



Figure 3. Calculating mean quality-adjusted duration by stochastic tree rollback.

What is the largest value of p you would accept? Using stochastic tree notation, we seek the value of p which produces the indifference




Call this indifference probability $p_{vaccine}$. The surprising fact is that if $p = c$, then both of these alternatives have mean lifetime $1/\mu$. Therefore according to the mean quality-adjusted lifetime criterion, one should take $p_{vaccine} = c$. If, for example, you are a 40-year-old white male, your annual mortality rate is approximately 2.5 per 1,000 (U.S. National Center for Health Statistics 1986). Would you take a 50% chance at immediate death to cut this mortality rate in half? Your answer would be yes if your sole criterion were mean quality-adjusted lifetime. However, most individuals would find this a ludicrous choice.

The obvious remedy is to replace quality-adjusted duration by a risk-averse utility function. However, it is unclear whether an arbitrary utility function possesses similar desirable recursive properties in stochastic trees. For example, suppose utility is given by


$$\text{Utility} = \left(\frac{\text{Quality adjusted}}{\text{duration}} \right)^\alpha$$

for $\alpha > 0$. How might one roll back a stochastic tree using this utility function? Here is a naive rollback procedure applied to the simple stochastic tree



in which sojourn times are certain. Assume state y has quality factor $v(y) = 1$. Obtain the final utility $U = (2 + p^{1/\alpha})^\alpha$ by adding the two-year sure duration in state y to the $p^{1/\alpha}$ -year certainty equivalent that follows, and then applying the utility function. Unfortunately, this procedure is incoherent: if one moves the chance node to beginning of

the tree and concatenates the two resulting one-year durations, one obtains the correct expected utility:



The lesson is that one cannot evaluate utility over stochastic trees in arbitrarily capricious ways.

Because recursive procedures offer computational advantages for stochastic trees which are large and/or cyclic, recursively computable utility functions would certainly be desirable. The purpose of this paper is to investigate whether recursively computable utility functions more general than quality-adjusted duration exist, and if so, how the recursion should be performed. The surprising answer given in the next section is that in spite of the example just given, recursive calculation may in fact be performed with any utility function. Kreps (1977) gives the first account of this of which we are aware (but in discrete time). The catch is that in general the recursion is too computationally demanding to be useful, unless there are convenient preference summaries (called state descriptors by Meyer 1976). Therefore, in the next section we also investigate recursive utility forms having natural and tractable preference summaries.

2. RECURSIVE UTILITY OVER STOCHASTIC TREES

Mathematical Overview

The stochastic tree is a continuous-time probability model. To convey a clearer picture of the mathematical contribution of this paper and relationships to previous work, we present in this section the discrete-time versions of our general recursive utility theorems. The nature of our results are clearer when presented by analogy in discrete time, where the mathematical and notational baggage is less burdensome. It is also easier to summarize relationships to prior contributions, since most of them were done in a discrete time context. Readers wishing to proceed immediately to our main results can skip this subsection.

Let X_1, X_2, X_3, \dots be a discrete-time stochastic process, and for each $t \geq 0$, let

$$X_{\leq t} = (X_1, \dots, X_t),$$

$$X_{> t} = (X_{t+1}, X_{t+2}, \dots),$$

be respectively the past at time t and the future at time t . Let the utility function $U(x_1, x_2, \dots)$ be defined over possible realizations of (X_1, X_2, \dots) . One way to evaluate

$E[U(X_1, X_2, \dots)]$ recursively without any assumption on the form of U is as follows. Let

$$\Phi_t(X_{\leq t}) = E_{X_{>t}}[U(X_{\leq t}, X_{>t})|X_{\leq t}],$$

be expected utility evaluated beyond stage t given history $X_{\leq t}$ up to time t . Then

$$\begin{aligned}\Phi_{t-1}(X_{\leq t-1}) &= E_{X_{>t-1}}[U(X_{\leq t-1}, X_{>t-1})|X_{\leq t-1}] \\ &= E_{X_{>t-1}}[U(X_{\leq t-1}, X_t, X_{>t})|X_{\leq t-1}] \\ &= E_{X_t}[E_{X_{>t}}[U(X_{\leq t-1}, X_t, X_{>t}) \\ &\quad \cdot |X_t, X_{\leq t-1}]|X_{\leq t-1}] \\ &= E_{X_t}[E_{X_{>t}}[U(X_{\leq t}, X_{>t})|X_{\leq t}]|X_{\leq t-1}] \\ &= E_{X_t}[\Phi_t(X_{\leq t})|X_{\leq t-1}].\end{aligned}$$

This recursion

$$\Phi_{t-1}(X_{\leq t-1}) = E_{X_t}[\Phi_t(X_{\leq t})|X_{\leq t-1}], \quad (1)$$

has been studied by Kreps (1978). It is computationally intractable without further simplifying assumptions, because the quantities $\Phi_t(x_{\leq t})$ must be evaluated for every possible history $x_{\leq t}$ up to time t . This is true even if the X_t are probabilistically independent.

The nature of the required simplifying assumptions can be seen when the recursion is expressed in terms of conditional utilities. Suppose $U(x_1^0, x_2^0, \dots) = 0$, and use the notation

$$U(x_{\leq t}, x_{>t}^0) = U(x_{\leq t}).$$

Let the conditional utility $U(x_{>t}|x_{\leq t})$ be strategically equivalent to $U(x_{\leq t}, x_{>t})$ with

$$U(x_{>t}^0|x_{\leq t}) = 0 \quad \text{for all } x_{\leq t}.$$

Then $U(x_{>t-1}|x_{\leq t-1})$ can be recursively expressed in terms of $U(x_{>t}|x_{\leq t})$. This is due to the fact that these two utility functions are strategically equivalent over $x_{>t}$:

$$\begin{aligned}U(x_{>t-1}|x_{\leq t-1}) &= U(x_t, x_{>t}|x_{\leq t-1}) \sim_{x_{>t}} U(x_{\leq t-1}, x_t, x_{>t}) \\ &= U(x_{\leq t}, x_{>t}) \sim_{x_{>t}} U(x_{>t}|x_{\leq t}).\end{aligned}$$

Therefore there is an affine relationship between $U(x_{>t-1}|x_{\leq t-1})$ and $U(x_{>t}|x_{\leq t})$, which may depend on the other variables $x_{\leq t-1}$ and x_t , that is, there are quantities $\Delta U(x_t|x_{\leq t-1}) > 0$ and $a(x_t|x_{\leq t-1})$ such that

$$\begin{aligned}U(x_{>t-1}|x_{\leq t-1}) &= a(x_t|x_{\leq t-1}) \\ &\quad + \Delta U(x_t|x_{\leq t-1})U(x_{>t}|x_{\leq t}).\end{aligned}$$

Set $x_{>t} = x_{>t}^0$ to conclude that

$$a(x_t|x_{\leq t-1}) = U(x_t, x_{>t}^0|x_{\leq t-1}) = U(x_t|x_{\leq t-1}),$$

and substitute into the previous relation to get

$$\begin{aligned}U(x_{>t-1}|x_{\leq t-1}) &= U(x_t|x_{\leq t-1}) \\ &\quad + \Delta U(x_t|x_{\leq t-1})U(x_{>t}|x_{\leq t}).\end{aligned} \quad (2)$$

The analog of the recursion (1) in terms of conditional utilities can be expressed if we let

$$\phi_t(X_{\leq t}) = E_{X_{>t}}[U(X_{>t}|X_{\leq t})|X_{\leq t}],$$

be the conditional utility of the future given history $X_{\leq t}$. Using (2) we obtain the conditional recursion

$$\begin{aligned}\phi_{t-1}(X_{\leq t-1}) &= E_{X_t}[U(X_t|X_{\leq t-1})|X_{\leq t-1}] \\ &\quad + E_{X_t}[\Delta U(X_t|X_{\leq t-1})\phi_t(X_{\leq t})|X_{\leq t-1}].\end{aligned} \quad (3)$$

This is the analog of Theorem 2.1 below.

Opportunities for computational simplification in (3) are more obvious. For example, suppose the X_t are probabilistically independent. Also suppose there is a low-dimensional preference summary $q = e(x_{\leq t-1})$ such that $U(x_t|x_{\leq t-1})$ and $\Delta U(x_t|x_{\leq t-1})$ depend on $x_{\leq t-1}$ only through q . If the preference summary is *updatable* in the sense that

$$e(x_{\leq t-1}, x_t) = \theta(x_t, e(x_{\leq t-1})),$$

for some function θ , then the recursion (3) simplifies to

$$\begin{aligned}\phi_{t-1}(q) &= E_{X_t}[U(X_t|q)] \\ &\quad + E_{X_t}[\Delta U(X_t|q)\phi_t(e(X_t, q))].\end{aligned}$$

This is the analog of Theorem 2.2 below.

The notion of preference summary was introduced by Meyer, who uses the term *state descriptor*. Updatability is also discussed by Meyer. Examples of updatable preference summaries are the *memoryless* summary $e(x_{\leq t}) = \emptyset$, in which conditional preference depends not at all on the past; and the *Markov* summary $e(x_{\leq t}) = x_t$, in which conditional preference depends only on the most recent state.

An important question is whether there are any utility functions at all which possess these preference summaries, and if so, what they are. For example, the existence of a memoryless preference summary is equivalent to the statement that the future is utility independent of the past. Meyer presents the most general class of utility functions having this property. Fishburn (1965) and Bell present utility functions with Markov preference summaries.

We investigate analogous existence questions below, but in the context of stochastic trees. We introduce the corresponding notions of memoryless, Markovian, and also semi-Markovian preference summaries appropriate for stochastic tree models, and identify the classes of utility functions over stochastic trees having these preference summaries.

Recursive utility has also been studied in the economics literature, where the motivation is descriptive rather than prescriptive, as here. Epstein and Zinn (1989) examine a recursive model in which there is a value function V assigning conditional utility $V(X_{>t}|x_{\leq t})$ to the stochastic process $X_{>t}$ given the realization $X_{\leq t} = x_{\leq t}$. V is postulated to obey the recursive rule

$$V(X_{>t}|x_{\leq t}) = W(x_t, \mu(V(X_{>t+1}|x_{\leq t}, X_{t+1}))). \quad (4)$$

Here μ is a *certainty equivalent* operator, assigning a certainty equivalent to the uncertain value $V(X_{>t+1}|x_{\leq t}, X_{t+1})$ (uncertain because X_{t+1} is uncertain); and W is an *aggregator*, combining current payoff x_t with certainty equivalent future payoffs. Overall utility is $V(X_{>0}) = V(X_{>0}|x_{\leq 0})$. Epstein and Zinn show the existence of solutions V to (2.4) under reasonable conditions. Duffie and Epstein (1992) extend (4) to continuous time.

For specific choices of W and μ , the recursion (4) includes the dynamic choice theory of Kreps and Porteus (1978), which in turn includes the expected utility recursion (1). The latter, to which we confine ourselves in this paper, has the simple aggregator $W(x, y) = U(y)$, where U is the utility function. Despite its suggestive name, the aggregator W has no connection to our notion of preference summary. Our results concerning utility functions with specific preference summaries have therefore no direct connection to the economics literature on recursive utility. Moreover, none of that work attempts to find structural preference conditions which speed recursive computations, which is our main focus.

Preliminaries

Let x, y, z denote *states* of a stochastic process. For a state x and a duration $t > 0$, let x^t denote the function defined by

$$x^t(\tau) = x \quad \text{if } 0 \leq \tau < t.$$

Then x^t is possible sample path of a stochastic process. If g is a sample path with domain $[0, s)$ and h is any other sample path, define the *concatenation* gh of g and h by

$$gh(t) = \begin{cases} g(t) & \text{if } 0 \leq t < s, \\ h(t - s) & \text{if } t \geq s. \end{cases}$$

Informally, gh is obtained by sliding h a distance s to the right and appending it to g . Note the visually familiar identity

$$x^s x^t = x^{s+t},$$

which is one motivation for the notation.

We follow the usual convention of designating random quantities by capital letters. So x^T denotes a sojourn of random duration T in state x , and X^T denotes a sojourn of random duration T in uncertain state X .

Stochastic Trees


The stochastic trees discussed in Section 1 were Markov processes with chance nodes added. In general we allow semi-Markov processes as well. Formally, we define a *stochastic tree diagram* to be a directed graph with a finite or countable number of nodes such that

- (a) each node in the diagram is either a *stochastic* node or a *chance* node;
- (b) arcs from stochastic nodes are labeled with absolutely continuous probability distributions over $[0, \infty)$;

- (c) arcs from chance nodes are labeled with probabilities summing to 1.

The *tree* terminology is meant to emphasize the analogy with decision trees. We do not formally require that the diagram be a tree, although it may always be converted to one by using phantom nodes, as illustrated in Section 1. A *stochastic tree* is the stochastic process suggested by the stochastic tree diagram. Formally, we replace each distribution in the diagram with an independent random duration having that distribution, and at each chance node we independently replace all probabilities by mutually exclusive, collectively exhaustive events having those probabilities. In the resulting process, transition from a stochastic node x proceeds along the arc from x whose duration is the minimum, and transition from a chance node y proceeds along the arc from y whose event occurs. Any stochastic tree is a continuous-duration semi-Markov process and any continuous-duration semi-Markov process can be represented by a stochastic tree diagram. (The requirement that all durations have densities is meant to insure that ties have probability zero in determining the minimum-duration arc out of a given node. The density requirement may be relaxed as long as the probability of a tie remains zero.) We shall be interested only in stochastic trees which are really stochastic, that is, which cannot visit only chance nodes. Formally, we require that every chance node leads with positive probability to some stochastic node.

