Work in Progress
Please do not quote without the authors' approval

Temporal Uncertainty in Cost-effectiveness Decision Models

Ronan Mahon*, Andrea Manca, Stephen Palmer
Centre for Health Economics, University of York

* Address correspondence to Ronan Mahon, Centre for Health Economic, University of York, Heslington, York, YO10 5DD; ronan.mahon@york.ac.uk

Abstract
One of the requirements of cost-effectiveness analysis is to apply an appropriate time horizon. The appropriate time horizon for the analysis is often far greater than the time horizon of the evidence available. As a result, temporal uncertainties could be said to exist in the analysis. This paper explores the consequences of temporal uncertainty in terms of the goals of cost-effectiveness analysis in health care. Through a motivating example, it is discussed where temporal uncertainty is expected to arise, when and why it is crucial to characterise this uncertainty, and what features any modelling solution ought to have. Finally an approach to incorporating temporal uncertainty into the analysis is outlined. In particular, it is suggested that the uncertain temporal trajectories of parameters are themselves parameterised. These temporal parameters allow expression of expectations and uncertainties regarding the temporal behaviour of parameters, while also facilitating the calculation of the value of obtaining further evidence that would reduce temporal uncertainty.

Contents

1. Introduction

1.1. A Problem in CEA

Cost-effectiveness analysis (CEA) is used in healthcare in order to aid decision-making regarding firstly, whether a health intervention should be endorsed given the evidence currently available, and secondly, whether further evidence should be sought (1). The type of evidence generally available when conducting a CEA tends to have certain characteristics (e.g. partial comparisons, lack of generalisability) which may be at odds with the requirements for an appropriate CEA (e.g. inclusion of all relevant comparators, relevance to the decision context (2)). To achieve the requirements of CEA given the limitations of the evidence available, a decision analytic framework (comprising evidence synthesis and decision analysis) is often employed (3).

The issue under discussion in this research arises as a result of one particular evidence limitation (a short-term time horizon) and two particular requirements of CEA (use of an appropriately long time horizon, and characterisation of uncertainty). The time horizons from the primary source(s) of evidence in a CEA (e.g. a randomised controlled trial) are usually short-term in nature, whereas the appropriate time horizon to be applied in the CEA (the time over which costs and effects are expected to differ between alternative health interventions) can often be long-term (lifetime when there are mortality effects involved). The existence of such a time horizon mismatch is clearly problematic for the CEA.

1.2. The Nature of this Problem

In the context of decision modelling, this problem can be thought of in several different ways. It is, perhaps primarily, a missing data problem as there indeed exists a significant evidence gap in the analysis. The literature on missing data in CEA however tends to refer to instances of data missing amongst datasets that generally cover the period of interest (4-5) which is not truly the concern here. Since what is required is an expression of our lack of knowledge regarding a period of time in the analysis, this is perhaps better thought of as an issue of uncertainty. The issue can be thought of as one of parameter uncertainty (6), as it is literally the values of parameters that are required. However, unlike how parameter uncertainty is usually addressed, this problem does not simply require distributions to be assigned to established point estimates, as even the point estimates are not known. The issue could also be characterised as one of structural uncertainty (7), where an assumption is required in the model concerning the temporal behaviour of one or more parameters. To a large extent, the structural uncertainty approach is what is currently used in health technology assessments (HTAs). It is common for a HTA to apply one temporal assumption as part of the decision model and then to explore the related uncertainty by applying alternative assumptions in a deterministic sensitivity analysis. While there may be use in taking this structural uncertainty approach (primarily insofar that methods, or variants of methods, used to address structural uncertainty can be helpfully employed here), there are two features of this issue that make it distinct from how a structural uncertainty issue is usually thought of:

- **The need to explicitly model time.** What it will be desirable to express is an increasing uncertainty over time. It may be necessary therefore to model time explicitly in the decision

model. Currently, a single value or distribution is often calculated for a parameter and no relationship with time is modelled. This is especially true of cost and utility estimates.

- **Discrete vs. continuous distributions.** Although depending on the context, this uncertainty may not be accurately represented by discrete alternative 'scenarios'. It may be more suitable to build a continuous distribution (as we would for parameter uncertainty) to convey our true uncertainty regarding the behaviour of parameters over time.

In reality, all of the above concepts are relevant here and it is probably not helpful to categorize this problem as any one 'type' of problem. Let us define the problem generally as 'temporal uncertainty'. That is, let us think of the problem as uncertainty regarding the behaviour/trajectory of parameter values and model structures over time given that there is a dearth of evidence that relates to longer-term time periods.

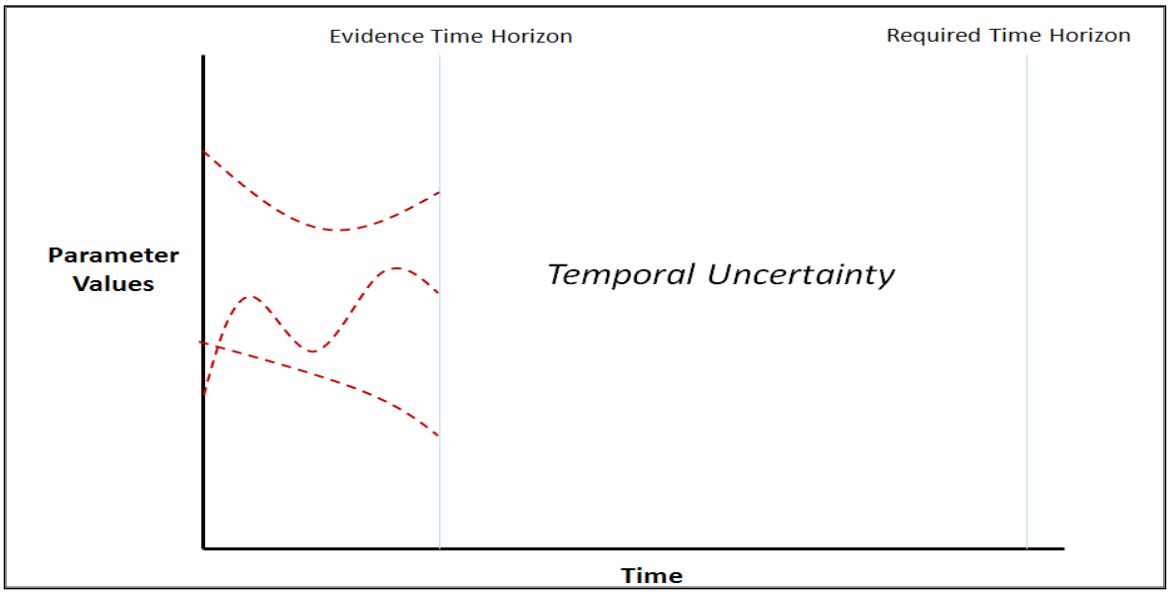

Figure 1: Temporal Uncertainty – uncertainty regarding the behaviour of parameters and the validity of structures over time.

1.3. Key Principles in Addressing Temporal Uncertainty

Whatever the approach ultimately taken to addressing temporal uncertainty, and whatever the precise decision context, there are some key principles/steps that ought to guide the approach to addressing temporal uncertainty in CEA.

I. Identify the components of the decision model that are, or may be, affected by temporal uncertainty.

II. Determine the need to characterise this temporal uncertainty for the purposes of producing useful cost-effectiveness results.

III. Characterise this uncertainty (to the required extent) by exploiting all relevant evidence and validating the resultant model.

IV. Calculate the value of further research that would reduce the temporal uncertainty.

1.4. Objectives of Research

With the above principles in mind, this research sets out to explore, though a motivating example, the task of addressing temporal uncertainty in CEA. The objectives of the paper then are to:

- Outline how and where temporal uncertainty arises in CEA.
- Explore the issues and challenges in appropriately accounting for temporal uncertainty in CEA.
- Convey the consequences of not accounting for temporal uncertainty.

