Use of Richardson's Extrapolation to Reduce Timing Errors in Markov Models

Pelham Barton
Health Economics Unit
Public Health Building
University of Birmingham
B15 2TT
p.m.barton@bham.ac.uk

Status: Work at an early stage of development.

Background: For computational efficiency, it is desirable to use as large a cycle length as possible in a Markov model. However, large cycle lengths incur the risk of timing errors. Methods such as half-cycle correction and Simpson's rule can account for timing errors in calculating total costs and outcomes from Markov models, under the assumption that the state probabilities have been calculated correctly. However, it is also possible that timing errors may have been made in estimating the transition probabilities. Particular cases where there is scope for such errors are the presence of competing risks and the possibility of more than one event happening within a cycle, in particular for a progressive disease where some individuals may progress by more than one stage within a cycle. In principle these problems can be addressed by careful calculation of the probabilities: an alternative approach is called Richardson's extrapolation, in which the model is constructed and run using two or more different cycle lengths. Comparing the results from different cycle lengths allows an estimate of the timing errors to be made and hence an estimate of the results without timing errors.

Aims: To explore the possibility of using Richardson's extrapolation in Markov models.

Methods: Hypothetical models involving multi-state progression and competing risks were constructed both as continuous time models and discrete time Markov models and the results compared.

1. Introduction

The enormous advance in computing power over the last few decades has enabled more and more sophisticated models to become computationally feasible. It has also enabled the use of complex evaluations of such models, including estimates of the expected value of perfect parameter information and expected value of sample information. Such methods are potentially highly sensitive to any inaccuracy in the representation of the "real world" within the model. Also, they require a very large number of model runs with different parameter sets. Accordingly, there is a greater need than ever for methods of implementation of models which allows efficient calculation without loss of accuracy.

A commonly used model type in economic evaluation of health care technologies is the discrete time Markov model. In this type of model, possible conditions of an individual patient are grouped into "health states" so that by definition a patient is only in one state at any given time. The model runs with a fixed time cycle: the probability that a patient is in a given state at the end of a cycle depends only on the state at the start of the cycle. This "memoryless" property is the essential feature indicated by the use of the name "Markov": a further assumption of any specific model is that each state is sufficiently homogeneous that the model represents a reasonable approximation to reality.

When a Markov model is used in cost-effectiveness analysis, total costs and QALYs for each treatment strategy are estimated. There are two potential sources of error in such models due to the use of discrete time. First, there may be errors in the estimated probabilities of patients occupying each state at the end of each cycle and second, there may be errors in aggregating across cycles. The second potential source of error is traditionally handled through half-cycle correction (Sonnenberg and Beck, 1993) although caution is recommended in the unthinking application of this technique (Barton, 2009).

This paper is principally concerned with the first type of error. Two illustrative models have been developed. They are described first as continuous time models and then various discrete time implementations are shown and the results compared. Of particular

interest is the way in which the discretisation error varies with the cycle length. If the model is implemented with two different cycle lengths, the results can be compared and a process known as Richardson's extrapolation used to give a more precise estimate of the true value (Richardson, 1910).

Section 2 of this paper introduces the two examples as continuous time models. Section 3 shows how the continuous time models can be converted into discrete time Markov models preserving the full accuracy of state changes. Section 4 considers simplified versions of the models where the transition probabilities are calculated directly from hazard rates and compares the results for implementation with a range of different cycle lengths. In Section 5, the use of Richardson's extrapolation is explored. A discussion follows in Section 6.

2. Illustrative continuous time models

We consider two examples of five state Markov models. In each case, a patient in any of the four live states is at risk of moving to a worse state. The risk is assumed to be at a constant hazard rate in continuous time, but may be modified by treatment. Then it is clear that a patient's expected quality of life (QoL) varies continuously with time. The same is not necessarily true for costs, which are often incurred at discrete times. Two options for the costing are considered here: one is where "one off" costs are incurred every 4 months, and the other is where costs are assumed to be incurred continuously.

In both models, the unit costs are assumed to vary by health state, and apply to the treatment arm only. The assumption in the "discrete costs" case is that treatment is issued at regular appointments which are booked in advance. If the patient changes health state between appointments, the new cost will be incurred from the next appointment, but it is assumed that the risk of progressing a second or subsequent time within the 4 month cycle will be the same as for a patient on treatment. If the patient dies between appointments, the full cost for that 4 month period has already been incurred, and no refund is possible. Note that in such cases, it is important to avoid what would best be described as "half-cycle incorrection" in calculating the costs (Barton, 2009).