Consider a stochastic tree diagram H' of the form



where x_1, \dots, x_n are states, $\lambda_1, \dots, \lambda_n$ are distributions and G' is an arbitrary stochastic tree diagram. The realization of H' as a stochastic tree is

$$H = x_1^{T_1} \dots x_n^{T_n} G = \mathbf{x}^T G,$$

where $\mathbf{T} = (T_1, \dots, T_n)$ are independent durations with distributions $\lambda = (\lambda_1, \dots, \lambda_n)$, and $\mathbf{x} = (x_1, \dots, x_n)$. We therefore write

$$H' = x_1^{\lambda_1} \dots x_n^{\lambda_n} G' = \mathbf{x}^\lambda G'.$$

We refer to H as a stochastic tree *with trunk* \mathbf{x}^T , and to H' as a stochastic tree diagram *with trunk* \mathbf{x}^λ . However, in the sequel, we will often for simplicity drop the distinction between stochastic tree *diagrams* and stochastic trees. We will sometimes refer to \mathbf{x}^λ as a *history*. We allow the empty history \emptyset , with the property that $\emptyset G = G$ for any stochastic tree G .

We denote the convolution of two distributions λ and μ by $\lambda * \mu$. If S and T are independent durations with respective distributions λ and μ , then $S + T$ has distribution

$\lambda * \mu$. Therefore the two stochastic tree diagrams $x^\lambda x^\mu$ and $x^{\lambda * \mu}$, which are distinct as diagrams, represent the same stochastic tree, and we simply write $x^\lambda x^\mu = x^{\lambda * \mu}$.

Preference over Stochastic Trees

We assume that preference over stochastic trees is represented by a utility function U , so that for two stochastic trees H and G

$$H > G \Leftrightarrow E[U(H)] > E[U(G)].$$

The domain of U is the set of all semi-Markov sample paths h .

For stochastic trees $x^\lambda G$ we write $U(x^\lambda G)$ to mean $E_{\mathbf{T}}[U(x^\lambda \mathbf{T}G)]$ where \mathbf{T} has distribution λ . In addition, define a conditional preference relation over stochastic trees G, H by

$$G > H \text{ given } x^\lambda \Leftrightarrow x^\lambda G > x^\lambda H.$$

In other words, conditional preference over stochastic trees H given x^λ coincides with unconditional preference over trees with trunk x^λ . It therefore has an expected utility representation, with utility function which we denote by $U(h|x^\lambda)$. The latter is by definition strategically equivalent (as a function of h) to $U(x^\lambda h)$. Conditional preference given the empty history \emptyset is the same as unconditional preference, and $U(h|\emptyset) = U(h)$.

We introduce a special sample path ϕ to which all conditional and unconditional utility functions $U(\cdot | x^\lambda)$ assign utility zero:

$$U(\phi|x^\lambda) = 0.$$

In medical models, it is convenient to take ϕ to be the sample path in which the dead state is occupied forever. However, any convenient sample path is acceptable. If y^μ is any history, we will write $U(y^\mu h|x^\lambda)$ to mean $E_S[U(y^\mu h|x^\lambda)]$, where S has distribution μ . We will abbreviate $U(y^\mu \phi)$ by $U(y^\mu)$. We say that preference over stochastic trees is *nontrivial* if for every history x^λ there is a sample path h such that $x^\lambda h > x^\lambda \phi$ or $x^\lambda h < x^\lambda \phi$.

Theorem 1. Assume nontriviality holds and $U(\phi|x^\lambda) = 0$ for all x^λ . Then there are functions $\Delta U(y^\mu|x^\lambda) > 0$ such that an affine restriction

$$\Delta U(y^\mu|x^\lambda)U(h|x^\lambda y^\mu) = \int_0^\infty \Delta U(y^s|x^\lambda)U(h|x^\lambda y^s)d\mu(s), \tag{5}$$

holds for all histories h , a concatenation restriction

$$\Delta U(y^{\mu * \nu}|x^\lambda) = \Delta U(y^\mu|x^\lambda)\Delta U(y^\nu|x^\lambda y^\mu), \tag{6}$$


holds, and a recursive equation

$$U(y^\mu h|x^\lambda) = U(y^\mu|x^\lambda) + \Delta U(y^\mu|x^\lambda)U(h|x^\lambda y^\mu), \tag{7}$$

holds for all sample paths h .


The recursive Equation (7) allows stochastic tree rollback for an arbitrary utility function U . To see how this

works for acyclic trees G with a single root node, consider an arbitrary stochastic node y in G :



where we let H be the subtree of G with root node y . Since G is acyclic with a single root, there is a unique path $x_1^{\lambda_1} \dots x_n^{\lambda_n} = x^\lambda$ from the root node x_1 to y . Attach expected utility $E_H[U(H|x^\lambda)]$ to node y . We show below how this expected utility can be calculated as a function of the expected utilities attached to the direct successors of y . At chance nodes z , calculate expected utility by averaging utilities at the direct successors of z , just as in a decision tree. These calculations repeated recursively from right to left in the tree results in $E[U(G)]$ at the root node of G . Here are the details:

We begin by converting the fork at y to the equivalent form



where the new states y_i represent the information “state y is occupied and fork i is taken.” If T_i are the random durations corresponding to the distributions λ_i , then $p_i = P(T_i = \min_j T_j)$ is the probability that the i th fork is taken, and μ_i is the conditional distribution of T_i given $T_i = \min_j T_j$. Then $H = Y^\delta K$ is a stochastic tree in which:

- (i) an uncertain state $Y = y_i$ is first visited with probability p_i ,
- (ii) an uncertain sojourn S in Y next occurs, where S has distribution μ_i given $Y = y_i$,
- (iii) stochastic tree K is entered, where K is identical in distribution to the tree K_i given $Y = y_i$ (so K is conditionally independent of S given Y).

This situation is depicted in Figure 4, along with a proposed rollback procedure.




Figure 4. Stochastic tree rollback for arbitrary U .

This rollback procedure can be derived algebraically as follows. We have

$$\begin{aligned}
 E_H[U(H|\mathbf{x}^\lambda)] &= E_Y[E_{S,K}[U(Y^S K|\mathbf{x}^\lambda)|Y]] \\
 &= \sum_i p_i E_{S,K}[U(y_i^S K|\mathbf{x}^\lambda)|Y = y_i] \\
 &= \sum_i p_i E_{K_i}[U(y_i^{\mu_i} K_i|\mathbf{x}^\lambda)] \\
 &\quad (\text{because } K, S \text{ are independent given } Y = y) \\
 &= \sum_i p_i E_{K_i}[U(y_i^{\mu_i}|\mathbf{x}^\lambda) \\
 &\quad + \Delta U(y_i^{\mu_i}|\mathbf{x}^\lambda)U(K_i|\mathbf{x}^\lambda y_i^{\mu_i})] \\
 &\quad (\text{recursive equation}) \\
 &= \sum_i p_i (U(y_i^{\mu_i}|\mathbf{x}^\lambda) \\
 &\quad + \Delta U(y_i^{\mu_i}|\mathbf{x}^\lambda) E_{K_i}[U(K_i|\mathbf{x}^\lambda y_i^{\mu_i})]).
 \end{aligned}$$

According to this last equation, the conditional expected utility $E_H[U(H|\mathbf{x}^\lambda)]$ at node y may be calculated recursively from the conditional expected utilities $E_{K_i}[U(K_i|\mathbf{x}^\lambda y_i^{\mu_i})]$. The latter are precisely the utilities attached to the successors of y . A rollback procedure for single-root acyclic trees G has therefore been demonstrated. For a single-root cyclic tree G , rollback may be performed by the usual successive approximation of G by finitely deep acyclic trees G_n . We present an example at the end of this section.

For the most general types of stochastic trees, this recursion may be intractable, because the probabilities p_i and distributions μ_i must be calculated from the distributions λ_i at each node. However, for special cases such as exponential distributions or Weibull distributions with common shape parameter, the calculation may be done in closed form. We assume that the λ_i have been chosen so as to facilitate this calculation.

The recursion will, however, still be intractable without further restrictions on the utility structure. The reason is that the recursion requires the calculation at state y of a conditional utility given every possible path through the stochastic tree leading to y , a computationally intractable task for cyclic stochastic trees, which are, in effect, infinitely deep. For acyclic trees, the recursion is equivalent to the direct calculation of the unconditional expected utility of the entire tree by successively applying conditional expectation from right to left. Therefore, although this calculation technically qualifies as a recursion, it carries with it none of the conceptual or computational advantages usually associated with recursive calculation.

Therefore, in answer to the fundamental question raised in Section 1, stochastic tree rollback may be performed with an arbitrary utility function. However, the rollback algorithm is not a practical one for cyclic trees. Nevertheless, in special cases the utility structure may allow computational savings. We discuss how this may occur next.




Figure 5. Stochastic tree rollback when there is an updatable preference summary q .

Preference Summaries

Potential computational simplifications are available in stochastic tree rollback if the quantities $U(\cdot|\mathbf{x}^\lambda)$ and $\Delta U(\cdot|\mathbf{x}^\lambda)$ depend on \mathbf{x}^λ only through a low dimensional preference summary $q = e(\mathbf{x}^\lambda)$ of \mathbf{x}^λ . To be useful in a recursion such as (7), the summary should be *updatable*, that is, if $q = e(\mathbf{x}^\lambda)$ then $e(\mathbf{x}^\lambda y^\mu)$ should be obtainable as some function of q and y^μ .

We therefore make the following definitions. A *preference summary* is a function which assigns *preference states* q to histories \mathbf{x}^λ in such a way that if $e(\mathbf{x}^\lambda) = e(\mathbf{z}^\nu)$ then conditional preferences given \mathbf{x}^λ and given \mathbf{z}^ν are identical. We say that e is an *updatable* preference summary if there is a function θ such that $e(\mathbf{x}^\lambda y^\mu) = \theta(e(\mathbf{x}^\lambda), y^\mu)$.

When there is an updatable preference summary, the results of Theorem 1 take the following form.

Theorem 2. Assume nontriviality holds and $U(\phi|\mathbf{x}^\lambda) = 0$ for all \mathbf{x}^λ . If U has an updatable preference summary (e, θ) , then there is a summary utility function $u(h|q)$ on sample paths h given preference states q such that $u(\phi|q) = 0$ and

$$u(\cdot|q) \sim U(\cdot|\mathbf{x}^\lambda) \text{ whenever } e(\mathbf{x}^\lambda) = q.$$

Moreover, there are summary discount factors $\Delta u(y^\mu|q) > 0$ satisfying the affine restriction

$$\begin{aligned}
 \Delta u(y^\mu|q)u(h|\theta(q, y^\mu)) \\
 = \int_0^\infty \Delta u(y^s|q)u(h|\theta(q, y^s))d\mu(s),
 \end{aligned} \tag{8}$$

for all sample paths h , the concatenation restriction

$$\Delta u(y^{\mu*\nu}|q) = \Delta u(y^\mu|q)\Delta u(y^\nu|\theta(q, y^\mu)), \tag{9}$$

and the recursive equation

$$u(y^\mu h|q) = u(y^\mu|q) + \Delta u(y^\mu|q)u(h|\theta(q, y^\mu)), \tag{10}$$

for all sample paths h .

Figure 5 illustrates rollback when there is an updatable preference summary. Computational savings can occur because conditional utilities need be calculated only for each of the preference summaries that could occur at a given state y , rather than each of the histories that could precede y . For cyclic trees the number of preceding histories can be

arbitrarily large. However, for properly chosen utility functions, the number of preceding preference summaries may be small even for cyclic trees (see the computational example in the concluding subsection below). We discuss such utility functions in the next three subsections.

Memoryless Preference Summaries

Stochastic tree rollback is particularly simple with quality-adjusted duration as the utility measure, as we saw in Section 1. What preference summary is associated with this utility function? The quality-adjusted duration expected utility measure is

$$U(\mathbf{x}^\lambda) = U(x_1^{\lambda_1} \dots x_n^{\lambda_n}) = \sum_i v(x_i) m_i,$$

where $m_i = \int_0^\infty t \cdot d\lambda_i(t)$ is the mean of the distribution λ_i . Then

$$U(y^s | \mathbf{x}^\lambda) \sim_y U(y^s \mathbf{x}^\lambda) = v(y) \cdot s + \sum_i v(x_i) m_i \sim_y v(y) \cdot s.$$

By inspection, we see that conditional utility given \mathbf{x}^λ depends not at all on \mathbf{x}^λ , that is, it is *memoryless*. Therefore, the preference summary can be taken to be the constant function $e(\mathbf{x}^\lambda) = \emptyset$. This preference summary is (trivially) updatable. The summary utility function $u(\cdot | q)$ of Theorem 2 is

$$u(y^s | \emptyset) = u(y^s) = v(y) \cdot s.$$

By inspection, the summary discount factor $\Delta u(\cdot | q)$ of Theorem 2.2 is $\Delta u(y^s | \emptyset) = \Delta u(y^s) = 1$.

Are there other utility functions over stochastic trees having memoryless preference summaries, and if so, what are they? In the following, we say that utility is *continuous in its duration arguments* if for all n , $U(x_1^{\lambda_1} \dots x_n^{\lambda_n})$ is continuous as a function of $t_1 > 0, \dots, t_n > 0$. We also use the notation $[\frac{\alpha}{\beta}]$ to represent the affine function with slope β and intercept α . In other words

$$\left[\frac{\alpha}{\beta} \right](t) = \alpha + \beta t.$$

Moreover, if f and g are functions, let the symbol $f \circ g$ denote functional composition of f and g , so that $(f \circ g)(t) = f(g(t))$.