2. Motivating Example (RITA3 Case Study)

This case study comprises a re-analysis of a cost-effectiveness decision model that sought to estimate the cost-effectiveness of an early interventional strategy in non-ST-elevation acute coronary syndrome (NSTE-ACS) based primarily on data from the Randomised Intervention Treatment of Angina (RITA3) trial.

2.1. Background and Motivation for the Cost-effectiveness Analysis

Patients with NSTE-ACS face a significant risk of mortality and cardiovascular events. It is expected that an early interventional strategy (routine angiography followed by revascularisation if clinically indicated) will represent a lower risk of death/cardiovascular event compared to a conservative strategy (ischaemia or symptom-driven angiography), but also a higher cost to the health system. There is uncertainty regarding whether implementing the early interventional strategy represents good value for money from a health system's point of view. The decision problem in the analysis therefore, concerns whether the early interventional strategy or the conservative strategy should be recommended for patients presenting with NSTE-ACS. The key result of the analysis is the mean incremental cost-effectiveness ratio (ICER).

2.2. Outline of the Model Structure

The original decision model was created by Henriksson et al in 2008 (8). The model comprises a short-term decision tree (assumed to be instantaneous in time) and a long-term Markov structure. A series of regression models (based on the RITA3 individual patient-level data) are used to estimate the transition probabilities between Markov states. Similarly, costs and QALYs per Markov state (and for the index hospitalisation period) are estimated using standard OLS regressions. While the bulk of the data comes from the RITA3 trial, additional information on standard mortality rates and treatment effect are incorporated from standard life-table and a meta-analysis of trials respectively. The model structure is depicted in the figure below.

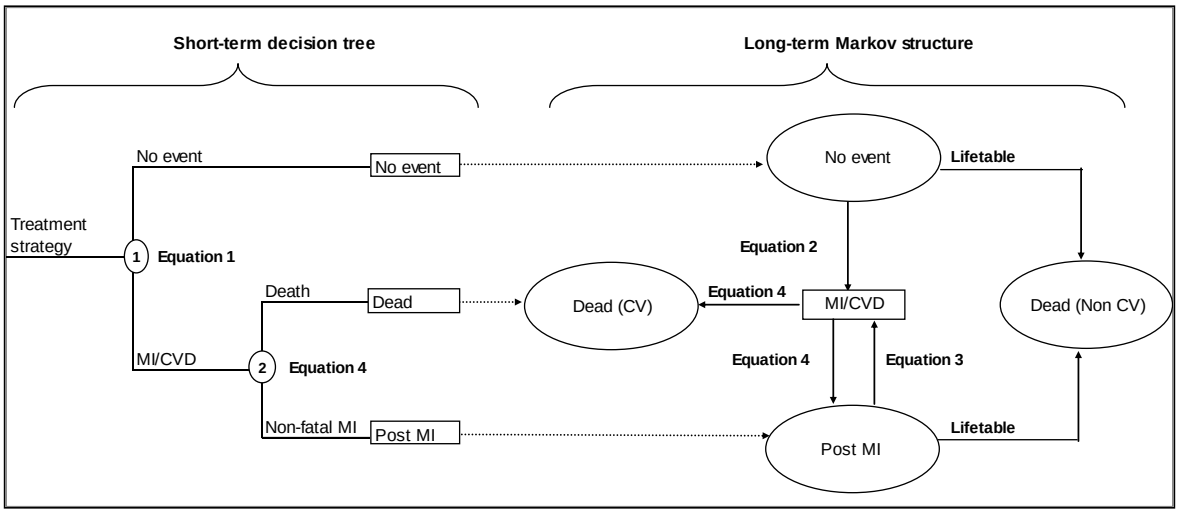

Figure 2: The structure of the RITA3 model

3. Temporal Uncertainty in the Model

3.1. Why Temporal Uncertainty Arises in this CEA

The bulk of the data that inform this model come in the form of individual patient-level data (IPD) from the RTIA3 trial. These data are 5 years in duration. The appropriate model time horizon (the time over which costs and effects are expected to differ between the two strategies) however is circa 60 years, after which effectively all patients are deceased. It is this time horizon mismatch that leads to temporal uncertainty within the model. For the purposes of this report, the 5 year evidence period will be referred to as the 'observed period' and the subsequent 55 year period as 'the unobserved period'.

3.2. Identifying Where There is Temporal Uncertainty

In short, all of the key components of the long-term Markov portion of the model are potentially affected by temporal uncertainty.

- The **transition probabilities** between Markov states are arguably the driving force behind the cost-effectiveness results. For the 'observed period', these probabilities are calculated in different ways (for each transition and for each arm) based on a combination of the data available and structural assumptions. Considerable decision uncertainty relates to how these transition probabilities evolve over time.

- **Utility values** attributed to each Markov states are often assumed to be constant. Under some circumstances, this may be a reasonable assumption. Indeed, the decline in utility is already represented to some extent by progression to 'worse' health states. There is still however temporal uncertainty to be considered when we apply a 60 year time horizon. The natural decline in utility associated with age, for example, is not catered for in the current version of the model.

- Similarly, constant **costs** are often assumed for Markov states. Costs are somewhat more tangible and may be more easily estimated compared to utility values. A significant source of temporal uncertainty related to costs is the possibility of some future event (see below).

- The **structures and other assumptions** imposed by the model are also subject to temporal uncertainty. What may be shown to be a valid assumption over the observed period may not be valid over the unobserved period. This includes the appropriateness of measuring the probability of experiencing a 'composite event' rather than considering myocardial infarction and cardio-vascular death separately, as well as appropriateness of the 'post MI' state (in fact 5 tunnel states are used in the original model). It is also assumed that when estimating transition probabilities, age can be modelled by updating the age group covariate every 10 years.

- Another consequence of this time horizon mismatch is the potential need to incorporate **other future events** which may affect parameters in the model, e.g. price shocks, new comparators emerging, etc. It is unlikely that such events will be pertinent to this case study and so it is assumed here that there is no need to incorporate them into the model. It is also assumed here that there is no other long-term clinical event to take account of, e.g. side-effects of a treatment.

3.3. Focus and Assumptions of this Case Study

Although, we highlight above all the potential areas of temporal uncertainty in this model, the scope for this case study is to re-analyse the temporal uncertainty pertaining only to the transition probabilities between the 'no event' state and the 'composite event' (myocardial infarction/cardio vascular death) state. The motivation for this is to focus on just one component in detail in order to demonstrate the modelling aspects of this issue. It is also expected that it is this component of the model where temporal uncertainty will have the greatest contribution towards the overall decision uncertainty.

Within this component of the model, there are two related parameters to be examined.

- Baseline risk (which represents the conservative arm of the trial)
- Treatment effect (which when applied to the baseline risk represents the early interventional arm of the trial)

Only one sub-group (of five in the original analysis) is being examined here and it is assumed that there is no unobserved heterogeneity. As mentioned above, in the original analysis it is assumed that age can be modelled by updating the age covariate (in the regression analysis pertaining to the observed period) every ten years. That assumption is also employed here and so what are being estimated are the probabilities of a first 'event' for every year after the intervention, for a constant age.

It is assumed that the ICER threshold is £20,000 per QALY.

4. Does this Temporal Uncertainty Matter?

We know that in order to achieve a complete picture of cost-effectiveness, we require full characterisation of all decision uncertainty. However, such thorough modelling may not be required.

4.1. Initial Indicators of the Extent of Temporal Uncertainty

Firstly, to get an idea of the expected impact of this temporal uncertainty on the cost-effectiveness results, we can contrast the duration of the observed period with the duration of the 'unobserved' period. The below graph illustrates the survival data we have available (in the form of Kaplan-Meier curves) in the context of the survival space we are expected to fill.

