

Methods for handling unobserved covariates in a Bayesian update of a cost-effectiveness model

Benjamin Thorpe, PhD*, Orlagh Carroll, BSc†, Linda Sharples, PhD‡

September 11, 2017

Abstract

Health economic decision models often involve a wide-ranging and complicated synthesis of evidence from a number of sources, making design and implementation of such models resource-heavy. When new data become available and reassessment of treatment recommendations is warranted, it may be more efficient to perform a Bayesian update of an existing model than to construct a new model. If the existing model depends on many, possibly correlated, covariates then an update may produce biased estimates of model parameters if some of these covariates are completely absent from the new data. Motivated by the need to update a cost-effectiveness analysis comparing diagnostic strategies for coronary heart disease, this study develops methods to overcome this obstacle by either introducing additional data or using results from previous studies. We outline a framework to handle unobserved covariates, and use our motivating example to illustrate both the flexibility of the proposed methods and some potential difficulties in applying them.

Introduction

Long-term health economic decision models have complicated mathematical structures, with model parameters informed by a range of data sources, including mortality rates, clinical outcomes, incidence of clinical events, costs and

*Leeds Institute of Clinical Trials Research, University of Leeds, UK

†Leeds Institute of Clinical Trials Research, University of Leeds, UK

‡London School of Hygiene and Tropical Medicine, UK

Address correspondence to: Benjamin Thorpe, PhD, Leeds Institute of Clinical Trials Research, University of Leeds, Worsley Building, Clarendon Way, Leeds, LS2 9NL, UK; telephone: 0113 343 2785; e-mail: b.thorpe@leeds.ac.uk

Financial support for this study was provided entirely by the National Institute for Health Research through a Research Methods Fellowship (BT) and a Research Methods Internship (OC). The funding agreement ensured the authors' independence in designing the study, interpreting the data, writing, and publishing the report.

utilities.¹ Given the resources involved in the design and implementation of these models, when reassessment of the treatment decisions is required due to accrual of new data, conducting a Bayesian update of the model parameters may be more efficient than designing a new model.

Many parameters within decision models are conditional on patient covariates, representing demographics, clinical histories and treatments. These covariates might be necessary to allow for changes in patients' characteristics as they transition between health states within the model, or they might be used to restrict the analysis to a specific patient population. If, for example, a decision model depends on prediction equations for risk of clinical events then the coefficients of the regression equations will have been adjusted for the other variables involved, and this should be reflected in the analysis.² Any Bayesian updating of these equations must therefore update all coefficients simultaneously in order to accommodate correlations between coefficients.³

Intermittently missing data are common in updates that require data on many different covariates. They may be handled by treating missing values as uncertain quantities and assigning distributions to them.^{4,5} A potentially greater obstacle is the complete absence of some necessary variables from the accrued data. Commonly-used missing data methods cannot be applied because they require at least some observations for each variable.

Our motivating example is an update of an existing cost-effectiveness study and its probabilistic sensitivity analysis (PSA). Assessment of decision uncertainty through PSA, by assigning distributions to uncertain model parameters and propagating that uncertainty through the model,⁶ is a requirement of decision makers such as The National Institute for Health and Care Excellence.⁷ It is also considered a hallmark of good practice by the ISPOR-SMDM Modeling Good Research Practices Task Force, who recommend incorporating all available data through evidence synthesis techniques.² In our example, individual patient data (IPD) had become available, which we hoped to incorporate in the PSA by performing a Bayesian update of the parameter distributions. However, a necessary covariate had not been observed.

Ignoring new data would not be good practice as it would ignore current, available and relevant evidence. Proceeding with an inappropriate synthesis, ignoring important correlations between observed and unobserved covariates, could produce biased results. Our goal in this paper is therefore to present simple methods to overcome the problem of unobserved covariates in a Bayesian update. These are similar to methods used for intermittently missing data,

with the important difference that distributions assigned to unobserved covariates must be informed by external evidence, such as estimated distributions from published research or IPD from another study.

Full details of our motivating example are given in the following section. In the subsequent section, we propose methods to handle unobserved covariates in a number of general situations. We apply these methods to our example, illustrating their flexibility and highlighting potential difficulties. In the final section, we discuss their advantages and limitations. Although we focus on updating the inputs of a cost-effectiveness model, rather than the model structure, our motivating example demonstrates how the structure and assumptions of the model are important considerations in applying these methods.