Theorem 3. *Let U be a utility function over stochastic trees which is continuous in its duration arguments, satisfies non-triviality and has $U(\phi | \mathbf{x}^\lambda) = 0$ for all \mathbf{x}^λ . Then U is memoryless if and only if there are functions $v(\cdot)$ and $a(\cdot)$ such that*

$$U(\mathbf{x}^t) \sim \left[\frac{u(x_1^{t_1})}{\Delta u(x_1^{t_1})} \right] \circ \dots \circ \left[\frac{u(x_n^{t_n})}{\Delta u(x_n^{t_n})} \right](0),$$

where

$$u(x^t) = \int_0^t v(x) e^{-a(x)s} ds, \quad \Delta u(x^t) = e^{-a(x)t}.$$

Theorem 4. *Under the assumptions of the previous theorem, if U has a memoryless preference summary, then its summary utility and discount factors are given for $t \geq 0$ by*

$$u(y^t | \emptyset) = u(y^t) = \int_0^t v(y) e^{-a(y)s} ds,$$

$$\Delta u(y^t | \emptyset) = \Delta u(y^t) = e^{-a(y)t}.$$

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

$$\Delta u(y^\mu) = \int_0^\infty \Delta u(y^s) d\mu(s) = \int_0^\infty e^{-a(y)s} d\mu(s).$$

For utility functions with memoryless preference summaries, the recursive equations (10) become

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

Memoryless utility is therefore identical to quality-adjusted duration with the future discounted at the state-dependent rate $a(y)$ when state y is occupied. When $a(y) = 0$ for all y , we recover the pure quality-adjusted duration form. When $a(y) \neq 0$ we obtain

$$u(y^t) = \int_0^t v(y) e^{-a(y)s} ds = \frac{v(y)}{a(y)} (1 - e^{-a(y)t}).$$

The summary u may therefore be regarded as a utility function having constant risk attitude $a(y)$ for durations spent in state y .

It is interesting to calculate $u(y^\lambda)$ when λ represents an exponential (λ) distribution convenient for stochastic tree modeling. Here λ is the rate of departure from state y . We obtain

$$u(y^\lambda) = \frac{v(y)}{a(y)} \int_0^\infty (1 - e^{-a(y)t}) \lambda e^{-\lambda t} dt = \frac{v(y)}{a(y) + \lambda}, \quad (11)$$

provided $a(y) + \lambda > 0$ (and equal to $+\infty$ if $a(y) + \lambda \leq 0$). In the risk-neutral case $a(y) = 0$, we obtain the quality-adjusted duration $u(y^\lambda) = v(y) \cdot \lambda^{-1}$. When $a(y) \neq 0$, $u(y^\lambda)$ may still be regarded as quality-adjusted duration, but with departure rate incremented by the subjective term $a(y)$. A risk averter in effect perceives the departure rate from y to be larger than its true value, and a risk seeker perceives it smaller.

One difficulty with the combination of memoryless utility and exponential durations is that $u(y^\lambda)$ diverges to $+\infty$ when $a(y) + \lambda \leq 0$. Therefore only moderately risk-seeking preference ($-\lambda \leq a(y) \leq 0$) may be portrayed.

Markovian Preference Summaries

The simplest preference summary which has some memory of its history remembers only the most recent state visited:

$$e(x_1^{\lambda_1} \dots x_n^{\lambda_n}) = x_n.$$

We call this the *Markovian* preference summary. For completeness, set $e(\emptyset) = \emptyset$. This is an updatable preference summary with

$$\theta(x, y^\mu) = y.$$

Call the utility functions having this preference summary the *Markovian* utility functions. Memoryless utility is Markovian because preferences depending on the null portion of the preceding history depend (vacuously) on the last state in the history. It follows that the class of Markovian utility functions includes the memoryless utility functions and possibly more. We describe exactly how much more in the following theorem.

Theorem 5. *Let U be a utility function over stochastic trees which is continuous in its duration arguments, satisfies nontriviality and has $U(\phi|x^\lambda) = 0$ for all x^λ . Then U has a Markovian preference summary if and only if there are functions $w(y|x)$, $\Delta w(y|x) > 0$, $v(y)$ and $a(y)$ such that*

$$U(x^t) \sim \left[\begin{array}{c} u(x_1^t | \emptyset) \\ \Delta u(x_1^t | \emptyset) \end{array} \right] \circ \left[\begin{array}{c} u(x_2^t | x_1) \\ \Delta u(x_2^t | x_1) \end{array} \right] \\ \circ \dots \circ \left[\begin{array}{c} u(x_n^t | x_{n-1}) \\ \Delta u(x_n^t | x_{n-1}) \end{array} \right] (0),$$

where

$$u(y^t|x) = w(y|x) + \Delta w(y|x) \int_0^t v(y)e^{-a(y)s} ds,$$

$$\Delta u(y^t|x) = \Delta w(y|x)e^{-a(y)t},$$

for some functions $w(y|x)$, $\Delta w(y|x) > 0$, $v(y)$, $a(y)$ with $w(y|y) = 0$, $\Delta w(y|y) = 1$.

Theorem 6. *Under the assumptions of the previous theorem, if U has a Markovian preference summary, then its summary utility and discount factors are given for $t \geq 0$ by*

$$u(y^t|x) = w(y|x) + \Delta w(y|x) \int_0^t v(y)e^{-a(y)s} ds,$$

$$\Delta u(y^t|x) = \Delta w(y|x)e^{-a(y)t}.$$

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

$$\Delta u(y^\mu|x) = \int_0^\infty \Delta u(y^s|x) d\mu(s) \\ = \Delta w(y|x) \int_0^\infty e^{-a(y)s} d\mu(s).$$

Markovian utility extends memoryless utility by adding the functions $w(y|x)$ and $\Delta w(y|x)$. The former may be interpreted as an instantaneous *toll* or *bonus* accrued in the transition from x to y . The latter is a *discount factor*, in that all utility subsequent to the transition from x to y is increased or decreased by that factor. (Roach et al. use such transition-induced discount factors to model morbidly effects.) The quantities $v(y)$ and $a(y)$ have the same interpretation as under the memoryless preference summary. In particular, risk attitude for durations spent in state y is constant as well under Markovian preference summaries.

Semi-Markovian Preference Summaries

One possible generalization of the Markovian preference summary would be a preference summary which remembers not only the last state visited, but also the associated duration (or distribution thereof). We call this the *semi-Markov* preference summary. It is formally defined as follows: For nonzero $\lambda_1, \lambda_2, \dots, \lambda_n$

$$e(x_1^{\lambda_1} \dots x_n^{\lambda_n}) = x_k^{\lambda_k} \dots \lambda_n \quad \text{if } x_{k-1} \neq x_k = \dots = x_n.$$

Note that the preference summary must keep track of whether previous states in the history are identical to the last one, in which case the distribution of time spent in that state is the convolution $\lambda_k * \dots * \lambda_n$ of the distributions λ_i of the previous identical states. For completeness, set $e(\emptyset) = \emptyset$. This is also an updatable preference summary, with

$$\theta(x^\lambda, y^\mu) = \begin{cases} y^\mu, & y \neq x, \\ x^{\lambda * \mu}, & y = x. \end{cases}$$

The corresponding class of semi-Markovian utility functions includes the Markovian utility functions but is considerably broader than that class, as the next two results show.

Theorem 7. *Assume nontriviality holds and $U(\phi|x^\lambda) = 0$ for all x^λ . Then U has a semi-Markovian preference summary if and only if there are functions $\Delta w(x^\lambda) > 0$, $w(y^\mu)$, and for $x \neq y$, functions $\Delta w(y|x^\lambda) > 0$, $w(y|x^\lambda)$ with*

$$w(y|\emptyset) = 0, \quad \Delta w(y|\emptyset) = 1,$$

such that the affine restrictions

$$w(y^\mu) = \int_0^\infty w(y^t) d\mu(t), \tag{12}$$

$$\Delta w(x^\lambda)w(y|x^\lambda) = \int_0^\infty \Delta w(x^t)w(y|x^t) d\lambda(t), \tag{13}$$

$$\Delta w(x^\lambda)\Delta w(y|x^\lambda) = \int_0^\infty \Delta w(x^t)\Delta w(y|x^t) d\lambda(t), \tag{14}$$

hold, and for $x_1 \neq x_2 \neq \dots \neq x_n$

$$U(x_1^{\lambda_1} x_2^{\lambda_2} \dots x_n^{\lambda_n}) = \left[\begin{array}{c} w(x_1 | \emptyset) \\ \Delta w(x_1 | \emptyset) \end{array} \right] \circ \left[\begin{array}{c} w(x_1^{\lambda_1}) \\ \Delta w(x_1^{\lambda_1}) \end{array} \right] \\ \circ \dots \circ \left[\begin{array}{c} w(x_n | x_{n-1}^{\lambda_{n-1}}) \\ \Delta w(x_n | x_{n-1}^{\lambda_{n-1}}) \end{array} \right] \circ \left[\begin{array}{c} w(x_n^{\lambda_n}) \\ \Delta w(x_n^{\lambda_n}) \end{array} \right] (0). \tag{15}$$

Theorem 8. *Under the assumptions of the previous theorem, if U has a semi-Markovian preference summary, then its summary utility and discount factors are given by*

$$u(y^\nu|x^\lambda) = \begin{cases} w(y|x^\lambda) + \Delta w(y|x^\lambda)w(y^\nu), & y \neq x, \\ \frac{w(y^{\lambda * \nu}) - w(y^\lambda)}{\Delta w(y^\lambda)}, & y = x, \end{cases} \tag{16}$$

and

$$\Delta u(y^\nu|x^\lambda) = \begin{cases} \Delta w(y|x^\lambda)\Delta w(y^\nu), & y \neq x, \\ \frac{\Delta w(y^{\lambda * \nu})}{\Delta w(y^\lambda)}, & y = x. \end{cases} \tag{17}$$

To gain some insight into the nature of semi-Markov utility, use the results of the last theorem in the recursive equation (10) to obtain for $y \neq x$

$$\begin{aligned} u(y^v h | x^\lambda) &= u(y^v | x^\lambda) + \Delta u(y^v | x^\lambda) u(h | y^v) \\ &= w(y | x^\lambda) + \Delta w(y | x^\lambda) w(y^v) \\ &\quad + \Delta w(y | x^\lambda) \Delta w(y^v) u(h | y^v) \\ &= \begin{bmatrix} w(y | x^\lambda) \\ \Delta w(y | x^\lambda) \end{bmatrix} \circ \begin{bmatrix} w(y^v) \\ \Delta w(y^v) \end{bmatrix} (u(h | y^v)). \end{aligned} \quad (18)$$

From this we see that a single semi-Markov recursion consists of two stages: first the subsequent utility $u(h | y^v)$ is discounted by the factor $\Delta w(y^v)$ and the result is incremented by the utility $w(y^v)$; second, this result is in turn discounted by the factor $\Delta w(y | x^\lambda)$, after which the bonus/toll $w(y | x^\lambda)$ is added. The first stage uses the state y and its duration v , but not the preference summary x^λ ; the second stage uses the state y without its duration v , and the preference summary x^λ .

On the other hand, when $y = x$, the recursive equation (10) becomes

$$\begin{aligned} u(y^v h | y^\lambda) &= u(y^v | y^\lambda) + \Delta u(y^v | y^\lambda) u(h | y^{\lambda * v}) \\ &= \frac{w(y^{\lambda * v}) - w(y^\lambda)}{\Delta w(y^\lambda)} + \frac{\Delta w(y^{\lambda * v})}{\Delta w(y^\lambda)} u(h | y^{\lambda * v}). \end{aligned}$$

Semi-Markovian utility must distinguish the cases $y \neq x$ and $y = x$ in order to keep utility recursion consistent with concatenation. In essence, this is the price paid for departure from exponential utility (the Markovian case), which is automatically consistent with concatenation.

A particularly convenient form of semi-Markov utility arises when it is assumed that the tolls $w(y | x^\lambda)$ and discount factors $\Delta w(y | x^\lambda)$ are *duration-independent*, that is, $w(y | x^\lambda) = w(y | x)$ and $\Delta w(y | x^\lambda) = \Delta w(y | x)$. In this case, these quantities cancel out of the affine restrictions (13) and (14), which reduce to

$$\Delta w(x^\lambda) = \int_0^\infty \Delta w(x^t) d\lambda(t). \quad (19)$$

Semi-Markov utility is in this case completely determined by $w(y | x)$, $\Delta w(y | x)$, and the quantities $w(y^t)$ and $\Delta w(y^t)$, for constant durations t .

Representations for Semi-Markovian Utility

The semi-Markovian class of utilities broadens the class of Markovian utilities by replacing exponential utility and exponential discount factors with utilities $w(y^t)$ and discount factors $\Delta w(y^t)$ of arbitrary forms. However, it is possible to represent semi-Markovian utility in a generalized exponential form which clarifies some of its properties. If $w(y^t)$ is positive and strictly increasing (or negative and strictly decreasing) in t , and differentiable, it may always be represented as

$$w(y^t) = \int_0^t v(y) e^{-A_y(s)} ds,$$

for some functions $v(y)$, $A_y(t)$. Moreover, because $\Delta w(y^t) > 0$, it may always be represented as

$$\Delta w(y^t) = e^{-B_y(t)},$$

for some function $B_y(t)$. When $A_y(t) = B_y(t) = a(y)t$ and duration independence holds, we recover the Markovian case. In addition, the risk aversion function $a_y(t)$ associated with $w(y^t)$ is the *derivative* (should it exist) of $A_y(t)$:

$$a_y(t) = \frac{d^2/dt^2 w(y^t)}{d/dt w(y^t)} = A'_y(t).$$

General formulas for evaluating an exponential(μ) sojourn in a state y are readily devised. By interchanging the order of integration one can show

$$\begin{aligned} w(y^\mu) &= \int_0^\infty \left(\int_0^t v(y) e^{-A_y(s)} ds \right) \mu e^{-\mu t} dt \\ &= v(y) \int_0^\infty e^{-A_y(s) - \mu s} ds. \end{aligned} \quad (20)$$

Moreover, in the duration-independent case

$$\Delta w(y^\mu) = \int_0^\infty e^{-B_y(t)} \mu e^{-\mu t} dt = \mu \int_0^\infty e^{-B_y(t) - \mu t} dt. \quad (21)$$

A convenient form for evaluating exponential sojourns is the power form for semi-Markov utility

$$w(y^t) = \frac{v(y)}{1 - \alpha_y} t^{1 - \alpha_y} \quad (22)$$

defined for $\alpha_y < 1$. The associated risk-aversion function is

$$a_y(t) = \frac{\alpha_y}{t}.$$

It is not difficult to show that for exponential (μ) sojourns in state y

$$w(y^\mu) = v(y) \left(\frac{1}{\mu} \right)^{1 - \alpha_y} \Gamma(1 - \alpha_y). \quad (23)$$

The advantage here is that in contrast to the case of exponential utility, $w(y^\mu)$ is finite for all values of μ and all permitted risk attitudes, both risk averse ($0 \leq \alpha_y < 1$) and risk seeking ($\alpha_y \leq 0$). The companion power form

$$\Delta w(y^t) = t^{-\beta_y},$$

is less realistic as a discount factor since its values can exceed 1 for $\beta_y \neq 0$. However for exponential sojourns, there is the convenient closed-form expression


$$\Delta w(y^\mu) = \mu^{\beta_y} \Gamma(1 - \beta_y), \quad (24)$$

the latter equation holding in the duration-independent case and for $\beta_y < 1$.