In the "continuous costs" case, it is assumed that the new level of costs is incurred immediately on progression to a more advanced state of disease, but equally that costs cease to be incurred immediately on the death of the patient. This might be a reasonable assumption if the costs are largely made up of nursing care, for example, or medication that is held by the provider, so that pills from a partially used pack could be given to another patient.

2.1 Multi-state progressive condition

In this model, the condition is assumed to be relentlessly progressive through four clearly defined stages: Mild, Moderate, Severe, and Terminal. It is assumed that progression in continuous time is through these four stages. Other cause deaths are excluded from the illustrative model for simplicity of exposition. Figure 1 shows the allowable transitions within the model and the necessary data inputs are in Table 1. Note that although progression in continuous time can only be one stage at a time, any forward progression is possible in the discrete time model.

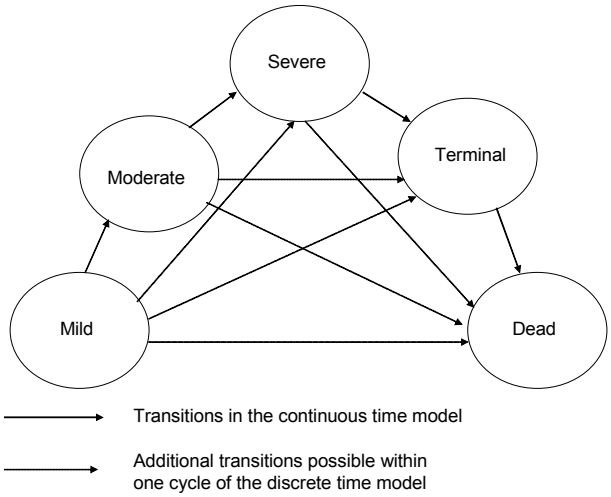

Figure 1 The "progressive condition" model

Table 1. Data inputs for the "progressive condition" model

			Hazard rate for progression (/month)	
State	Cost	QoL	No treatment	With treatment
Mild	600	0.8	0.2	0.15
Moderate	700	0.6	0.25	0.2
Severe	810	0.4	0.3	0.25
Terminal	930	0.2	0.35	0.3

Costs are either "one off" costs incurred every 4 months or the cost per 4 months in continuous time. QoL scores are on the usual scale where 1 = full health and 0 = equivalent to dead. Progression in the continuous time model is to the next state only (from Terminal to Dead).

2.2 The "competing risks" model

In this model, we assume that the initial ill state (called "progression free") is subject to two separate forms of progression, A and B. In this model, death is possible from any state. The treatment reduces the mortality from any state, and reduced the risk of progression factor A. In this case, the treatment has no effect on progression factor B, but indirectly reduces the risk of progression factor B by delaying progression factor A. The idea behind this selection is that factor A is an exacerbation of the original disease, while factor B is a comorbidity: abstract descriptions are used to avoid unnecessary discussion of features of a particular disease. Figure 2 shows the transitions possible, with the data inputs in Table 2. Again, any worsening is possible in the discrete time model. Note that it is never possible to move between the states "Progression A" and "Progression B". To clarify the progressions in factors A and B, the monthly hazard for "Progression free" to "Progression B" or for "Progression A" to "Progressions A and B" is 0.2 without treatment, reduced to 0.1 with treatment. The hazard for "Progression free" to "Progression B" is 0.15 regardless of treatment, but the hazard for "Progression A" to "Progressions A and B" is 0.3 regardless of treatment, so slowing progression factor A indirectly slows progression factor B in this model.

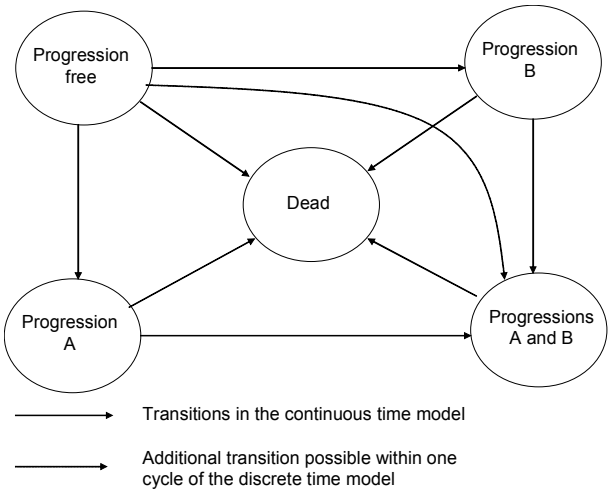

Figure 2. The "competing risks" model

Table 2. Data inputs for the "competing risks" model.