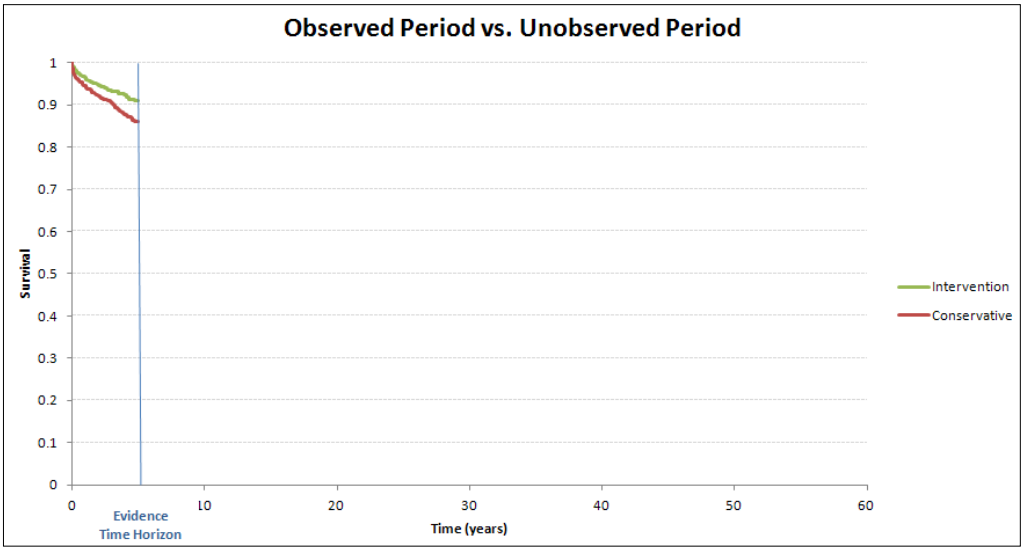

Figure 3: Available data against full model time horizon

It can be seen from this graph that the survival data are not very 'mature'. Specifically, the data cover one twelfth of the model time period and only ten percent of the overall cohort have experienced the event. However, it must also be noted that outcomes attributable to earlier periods are of more significance. In this model, approximately 31% of all costs and QALYs are attributable to the observed period (using the basecase assumptions in the original version of the model). The reason for this is that there are more patients alive or in better health states over earlier periods and costs and QALYs accrued are discounted to a lesser degree.

Although this time horizon mismatch is an indicator of the extent of the temporal uncertainty that may exist in relation to these parameters, what ultimately is of consequence is the impact that this temporal uncertainty has on the decisions the CEA is designed to inform, i.e. whether or not to adopt a new health intervention, and whether or not to seek further evidence.

The below figure outlines (in a stylised simple world of decision making) when and why modelling of temporal uncertainty is expected to be a necessity, as well as occasions where such modelling arguably is not a necessity.

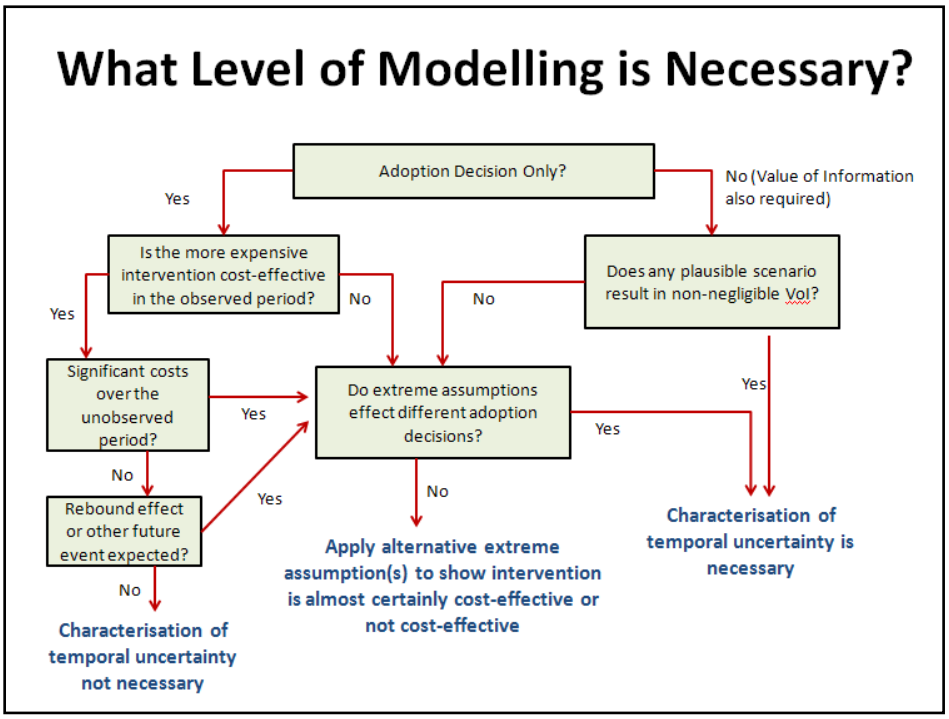

Figure 4: Flow chart showing what level of modelling regarding temporal uncertainty may be necessary under certain scenarios. It is assumed for simplicity that there are two competing interventions

It must be noted however that in reality there is always motivation to fully characterise all sources of decision uncertainty for the purposes of incentivising manufacturers to lower prices or provide further information (9). The above schematic is simply intended to demonstrate the relationship between temporal uncertainty and the goals of cost-effectiveness analysis. In practice, the above decision contexts rather than indicate the need for characterisation of temporal uncertainty, perhaps indicate the required complexity of the modelling. Nonetheless, let us consider these questions in the case study, in order to gauge the requirement for characterisation of temporal uncertainty.

4.2. Is the Adoption Recommendation Expected to Change over the Unobserved Period?

One could imagine a scenario (depicted in the flow chart above) where the most costly health intervention is shown to be cost-effective within the observed period (implying that significant health gains are delivered). If this was the case, and no significant further costs were expected over the unobserved period, no rebound effect or other further effect was expected, and the CEA existed only to inform the adoption decision, then one could argue that there is no need to model the unobserved period. We would be confident that this health intervention will only continue to accrue more QALYs relative to its comparators and at no significant extra cost and therefore remain cost-effective.

In this case study (for now, using the temporal assumptions implemented in the original analysis), if we calculate the cumulative incremental net health benefit over time, it can be observed that the adoption decision could plausibly (and in fact is expected to) change over the unobserved period, i.e. for the early intervention treatment (the more costly treatment) to be cost-effective it relies on the QALYs accrued over the unobserved period and so it may be that the assumptions imposed regarding the temporal behaviour of parameters will affect the adoption decision.