Stochastic Tree Rollback

The results presented above may be used to derive special rollback formulas for stochastic trees having only exponential durations. That is the purpose of this section. We begin by considering the case of semi-Markovian utility, since both Markovian and memoryless utility are special cases.

Consider a schematic representation of a stochastic fork



formed by competing rates $\mu_i > 0$ at state y , where $\mu = \sum_i \mu_i$ and $p_i = \mu_i/\mu$. Here H_i is itself a stochastic subtree for which the utility $E[u(H_i|y^\mu)]$ has already been calculated. Supposing that the preference summary at G is x^λ , we wish to derive a useful recursive formula for $E[u(G|x^\lambda)]$. Because self-transitions may be eliminated when durations are exponential, we assume that $x \neq y$. Denote by H the stochastic subtree consisting of a chance p_i of subtree H_i . Then $G = y^\mu H$, and according to (18)

$$\begin{aligned} E_G[u(G|x^\lambda)] &= E_H[u(y^\mu H|x^\lambda)] \\ &= E_H \left[\left[\frac{w(y|x^\lambda)}{\Delta w(y|x^\lambda)} \right] \circ \left[\frac{w(y^\mu)}{\Delta w(y^\mu)} \right] (u(H|y^\mu)) \right] \\ &= \left[\frac{w(y|x^\lambda)}{\Delta w(y|x^\lambda)} \right] \circ \left[\frac{w(y^\mu)}{\Delta w(y^\mu)} \right] \left(\sum_i p_i E[u(H_i|y^\mu)] \right). \end{aligned}$$

When the semi-Markov representations (20), (21) are substituted, the result in the duration-independent case is

$$\begin{aligned} E_G[u(G|x^\lambda)] &= \left[\frac{w(y|x)}{\Delta w(y|x)} \right] \circ \left[\frac{v(y) \int_0^\infty e^{-A_y(s) - \mu s} ds}{\mu \int_0^\infty e^{-B_y(s) - \mu s} ds} \right] \\ &\quad \left(\sum_i p_i E[u(H_i|y^\mu)] \right) \\ &= \left[\frac{w(y|x)}{\Delta w(y|x)} \right] \circ \left[\frac{v(y) \int_0^\infty e^{-A_y(s) - \mu s} ds}{\int_0^\infty e^{-B_y(s) - \mu s} ds} \right] \\ &\quad \left(\sum_i \mu_i E[u(H_i|y^\mu)] \right). \end{aligned}$$

In the special case of the semi-Markov power forms (23), (24), we obtain

$$E_G[u(G|x^\lambda)] = \left[\frac{w(y|x)}{\Delta w(y|x)} \right] \circ \left[\frac{v(y) \left(\frac{1}{\mu}\right)^{1-\alpha_y} \Gamma(1-\alpha_y)}{\left(\frac{1}{\mu}\right)^{1-\beta_y} \Gamma(1-\beta_y)} \right] \left(\sum_i \mu_i E[u(H_i|y^\mu)] \right), \quad (25)$$

valid for $\alpha_y < 1$ and $\beta_y < 1$. Finally, in the statewise exponential case giving rise to Markovian utility, we obtain


$$\begin{aligned} E_G[u(G|x)] &= \left[\frac{w(y|x)}{\Delta w(y|x)} \right] \circ \left[\frac{v(y)/(a(y) + \mu)}{1/(a(y) + \mu)} \right] \\ &\quad \left(\sum_i \mu_i E[u(H_i|y)] \right) \\ &= \left[\frac{w(y|x)}{\Delta w(y|x)} \right] \left(\frac{v(y) + \sum_i \mu_i E[u(H_i|y)]}{a(y) + \mu} \right). \end{aligned} \quad (26)$$

valid whenever $a(y) + \mu > 0$.

Cyclic Stochastic Tree Rollback

The utility rollback algorithm for cyclic stochastic trees is an intuitive extension of the acyclic rollback algorithm, and indeed is merely a version of the method of successive approximations (value iteration) from infinite horizon dynamic programming (e.g., Bertsekas, Ross). We give here a simple example to illustrate the technique. The example is also meant to concretely illustrate rollback using preference summaries, with its associated computational simplifications.

Consider the hypothetical stochastic tree



having three possible states which we label 0, 1, 2. The transition rate from x to y is λ_{xy} . A transition with rate μ_x having no destination state indicates subsequent sample path ϕ . The rates are given by

$$\lambda = (\lambda_{xy}) = \begin{bmatrix} 0 & 10 & 5 \\ 20 & 0 & 20 \\ 40 & 0 & 0 \end{bmatrix}, \quad \mu = (\mu_x) = \begin{bmatrix} 10 \\ 10 \\ 5 \end{bmatrix}.$$

Table I
Recursive Calculations for a Cyclic Stochastic Tree

n	Eu ($G_{1n} 0$)	Eu ($G_{2n} 0$)	Eu ($G_{2n} 1$)	Eu (G_{0n})	Eu ($G_{0n} 1$)	Eu ($G_{0n} 2$)
0	0	0	0	0	0	0
1	0.15	0.053	0.053	0.425	0.283	0.013
2	0.285	0.064	0.064	0.481	0.333	0.041
3	0.309	0.089	0.089	0.496	0.346	0.048
4	0.325	0.095	0.095	0.503	0.353	0.052
5	0.33	0.099	0.099	0.506	0.356	0.053
6	0.332	0.1	0.1	0.507	0.357	0.054
7	0.333	0.101	0.101	0.508	0.357	0.054
8	0.333	0.101	0.101	0.508	0.357	0.054
9	0.334	0.101	0.101	0.508	0.357	0.054
10	0.334	0.101	0.101	0.508	0.357	0.054

We evaluate the tree using the power form of semi-Markov utility with $\beta_y = 0$, $\alpha_y = 0.5$ for all states y , and

$$v = (v(y)) = \begin{bmatrix} 1.0 \\ 0.6 \\ 0.2 \end{bmatrix}.$$

We set all bonuses $w(y|x)$ to zero and all discount factors $\Delta w(y|x)$ to one, with the exceptions

$$w(0|1) = -0.10, \quad \Delta w(0|1) = 0.90, \\ w(0|2) = -0.20, \quad \Delta w(0|2) = 0.50.$$

In other words, a transition from 1 to 0 is accompanied by a toll of 0.10, and all subsequent utility is discounted by a factor of 0.90; for the transition from 2 to 0, the toll is 0.20, and subsequent utility is discounted by a factor of 0.50.

We use the recursion formula (25) moving from right to left in the tree. We begin by setting the utility of all terminal nodes in the tree to zero, then applying (25) from right to left until all nodes are assigned utilities. These newly computed values are assigned to the terminal nodes, and the process is repeated. Eventually the utilities converge to the true values for the cyclic tree. A numerical summary is presented in Table I.

In Table I, the expression $Eu(G_{yn}|x)$ denotes the expected utility, given preference summary (that is, preceding state) x , of the stochastic subtree G_{yn} which begins at state y and is truncated after n repeated levels. Notice that for each y , $Eu(G_{yn}|x)$ must be calculated for each x which could precede y , that is, for each preference summary which could occur at y . So state 0 is tagged with three utility values (predecessors $\emptyset, 1, 2$), state 1 with one (predecessor 0), and state 2 by two (predecessors 0, 1). Contrast this with the most general form of recursion (Theorem 1), in which the preference summary at y is in effect the entire history preceding y . In that case $Eu(G_{yn}|x^\lambda)$ would be required for every possible history x^λ preceding y . The computational simplification due to an available preference summary is dramatic.

We assert without proof that the expected utilities $Eu(G_{yn}|x)$ converge to the corresponding infinite-horizon values $Eu(G_y|x)$ as $n \rightarrow \infty$. Convergence questions will not be treated here. The overall expected utility of the cyclic

tree is therefore the approximate value $Eu(G_{0n}) = 0.508$ with $n = 10$.


3. ASSESSMENT AND APPLICATION

Utility Assessment Procedures

Space does not permit a complete discussion of assessment procedures for the classes of utility functions we have introduced here. We will confine the discussion to the assessment of the risk attitude parameters for the exponential and power forms described above. A possible assessment setting is a generalization of the vaccine scenario discussed in Section 1, which an assessor might present as follows:

You have been exposed to a virus which might reduce your life-span if left untreated. Your untreated lifetime has distribution μ_1 with mean $c\%$ lower than before. There is a treatment which will restore your life-span to its baseline distribution μ_0 but there is a probability p of fatal side effects. What is the largest value of p for which you would accept treatment?

Using stochastic tree notation, we require a value $p = p_{\text{vaccine}}$ which produces the indifference



A convenient aspect of this scenario is that the risk-neutral response is $p_{\text{vaccine}} = c$, so that lower values of p_{vaccine} correspond to risk aversion.

Ideally, the distributions μ_0 and μ_1 should be realistic lifetime distributions whose expected utilities have convenient closed forms. A class of distributions compatible with exponential utility is the gamma (m, λ) family (the distribution of the sum of m independent exponential (λ) durations). The gamma parameters can be chosen to reasonably approximate human survival curves (Hazen 1992). If μ is gamma (m, λ) and utility has the exponential form of Theorem 4, then

$$u(y^\mu) = \frac{v(y)}{a(y)} \left(1 - \left(\frac{\lambda}{\lambda + a(y)} \right)^m \right).$$

When μ_0 is gamma (m, λ) and μ_1 is gamma ($m, \lambda/(1 - c)$), the indifference (27) forces the equality

$$p + (1 - p) \left(\frac{\lambda}{\lambda + a(y)} \right)^m = \left(\frac{\lambda}{\lambda + (1 - c)a(y)} \right)^m,$$

which can be solved numerically for $a(y)$. When $m = 1$ (exponential lifetimes), the solution is

$$a(y) = p \frac{c - p}{1 - c} \lambda.$$

A flexible class of distributions compatible with power utility is the Weibull family with parameters m and θ , having hazard rate function

$$\mu(t) = \frac{m}{\theta} \left(\frac{t}{\theta}\right)^{m-1},$$

and mean $\theta\Gamma(1 + 1/m)$. When semi-Markov utility takes on the power form (22), then

$$w(y^\mu) = v(y) \frac{\theta^{1-\alpha_y}}{m} \Gamma\left(\frac{1-\alpha_y}{m}\right).$$

When μ_0 is Weibull(m, θ) and μ_1 is Weibull($m, (1 - c)\theta$), the indifference (28) forces an equality which can be solved to yield

$$\alpha_y = 1 - \frac{\ln(1 - p)}{\ln(1 - c)}$$

(independent of m and θ).

We advocate the use of *continuous-risk* assessment scenarios such as (27) involving small risks present in continuous time because such scenarios more closely approximate the decisions to which the assessed utility functions will be applied. Although continuous-risk assessment scenarios are cognitively demanding, we have found that subjects can provide answers when options are portrayed as survival curves using an interactive graphical computer interface. The resulting utility functions seem to more accurately portray risk attitudes than do utilities assessed with standard binary gamble/sure thing lotteries (Hazen et al. 1990, Pellissier and Hazen 1993).

An Analysis of the Decision to Perform Carotid Endarterectomy

The material in this section is based on Matchar and Pauker, who conduct a decision analysis of whether to perform a carotid endarterectomy (a surgical clearing of the carotid arteries) on a patient experiencing transient ischemic attacks (temporary symptoms such as lightheadedness and garbled speech, due to impaired blood supply). Matchar and Pauker report that “despite controversy regarding their value, approximately 85,000 carotid endarterectomies are performed annually.” Their analysis, which was based on a discrete-time Markov chain model, found the decision to be a toss-up, that is, there was no clinically significant net benefit to performing carotid endarterectomy. The evaluation criterion was mean quality-adjusted duration. Here we present the equivalent stochastic tree model, but evaluate the alternatives using Markovian and semi-Markovian utility functions. We can therefore examine the impact of patient risk attitude on the preferred decision.

Henceforth let the *surgical option* be to perform endarterectomy, and the *nonsurgical option* to forego it. The stochastic tree for the nonsurgical option has been presented in Figure 1 of Section 1. The complete stochastic

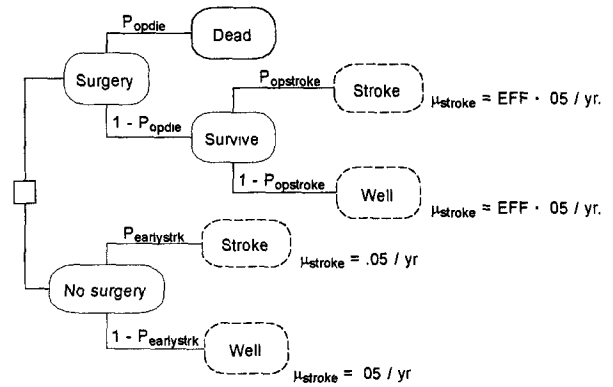


Figure 6. The decision to perform carotid endarterectomy.

tree for this decision is depicted in Figure 6, which shows the initial decision as to whether to perform surgery, as well as chances of mortality and stroke following surgery, and the 30-day chance of stroke for transient ischemic patients in the absence of surgery. As an approximation, all states depicted in Figure 6 are assumed instantaneous. (The true elapsed time of one month or less is small compared to other durations in the model.) This portion of the stochastic tree is therefore identical in structure and function to a decision tree. Patients foregoing surgery enter the stochastic tree of Figure 1 at the appropriate state (Well or Stroke, depending on whether a stroke has occurred in the short term) with unreduced subsequent stroke rate $\mu_{stroke} = 0.05/yr$. Patients undergoing and surviving surgery enter the same tree at the appropriate state with subsequent stroke rate $\mu_{stroke} = 0.05/yr$, reduced by the efficacy factor EFF. Matchar and Pauker use $EFF = 0.50$. The values for the remaining parameters in the stochastic tree are summarized in Table II.