			Hazard rates (/month) for progression					
			Factor A		Factor B		To Death	
State	Cost	QoL	No Rx	Rx	No Rx	Rx	No Rx	Rx
Prog free	900	0.9	0.2	0.1	0.15	0.15	0.05	0.04
Prog A	1020	0.5			0.3	0.3	0.1	0.08
Prog B	1180	0.8	0.2	0.1			0.2	0.16
Prog AB	1350	0.2					0.4	0.32

Rx = treatment Prog = Progression(s). Costs are either "one off" costs incurred every 4 months or the cost per 4 months in continuous time. QoL scores are on the usual scale where 1 = full health and 0 = equivalent to dead.

2.3 Running the continuous time models

The continuous time models can be run using numerical methods which are beyond the scope of this paper. Costs and QALYs accumulated are discounted at 3.5% per annum using a continuous discounting function. Sufficient computational effort has been made to ensure that the results may be regarded as an exact representation of the continuous time model. The time horizon for the "progressive condition" model was set at 72 months, and for the "competing risks" model 40 months, by which time the probability of death on treatment exceeds 99.96% in each case and the probability of death without

treatment exceeds 99.99%. The results are shown in Table 3. Suspicions that the costs in the various states have been carefully selected to produce ICERs (for the discrete cost case) just below £20,000/QALY are entirely justified. The results here are taken as the "exact" results for the reality being modelled. They are given to greater precision than would normally be justified for a model to show the differences between the various forms for the discrete time models.

Table 3. Results for continuous time models

	Progressive condition model	Competing risks model
Costs discrete	3631.54	2498.02
Costs continous	3336.17	2053.48
QALYs with treatment	0.858046	0.446974
QALYs without treatment	0.676434	0.322064
Difference in QALYs	0.181612	0.124910
ICER discrete costs	19,996	19,999
ICER continuous costs	18,370	16,440

3. Discrete time Markov models

To build a discrete time Markov model with a fixed time cycle T involves estimating the transition probabilities. For any pair of states X and Y, we need to estimate the probability that an individual in state X at any given time will be in state Y after a further time T. Given the continuous time model, these transition probabilities can be found by successive use of starting populations with a probability of 1 of being in each of the states. The resulting transition matrices with a time cycle of 4 months are shown in Table 4.

Table 4. Transition matrices for exact Markov models with 4 month cycle time

(*a*) Progressive condition model with no treatment

from \ to	Mild	Moderate	Severe	Terminal	Dead
Mild	0.449329	0.325798	0.147643	0.053526	0.023704
Moderate	0	0.367880	0.333425	0.181320	0.117375
Severe	0	0	0.301195	0.327583	0.371223
Terminal	0	0	0	0.246598	0.753402

(*b*) Progressive condition model with treatment

from \ to	Mild	Moderate	Severe	Terminal	Dead
Mild	0.548812	0.298448	0.108199	0.032689	0.011853
Moderate	0	0.449329	0.325798	0.147643	0.077230
Severe	0	0	0.367880	0.333425	0.298695
Terminal	0	0	0	0.301195	0.698805

(*c*) Competing risks model with no treatment

from \ to	Prog free	Prog A	Prog B	Prog AB	Dead
Prog free	0.201897	0.161517	0.121137	0.145365	0.370083
Prog A	0	0.201897	0	0.242275	0.555828
Prog B	0	0	0.201897	0.161517	0.636586
Prog AB	0	0	0	0.201897	0.798103

(*d*) Competing risks model with treatment

from \ to	Prog free	Prog A	Prog B	Prog AB	Dead
Prog free	0.313487	0.105304	0.199842	0.101954	0.279413
Prog A	0	0.218713	0	0.296626	0.484661
Prog B	0	0	0.353455	0.125695	0.520850
Prog AB	0	0	0	0.278038	0.721962

In each case, "Dead" is the unique absorbing state: the transition probability from "Dead" to "Dead" is 1; from "Dead" to any other state, 0.

When a standard cohort simulation is run using any of the transition matrices in Table 4, the probability distribution at the end of each cycle is exactly the same as for the continuous time model (to within computer accuracy). When costs are assumed to be incurred at discrete times, simply adding the (discounted) expected cost using the

distribution at the start of each cycle also gives exactly the same answer as for the continuous model.