Motivating example

Coronary heart disease (CHD) is the stenosis (narrowing) of coronary arteries. The extent of stenosis, and thus a patient's need for revascularisation, can be assessed through diagnostic imaging techniques, including cardiovascular magnetic resonance imaging (MRI), single-photon emission computed tomography (SPECT) and X-ray coronary angiography (CA).

The CE-MARC study^{8,9} compared the accuracy of SPECT and MRI in the diagnosis of CHD with *significant* stenosis (i.e. stenosis sufficient for revascularisation). Eligible patients were referrals to cardiologists with suspected angina. Participants were scheduled to undergo both MRI and SPECT. Each participant was also scheduled to have the true extent of their stenosis classified by CA (Table 1), which was assumed to be the gold standard (perfect sensitivity and specificity).

CHD with significant stenosis	At least 70% stenosis of a first order coronary artery at least 2mm in diameter, or left main stem stenosis of at least 50%.
CHD without significant stenosis	No significant stenosis, but at least 10% stenosis of at least one first order coronary artery at least 2mm in diameter, or at least 10% stenosis of the left main stem.
No CHD	Less than 10% stenosis of all first order coronary arteries at least 2mm in diameter, and less than 10% stenosis of the left main stem.

Table 1: Classification of CHD by angiography. The CE-MARC study considered diagnosis of CHD with significant stenosis. Further classification of patients without significant stenosis into those with and without CHD was introduced later, in the cost-effectiveness study.

Alongside the CE-MARC trial analysis, a cost-effectiveness study was undertaken.¹⁰ This study considered the subset of the CE-MARC population that was eligible for exercise tolerance testing (ETT) and had not previously undergone revascularisation. Eight diagnostic strategies were compared, including combinations of ETT, SPECT, MRI and CA. For example, a patient might undergo an MRI-CA sequence. A patient testing negative in MRI would not undergo the subsequent CA or revascularisation. A patient testing positive in MRI would then undergo CA and thus be correctly diagnosed and revascularised if necessary.

Using a decision analytic model (Figure 1), expected costs and quality-adjusted life years (QALYs) were estimated for each strategy over a 50-year time horizon. Each strategy would allocate patients to one of four states: true positive (TP), true negative (TN), false negative (FN), or death (caused by CA). This was modelled using decision trees. Positive test sequences would always end with CA and therefore false positives were not considered possible. For patients with and without CHD, separate Markov models were used to model pathways following testing.

Figure 1: The cost-effectiveness model used in CE-MARC, illustrated for the MRI-CA testing strategy. Health states are represented by ellipses, events by rectangles, and transitions by arrows. Only patients with significant stenosis, confirmed by CA, would undergo revascularisation and thus enter the TP state. Transitions from FN to TP would require identification via CA, followed by revascularisation.

For patients with CHD, transitions involving CV events (defined as myocardial infarction, cardiac arrest or cardiovascular death) were explicitly modelled, with transition probabilities depending on the incidence rate and mortality of these events. Before long-term follow-up data from CE-MARC became available, the required estimates were taken from published results of the EUROPA study.^{11,12}

For example, a patient’s time to first CV event, which we denote here by y , was assumed to be exponentially distributed, with event rate λ depending on 17 patient-level covariates in a proportional hazards model:

$$y \mid x, \beta \sim \text{Exp}(\lambda), \quad \lambda = \exp(\beta_0 + \beta_1 x_1 + \dots + \beta_{17} x_{17}),$$

where $\beta = (\beta_0, \dots, \beta_{17})$ denotes a vector of parameters and $x = (x_1, \dots, x_{17})$ includes baseline characteristics such as previous myocardial infarction and previous revascularisation. In the CE-MARC cost-effectiveness study, β was assigned a multivariate normal distribution, with mean and covariance matrix as estimated in EUROPA.

However, EUROPA recruited patients with evidence of CHD, such as previous myocardial infarction (64% of patients recruited), previous revascularisation (55%) or angiographic evidence of significant stenosis (61%), and a few with only a positive nuclear stress test (5%). Therefore estimates from EUROPA might require adjustment to accurately reflect the population of the CE-MARC cost-effectiveness analysis. In particular, patients with CHD in the TN state of Figure 1 do not have significant stenosis and therefore do not necessarily meet the main inclusion criteria of EUROPA. Despite this, the original CE-MARC cost-effectiveness model used the same estimated incidence rates from EUROPA in all three initial states (TN, FN and TP). This required an implicit assumption that, in patients with CHD, time to first CV event is independent of the extent of stenosis, given the other patient characteristics adjusted for in EUROPA.