We begin by evaluating the stochastic tree under Markovian utility. For simplicity and consistency with Matchar and Pauker, we use no tolls or discount factors (that is, we assume $w(y|x) = 0$ and $\Delta w(y|x) = 1$ for all x ,

Table II
Variables Used in the Matchar-Pauker Analysis

Description	Variable Name	Value
Operative mortality probability (30 days)	P_{opdie}	0.06
Operative stroke probability (30 days)	$P_{opstroke}$	0.08
30-day stroke probability for nonsurgical patients	$P_{earlystrk}$	0.04
Annual stroke rate for nonsurgical patients	μ_{stroke}	0.05/yr.
First-month death probability for big-stroke victims	$P_{earlydie}$	0.38
Proportion of strokes having significant mortality/morbidity	P_{big}	2/3
Excess risk of nonstroke death	μ_{excess}	0.065/yr.
Mortality due to other causes (58-yr-old white male)	μ_0	0.01106/yr.

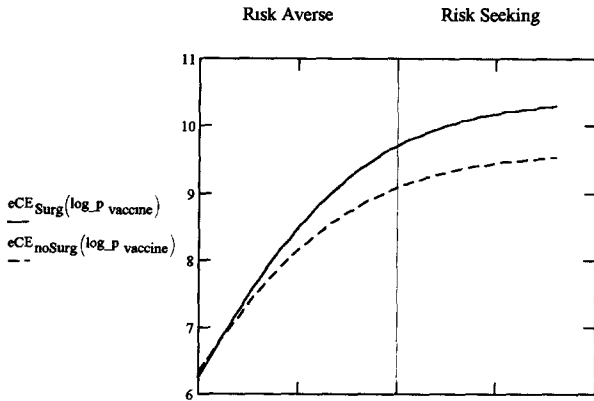


Figure 7. Certainty equivalent lifetimes (in years) as a function of p_{vaccine} for exponential utility.

y). We assign the same quality factors $v(y)$ as Matchar and Pauker, namely

y	$v(y)$
Well	1.0
Post Small Stroke	0.8
Post Big Stroke	0.2
Dead	0

We suppose the risk aversion parameters $a(y)$ are determined using the vaccine scenario discussed in the previous subsection. Specifically, let $a(y)$ be the parameter for which an immediate mortality probability p_{vaccine} would be just acceptable in order to eliminate a $c = 10\%$ reduction in mean (exponentially distributed) lifetime. Recall that the risk neutral response would be $p_{\text{vaccine}} = c = 10^{-1}$. For simplicity we assume that p_{vaccine} (and therefore $a(y)$) is the same regardless of the state y .

Stochastic tree rollback proceeds recursively by invoking the exponential utility formula (26) at stochastic forks, and averaging utilities at chance forks. Rollback proceeds from right to left in the stochastic tree, until utilities are calculated for the Surgery and No Surgery nodes in Figure 6. As an aid to interpreting the utility values, one may calculate *certainty equivalent lifetimes* at each node, that is, constant durations t spent in the state Well such that $u(\text{Well}^t)$ is equal to the expected utility at the node in question.

These certainty equivalent lifetimes are given as a function of $\log_{10}(p_{\text{vaccine}})$ in Figure 7. Values of p_{vaccine} less than 10^{-1} yield risk aversion ($a(y) > 0$), and those exceeding 10^{-1} yield risk seeking ($a(y) < 0$). The risk-neutral certainty equivalent lifetimes are $eCE_{\text{Surg}} = 9.705$ yr., $eCE_{\text{noSurg}} = 9.093$ yr., slightly favoring surgery. Matchar and Pauker obtain the respective values 8.00 yr. and 7.76 yr. We attribute the discrepancy to continuous- versus discrete-time model structure.

Matchar and Pauker do attempt to model risk aversion by including a discount rate for future life-years. Because our risk attitude parameter $a(y)$ may be regarded as a discount rate, their approach would be the discrete-time equivalent of ours. They found that the surgical and nonsurgical




Figure 8. Certainty equivalent lifetimes (in years) as a function of p_{vaccine} for power utility.

options become essentially equivalent as discount rate increases beyond 5%. This is in agreement with the results of Figure 7 for $\log(p_{\text{vaccine}}) < -1.8$, or equivalently $a(y) > 6\%$.

For other utility functions, however, results can be quite different. Figure 8 shows certainty equivalent lifetimes as a function of p_{vaccine} when a particular version of the power form of semi-Markov utility is used. Here we have assumed the same quality factors $v(y)$ as above, but have eliminated any discounting effects by setting $\beta_y = 0$, and have let the remaining parameters α_y be determined by assessing the immediate mortality probability p_{vaccine} which would be just acceptable in order to eliminate a $c = 10\%$ reduction in mean (Weibull distributed) lifetime, as described in the previous section. Once again, $p_{\text{vaccine}} = 10^{-1}$ is the risk-neutral response. For simplicity we suppose α_y is independent of y . Rollback proceeds using (25) at stochastic forks, and utility averaging at chance forks. As Figure 8 shows, the nonsurgical option is increasingly preferred as risk aversion increases. This is not surprising, due to the form $a_y(t) = \alpha_y/t$ of the coefficient of risk aversion for power utility, which becomes arbitrarily risk averse as t approaches zero. Risks of immediate death are therefore increasingly undesirable, and options which avoid such near-term risks are increasingly desirable as risk aversion increases.

The lessons here are threefold: First, risk attitude as captured by utility functions over stochastic trees can play an important role in medical treatment choice. Second, there are aspects of risk aversion, such as aversion to near-term risks, which are not adequately modeled by time discounting, but which can be captured by the semi-Markovian class of utility functions. Third, it may be beneficial to portray risk attitude not just as an abstract parameter, but also as the indifference level of a proxy quantity such as p_{vaccine} which plays a meaningful role in some simple scenario resembling the decision problem of interest. The relation between treatment choice and a meaningful proxy measure of risk attitude can provide a convincing argument for the optimality of the option in question.

CONCLUSION

What utility functions over stochastic trees can be evaluated recursively? We have shown that any utility function can be so evaluated, but that recursive evaluation is computationally worthwhile only for utility functions having preference summaries of low dimension. For three specific low-dimensional summaries, the memoryless, the Markovian, and the semi-Markovian, we have identified the respective classes of utility functions having those summaries. We have illustrated how such utility functions may be used to model risk attitude in the medical decision of whether to perform a carotid endarterectomy.

APPENDIX

Proof of Theorem 1. For fixed \mathbf{x}^λ and y^μ , the utility functions $U(h|\mathbf{x}^\lambda y^\mu)$ and $U(y^\mu h|\mathbf{x}^\lambda)$ both represent the same preference relation over stochastic trees of the form $\mathbf{x}^\lambda y^\mu h$. They must therefore be strategically equivalent, that is, there must be constants $\Delta U(y^\mu|\mathbf{x}^\lambda) > 0$ and $a(y^\mu|\mathbf{x}^\lambda)$ such that

$$U(y^\mu h|\mathbf{x}^\lambda) = a(y^\mu|\mathbf{x}^\lambda) + \Delta U(y^\mu|\mathbf{x}^\lambda)U(h|\mathbf{x}^\lambda y^\mu).$$

Set $h = \phi$ to conclude that $a(y^\mu|\mathbf{x}^\lambda) = U(y^\mu|\mathbf{x}^\lambda)$, from which the recursive equation (7) follows.

To show the affine restriction (5), note that by definition we have

$$\begin{aligned} U(y^\mu h|\mathbf{x}^\lambda) &= \int_0^\infty (U(y^s|\mathbf{x}^\lambda) \\ &\quad + \Delta U(y^s|\mathbf{x}^\lambda)U(h|\mathbf{x}^\lambda y^s))d\mu(s) \\ &= U(y^\mu|\mathbf{x}^\lambda) \\ &\quad + \int_0^\infty \Delta U(y^s|\mathbf{x}^\lambda)U(h|\mathbf{x}^\lambda y^s)d\mu(s). \end{aligned}$$

The affine restriction (5) is an immediate consequence when the last equation is compared to (7).

To show the concatenation restriction (6), note first that

$$\mathbf{x}^\lambda y^\mu y^\nu h = \mathbf{x}^\lambda y^{\mu\nu} h.$$

The utility of the left side of this expression is, using the recursion (7),

$$\begin{aligned} U(y^\mu y^\nu h|\mathbf{x}^\lambda) &= U(y^\mu|\mathbf{x}^\lambda) + \Delta U(y^\mu|\mathbf{x}^\lambda)U(y^\nu h|\mathbf{x}^\lambda y^\mu) \\ &= U(y^\mu|\mathbf{x}^\lambda) + \Delta U(y^\mu|\mathbf{x}^\lambda)(U(y^\nu|\mathbf{x}^\lambda y^\mu) \\ &\quad + \Delta U(y^\nu|\mathbf{x}^\lambda y^\mu)U(h|\mathbf{x}^\lambda y^\mu y^\nu)), \end{aligned}$$

and the utility of the right side is

$$U(y^{\mu\nu} h|\mathbf{x}^\lambda) = U(y^{\mu\nu}|\mathbf{x}^\lambda) + \Delta U(y^{\mu\nu}|\mathbf{x}^\lambda)U(h|\mathbf{x}^\lambda y^{\mu\nu}).$$

Set $h = \phi$ in each of these last two equations and equate the results to obtain

$$U(y^{\mu\nu}|\mathbf{x}^\lambda) = U(y^\mu|\mathbf{x}^\lambda) + \Delta U(y^\mu|\mathbf{x}^\lambda)U(y^\nu|\mathbf{x}^\lambda y^\mu).$$

Substitute this into the right side of the preceding equation, equate the result to the right side of its predecessor, and cancel terms to conclude

$$\begin{aligned} \Delta U(y^\mu|\mathbf{x}^\lambda)\Delta U(y^\nu|\mathbf{x}^\lambda y^\mu)U(h|\mathbf{x}^\lambda y^\mu y^\nu) \\ = \Delta U(y^{\mu\nu}|\mathbf{x}^\lambda)U(h|\mathbf{x}^\lambda y^{\mu\nu}). \end{aligned}$$

By nontriviality there is a sample path h with $U(h|\mathbf{x}^\lambda y^\mu y^\nu) = U(h|\mathbf{x}^\lambda y^{\mu\nu}) \neq 0$. Therefore this utility term can be cancelled from the last equation, and (6) results. QED.

Proof of Theorem 2. Given a sample path h and preference state q , we define $u(h|q)$ by selecting an arbitrary \mathbf{z}^ν with $e(\mathbf{z}^\nu) = q$, and letting $u(h|q) = U(h|\mathbf{z}^\nu)$. If $e(\mathbf{x}^\lambda) = q = e(\mathbf{z}^\nu)$, then because e is a preference summary, conditional preferences given \mathbf{x}^λ and given \mathbf{z}^ν are identical, so $U(\cdot|\mathbf{x}^\lambda) \sim U(\cdot|\mathbf{z}^\nu) = u(\cdot|q)$, as claimed. In fact, because $U(\phi|\mathbf{x}^\lambda) = 0$ for all \mathbf{x}^λ , it follows that

$$U(\cdot|\mathbf{x}^\lambda) = a(\mathbf{x}^\lambda)u(\cdot|q),$$

for some function $a(\mathbf{x}^\lambda) > 0$. Substitute this equation into (1) and use preference updatability to conclude

$$\begin{aligned} a(\mathbf{x}^\lambda)u(y^\mu h|q) &= a(\mathbf{x}^\lambda)u(y^\mu|q) \\ &\quad + \Delta U(y^\mu|\mathbf{x}^\lambda)a(\mathbf{x}^\lambda y^\mu)u(h|\theta(q, y^\mu)). \end{aligned}$$

Then

$$\frac{\Delta U(y^\mu|\mathbf{x}^\lambda)a(\mathbf{x}^\lambda y^\mu)}{a(\mathbf{x}^\lambda)} = \frac{u(y^\mu h|q) - u(y^\mu|q)}{u(h|\theta(q, y^\mu))} \equiv \Delta u(y^\mu|q).$$

Equation (10) is an immediate consequence of this definition of $\Delta u(y^\mu|q)$. To show (9), use the definition of Δu and (6) to write

$$\begin{aligned} \Delta u(y^\mu|q)\Delta u(y^\nu|\theta(q, y^\mu)) \\ = \Delta U(y^\mu|\mathbf{x}^\lambda)\Delta U(y^\nu|\mathbf{x}^\lambda y^\mu) \frac{a(\mathbf{x}^\lambda y^\mu)}{a(\mathbf{x}^\lambda)} \frac{a(\mathbf{x}^\lambda y^\mu y^\nu)}{a(\mathbf{x}^\lambda y^\mu)} \\ = \Delta U(y^{\mu\nu}|\mathbf{x}^\lambda) \frac{a(\mathbf{x}^\lambda y^\mu y^\nu)}{a(\mathbf{x}^\lambda)} \\ = \Delta u(y^{\mu\nu}|q), \end{aligned}$$

which is the desired result. Finally, Equation (8) follows from (10) because

$$\begin{aligned} \Delta u(y^\mu|q)u(h|\theta(q, y^\mu)) \\ = u(y^\mu h|q) - u(y^\mu|q) \\ = \int_0^\infty (u(y^s h|q) - u(y^s|q))ds \\ = \int_0^\infty \Delta u(y^s|q)u(h|\theta(q, y^s))ds. \quad \text{QED} \end{aligned}$$

Proof of Theorems 3 and 4. We require the following lemma.