If costs are assumed to be incurred continuously, and to estimate (discounted) QALYs in any case, it is a question of estimating the area under a smooth curve given the heights of certain points on that curve. This is usually done by a method called "half-cycle correction" (see Sonnenberg and Beck, 1993, for the method), but is better done by Simpson's rule (Barton, 2009: see Appendix). Similar remarks apply for costs under the "continuous costs" assumption. Table 5 gives the results of applying these methods to the Markov model with these transition matrices: they can be seen to be a very good approximation to the continuous model, even with a cycle length as long as 4 months.

Table 5 Comparing continuous time model with discrete time Markov model based on exact transition matrices

Conceptual model	Progressive condition		Competing risks	
Implementation	continuous	discrete	continuous	discrete
Costs discrete	3631.54	3631.54	2498.02	2498.02
Costs continous	3336.17	3337.06	2053.48	2050.87
QALYs with Rx	0.858046	0.858068	0.446974	0.445748
QALYs without Rx	0.676434	0.676462	0.322064	0.320292
Difference in QALYs	0.181612	0.181606	0.124910	0.125456
ICER discrete costs	19,996	19,997	19,999	19,911
ICER continuous costs	18,370	18,375	16,440	16,347

4. Approximating the Markov transition matrices from the hazard rates

The transition matrices in Table 4 take proper account of the possibility of an individual making more than one change of state during a cycle of the discrete time model, as shown by the non-zero entries in the positions corresponding to the dotted arrows in Figures 1 and 2. Unfortunately, there is in general no closed algebraic form for these entries in terms of the underlying risks. Therefore, if the effort of constructing the continuous time model is to be avoided, some assumption must be made. The simplest assumption is that at most one progression can happen in each cycle. This assumption

becomes more reasonable the shorter the cycle length. While this argument works in favour of short cycle lengths, it goes against the desire for speed in evaluating the model. In either model, the probability of not progressing within a cycle is easily calculated to be $e^{-\lambda T}$, where λ is the sum of the hazard rates for progressing and T is the cycle length. For the "progressive condition" model, this is sufficient to determine the transition matrix for the simplified model with any cycle length. For the "competing risks" model, matters are more complicated. It is therefore convenient at this stage to consider the two models separately.

4.1 The "progressive condition" model

Table 6 shows the transition matrices for the simplified version of this model with a 4 month cycle. Comparing with Table 4, all of the probabilities to the right of the diagonal on any row have been combined into the probability of progression by a single step, and effectively the dotted arrows in Figure 1 have been removed from the model. Clearly these transition matrices are quite different, and the assumption of only one progression within 4 months is not reasonable.

Table 6. Transition matrices for the simplified "progressive condition" model with a 4 month cycle

(*a*) With no treatment

from \ to	Mild	Moderate	Severe	Terminal	Dead
Mild	0.449329	0.550671	0	0	0
Moderate	0	0.367879	0.632121	0	0
Severe	0	0	0.301194	0.698806	0
Terminal	0	0	0	0.246597	0.753403

(*b*) With treatment

from \ to	Mild	Moderate	Severe	Terminal	Dead
Mild	0.548812	0.451188	0	0	0
Moderate	0	0.449329	0.550671	0	0
Severe	0	0	0.367879	0.632121	0
Terminal	0	0	0	0.301194	0.698806

Similar methods may be used to derive transition matrices for shorter cycle times. Evaluating accumulated QALYs follows the same principles regardless of cycle length, and this is also true for costs under the assumption that costs are incurred continuously. For costs incurred at discrete times, it is convenient to consider the example of a 1 month cycle time. In this case, costs are incurred at the beginning of the model, and then at the beginning of every fourth cycle thereafter. Similar principles apply to other cycle times which divide exactly into the 4 month interval between cost-incurring events. Some form of interpolation method would be needed if the cycle length does not divide exactly into 4 months: it would be silly to use such a cycle length when there is no need to do so.

Table 7 shows the results from the simplified model with a range of cycle lengths, compared to the exact results shown earlier. Discussion of these results appears in Section 4.3, after consideration of the "competing risks" model.