Annual follow-ups for cardiac events were planned for each CE-MARC patient over a period of at least five years.¹³ With these data, the cost-effectiveness

analysis could be updated by obtaining posterior distributions for model parameters related to the rate of CV events. Furthermore, information on extent of stenosis, from the detailed angiography results of each patient, could be used to investigate the assumption described above. Our goal was therefore to use the newly available IPD both to update model parameters and to adjust for an additional variable indicating severity of stenosis. CV events are associated with high risk of death, high treatment costs and reduced quality of life, and therefore such an update could have a significant impact on the cost-effectiveness results.

A major obstacle to this update was the complete absence of a necessary variable, symptomatic angina, from the CE-MARC data. This was a covariate in the proportional hazards model estimated in EUROPA, and was also used to specify scenario analyses in the CE-MARC cost-effectiveness study. It was therefore important that we did not ignore this variable in our update.

Methods

General problem

Suppose we wish to update model parameter estimates from an initial study (Study 1) using data from a new study (Study 2). We assume that we only have summary data from Study 1, but that IPD are available from Study 2.

We are interested in the situation where a conditional probability distribution $p(y | x, \beta)$ has been estimated in Study 1; this represents the distribution of an outcome y (e.g. time to CV event) conditional on covariates x and regression coefficients β . A joint distribution $p(\beta)$ describes the current uncertainty in the parameter vector β , prior to observing data from Study 2.

We denote by $y_{(k)i}$ and $x_{(k)i}$ the values of the outcome variable y and covariate vector x for the i -th individual of Study k , where $i = 1, 2, \dots, n_{(k)}$. Suppose that the vector of covariates can be decomposed as $x = (x', x'')$, such that the vector $x''_{(2)i}$ was observed for each individual in Study 2 but the vector $x'_{(2)i}$ was not. Our general problem reduces to estimating the posterior distribution of β conditional on all $n_{(2)}$ observations $(y_{(2)i}, x''_{(2)i})$ from Study 2.

Updating when all covariates are available in Study 2

Bayesian updating is most straightforward when we observe all variables (y, x) for each individual in Study 2. We only need to specify the prior $p(\beta)$ from Study 1 and the likelihood $p(y | x, \beta)$ in order to obtain the posterior distribution of β , conditional on all $n_{(2)}$ observations $(y_{(2)i}, x_{(2)i})$, via Bayes' theorem.

The necessary relationships between observed data and model parameters can be specified in the directed acyclic graph (DAG)¹⁴ shown in Figure 2. We have the observed data $y_{(2)i}$ and $x_{(2)i}$ from Study 2 (rectangular nodes), and an uncertain parameter vector β (circular). Arrows into $y_{(2)i}$ show its dependence on $x_{(2)i}$ and β through the specified likelihood. The plate (dashed rectangle) indicates that the same relationship holds for all $i = 1, 2, \dots, n_{(2)}$.

Figure 2: DAG for updating $p(\beta)$ when all covariates are observed in Study 2.

In many Bayesian analyses, particularly those involving many parameters, mathematical derivations of posterior distributions are unavailable. Markov Chain Monte Carlo (MCMC) methods, such as Gibbs sampling, are therefore used to sample from the posterior distribution of the model parameters.¹⁵

Unobserved covariates in Study 2

Suppose that a subset x' of the covariates $x = (x', x'')$ was not observed in Study 2, so that the observed data in Study 2 comprises $(y_{(2)i}, x''_{(2)i})$ for $i = 1, 2, \dots, n_{(2)}$. Since our likelihood is $p(y | x', x'', \beta)$, without data on $x'_{(2)i}$ we cannot directly use Bayes' theorem to obtain the posterior distribution of β . However, if we can assign a distribution to each $x'_{(2)i}$, conditional on $x''_{(2)i}$, we might proceed by treating each $x'_{(2)i}$ as an uncertain parameter.

We need to identify a distribution, $p(x' | x'', \gamma)$ (with parameter vector γ), relating the unobserved covariates and the observed covariates. For example, this distribution might have been estimated in a previous study (Study 3). This leads to the model in Figure 3.

Figure 3: DAG for updating β when there are unobserved covariates x' in Study 2 imputed from observed covariates x'' .

When performed via Gibbs sampling, this method of handling unobserved covariates is analogous to multiple imputation of missing data. All uncertain parameters, including each component of the vector $x'_{(2)i}$, are sampled sequentially from their full conditional distributions (conditional on all other parameters and the observed data). In each iteration, values of β are therefore sampled conditional on the previously sampled values of $x'_{(2)i}$.