Lemma 1. *The only continuous functions $g(t)$ satisfying $g(s + t) = g(s) + e^{-as}g(t)$,*

in some interval of values t are

$$g(t) = k \int_0^t e^{-au} du,$$

for real k

Proof. If $a = 0$, the functional equation becomes

$$g(s + t) = g(s) + g(t),$$

whose only continuous solution is the linear function $g(t) = kt = k \int_0^t 1 \cdot ds = k \int_0^t e^{-as} ds$, as claimed.

Suppose $a \neq 0$. Equate $g(s + t)$ with $g(t + s)$ to obtain

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

or equivalently

$$\frac{g(s)}{1 - e^{-as}} = \frac{g(t)}{1 - e^{-at}},$$

holding for all nonzero s, t . It follows that both sides must be constant, that is,

$$\frac{g(t)}{1 - e^{-at}} = \frac{k}{a},$$

for all nonzero t (and by continuity, for $t = 0$ as well, if 0 is in the interval of interest). So

$$g(t) = \frac{k}{a} \cdot (1 - e^{-at}) = k \int_0^t e^{-as} ds.$$

The latter function has the property that

$$\begin{aligned} g(s + t) &= k \int_0^{s+t} e^{-au} du \\ &= k \left(\int_0^s e^{-au} du + \int_s^{s+t} e^{-au} du \right) \\ &= k \left(\int_0^s e^{-au} du + e^{-as} \int_0^t e^{-au} du \right) \\ &= g(s) + e^{-as}g(t), \end{aligned}$$

as claimed. QED.

To begin the proof of Theorem 4, invoke the concatenation restriction (9) under memoryless utility to obtain

$$\Delta u(y^{\mu * \nu}) = \Delta u(y^\mu) \Delta u(y^\nu).$$

Letting μ and ν be degenerate distributions at s and t , respectively, we conclude that the function $g(t) = \Delta u(y^t)$ satisfies the functional equation $g(s + t) = g(s)g(t)$ for all $s, t > 0$. Since Δu may be expressed in terms of U , which is continuous in its duration arguments, it follows that the function g is continuous. Therefore $g(t) = e^{-a(y)t} = \Delta u(y^t)$ for some constant $a(y)$ which may depend on the state y (e.g., Section 1 in Aczel 1987).

Next, invoke the recursive equation (10) under memoryless utility to obtain

$$u(y^{\mu * \nu}) = u(y^\mu) + \Delta u(y^\mu)u(y^\nu).$$

Let μ and ν be degenerate respectively at s and t to obtain

$$u(y^{s+t}) = u(y^s) + \Delta u(y^s)u(y^t) = u(y^s) + e^{-a(y)s}u(y^t).$$

The function $g(t) = u(y^t)$ therefore satisfies the functional equation of Lemma 1 and is continuous in t . Therefore

$$u(y^t) = v(y) \int_0^t e^{-a(y)s} ds,$$

as the theorem claims.

Finally, invoke the affine restriction (8) under memoryless utility to obtain

$$\Delta u(y^\mu)u(h) = \int_0^\infty \Delta u(y^s)u(h)d\mu(s).$$

Cancel $u(h)$ to obtain the last claim of the theorem.

Next we turn to Theorem 3. By Theorem 2 under memoryless utility, $U(\cdot) \sim u(\cdot | \emptyset) = u(\cdot)$. To show necessity, invoke the recursive equation (10) for memoryless utility to obtain

$$\begin{aligned} u(x^t y^s) &= u(x^t) + \Delta u(x^t)u(y^s) \\ &= \left[\frac{u(x^t)}{\Delta u(x^t)} \right] u(y^s). \end{aligned}$$

Now apply the same procedure to $u(y^s)$, and so on, to obtain the claim of the theorem.

To show sufficiency it must be demonstrated that U has a memoryless preference summary. As a preliminary step, note that

$$\begin{aligned} U(\mathbf{x}^\lambda y^s) &= E[U(\mathbf{x}^T y^s)] \\ &= a + b \cdot E \left[\left[\frac{u(x_1^{T_1})}{\Delta u(x_1^{T_1})} \right] \circ \dots \circ \left[\frac{u(x_n^{T_n})}{\Delta u(x_n^{T_n})} \right] \right. \\ &\quad \left. (u(y^s)) \right] \\ &= a + b \cdot \left[\frac{u(x_1^{\lambda_1})}{\Delta u(x_1^{\lambda_1})} \right] \circ \dots \circ \left[\frac{u(x_n^{\lambda_n})}{\Delta u(x_n^{\lambda_n})} \right] (u(y^s)), \end{aligned}$$

for some constants $b > 0, a$. In the same way

$$U(\mathbf{x}^\lambda) = a + b \cdot \left[\frac{u(x_1^{\lambda_1})}{\Delta u(x_1^{\lambda_1})} \right] \circ \dots \circ \left[\frac{u(x_n^{\lambda_n})}{\Delta u(x_n^{\lambda_n})} \right] (0).$$

Then

$$\begin{aligned} U(y^s | \mathbf{x}^\lambda) &= c \cdot (U(\mathbf{x}^\lambda y^s) - U(\mathbf{x}^\lambda)), \quad c > 0. \\ &= c \cdot b \cdot \Delta u(x_1^{\lambda_1}) \cdot \dots \cdot \Delta u(x_n^{\lambda_n}) u(y^s), \end{aligned}$$

which is strategically equivalent (as a function of y^s) to $u(y^s)$. The associated preference summary therefore does not depend on \mathbf{x}^λ , hence is memoryless. QED.

Proof of Theorems 5 and 6. We require the following extension of Lemma 1.

Lemma 2. *The only continuous functions $g(t), f(t)$ satisfying*

$$g(s + t) = g(s) + e^{-as}f(t), \tag{A.1}$$

for $s, t > 0$ are

$$f(t) = k \int_0^t e^{-au} du,$$

$$g(t) = c + k \int_0^t e^{-au} du,$$

for real k, c .

Proof. Set $s = 1$ in (A.1) to get

$$g(1 + t) = g(1) + e^{-a}f(t)$$

$$f(t) = e^a(g(t + 1) - g(1)).$$

Resubstitute into (A.1) to get

$$g(s + t) = g(s) + e^{-a(s-1)}(g(t + 1) - g(1)).$$

Replace s by $s + 1$ and subtract $g(1)$ to get

$$g(s + t + 1) - g(1) = g(s + 1) - g(1) + e^{-as}(g(t + 1) - g(1)).$$

Now apply Lemma A.1 to the function $g(t + 1) - g(1)$ to conclude

$$g(t + 1) = g(1) + k_1 \int_0^t e^{-au} du,$$

and

$$f(t) = e^a(g(t + 1) - g(1)) = k \int_0^t e^{-au} du,$$

where $k = k_1 e^a$. Replace t by $t - 1$ in the penultimate equation to get

$$\begin{aligned} g(t) &= g(1) + k_1 \int_0^{t-1} e^{-au} du = g(1) \\ &\quad + k_1 \int_1^t e^{-a(v-1)} dv = k_1 e^a \int_1^t e^{-av} dv \\ &= g(1) + k_1 e^a \left(- \int_0^1 e^{-av} dv + \int_0^t e^{-av} dv \right) \\ &= c + k \int_0^t e^{-av} dv, \end{aligned}$$

as claimed. This pair $g(t), f(t)$ does satisfy A.1:

$$\begin{aligned} g(s + t) &= c + k \int_0^{s+t} e^{-au} du \\ &= c + k \left(\int_0^s e^{-au} du + \int_s^{s+t} e^{-au} du \right) \\ &= c + k \left(\int_0^s e^{-au} du + e^{-as} \int_0^t e^{-au} du \right) \\ &= g(s) + e^{-as}f(t), \end{aligned}$$

as desired. QED.

We now turn to Theorem 6. From the concatenation restriction (9) we get

$$\Delta u(y^{s+t}|x) = \Delta u(y^s|x)\Delta u(y^t|y).$$

Take natural logs and apply Lemma 2 with $a = 0$ to get

$$\Delta u(y^s|x) = \exp(k_{xy}s + c_{xy}),$$

$$\Delta u(y^t|y) = \exp(k_{xy}t).$$

The last equation forces k_{xy} to be independent of x . So let $k_{xy} = -a(y)$, $\Delta w(y|x) = \exp(c_{xy})$ for $x \neq y$ and $\Delta w(y|y) = 1$ to get

$$\Delta u(y^t|x) = \Delta w(y|x)e^{-a(y)t},$$

for all x, y , as desired.

According to the recursive equation (10) we have

$$\begin{aligned} u(y^{s+t}|x) &= u(y^s|x) + \Delta u(y^s|x)u(y^t|y) \\ &= u(y^s|x) + \Delta w(y|x)e^{-a(y)s}u(y^t|y). \end{aligned}$$

Invoking Lemma 2 we obtain

$$\Delta w(y|x)u(y^t|y) = k_{xy} \int_0^t e^{-a(y)s} ds, \tag{A.2}$$

$$u(y^t|x) = c_{xy} + k_{xy} \int_0^t e^{-a(y)s} ds, \tag{A.3}$$

for new constants k_{xy} and c_{xy} . Set $x = y$ in (A.2) and define $v(y) = k_{yy}$ to get

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

Resubstitute into (A.2) to get

$$\Delta w(y|x)v(y) = k_{xy}.$$

Letting $w(y|x) = c_{xy}$, we have according to (A.3)

$$u(y^t|x) = w(y|x) + \Delta w(y|x)v(y) \int_0^t e^{-a(y)s} ds.$$

Substitute $x = y$ here and compare to the equation for $u(y^t|y)$ to conclude that $w(y|y) = 0$.

This establishes all but the last claim of Theorem 6. To derive that, substitute the Markovian summary into the affine restriction (8) to get

$$\Delta u(y^\mu|x)u(h|y) = \int_0^\infty \Delta u(y^s|x)u(h|y)d\mu(s),$$

and cancel the factor $u(h|y)$.

To show necessity in Theorem 5, invoke the recursive equation (10) to obtain

$$\begin{aligned} u(x^t|\emptyset) &= u(x_1^{t_1}|\emptyset) + \Delta u(x_1^{t_1}|\emptyset)u(x_2^{t_2} \dots x_n^{t_n}|x_1) \\ &= \left[u(x_1^{t_1}|\emptyset) \right] (u(x_2^{t_2} \dots x_n^{t_n}|x_1)). \end{aligned}$$

Apply (10) again to obtain

$$\begin{aligned}
 u(\mathbf{x}^t|\emptyset) &= \begin{bmatrix} u(x_1^{t_1}|\emptyset) \\ \Delta u(x_1^{t_1}|\emptyset) \end{bmatrix} (u(x_2^{t_2}|x_1) \\
 &\quad + \Delta u(x_2^{t_2}|x_1)u(x_3^{t_3} \dots x_n^{t_n}|x_2)) \\
 &= \begin{bmatrix} u(x_1^{t_1}|\emptyset) \\ \Delta u(x_1^{t_1}|\emptyset) \end{bmatrix} \circ \begin{bmatrix} u(x_2^{t_2}|x_1) \\ \Delta u(x_2^{t_2}|x_1) \end{bmatrix} \\
 &\quad (u(x_3^{t_3} \dots x_n^{t_n}|x_2)).
 \end{aligned}$$

Repeat until one obtains

$$\begin{aligned}
 u(\mathbf{x}^t|\emptyset) &= \begin{bmatrix} u(x_1^{t_1}|\emptyset) \\ \Delta u(x_1^{t_1}|\emptyset) \end{bmatrix} \circ \begin{bmatrix} u(x_2^{t_2}|x_1) \\ \Delta u(x_2^{t_2}|x_1) \end{bmatrix} \\
 &\quad \circ \dots \circ \begin{bmatrix} u(x_n^{t_n}|x_{n-1}) \\ \Delta u(x_n^{t_n}|x_{n-1}) \end{bmatrix} (0).
 \end{aligned}$$

Because $U(\mathbf{x}^t) \sim u(\mathbf{x}^t|\phi)$, the conclusion of the theorem follows.

To show sufficiency in Theorem 5, we must show that the given U has a Markovian preference summary. We have

$$\begin{aligned}
 U(\mathbf{x}^\lambda y^s) &= E[U(\mathbf{x}^T y^s)] \\
 &= E \left[\begin{bmatrix} u(x_1^{T_1}|\emptyset) \\ \Delta u(x_1^{T_1}|\emptyset) \end{bmatrix} \circ \begin{bmatrix} u(x_2^{T_2}|x_1) \\ \Delta u(x_2^{T_2}|x_1) \end{bmatrix} \right. \\
 &\quad \left. \circ \dots \circ \begin{bmatrix} u(x_n^{T_n}|x_{n-1}) \\ \Delta u(x_n^{T_n}|x_{n-1}) \end{bmatrix} (u(y^s|x_n)) \right] \\
 &= \begin{bmatrix} u(x_1^{\lambda_1}|\emptyset) \\ \Delta u(x_1^{\lambda_1}|\emptyset) \end{bmatrix} \circ \begin{bmatrix} u(x_2^{\lambda_2}|x_1) \\ \Delta u(x_2^{\lambda_2}|x_1) \end{bmatrix} \\
 &\quad \circ \dots \circ \begin{bmatrix} u(x_n^{\lambda_n}|x_{n-1}) \\ \Delta u(x_n^{\lambda_n}|x_{n-1}) \end{bmatrix} (u(y^s|x_n)),
 \end{aligned}$$

and similarly

$$\begin{aligned}
 U(\mathbf{x}^\lambda) &= \begin{bmatrix} u(x_1^{\lambda_1}|\emptyset) \\ \Delta u(x_1^{\lambda_1}|\emptyset) \end{bmatrix} \circ \begin{bmatrix} u(x_2^{\lambda_2}|x_1) \\ \Delta u(x_2^{\lambda_2}|x_1) \end{bmatrix} \\
 &\quad \circ \dots \circ \begin{bmatrix} u(x_n^{\lambda_n}|x_{n-1}) \\ \Delta u(x_n^{\lambda_n}|x_{n-1}) \end{bmatrix} (0).
 \end{aligned}$$

Therefore

$$\begin{aligned}
 U(y^s|\mathbf{x}^\lambda) &= U(\mathbf{x}^\lambda y^s) - U(\mathbf{x}^\lambda) \\
 &= \Delta u(x_1^{\lambda_1}|\emptyset) \Delta u(x_2^{\lambda_2}|x_1) \\
 &\quad \cdot \dots \Delta u(x_n^{\lambda_n}|x_{n-1}) u(y^s|x_n),
 \end{aligned}$$

from which it may be seen that $U(y^s|\mathbf{x}^\lambda)$ is strategically equivalent (as a function of y^s) to $u(y^s|x_n)$. Thus U has a Markovian preference summary. QED.