Table 7 Results for the "progressive condition" model with a variety of cycle lengths

Model	continuous	Simplified (discrete time Markov)			
Cycle length	(none)	4 months	2 months	1 month	0.5 month
Costs discrete	3631.54	4996.02	4262.41	3934.14	3779.50
Costs continous	3336.17	4703.83	3966.89	3638.10	3483.74
QALYs with Rx	0.858046	1.086993	0.963008	0.908123	0.882484
QALYs without Rx	0.676434	0.917434	0.784907	0.727625	0.701269
Difference in QALYs	0.181612	0.169559	0.178102	0.180498	0.181215
ICER discrete costs	19,996	29,465	23,932	21,796	20,856
ICER continuous costs	18,370	27,742	22,273	20,156	19,224

4.2 The "competing risks" model

Where direct transition to more than one different state is possible from a given starting state *X*, the probability of remaining in that state for a time *T* is still $e^{-\lambda T}$, where λ is the sum of the hazard rates for progressing. Suppose that it is possible to progress to three different states Y_1, Y_2, and Y_3, and that the hazard rate for progression to state Y_i is a

constant λ_i, so that $\lambda = \lambda_1 + \lambda_2 + \lambda_3$. Then it is not difficult to see that the probability of transition to state Y_i conditional upon some transition being made at a time $t < T$ is simply $\frac{\lambda_i}{\lambda}$. Therefore the transition probability from X to Y_i required for the simplified version of the model is $\frac{\lambda_i}{\lambda} e^{-\lambda T}$. These transition probabilities would be exactly correct if all three states Y_i were absorbing states. Using these principles, the transition matrices for the simplified model with a 4 month cycle length appear in Table 8. These matrices may be compared with those in Table 4(*c*) and (*d*).

Table 8 Transition matrices for the simplified "competing risks" model with a 4 month cycle

(*a*) With no treatment

from \ to	Prog free	Prog A	Prog B	Prog AB	Dead
Prog free	0.201897	0.399052	0.299289	0	0.099763
Prog A	0	0.201897	0	0.598578	0.199526
Prog B	0	0	0.201897	0.399052	0.399052
Prog AB	0	0	0	0.201897	0.798103

(*b*) With treatment

from \ to	Prog free	Prog A	Prog B	Prog AB	Dead
Prog free	0.313486	0.236729	0.355093	0	0.094692
Prog A	0	0.218712	0	0.616806	0.164482
Prog B	0	0	0.353455	0.248671	0.397874
Prog AB	0	0	0	0.278037	0.721963

As before, the diagonal entries are unchanged, and the entry corresponding to the dotted arrow in Figure 2 has become 0. However, this time the probabilities have also changed even when remaining non-zero. For example, in the "exact" model with no treatment, there is a probability of nearly 64 percent of dying within 4 months if in the state "Prog AB": see Table 4(*c*). The corresponding entry in Table 8(*a*) is just under 40 percent: the

remaining 24 percent is accounted for by people who (in the assumed reality) progress to "Prog AB" and then die within the 4 months.

Again, this model has been run with a range of cycle lengths. The results are shown in Table 9, and discussion of them follows in Section 4.3.

Table 9 Results for the "competing risks" model with a variety of cycle lengths

Model	continuous	Simplified (discrete time Markov)			
Cycle length	(none)	4 months	2 months	1 month	0.5 month
Costs discrete	2498.02	3515.70	2947.32	2704.71	2596.94
Costs continous	2053.48	3066.45	2504.40	2263.62	2154.64
QALYs with Rx	0.446974	0.604255	0.513496	0.477358	0.461486
QALYs without Rx	0.322064	0.491759	0.391116	0.352735	0.336493
Difference in QALYs	0.124910	0.112496	0.122380	0.124623	0.124993
ICER discrete costs	19,999	31,252	24,083	21,703	20,777
ICER continuous costs	16,440	27,258	20,464	18,164	17,238

4.3 Comparison of model results with different cycle lengths

In Tables 7 and 9, it can clearly be seen that the simplified Markov model overestimates the costs and QALYs gained with or without treatment. This makes sense because the simplified model is in effect slowing down progression. It can also be seen that the errors reduce with cycle length. The modelling errors are shown in Table 10 for each of these outcomes. It can be seen that the errors approximately halve as the cycle length halves. This effect can be seen more clearly from the figures in Table 11, which give the ratio of modelling errors between successive columns in Table 10. The ratios are all slightly over 2, and themselves become closer to 2 as the cycle length reduces.