Properties of this method are also comparable to the usual recommendations¹⁶ for imputation of missing covariates: each component of $x'_{(2)i}$ is drawn

from a distribution that conditions on all other variables in the model of interest (including $y_{(2)i}$) and additional parameter uncertainty is accounted for in the prior $p(\gamma)$.

We have not assumed that we have IPD from the distribution of x' , only that the parameters and nature of their distribution are available from an existing study. However, if we had IPD from Study 3, we could estimate $p(x' | x'', \gamma)$ ourselves and (for example) determine $p(\gamma)$ using asymptotic properties of maximum likelihood estimators.

Additional assumptions required for the case of unobserved covariates in Study 3

Where the number of covariates in x'' is large, it is possible, even likely, that some of these will not have been observed in Study 3 (or simply not included when estimating the distribution of x'). Thus the covariate vector x can be decomposed into $x = (x', x'', x''')$, where x' was observed in Study 1 and Study 3 only, x'' was observed in all three studies, and x''' was observed in Study 1 and Study 2 only.

If we have $p(x' | x'', \gamma)$ from Study 3, we can use this to update β , provided that we can make an additional assumption that x' is independent of x''' given x'' and γ ; we denote this by $x' \perp\!\!\!\perp x''' | (x'', \gamma)$. Figure 4 shows the DAG corresponding to these assumptions.

Figure 4: DAG for updating β when there are unobserved covariates x' in Study 2 imputed from a subset x'' of the observed covariates.

Additional covariates in Study 3

Another potential difficulty is that the distribution estimated for x' in Study 3 might include additional, unwanted covariates. Suppose that we have an estimated distribution $p(x' | x'', z, \gamma)$ from Study 3, where z is a vector of covariates not in x . Suppose also that we do not have IPD from Study 3. Our update can still be performed if z was observed in Study 2 and if we can assume both $x' \perp\!\!\!\perp x''' | (x'', z, \gamma)$ and $y \perp\!\!\!\perp z | (x, \beta)$. See Figure 5.

Figure 5: DAG for updating β when there are unobserved covariates x' in Study 2 imputed from observed covariates x'' and additional observed covariates z .

Additional covariates can sometimes be used to our advantage. For example, if the IPD from Study 3 are available, we might incorporate additional covariates to strengthen confidence in our independence assumptions. We would choose variables z that are not in x and then estimate $p(x' | x'', z, \gamma)$ and $p(\gamma)$. For example, we might choose z such that an assumption $x' \perp\!\!\!\perp x''' | (x'', z, \gamma)$ is more plausible than an assumption $x' \perp\!\!\!\perp x''' | (x'', \gamma)$. This is analogous to inclusion of auxiliary variables in a multiple imputation model to make a ‘missing at random’ assumption more plausible.

This two-stage approach (estimation of $p(x' | x'', z, \gamma)$, followed by updating β) could be performed in a single step by including the Study 3 data directly in our Bayesian model, as illustrated in Figure 6.

Figure 6: DAG for updating β and γ simultaneously using data sets from both Study 2 and Study 3.

Handling unobserved covariates in multiple studies in a one-stage analysis

If Study 2 and Study 3 do not provide covariates z for which assumptions $x' \perp\!\!\!\perp x''' | (x'', z, \gamma)$ and $y \perp\!\!\!\perp z | (x, \beta)$ are plausible then the model in Figure 6 is not tenable. One option is to use the covariates x'' , available in both Study 2 and Study 3, to impute x''' in Study 3, as illustrated in Figure 7. This requires additional distributional assumptions: $p(x''' | x'', \theta)$ with parameter vector θ , relating observed and unobserved covariates, and a prior $p(\theta)$. An assumption $x' \perp\!\!\!\perp x''' | (x'', \gamma)$ is not required.

Figure 7: DAG for updating β , using data sets from both Study 2 and Study 3, with imputation of unobserved covariates x' in Study 2 and x''' in Study 3.

The models presented so far apply to a broad range of problems and can be extended to more specific problems after careful specification of additional distributions and conditional independence assumptions. For example, the model in Figure 7 could be extended to include auxiliary variables z in the imputation of x' and x''' .

Application

We now describe the application of the proposed methods to our example of updating the CE-MARC cost-effectiveness analysis. We focus on updating parameters related to the rate of CV events for patients with CHD in the three

initial states (FN, TN and TP) of Figure 1.