Proof of Theorems 7 and 8. Invoke the concatenation equations (9) under the semi-Markov preference summary to obtain

$$\Delta u(y^{\mu * \nu}|y^\lambda) = \Delta u(y^\mu|y^\lambda) \Delta u(y^\nu|y^{\lambda * \mu}), \tag{A.4}$$

and for $y \neq x$

$$\Delta u(y^{\mu * \nu}|x^\lambda) = \Delta u(y^\mu|x^\lambda) \Delta u(y^\nu|y^\mu). \tag{A.5}$$

The last equation when x^λ is the null preference summary is

$$\Delta u(y^{\mu * \nu}|\emptyset) = \Delta u(y^\mu|\emptyset) \Delta u(y^\nu|y^\mu).$$

Define $\Delta w(y^\mu) = \Delta u(y^\mu|\emptyset)$, and solve the last equation for $\Delta u(y^\nu|y^\mu)$ to get

$$\Delta u(y^\nu|y^\mu) = \frac{\Delta w(y^{\mu * \nu})}{\Delta w(y^\mu)}. \tag{A.6}$$

Substitute this back into (A.5) to get

$$\frac{\Delta u(y^{\mu * \nu}|x^\lambda)}{\Delta w(y^{\mu * \nu})} = \frac{\Delta u(y^\mu|x^\lambda)}{\Delta w(y^\mu)}, \quad x \neq y.$$

Replace $\mu * \nu$ and μ by constant durations $s > 0, t > 0$ to conclude that the ratio $\Delta u(y^s|x^\lambda)/\Delta w(y^s)$ does not depend on s . Let $\Delta w(y|x^\lambda)$ be this ratio. We therefore have

$$\Delta u(y^s|x^\lambda) = \Delta w(y|x^\lambda) \Delta w(y^s). \tag{A.7}$$

Note that

$$\Delta w(y|\emptyset) = \frac{\Delta u(y^s|\emptyset)}{\Delta w(y^s)} = \frac{\Delta u(y^s|\emptyset)}{\Delta u(y^s|\emptyset)} = 1.$$

We show that (A.7) remains valid when a distribution μ replaces s . Invoke the affine restriction (8) under the semi-Markov preference summary to get for $x \neq y$

$$\Delta u(y^\mu|x^\lambda) u(h|y^\mu) = \int \Delta u(y^s|x^\lambda) u(h|y^s) d\mu(s).$$

When x^λ is the null summary this becomes

$$\Delta w(y^\mu) u(h|y^\mu) = \int \Delta w(y^s) u(h|y^s) d\mu(s).$$

Use (A.7) in the penultimate equation to get

$$\begin{aligned}
 \Delta u(y^\mu|x^\lambda) u(h|y^\mu) &= \int_0^\infty \Delta w(y|x^\lambda) \Delta w(y^s) u(h|y^s) d\mu(s) \\
 &= \Delta w(y|x^\lambda) \int_0^\infty \Delta w(y^s) u(h|y^s) d\mu(s) \\
 &= \Delta w(y|x^\lambda) \Delta w(y^\mu) u(h|y^\mu).
 \end{aligned}$$

Cancel the factor $u(h|y^\mu)$, which by nontriviality can be considered nonzero, to obtain

$$\Delta u(y^\mu|x^\lambda) = \Delta w(y|x^\lambda) \Delta w(y^\mu). \tag{A.8}$$

The desired analog of (A.7), (A.6), and (A.8) together establish the second equation (17) of Theorem 8.

Now consider the recursive equation (10), which becomes under a semi-Markov preference summary:

$$\begin{aligned} u(y^\mu h|y^\lambda) &= u(y^\mu|y^\lambda) + \Delta u(y^\mu|y^\lambda)u(h|y^{\lambda*\mu}) \\ &= u(y^\mu|y^\lambda) + \frac{\Delta w(y^{\lambda*\mu})}{\Delta w(y^\lambda)} u(h|y^{\lambda*\mu}), \end{aligned}$$

and for $y \neq x$,

$$\begin{aligned} u(y^\mu h|x^\lambda) &= u(y^\mu|x^\lambda) + \Delta u(y^\mu|x^\lambda)u(h|y^\mu) \\ &= u(y^\mu|x^\lambda) + \Delta w(y|x^\lambda)\Delta w(y^\mu)u(h|y^\mu). \end{aligned}$$

Set $h = y^T$ in the last two equations, where T has distribution ν , and then take expectations to get

$$u(y^{\mu*\nu}|y^\lambda) = u(y^\mu|y^\lambda) + \frac{\Delta w(y^{\lambda*\mu})}{\Delta w(y^\lambda)} u(y^\nu|y^{\lambda*\mu}), \quad (\text{A.9})$$

and for $y \neq x$

$$u(y^{\mu*\nu}|x^\lambda) = u(y^\mu|x^\lambda) + \Delta w(y|x^\lambda)\Delta w(y^\mu)u(y^\nu|y^\mu). \quad (\text{A.10})$$

When x^λ is the null summary, the last equation becomes

$$w(y^{\mu*\nu}) = w(y^\mu) + \Delta w(y^\mu)u(y^\nu|y^\mu),$$

where we define $w(y^\mu) = u(y^\mu|\emptyset)$. Then w is affine in μ because $u(\cdot|y^0)$ is a utility function, so the first affine restriction (12) holds. Solve to obtain

$$u(y^\nu|y^\mu) = \frac{w(y^{\mu*\nu}) - w(y^\mu)}{\Delta w(y^\mu)},$$

and substitute back into (A.10) to get

$$u(y^{\mu*\nu}|x^\lambda) = u(y^\mu|x^\lambda) + \Delta w(y|x^\lambda)(w(y^{\mu*\nu}) - w(y^\mu)).$$

The constant-duration version of this last equation

$$u(y^{s+t}|x^\lambda) - u(y^t|x^\lambda) = \Delta w(y|x^\lambda)(w(y^{s+t}) - w(y^t)),$$

implies that $u(y^t|x^\lambda)$ and $w(y^t)$ are related by an affine transformation with slope $\Delta w(y|x^\lambda)$. Letting $w(y|x^\lambda)$ be the intercept, we have

$$u(y^t|x^\lambda) = w(y|x^\lambda) + \Delta w(y|x^\lambda)w(y^t).$$

Let t have distribution ν and take expectations to get

$$u(y^\nu|x^\lambda) = w(y|x^\lambda) + \Delta w(y|x^\lambda)w(y^\nu).$$

This establishes (16) and therefore completes the proof of Theorem 8.

Turning to the remaining necessary conditions in Theorem 7, note first that the affine restriction (8) requires that $\Delta u(y^\mu|q)u(h|\theta(q, y^\mu))$ be an affine function of the measure μ . For the semi-Markov preference summary, this restriction forces

$$\Delta u(y^\mu|x^\lambda)u(h|y^\mu) \quad \text{affine in } \mu,$$

$$\Delta u(y^\mu|y^\lambda)u(h|y^{\lambda*\mu}) \quad \text{affine in } \mu.$$

Using (A.8) we conclude

$$\Delta w(y^\mu)u(h|y^\mu) \quad \text{affine in } \mu.$$

For $z \neq y$, set $h = z^\nu\phi$ and then $h = z^\nu k$, where (by nontriviality) $w(z^\nu) \neq w(z^\nu k)$, to conclude using Theorem 8

$$\Delta w(y^\mu)(w(z|y^\mu) + \Delta w(z|y^\mu)w(z^\nu)) \quad \text{affine in } \mu;$$

$$\Delta w(y^\mu)(w(z|y^\mu) + \Delta w(z|y^\mu)w(z^\nu k)) \quad \text{affine in } \mu.$$

The difference of these two expressions

$$\Delta w(y^\mu)\Delta w(z|y^\mu)(w(z^\nu k) - w(z^\nu)),$$

is affine in μ , and is proportional to $\Delta w(y^\mu)\Delta w(z|y^\mu)$, so the latter is affine in μ as well. Subtract the affine $\Delta w(y^\mu)\Delta w(z|y^\mu)w(z^\nu)$ from the first affine expression to conclude that the result $\Delta w(y^\mu)w(z|y^\mu)$ is also affine in μ . This establishes the affine restrictions (13), (14).

Finally, to establish (15), suppose $x_1 \neq x_2 \neq \dots \neq x_n$, and let $\mathbf{x}_k = (x_k, \dots, x_n)$, $\mathbf{t}_k = (t_k, \dots, t_n)$. According to Theorem 8 and the recursive equations (10), we have

$$\begin{aligned} u(\mathbf{x}^t|\emptyset) &= u(\mathbf{x}^{t_1}|\emptyset) = u(x_1^{t_1}|\emptyset) + \Delta u(x_1^{t_1}|\emptyset)u(\mathbf{x}_2^{t_2}|x_1^{t_1}) \\ &= \left[\frac{w(x_1|\emptyset)}{\Delta w(x_1|\emptyset)} \right] \circ \left[\frac{w(x_1^{t_1})}{\Delta w(x_1^{t_1})} \right] u(\mathbf{x}_2^{t_2}|x_1^{t_1}). \end{aligned}$$

In the same way,

$$u(\mathbf{x}_2^{t_2}|x_1^{t_1}) = \left[\frac{w(x_2|x_1^{t_1})}{\Delta w(x_2|x_1^{t_1})} \right] \circ \left[\frac{w(x_2^{t_2})}{\Delta w(x_2^{t_2})} \right] u(\mathbf{x}_3^{t_3}|x_2^{t_2}),$$

and so on until

$$\begin{aligned} u(\mathbf{x}_n^{t_n}|x_{n-1}^{t_{n-1}}) &= \left[\frac{w(x_n|x_{n-1}^{t_{n-1}})}{\Delta w(x_n|x_{n-1}^{t_{n-1}})} \right] \circ \left[\frac{w(x_n^{t_n})}{\Delta w(x_n^{t_n})} \right] u(\phi|x_n^{t_n}) \\ &= \left[\frac{w(x_n|x_{n-1}^{t_{n-1}})}{\Delta w(x_n|x_{n-1}^{t_{n-1}})} \right] \circ \left[\frac{w(x_n^{t_n})}{\Delta w(x_n^{t_n})} \right] (0). \end{aligned}$$

Combine these to obtain $u(\mathbf{x}^t|\emptyset)$. Then (15) follows because $U(\mathbf{x}^t) \sim u(\mathbf{x}^t|\emptyset)$. This completes the proof of necessity in Theorem 7.

To show sufficiency, suppose U is specified by (15), where the affine restrictions (12) through (14) hold. To show U has a semi-Markov preference summary, we examine conditional preference over sample paths h given $\mathbf{x}'^{\lambda'}$, where $\mathbf{x}' = (x'_1, \dots, x'_n)$, $\lambda' = (\lambda'_1, \dots, \lambda'_n)$. If \mathbf{x}' has adjacent states $x'_{i+1} = x'_i$ which are the same, then we can concatenate $x'_i \lambda'_i x'_{i+1} \lambda'_{i+1}$ into $x'_i \lambda'_i \lambda'_{i+1}$, and we may continue concatenation until all adjacent states are distinct. Let \mathbf{x}^λ denote the result of this concatenation process, in which $x_1 \neq x_2 \neq \dots \neq x_n$. Note that the semi-Markov preference summary $e(\mathbf{x}'^{\lambda'})$ of $\mathbf{x}'^{\lambda'}$ is $e(\mathbf{x}'^{\lambda'}) = x_n^{\lambda_n}$. We intend to show that under the utility function U of (15), conditional preference given $\mathbf{x}'^{\lambda'}$ depends only on the semi-Markov summary $x_n^{\lambda_n}$.