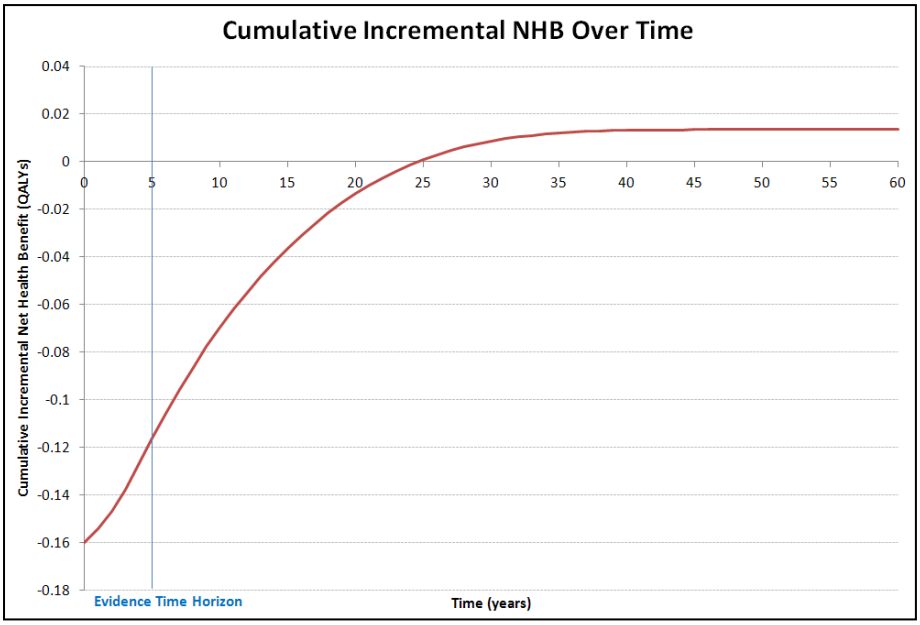

Figure 5: Cumulative Incremental Net Health Benefit over time

4.3. Do Extreme But Plausible Assumptions Alter the Adoption Recommendation?

To confirm that modelling of temporal uncertainty over the unobserved period is truly necessary to make an informed adoption recommendation, we could impose extreme (but plausible) assumptions in order to observe the effect on the mean ICER and the resulting adoption recommendation. If it turned out that no plausible assumption resulted in a different adoption recommendation, then one could argue that there is no need to fully characterise the temporal uncertainty. All that would be necessary would be to demonstrate with these extreme assumptions (or perhaps with a single 'conservative' assumption) that there is negligible chance of the recommendation being incorrect. Conversely, if the extreme assumptions did represent different adoption recommendations, we would be compelled to investigate further and characterise, to at least some extent, this temporal uncertainty.

Of course, what is extreme but plausible is a subjective concept and some validation may be necessary for this task. Taking the original analysis, let us impose some bounds of plausibility in relation to (i) the treatment effect over the unobserved period and (ii) the baseline risk over the unobserved period, in order to observe the resultant mean ICERs. The below table outlines the bounds of plausibility that were judged to be reasonable.

Parameter	Extreme Assumption
Treatment Effect	Average hazard ratio (~0.6) in observed period holds over lifetime of model
Treatment Effect	No treatment effect over unobserved period (hazard ratio = 1)
Baseline Risk	Hazard fixed at zero over unobserved period
Baseline Risk	Hazard fixed at highest average calculated in observed period (0.05)

Table 1: Extreme but plausible assumptions for the two parameters of interest

The assumption made in relation to baseline risk was found not to itself have a significant impact on the mean ICER (i.e. when the treatment effect assumption was held constant). The impact of the baseline assumption is best described as dictating the extent of the impact of the treatment effect assumption. This is intuitive as the greater the baseline risk, the greater the scope there is for the treatment effect to

have an impact. The below graph illustrates the extent of the impact on the mean ICER of the (plausible) treatment effect assumptions when the assumption imposed with regard to baseline risk is such that the treatment effect has most impact (i.e. baseline risk is assumed to hold at 0.05).

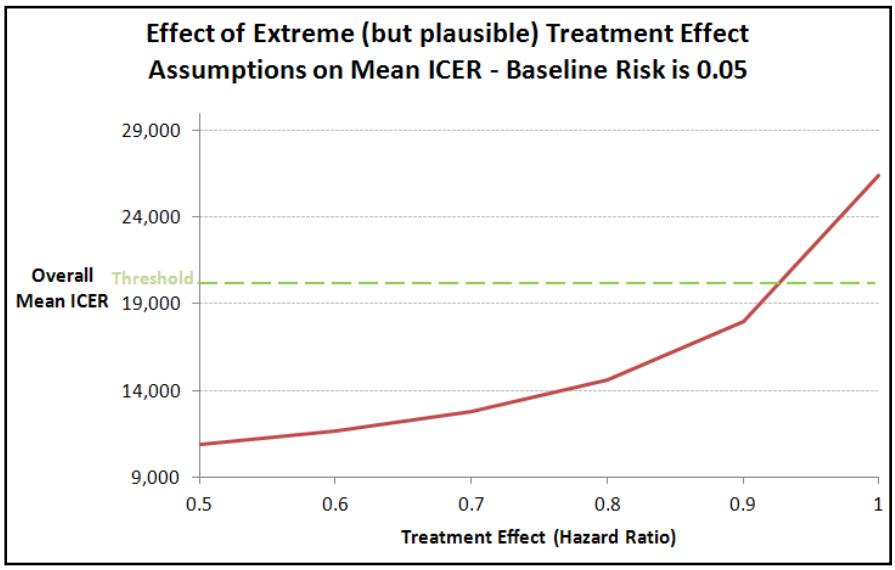

Figure 6: Resultant mean ICERs when extreme assumptions regarding the treatment effect are applied and a constant baseline risk of 0.05 is assumed

Since this range of plausible assumptions result in ICERs that fall either side of the ICER threshold, let us conclude that it is desirable to compute the true expected long-term outcomes and fully quantify the uncertainty.

Of course, if a recommendation is desired in relation to gathering further evidence then all sources of uncertainty should be quantified a fully as possible. Even in this decision context however, it could be that the expected ICER is so high or so low that it would be concluded that there is negligible value in collecting further information and we could again simply employ conservative assumptions to aid the adoption decision. Let us assume however, that in this case it is desirable to fully quantify the temporal uncertainty.

5. Requirements for Appropriate Characterisation of Temporal Uncertainty

5.1. Incorporating Temporal Uncertainty into the Decision Model

Assuming the necessity to address temporal uncertainty in a CEA, it is not enough to consider this uncertainty outside of the decision model (e.g. through a deterministic sensitivity analysis) as this may lead to implicit weighting of scenarios when the adoption recommendation is made. It is desirable to incorporate this uncertainty into the decision model itself. Even if it doesn't substantially change the expected cost-effectiveness (mean ICER), quantifying this uncertainty will allow a better estimate of the probability of cost-effectiveness and will also allow estimation of how valuable it would be to collect further information in order to reduce the uncertainty. It is also required that all relevant evidence is considered as the uncertainty is characterised and that we are explicit regarding what assumptions we are imposing for each parameter.

5.2. Contextual Factors

Further to these general requirements, we suggest that there are a number of factors in a HTA that may give rise to some additional requirements for an appropriate characterisation of temporal uncertainty. The table below outlines the factors, some of which have been discussed already, that could inform what is desired from a modelling approach.

Decision Context

What is the model time horizon	60 years
Are we interested in value of information?	Yes
How many comparators?	2
Do we need to model any future 'events' (e.g. price shock, new comparators)?	No
Is the adoption recommendation expected to change over the unobserved period?	Yes
Do we need to consider sub-groups?	No

Clinical Context

Is the disease chronic or acute?	Chronic
What is the disease area?	Cardiovascular
Is the intervention a one-off at the start or continued over the unobserved period?	One-off at start
Do we need to account for treatment switchover or withdrawal?	No

Data/Modelling Context

In what component of the model is there temporal uncertainty?	Transition probabilities
What is the nature of the evidence available?	IPD and AD from RCTs
What is the duration of the evidence (observed) period?	5 years

Table 2: Contextual factors pertinent to the characterisation of temporal uncertainty

From this, we can conclude that we need not consider long-term effects on continued intervention, nor sudden shifts in parameters associated with the discontinuation of an intervention. We need not consider heterogeneity (for now). However the modelling of temporal uncertainty should incorporate (and ideally fully exploit) the individual patient-level data from the RITA3 trial and the odds ratios from the other trials. We also need to interpret this evidence as transition probabilities.

5.3. Choosing a Modelling Approach for this Component of the Decision Model

With these requirements and circumstances in mind, let us consider some general modelling approaches that could be taken in relation to this component of the decision model and choose an approach that is conducive to effective characterisation of the temporal uncertainty.