Table 10 Modelling errors for the discrete time Markov models

(*a*) the "progressive condition" model

Cycle length	4 months	2 months	1 month	0.5 month
Costs discrete	1364.48	630.87	302.61	147.96
Costs continous	1367.67	630.72	301.93	147.57
QALYs with Rx	0.228947	0.104962	0.050077	0.024437
QALYs without Rx	0.241000	0.108472	0.051191	0.024834

(*b*) the "competing risks" model

Cycle length	4 months	2 months	1 month	0.5 month
Costs discrete	1017.68	449.29	206.69	98.92
Costs continous	1012.97	450.92	210.14	101.16
QALYs with Rx	0.157281	0.066522	0.030384	0.014511
QALYs without Rx	0.169694	0.069052	0.030671	0.014428

Table 11. Ratio of modelling errors between discrete time Markov models when the greater cycle length is double the lesser cycle length

Model	Progressive condition			Competing risks		
Greater cycle length	4 mo	2 mo	1 mo	4 mo	2 mo	1 mo
Costs discrete	2.163	2.085	2.045	2.265	2.174	2.090
Costs continous	2.168	2.089	2.046	2.246	2.146	2.077
QALYs with Rx	2.181	2.096	2.049	2.364	2.189	2.094
QALYs without Rx	2.222	2.119	2.061	2.458	2.251	2.126

The evidence in Tables 10 and 11 supports the view that the results from the discrete time Markov models converge to the results from the continuous time model as the cycle time becomes shorter. This must therefore also be the case for the outcomes in terms of difference in QALYs and ICER. There appears to be a slight anomaly in Table 9 with regard to the difference in QALYs, where the result for the 0.5 month cycle length is higher than the result for the continuous model. This model has also been run for cycle lengths 0.25 months and 0.125 months. The full results are not shown, but the difference in QALYs was calculated as 0.125004 and 0.124970 respectively, supporting the view that the values do indeed converge to the "correct" value of 0.124910. While the turning

point in the function mapping cycle length to modelled difference in QALYs was unexpected, there is no good reason why it should not happen, and indeed the general pattern was maintained when changes were made to the assumptions about hazard rates.

5. Use of Richardson's extrapolation

We now have sufficient background to reach the main point of this paper. Given the way the simplified discrete time Markov models were constructed, it is reasonable to expect that any modelled outcome measure *M* can be expressed as what is known as an analytic function of the cycle length *h*, by which is meant a function which can be expressed as a convergent Taylor series

$$M(h) = M_0 + M_1 h + M_2 h^2 + M_3 h^3 + \ldots \qquad (*)$$

for sufficiently small *h*. Here, M_0 is the "true" value of the outcome measure. The fact that the numbers in Table 10 are roughly proportional to the cycle length tells us both that M_1 is non-zero and that the term $M_1 h$ accounts for most of the error. The fact that the fraction parts of the numbers in Table 11 are themselves roughly proportional to the cycle length tells us that M_2 is non-zero and that the term $M_2 h^2$ accounts for most of the remaining error after accounting for the term $M_1 h$; the fact that the proportionality is only approximate tells us that the "higher" terms in the series are also present.

Note that if a Taylor series of this sort applies separately to the estimated QALYs with and without treatment, then it must also apply to the difference in QALYs, simply by subtracting the coefficients of the two individual series: the anomalous pattern in the case of the "competing risks" model tells us that the leading terms in the series are not the largest terms until *h* is considerably less than 1 month. It is also true that the ICERs can be expressed as a Taylor series, although the relationship between the coefficient of the ICER series and the series for differences in costs and QALYs is more complicated.

Given a complete Taylor series of the form (*), one can then calculate

$$R(h) = 2M(h) - M(2h) = M_0 - 2M_2 h^2 - 6M_3 h^3 + \ldots,$$

and if the term $M_1 h$ accounts for most of the error in estimating *M*, then $R(h)$ is likely to be a much improved estimate compared to $M(h)$. This idea was first proposed by LF Richardson (1910) and hence is known as "Richardson's extrapolation". In fact, in Richardson's case, the Taylor series for his model only contained even powers of *h*, so he could go straight to $R(h) = \frac{4}{3}M(h) - \frac{1}{3}M(2h) = M_0 - 4M_4 h^4 + \ldots$.

The potential value in the evaluation of Markov models is that the total computational effort in running a model with (say) a cycle time of 4 months and then again with a cycle time of 2 months is likely to be less than the effort required to run the model with a cycle time of 1 month, and much less than the effort required to run the model with a shorter cycle time. Applying Richardson's extrapolation to the results of such a procedure may therefore allow an improvement in the efficiency of modelling without loss of accuracy. This is effectively applying to discrete time Markov models an observation that has been made in the similar context of numerical evaluation of continuous time models (Barton and Tobias, 1998).