Let h be a sample path which after concatenation of successive identical states, becomes \mathbf{y}^s , where $\mathbf{y} = (y_1, \dots, y_m)$, $\mathbf{s} = (s_1, \dots, s_m)$, and $y_1 \neq y_2 \neq \dots \neq y_m$. Consider first the case in which $x_n \neq y_1$. According to (2.15), if $\mathbf{T} = (T_1, \dots, T_n)$ are independent durations with distributions λ then

$$\begin{aligned}
 U(\mathbf{x}^\lambda \mathbf{y}^s) &= E_T[U(\mathbf{x}^T \mathbf{y}^s)] \\
 &= E_T \left[\begin{matrix} \left[\begin{matrix} w(x_1 | \emptyset) \\ \Delta w(x_1 | \emptyset) \end{matrix} \right] \circ \left[\begin{matrix} w(x_1^{T_1}) \\ \Delta w(x_1^{T_1}) \end{matrix} \right] \\ \circ \left[\begin{matrix} w(y_1 | x_n^{T_n}) \\ \Delta w(y_1 | x_n^{T_n}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_1^{s_1}) \\ \Delta w(y_1^{s_1}) \end{matrix} \right] \\ \circ \dots \circ \left[\begin{matrix} w(x_n | x_{n-1}^{T_{n-1}}) \\ \Delta w(x_n | x_{n-1}^{T_{n-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(x_n^{T_n}) \\ \Delta w(x_n^{T_n}) \end{matrix} \right] \\ \circ \dots \circ \left[\begin{matrix} w(y_m | y_{m-1}^{s_{m-1}}) \\ \Delta w(y_m | y_{m-1}^{s_{m-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_m^{s_m}) \\ \Delta w(y_m^{s_m}) \end{matrix} \right] \end{matrix} \right] \quad (0) \\
 &= \left[\begin{matrix} w(x_1 | \emptyset) \\ \Delta w(x_1 | \emptyset) \end{matrix} \right] \circ \left[\begin{matrix} w(x_1^{\lambda_1}) \\ \Delta w(x_1^{\lambda_1}) \end{matrix} \right] \\
 &\quad \circ \left[\begin{matrix} w(y_1 | x_n^{\lambda_n}) \\ \Delta w(y_1 | x_n^{\lambda_n}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_1^{s_1}) \\ \Delta w(y_1^{s_1}) \end{matrix} \right] \\
 &\quad \circ \dots \circ \left[\begin{matrix} w(x_n | x_{n-1}^{\lambda_{n-1}}) \\ \Delta w(x_n | x_{n-1}^{\lambda_{n-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(x_n^{\lambda_n}) \\ \Delta w(x_n^{\lambda_n}) \end{matrix} \right] \\
 &\quad \circ \dots \circ \left[\begin{matrix} w(y_m | y_{m-1}^{s_{m-1}}) \\ \Delta w(y_m | y_{m-1}^{s_{m-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_m^{s_m}) \\ \Delta w(y_m^{s_m}) \end{matrix} \right] \quad (0). \quad (A.11)
 \end{aligned}$$

The last equality results from n applications of the affine restrictions (12)–(14), which at the i th application looks like

$$\begin{aligned}
 E_{T_i} \left[\dots \circ \left[\begin{matrix} w(x_i^{T_i}) \\ \Delta w(x_i^{T_i}) \end{matrix} \right] \circ \left[\begin{matrix} w(x_{i+1} | x_i^{T_i}) \\ \Delta w(x_{i+1} | x_i^{T_i}) \end{matrix} \right] \circ \dots \right] \\
 &= E_{T_i} [\dots (w(x_i^{T_i}) + \Delta w(x_i^{T_i})(w(x_{i+1} | x_i^{T_i}) \\
 &\quad + \Delta w(x_{i+1} | x_i^{T_i})(\dots))) \\
 &= \dots (E_{T_i} [w(x_i^{T_i})] + (E_{T_i} [\Delta w(x_i^{T_i})w(x_{i+1} | x_i^{T_i})] \\
 &\quad + E_{T_i} [\Delta w(x_i^{T_i})\Delta w(x_{i+1} | x_i^{T_i})(\dots)])) \\
 &= \dots (w(x_i^{\lambda_i}) + (\Delta w(x_i^{\lambda_i})w(x_{i+1} | x_i^{\lambda_i}) \\
 &\quad + \Delta w(x_i^{\lambda_i})\Delta w(x_{i+1} | x_i^{\lambda_i})(\dots))) \\
 &= \dots \circ \left[\begin{matrix} w(x_i^{\lambda_i}) \\ \Delta w(x_i^{\lambda_i}) \end{matrix} \right] \circ \left[\begin{matrix} w(x_{i+1} | x_i^{\lambda_i}) \\ \Delta w(x_{i+1} | x_i^{\lambda_i}) \end{matrix} \right] \circ \dots
 \end{aligned}$$

The case where $x_n = y_1$ is similar. In $\mathbf{x}^\lambda \mathbf{y}^s$, concatenate $x_n^{\lambda_n} y_1^{s_1}$ into $x_n^{\lambda_n s_1}$. The analog of (A.11) is then

$$\begin{aligned}
 U(\mathbf{x}^\lambda \mathbf{y}^s) &= \left[\begin{matrix} w(x_1 | \emptyset) \\ \Delta w(x_1 | \emptyset) \end{matrix} \right] \circ \left[\begin{matrix} w(x_1^{\lambda_1}) \\ \Delta w(x_1^{\lambda_1}) \end{matrix} \right] \\
 &\quad \circ \left[\begin{matrix} w(y_2 | x_n^{\lambda_n} y_1^{s_1}) \\ \Delta w(y_2 | x_n^{\lambda_n} y_1^{s_1}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_2^{s_2}) \\ \Delta w(y_2^{s_2}) \end{matrix} \right] \\
 &\quad \circ \dots \circ \left[\begin{matrix} w(x_n | x_{n-1}^{\lambda_{n-1}}) \\ \Delta w(x_n | x_{n-1}^{\lambda_{n-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(x_n^{\lambda_n} y_1^{s_1}) \\ \Delta w(x_n^{\lambda_n} y_1^{s_1}) \end{matrix} \right] \\
 &\quad \circ \dots \circ \left[\begin{matrix} w(y_m | y_{m-1}^{s_{m-1}}) \\ \Delta w(y_m | y_{m-1}^{s_{m-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_m^{s_m}) \\ \Delta w(y_m^{s_m}) \end{matrix} \right] \quad (0). \quad (A.12)
 \end{aligned}$$

Now consider $U(\mathbf{x}^\lambda \mathbf{y}^s)$ as a utility function over \mathbf{y}^s with \mathbf{x}^λ fixed. Then in (A.11) and (A.12) all terms not involving \mathbf{y}^s are effectively constants. By removing all initial additive constant terms and initial positive constant factors, we see that as a function of \mathbf{y}^s

$$\begin{aligned}
 U(\mathbf{x}^\lambda \mathbf{y}^s) &\sim \left\{ \begin{matrix} \left[\begin{matrix} w(x_n^{\lambda_n}) \\ \Delta w(x_n^{\lambda_n}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_1 | x_n^{\lambda_n}) \\ \Delta w(y_1 | x_n^{\lambda_n}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_1^{s_1}) \\ \Delta w(y_1^{s_1}) \end{matrix} \right] \\ \left[\begin{matrix} w(x_n^{\lambda_n} y_1^{s_1}) \\ \Delta w(x_n^{\lambda_n} y_1^{s_1}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_2 | x_n^{\lambda_n} y_1^{s_1}) \\ \Delta w(y_2 | x_n^{\lambda_n} y_1^{s_1}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_2^{s_2}) \\ \Delta w(y_2^{s_2}) \end{matrix} \right] \\ \dots \circ \left[\begin{matrix} w(y_m | y_{m-1}^{s_{m-1}}) \\ \Delta w(y_m | y_{m-1}^{s_{m-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_m^{s_m}) \\ \Delta w(y_m^{s_m}) \end{matrix} \right] \quad (0) \text{ if } y_1 \neq x_n, \\ \dots \circ \left[\begin{matrix} w(y_m | y_{m-1}^{s_{m-1}}) \\ \Delta w(y_m | y_{m-1}^{s_{m-1}}) \end{matrix} \right] \circ \left[\begin{matrix} w(y_m^{s_m}) \\ \Delta w(y_m^{s_m}) \end{matrix} \right] \quad (0) \text{ if } y_1 = x_n. \end{matrix} \right.
 \end{aligned}$$

It follows that conditional preference given \mathbf{x}^λ depends only on the semi-Markov summary $e(\mathbf{x}^\lambda) = x_n^{\lambda_n}$, as claimed. QED.

ACKNOWLEDGMENT

This research was supported by grant number SES-9112854 from the National Science Foundation. Thanks to Chelsea White for pointing out the previous work by David Kreps.

REFERENCES

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

BARRAGER, S. M. 1980. Assessment of Simple Joint Time-risk Preference Functions. *Mgmt. Sci.* **26**, 620–632.

BECK, J. R. AND S. G. PAUKER. 1981. Anticoagulation and Atrial Fibrillation in the Brachycardia-tachycardia syndrome. *Med. Dec. Making*, **1**, 285–301.

BELL, D. E. 1977. A Utility Function for Time Streams Having Inter-period Dependencies. *Opns. Res.* **25**, 448–458.

BERTSEKAS, D. 1976. *Dynamic Programming and Stochastic Control*. Academic Press, New York.

BODILY, S. E. 1980. Analysis of Risks to Life and Limb. *Opns. Res.* **28**, 156–175.

CHANG, R. W., J. M. PELLISSIER, AND G. B. HAZEN. 1993. A Cost-effectiveness Analysis of Total Hip Arthroplasty for Osteoarthritis of the Hip (Abstract). *Clinical Res.* **41**, 235A, National Meeting, Washington DC.

CHUNG, K.-J. AND M. J. SOBEL. 1987. Discounted MDP's: Distribution Functions and Exponential Utility Maximization. *SIAM J. Control and Optimization*, **25**, 49–62.

CINLAR, E. 1975. *Introduction to Stochastic Processes*. Prentice-Hall, Englewood Cliffs, New Jersey.

DUFFIE, D. AND L. G. EPSTEIN. 1992. Stochastic Differential Utility. *Econometrica*, **60**, 353–394.

EPSTEIN, L. G. AND S. E. ZIN. 1989. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, **57**, 937–969.

- FISHBURN, P. C. 1965. Markovian Dependence in Utility Theory with Whole Product Sets. *Opns. Res.* **13**, 238–257.
- FISHBURN, P. C. 1978. Ordinal Preferences and Uncertain Lifetimes. *Econometrica*, **46**, 817–833.
- HARVEY, C. M. 1988. Utility Functions for Infinite-period Planning. *Mgmt. Sci.* **34**, 645–665.
- HAZEN, G. B. 1992. Stochastic trees: A New Technique for Temporal Medical Decision Modeling. *Med. Dec. Making*, **12**, 163–178.
- HAZEN, G. B. 1993. Factored Stochastic Trees: A Tool for Solving Complex Temporal Medical Decision Problems. *Med. Dec. Making*, **13**, 227–236.
- HAZEN, G. B., W. J. HOPP, AND J. M. PELLISSIER. 1991. Continuous-risk Utility Assessment in Medical Decision Making. *Med. Dec. Making*, **11**, 294–304.
- HILLNER, B. E., J. P. HOLLENBERG, AND S. G. PAUKER. 1986. Postmenopausal Estrogens in Prevention of Osteoporosis. *Amer. J. Medicine*, **80**, 1115–1127.
- HOWARD, R. A. AND J. E. MATHESON. 1972. Risk-sensitive Markov Decision Processes. *Mgmt. Sci.* **18**, 356–369.
- KEENEY, R. L. AND H. RAIFFA. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York.
- KREPS, D. M. 1977a. Decision Problems with Expected Utility Criteria, I: Upper and Lower Convergent Utility. *Math. Opns. Res.* **2**, 45–53.
- KREPS, D. M. 1977b. Decision Problems with Expected Utility Criteria, II: Stationarity. *Math. Opns. Res.* **2**, 266–274.
- KREPS, D. M. 1978. Sequential Decision Problems with Expected Utility Criteria, III: Upper and Lower Transience. *SIAM J. Control and Optimization*, **16**, 420–428.
- KREPS, D. M. AND E. L. PORTEUS. 1978. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, **46**, 185–200.
- MATCHAR, D. B. AND S. G. PAUKER. 1986. Transient Ischemic Attacks in a Man with Coronary Artery Disease: Two Strategies Neck and Neck. *Med. Dec. Making*, **6**, 239–249.
- MEYER, R. F. 1976. Preferences over Time. In *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Chapter 9. Wiley, New York.
- MITTEN, L. G. 1971. Preference Order Dynamic Programming. *Mgmt. Sci.* **21**, 43–46.
- MOONEY, C., A. I. MUSHLIN, AND C. E. PHELPS. 1990. Targeting Assessments of Magnetic Resonance Imaging in Suspected Multiple Sclerosis. *Med. Dec. Making*, **10**, 77–94.
- PELLISSIER, J. M. AND G. B. HAZEN. 1994. Implementation of Continuous-risk Utility Assessment: The Total Hip Replacement Decision. *Socio-Economic Planning Sci.* **28**, 251–276.
- PLANTE, D. A., J. F. PICCIRILLO, AND R. A. SOFFERMAN. 1987. Decision Analysis of Treatment Options in Pyriform Sinus Carcinoma. *Med. Dec. Making*, **7**, 74–83.
- PLISKIN, J. S., D. S. SHEPARD, AND M. C. WEINSTEIN. 1980. Utility Functions for Life Years and Health Status. *Opns. Res.* **28**, 206–224.
- ROACH, P. J., C. FLEMING, M. D. HAGEN, AND S. G. PAUKER. 1988. Prostatic Cancer in a Patient with Asymptomatic HIV Infection: Are Some Lives More Equal Than Others? *Med. Dec. Making*, **8**, 132–144.
- ROSS, S. 1983. *Introduction to Stochastic Dynamic Programming*. Academic Press, New York.
- SNIEDOVICH, M. 1981. Some Comments on Preference Order Dynamic Programming Models. *J. Math. Analysis and Applications*, **79**, 489–501.
- SOBEL, M. J. 1975. Ordinal Dynamic Programming. *Mgmt. Sci.* **21**, 967–975.
- TORRENCE, G. W. 1976. Health Status Index Models: A Unified Mathematical View. *Mgmt. Sci.* **22**, 990–1001.
- TORRENCE, G. W., M. H. BOYLE, AND S. P. HORWOOD. 1982. Application of Multi-attribute Utility Theory to Measure Social Preference for Health States. *Opns. Res.* **30**, 1043–1069.
- WEINSTEIN, M. C. AND W. B. STASON. 1976. *Hypertension: A Policy Perspective*. Harvard University Press, Cambridge MA.

Copyright 1996, by INFORMS, all rights reserved. Copyright of Operations Research is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.