A common approach is to assume proportional hazards (i.e. assume that the effect of treatment is multiplicative with respect to baseline risk and that this effect is independent of time). A proportional hazards (PH) survival function can be fit to the RITA 3 data in order to estimate the baseline hazards over the trial period and a single hazard ratio pertaining to treatment effect for the trial period. Cumulative hazards can be calculated for specific time intervals and transition probabilities can be produced form these (10). Separate judgements can then be made regarding the behaviour of both the baseline hazard and the treatment effect after the unobserved period. This is what was done in the original version of the

model where a Weibull PH model was fit. An assumption of constant hazards was then applied to the baseline (i.e. the hazard observed at the end of the trial period is retained), and an assumption of no further treatment effect was made. The treatment effect applied to the observed period was a pooled treatment effect comprised of the hazard ratio from the Weibull fit and the odds ratios from the external data. Although this approach is straightforward, it perhaps somewhat restrictive and does not make best use of the short-term evidence in the context of characterising temporal uncertainty.

An alternative approach that might be less restrictive and more conducive to extrapolating evidence over time would be to fit separate survival distributions to each arm of the trial data. No relationship would then be assumed between the arms. Each arm could then be independently extrapolated over the unobserved period, either assuming a continuation of the survival distributions, or making an alternative assumption regarding the behaviour of the hazards, perhaps based on their temporal behaviour in the observed period. Since this approach does not consider a relative risk, but a separate absolute risk for each arm, it is not necessary to only select distributions that use a proportional hazards metric. However this approach also implies that we will not use the external data available as they pertain to treatment effect specifically. Furthermore, this approach involves not making an explicit judgement regarding treatment effect, rather the assumption is implicit in the extrapolation of the independent arms. We should be explicit regarding what is being assumed regarding treatment effect, for the sake of clarity, but also because it is scenarios regarding future treatment effect that is highlighted in the current NICE guidance (11).

A compromise approach would be to focus on the two parameters separately (i.e. baseline risk and treatment effect – this as opposed the two arms of the trial). Model the baseline risk (i.e. the conservative arm of the trial) in relation to the observed period and then generate alternative long-term scenarios based on the temporal trend observed in the trial period, or on external evidence relating to natural history, or on other clinically plausible scenarios. Then model the treatment effect, firstly by analysing the change in hazard ratio over time in the RITA 3 data and the hazard ratios observed at different time points in the external data. Temporal behaviour can then potentially to some extent be based on the temporal trends observed. This would incorporate, and seem to best exploit, the relevant evidence available while also allowing explicit assumptions regarding the both parameters. We suggest here that given the requirements outlined earlier, this last modelling approach best facilitates the appropriate characterisation of temporal uncertainty in this particular circumstance.

6. Characterising Temporal Uncertainty by Parameterising the Temporal Trajectory of the Parameters of Interest

Now it is explored how one might characterise expectations and uncertainties regarding the temporal trajectory of parameters given the evidence available. It is probable that in this case study there are some

other sources of evidence that may to some extent aid the long-term modelling of these parameters[1], but for now let us assume that the short-term evidence is all that is available.

There are a number of approaches that could be taken to convey increasing uncertainty over time regarding the value of a parameter. Taking the 'parameter uncertainty' approach, we could split the unobserved period into several time intervals and apply distributions with wider and wider confidence intervals to the parameter in question. This would seem to be a cumbersome approach however as we effectively have to estimate distributions for several parameters. Taking the 'structural uncertainty' approach we could incorporate alternative assumptions regarding the behaviour of each parameter over time. However this may be too restrictive to convey the true extent of our uncertainty.

It has been suggested by Jackson et al. that all structural uncertainties can be expressed as parameters in a decision model (13). Here we apply that principle to temporal uncertainty and so we parameterise the temporal trajectories of parameters whose short-term values have already been estimated in the model. The characteristics of these temporal parameters will represent our expectations and uncertainties regarding the behaviour over the time of the underlying parameters.

We need to acknowledge the limitations of what short-term evidence can tell us about long-term parameter values. We therefore avoid relying on 'functional fits' to the short-term data to characterise uncertainty over the long-term. But of course, the fundamental problem here is that there are no (or few) data pertaining to the unobserved period. We cannot convey a complete lack of knowledge (even a uniform distribution implies an expected value). What we can do however is apply a very flat distribution around the temporal parameter and bound the value of the underlying parameter by what is deemed to be plausible. In terms of what mean is applied, we suggest that in the absolute absence of evidence pertaining to the long-term, it is reasonable to base the expected long-term parameter trajectories on the extrapolated short-term evidence, but we ought to apply this extra temporal uncertainty to reflect our lack of true knowledge about the behaviour of parameters over the long-term.

6.1. Baseline Risk

The extrapolated baseline hazards based on a Weibull fit to the RITA3 data are as illustrated in the below graph.

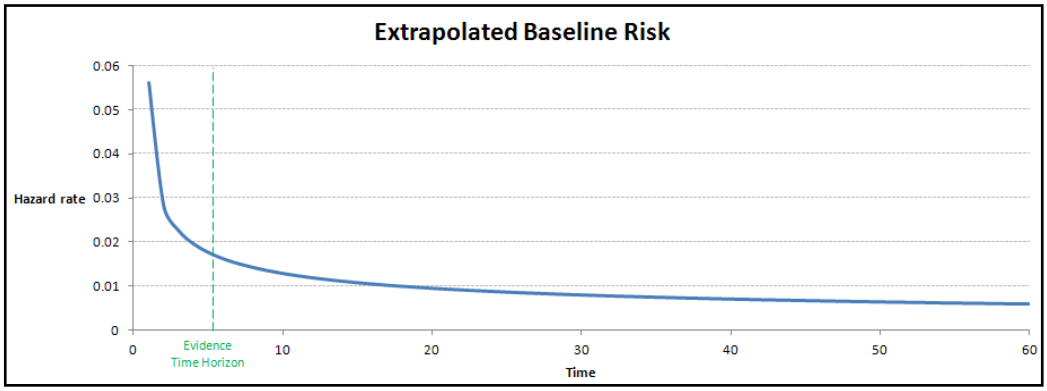

Figure 10: The expected temporal trajectory of the baseline risk parameter as per the Weibull extrapolation

[1] Other studies suggest that the risk of a composite event reduces over the long-term for this patient group (12)

Let us assume for now that this characterisation of the observed period is reasonable. Note that the shape and scale parameters of the Weibull function are applied probabilistically in the model but this does not convey any extra uncertainty related to time. We would like to incorporate into the model (through a probabilistic temporal parameter) scenarios where the baseline risk does not simply continue the trend of the Weibull distribution, but instead increases after the observed period, or decreases at a greater rate. Let us assume for simplicity that the temporal trajectory of baseline risk over the unobserved period is approximately linear (though of course it is perfectly feasible, and probably more desirable, to impose a non-linear trajectory).

Now we define the value of the baseline risk parameter over the unobserved period as a function of baseline risk at 5 years, the 'slope' of the curve for baseline risk, and time. We also restrain the values of the baseline parameter to be within our bounds of plausibility i.e. [0, average value in observed period] In particular:

$$b_t = min[\alpha, max[0, b_5 + slope * (t-5)]]$$

where: b_t = the value of baseline risk at time t

α = the average value of the hazard rate in the observed period (for that probabilistic draw)

We now also make the 'slope' parameter probabilistic in order to express our uncertainty regarding the long-term temporal trajectory of baseline risk.