Returning to our examples, Table 12 shows the results of applying Richardson's extrapolation to the outcomes recorded in Tables 7 and 9. Again, the results from the continuous model are included for comparison. Two methods for calculating the ICER have been shown: the first is to use the costs and difference in QALYs from the extrapolated models and the second is to apply extrapolation directly to the modelled ICERs.

Comparing Table 12 with Tables 7 and 9, the results from Richardson's extrapolation appear to compare reasonably favourably with the results from single runs of the model involving equivalent computational effort. This point is discussed further in Section 6.

Table 12. Results of applying Richardson's extrapolation to the results from the simplified discrete time Markov models

(*a*) the "progressive condition" model

Greater cycle length	4 months	2 months	1 month	continuous
Costs discrete	3528.80	3605.88	3624.85	3631.54
Costs continous	3229.94	3309.31	3329.38	3336.17
QALYs with Rx	0.839023	0.853238	0.856844	0.858046
QALYs without Rx	0.652379	0.670344	0.674912	0.676434
Difference in QALYs	0.186644	0.182894	0.181932	0.181612
ICER discrete costs	18,907	19,716	19,924	19,996
ICER continuous costs	17,305	18,094	18,300	18,370
Richardson's extrapolation applied directly to modelled ICERs:				
ICER discrete costs	18,400	19,660	19,917	19,996
ICER continuous costs	16,805	18,039	18,293	18,370

(*b*) the "competing risks" model

Greater cycle length	4 months	2 months	1 month	continuous
Costs discrete	2378.93	2462.11	2489.17	2498.02
Costs continous	1942.34	2022.84	2045.65	2053.48
QALYs with Rx	0.422737	0.441221	0.445613	0.446974
QALYs without Rx	0.290473	0.314355	0.320250	0.322064
Difference in QALYs	0.132263	0.126866	0.125363	0.124910
ICER discrete costs	17,986	19,407	19,856	19,999
ICER continuous costs	14,685	15,945	16,318	16,440
Richardson's extrapolation applied directly to modelled ICERs:				
ICER discrete costs	16,915	19,323	19,850	19,999
ICER continuous costs	13,670	15,863	16,312	16,440

6. Discussion

The astute reader will have noticed that, although the style of this paper so far has generally followed the usual principle of a draft journal article, the occasional sentence has been included which is only suitable "among friends" in an HESG paper. At this

point, I abandon all pretence of drafting an article, and use my remaining space to address myself directly to an HESG audience.

The underlying concept behind Richardson's extrapolation is one that I thought of for myself while working towards a PhD in simulation modelling from 1995 to 1998. I had a draft paper almost ready for submission to the journal *Simulation* when I found that the idea was far from new, and had to rewrite the paper in the form of the application of an existing idea to a new area. This is the paper that was eventually published in *System Dynamics Review* and appears in the reference list. It was my first published paper.

It has long been at the back of my mind to try the application of the same ideas to Markov models as used in Health Economics. Submitting an abstract to HESG is the best way of creating an obligation to do the work, and I have found a gap in my funded work to produce the results shown here. All of the calculations were done in the week before the submission date for this paper.

Things I would have liked to try, but for which I have not yet found the time, include even greater simplification in the construction of models, for example calculating transition probabilities by multiplying the hazard rates directly by the cycle length, and abandoning Simpson's rule (or even half-cycle correction) in favour of what may be called the "rectangle rule" for estimating total costs and QALYs. When an extrapolation method is used, it does not matter that the individual models are at all reasonable: what matters is that they are part of a sequence which converges smoothly to an accurate representation of the underlying conceptual model.

Compared to the results in my 1998 paper, the results here are somewhat disappointing in terms of the potential saving in computational effort without loss of accuracy. This may be because using Simpson's rule and correct conversion of rates to probabilities is already nearly good enough.

I have not had the opportunity before submitting this paper to test how robust the general pattern of my conclusions is to variations in the model inputs. Something else I would

like to consider is applying Richardson's extrapolation directly to the transition matrices. One possible approach here is to square the (simplified) transition matrix for a 2 month cycle and compare it with the directly calculated matrix for a 4 month cycle. I may do some of this before the meeting.