As said above, in the absence of other evidence, let the expected value of the slope be determined by the extrapolated (Weibull) curve based on the short-term evidence. Based on this, the mean baseline hazard at 5 years is 0.0176 and at 60 years is 0.006. Therefore the mean slope (assuming linearity) is (0.006-0.0176)/(60-5) = -0.00021

We would not expect the extent of the uncertainty imposed to have a significant impact on the mean ICER (all else held constant). It is only because of the non-linearity of the model that the mean ICER would change at all (14). Let us assume that the slope parameter follows a Normal distribution (again, we do this for simplicity, a Log-Normal or Gamma might be more suitable). We can assess the impact of applying increasing magnitudes of temporal uncertainty, but for now, let's say the standard deviation is 0.0003. The slope parameter is then computed as follows:

$$slope \sim Normal(-0.00021, 0.0003)$$

With this probabilistic temporal trajectory applied to the baseline risk parameter (along with probabilistic modelling of the initial Weibull curve) we would seem better able to express our full uncertainty with regard to the behaviour of the baseline risk parameter.

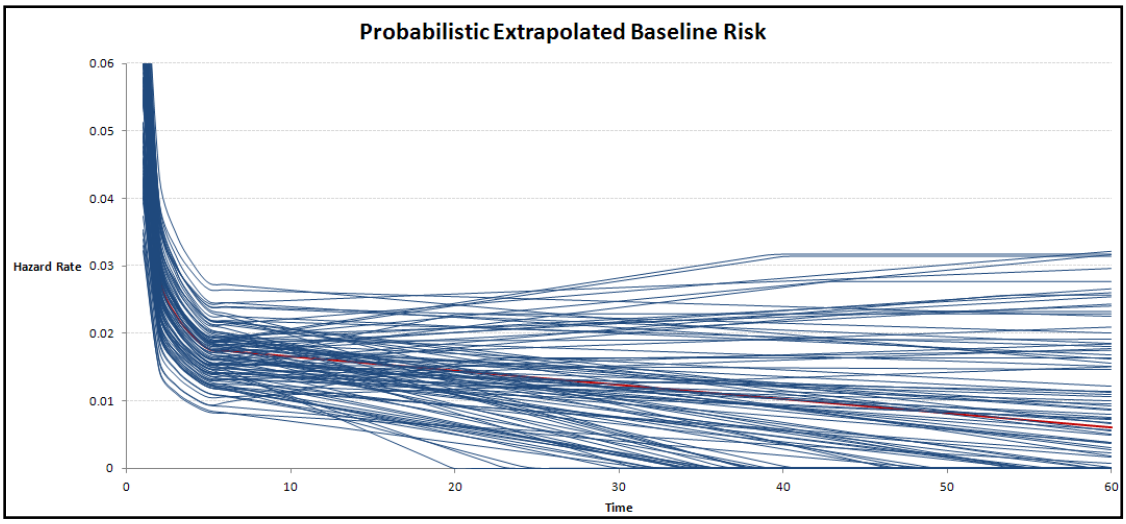

Figure 11: Baseline risk when temporal uncertainty is accounted for

6.2 Treatment Effect

To try to exploit the short-term evidence regarding treatment effect, we can examine the short-term data to look for any temporal trend in the hazard ratio (as opposed to assuming proportional hazards and obtaining one hazard ratio for the observed period). The below graph shows smoothed empirical hazard curves for the observed period based on the data from RITA3.

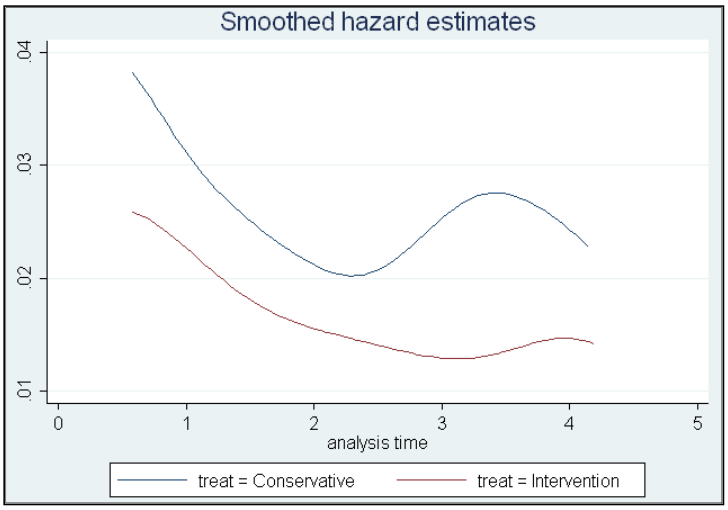

Figure 12: Empirical hazard curves showing treatment effect over the observed period

The graph indicates a marginal decline in treatment effect for the first two and a half years but then a jump in the 3rd/4th years followed by another decline. Our uncertainty regarding whether the jump is an anomaly or an indication of a more substantial treatment effect over the unobserved period is representative of the general limitations of only possessing short-term data. It is possible that the parameter has yet to reach a 'steady state'. To explore further what the short-term data indicates regarding the temporal behaviour of treatment effect, we could implement a Piecewise Cox Regression, in order to compute hazard ratios for different time periods[2]. The results concur with the graph above, with the jump in treatment effect captured in the 4th time period (note that a hazard ratio of 1 would suggest no treatment effect).

[2] This was done using the stsplit command in Stata (15)

Time Period (Year)	Estimated Hazard Ratio
1	0.647
2	0.671
3	0.96
4	0.444
5	0.604

Table 4: Results of Piecewise Cox Regression

Regarding the other evidence available pertaining to treatment effect, most of these data come in the form of a single Hazard or Odds ratio[3] for a short-term period. One report on the FRISC-II trial however, describes a fall in treatment effect from the first two years to the following three years (16)

Let us say that our uncertainty here is such that there are a discrete number of competing beliefs about the nature of treatment effect over the unobserved period, and relatedly the meaning of the jump in treatment effect during the observed period. Let us therefore take a different approach (to the one used for baseline risk) in characterising the temporal trajectory of the treatment effect parameter. In this case we will express our uncertainty by incorporating into the model alternative discrete scenarios (as opposed to characterising temporal trajectories with a continuous distribution).

Scenario 1: The jump was an anomaly and the true hazard ratio tends towards one quickly after the observed period. This scenario is represented by an assumption of no further treatment effect after the observed period (this was the basecase assumption in the original analysis).
Scenario 2: There is substantial treatment effect over the unobserved period and the hazard ratio will continue to jump up and down over the unobserved period as it did over the observed period. This scenario is represented by assuming the average hazard ratio over the observed period holds for the duration of the unobserved period.
Scenario 3: The treatment effect slowly dissipates over the unobserved period, as suggested by the report on the FRISC-II trial. This scenario is represented by assuming the treatment effect declines as per the rate indicated by that report.

The weights ascribed to each scenario are somewhat arbitrary at this stage of this research. A declining treatment effect is for now judged to be the most reasonable scenario.

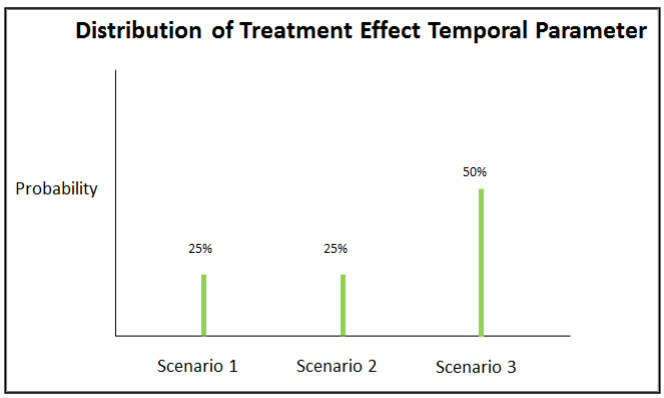

Figure 13: Distribution of Treatment Effect Temporal Parameter

[3] In this case study, an odds ratio can be used to approximate a hazard ratio because the risk of the event in question is very low and a hazard ratio tends towards an odds ratio as the underlying probability tends to zero.

The probabilistic hazard rates pertaining to the treatment arm of the trial are illustrated in the graph below.

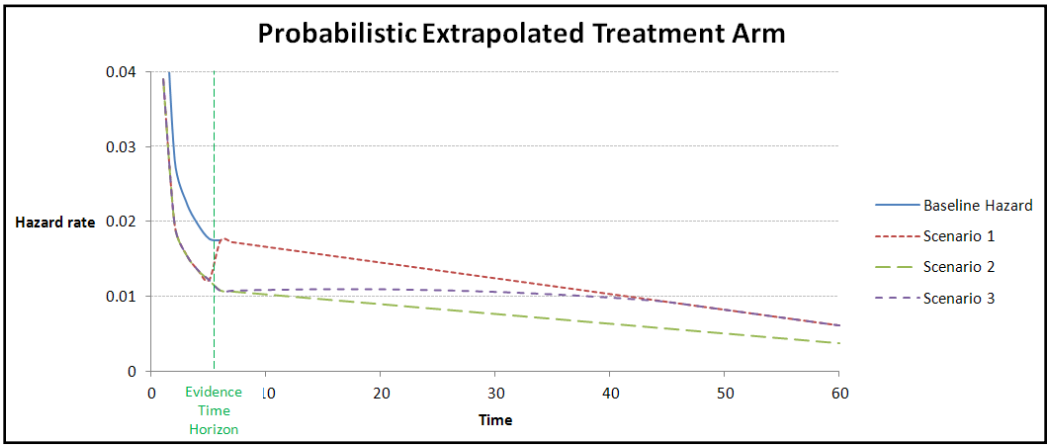

Figure 14: The realisations of extrapolated treatment arm hazards when the three treatment effect scenarios are applied

7. Cost-effectiveness Results

The cost-effectiveness results are outlined by comparing the decision model incorporating temporal uncertainty with the decision model not incorporating temporal uncertainty (i.e. just the expected values for each temporal parameter are applied).

The difference in mean ICER is minimal as expected. Any difference should simply relate to the non-linearity of the model.

Mean ICERs	
Without Temporal Uncertainty	With Temporal Uncertainty
£12,059	£12,772

The distributions of possible ICERs however indicate that there is a non-negligible number of possible ICERs that represent a different adoption recommendation when temporal uncertainty is accounted for. This is illustrated in the cost-effectiveness planes below.

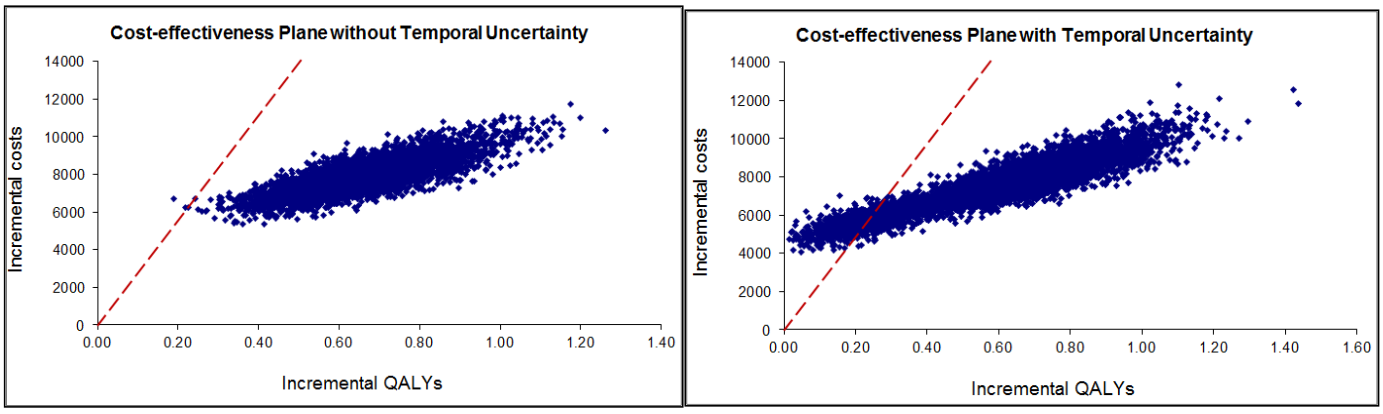

Figure 15: The cost-effectiveness planes for the decision model without and with temporal uncertainty

The consequences of this more elongated distribution can be seen when the probability of cost-effectiveness is considered. The graphs below illustrate the cost-effectiveness acceptability curves for each version of the decision model. When temporal uncertainty is not incorporated, the probability that

the early interventional treatment is cost-effective is 0.996. When temporal uncertainty is incorporated, this probability is 0.825.

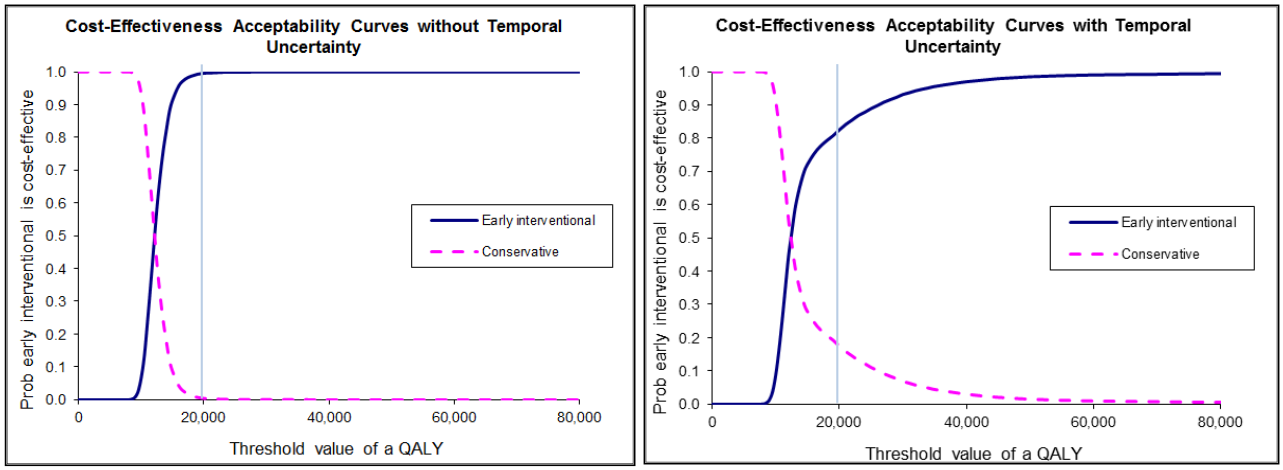

Figure 16: Cost-effectiveness acceptability curves for the decision model without and with temporal uncertainty

8. Value of Information

If the adoption decision is to be made based solely on expected cost-effectiveness, then the consideration of temporal uncertainty has minimal significance. It is more likely however that a decision maker will want to take into account the error probability, i.e. the chance that the recommendation will turn out to be 'wrong'. There are, at least in principle, other options available to the decision maker after CEA is complete. Recommendations of adoption subject to further research, or rejection with a view to conducting further research could be made. What is desirable therefore is to quantify the expected value of collecting further information (value that becomes manifest at times when new evidence compels decision maker to change his/her mind and therefore avoid incurring costs on the health system).

A probabilistic decision model, like the one constructed in this case study, already has the features required to conduct some basic value of information analysis. Here, the expected value of perfect information[4] is computed for the decision models with and without temporal uncertainty. The results are shown on the graphs below.