The most important question to raise is whether the ideas are worth pursuing for real life modelling. I have made no attempt at this stage to tackle the issue of how to apply these methods from observed data, let alone combine them with the techniques of probabilistic sensitivity analysis. One thought is to create simulated samples from an underlying continuous model, and then use a variety of methods to construct transition matrices with distributions around the parameters from which to generate such things as CEACs and 95 percent credible intervals for model outputs, and then see how often the true value came within the credible interval. One point to note here for the two underlying models is that there are more degrees of freedom in the transition matrices than in the continuous models.

Appendix – Simpson's rule

Suppose a discrete time Markov model is run with cycle length h and it is desired to find the area under a curve with height q_i after i cycles. (Here q can represent either expected rate of continuously incurring costs or expected quality of life, with or without discounting.) Then standard half-cycle correction estimates the total across n cycles by

$$T(h) = h\left(\tfrac{1}{2}q_0 + q_1 + q_2 + \ldots + q_{n-1} + \tfrac{1}{2}q_n\right).$$

It can easily be shown that this is equivalent to the trapezium rule (see any good textbook on numerical analysis). The trapezium rule assumes that changes in health state occur on average half way through a cycle, or estimates the area under a curve by joining the points with straight lines. Under the very slightly restrictive assumption that n is an even number, a better formula is Simpson's rule

$$S(h) = h\left(\tfrac{1}{3}q_0 + \tfrac{4}{3}q_1 + \tfrac{2}{3}q_2 + \tfrac{4}{3}q_3 + \tfrac{2}{3}q_4 + \ldots + \tfrac{2}{3}q_{n-2} + \tfrac{4}{3}q_{n-1} + \tfrac{1}{3}q_n\right).$$

Simpson's rule adjusts for the main bias in the trapezium rule, and is equivalent to fitting a parabola across a pair of gaps between points. See Figure 3.

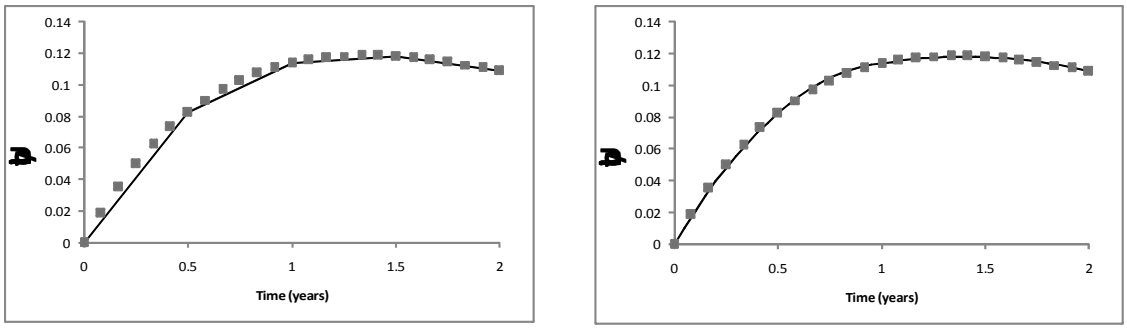

Figure 3. Comparison between the trapezium rule (half-cycle correction) and Simpson's rule. The example is taken from Barton (2009) using a different model from any of the ones defined in this paper. In each case, the rule is fitted to points with a cycle length of 6 months, but values of the underlying model are shown every month.

It can also be noted that the trapezium rule already corrects for first-order error terms, so it is a reasonable approximation that $T(2h)$ is four times as far away from the correct value as $T(h)$, in the same direction, and that Simpson's rule can be found as $S(h) = \frac{4}{3}T(h) - \frac{1}{3}T(2h)$, applying Richardson's extrapolation to the trapezium rule.

References

Barton PM (2009) The irrelevance of half-cycle correction in Markov models. Presentation at the 31st Annual Meeting of the Society for Medical Decision Making (abstract available at http://smdm.confex.com/smdm/2009ca/webprogram/Paper4912.html accessed 1 December 2009, and paper submitted to *Medical Decision Making*).

Barton PM and Tobias AM (1998) Accurate estimation of performance measures for system dynamic models. *System Dynamics Review* 14:85-94.

Richardson LF (1910) The approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stress in a masonry dam. *Philosphical Transactions of the Royal Society of London Series A* 210:307-357.

Sonnenberg FA, Beck JR (1993) Markov models in medical decision making: a practical guide. *Medical Decision Making* 13:322-338.