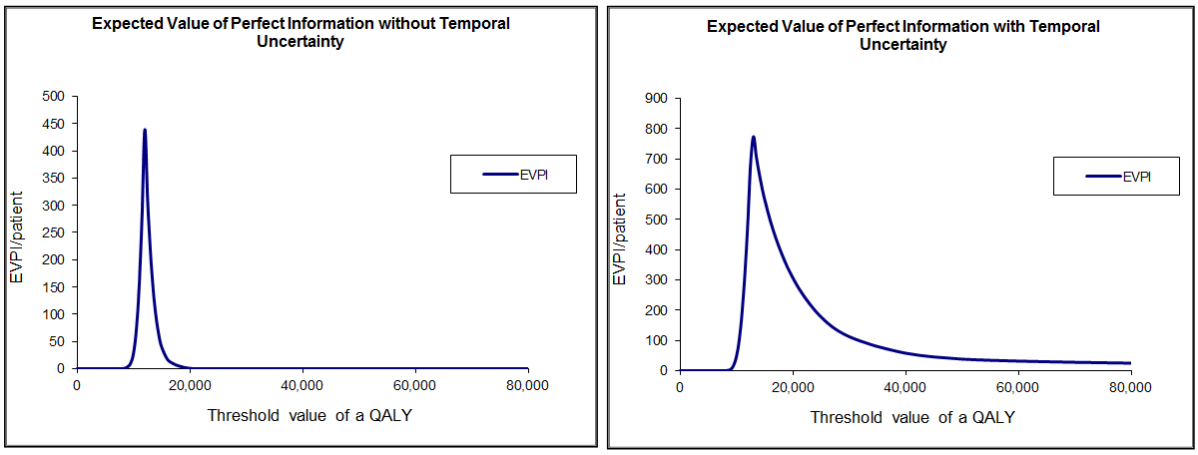

Figure 17: expected value of perfect information for the decision model without and with temporal uncertainty

[4] EVPI is an upper-bound on the worth of collecting further evidence. It is impossible to have perfect information, so if EVPI is less than the expected cost of collecting further information, then such an undertaking would not be worthwhile.

Intuitively, since we have incorporated more uncertainty into the model, the version with temporal uncertainty will show generally greater EVPI. What is of most consequence however is the position of the EVPI curve in relation to the ICER threshold (here £20,000). In this case study, there is significantly greater EVPI at a threshold of £20,000 due to the inclusion of temporal uncertainty. Although this approximates the value of collecting information on the temporal parameters in particular, we could more formally calculate the expected value of perfect information for particular parameters.

9. Conclusions and Discussion

While cost-effectiveness analysis strives to inform decisions regarding the adoption of health interventions and the value of further research based on the best available evidence, there are many analytic challenges involved. Dealing with temporal uncertainty is one of these challenges, a challenge that has not been explored in detail to date and for which there is not currently substantial guidance. This research highlights the extent to which temporal uncertainty may pervade cost-effectiveness decision models, the circumstances under which temporal uncertainty is likely to have an impact, and the potential magnitude of that impact. It is suggested here that it is important to be explicit about the temporal assumptions that are imposed on parameters and as such, it is often preferable that parameters subject to temporal uncertainty are considered individually and independently (although ideally correlation between parameters would also be accounted for). It is also suggested that parameterising the temporal trajectories may be a useful way of representing temporal uncertainties in a decision model.

While this research seeks to highlight the presence of and outline the meaning of temporal uncertainty in cost-effectiveness analysis, what must be explored further are methods to exploit the (mainly short-term) evidence available to inform, to the greatest extent possible, the long-term temporal trajectories of parameters (i.e. inform the distributions of the temporal parameters). In fact, this issue of temporal uncertainty is usually discussed in terms of methods to use short-term evidence to estimate long-term values, i.e. extrapolation of evidence over time. While it is important to be mindful of the limitations of using short-term evidence to estimate long-term values, it is also necessary to consider how the available evidence might be best interpreted, given that some estimation of expected values must be made now. In this research some simplistic and possibly unrealistic assumptions were applied such as the linear temporal trajectory of the baseline risk parameter, the normally distributed 'slope' parameter, the weights attributed to the treatment effect scenarios and the assumption of no unobserved heterogeneity. These assumptions, in practice, would need greater validation.

It is possible that an analysis (given the current evidence) suggests that more evidence is required to make an adoption decision. It is worth considering therefore what kind of evidence may be useful and how best to incorporate supplementary evidence into the decision model. In the case of temporal uncertainty, it is likely that expert elicitation would be suitable since there would not exist any other data that pertain to, for example, the future effects of a new drug. Also, the opportunity cost of waiting for

more trial data is likely to be high. The use of expert elicitation will be explored in the next stage of this research.

It is hoped that this and further research will contribute to greater recognition of the issue of temporal uncertainty and lead to guidance ensuring sound and consistent methods of addressing this issue.

References

1. Claxton K, Sculpher M, Drummond M. A rational framework for decision making by the National Institute For Clinical Excellence (NICE). Lancet. 2002 Aug 31;360(9334):711-5.
2. Sculpher MJ, Claxton K, Drummond M, McCabe C. Whither trial-based economic evaluation for health care decision making? Health Econ. [Research Support, Non-U.S. Gov't]. 2006 Jul;15(7):677-87.
3. Briggs A, Claxton K, Sculpher MJ. Decision modelling for health economic evaluation. Oxford: Oxford University Press; 2006.
4. Briggs A, Clark T, Wolstenholme J, Clarke P. Missing... presumed at random: cost-analysis of incomplete data. Health Econ. 2003 May;12(5):377-92.
5. Burton A, Altman DG. Missing covariate data within cancer prognostic studies: a review of current reporting and proposed guidelines. Br J Cancer. 2004 Jul 5;91(1):4-8.
6. Claxton K. Exploring uncertainty in cost-effectiveness analysis. Pharmacoeconomics. [Research Support, Non-U.S. Gov't]. 2008;26(9):781-98.
7. Bojke L, Claxton K, Sculpher M, Palmer S. Characterizing Structural Uncertainty in Decision Analytic Models: A Review and Application of Methods. Value Health. 2009 Jan 9.
8. Henriksson M, Epstein DM, Palmer SJ, Sculpher MJ, Clayton TC, Pocock SJ, et al. The cost-effectiveness of an early interventional strategy in non-ST-elevation acute coronary syndrome based on the RITA 3 trial. Heart. 2008 Jun;94(6):717-23.
9. Griffin SC, Claxton KP, Palmer SJ, Sculpher MJ. Dangerous omissions: the consequences of ignoring decision uncertainty. Health Econ. 2011 Feb;20(2):212-24.
10. Collett D. Modelling survival data in medical research. 2nd ed. Boca Raton, Fla. ; London: Chapman & Hall/CRC; 2003.
11. NICE. National Institute for Health and Clinical Excellence. Guide to the methods of thechnological appraisal. 2008.
12. van Domburg RT, van Miltenburg-van Zijl AJ, Veerhoek RJ, Simoons ML. Unstable angina: good long-term outcome after a complicated early course. J Am Coll Cardiol. 1998 Jun;31(7):1534-9.
13. Jackson CH, Bojke L, Thompson SG, Claxton K, Sharples LD. A Framework for Addressing Structural Uncertainty in Decision Models. Med Decis Making. 2011 May 20.
14. Rice JA. Mathematical statistics and data analysis. 2nd. ed. Belmont, Calif.: Duxbury Press; 1995.
15. Cleves MA. An introduction to survival analysis using Stata. 3rd ed. College Station, Tex.: Stata Press; 2010.
16. Lagerqvist B, Husted S, Kontny F, Stahle E, Swahn E, Wallentin L. 5-year outcomes in the FRISC-II randomised trial of an invasive versus a non-invasive strategy in non-ST-elevation acute coronary syndrome: a follow-up study. Lancet. 2006 Sep 16;368(9540):998-1004.