Recall that the distribution of time to first CV event, y , was estimated in the EUROPA study using a parametric proportional hazards model with 17 baseline characteristics as covariates. In the CE-MARC cost-effectiveness model, the population consisted of patients with no history of revascularisation, so the likelihood and prior information from EUROPA reduced to:

$$y \mid x, \beta \sim \text{Exp}(\lambda), \quad \lambda = \exp\left(\beta_0 + \sum_{i=1}^{16} \beta_i x_i\right), \quad \beta = (\beta_0, \dots, \beta_{16}) \sim N_{17}(\mu, \Sigma),$$

where λ now denotes the CV event rate for a patient with baseline characteristics $x = (x_1, \dots, x_{16})$ who has not undergone revascularisation prior to testing. Given the EUROPA eligibility criteria, these distributions were considered to be directly applicable to patients with significant stenosis, and therefore to transitions from the FN and TP states. However, as noted, their application to transitions from the TN state required an assumption that CV event rate does not depend on the extent of stenosis, given its existence and given the values of other covariates in the model. This assumption was not made on the basis of any established evidence. For CE-MARC participants, we had data on both time to first CV event (from follow-up) and stenosis (from angiography), and could therefore investigate the validity of this assumption.

We extended the proportional hazards model to include x_{17} , an indicator for insignificant stenosis, and β_{17} , an additional regression coefficient:

$$y \mid x, \beta \sim \text{Exp}(\lambda), \quad \lambda = \exp\left(\beta_0 + \sum_{i=1}^{17} \beta_i x_i\right),$$

where $x = (x_1, \dots, x_{17})$ and $\beta = (\beta_0, \dots, \beta_{17})$. The additional hazard ratio ($e^{\beta_{17}}$) would apply to patients in the TN state, since these patients would have only insignificant stenosis ($x_{17} = 1$). Patients in the TP or FN states would all have significant stenosis at the time of testing ($x_{17} = 0$), and therefore the model for y would reduce to the one specified in EUROPA.

We took the prior $(\beta_0, \dots, \beta_{16}) \sim N_{17}(\mu, \Sigma)$, as estimated in EUROPA. We took a weakly informative, independent normal prior for β_{17} , with mean 0 and standard deviation 0.97, so that $e^{\beta_{17}}$ would have median 1 and around 50% probability of lying between 1/2 and 2. Our primary goal was therefore to update $p(\beta)$, the multivariate normal distribution taken as the prior of β , using IPD from CE-MARC.

We note that other cost-effectiveness analyses, including some with the same decision model structure as CE-MARC,¹⁷ have similarly allowed event rates to differ between states by multiplying them by appropriate factors. These factors

are often taken from previously published studies, without any adjustment. Possible correlations with other parameters in the decision model are therefore overlooked, and potentially related parameters are varied independently in PSA. A strength of our Bayesian approach is that, by using IPD and specifying only weak prior information for the additional parameter, these potential correlations are accommodated in the posterior.

Data

CE-MARC recruited 752 patients, of whom 296 were eligible for ETT, had not previously undergone revascularisation, and had angiographic evidence of CHD with either significant or insignificant stenosis. To avoid overcomplicating presentation of our methods, we excluded a further 24 patients with intermittently missing data in the observed covariates. We therefore present a complete case analysis using IPD from 272 CE-MARC patients.

Symptomatic angina, defined in EUROPA as a Canadian Cardiovascular Society (CCS) grading of II or higher (chest pain during ordinary physical activities), was not recorded in CE-MARC. Because of this, straightforward updating (Figure 2) of $p(\beta)$ was not possible.

We therefore introduced a third study into our analysis, CECaT,¹⁸ from which had access to a subset of the baseline data, including CCS grading. CECaT compared diagnostic strategies for CHD and recruited 898 patients who were eligible for ETT and not recently revascularised. Of these patients, 218 were randomised to and underwent angiography. All but one had CCS grading recorded at baseline. We selected these 217 CECaT patients as IPD for our analysis.

In our terminology, EUROPA is Study 1, CE-MARC is Study 2, CECaT is Study 3, and symptomatic angina at baseline is x' . We denote by x'' the vector of covariates available from both CE-MARC and CECaT, and by x''' the vector of covariates not available from CECaT (Table 2).

Baseline characteristics (x)		EUROPA	CE-MARC	CECaT
x' :	Symptomatic angina (CCS grade II – IV)	✓		✓
x'' :	Insignificant stenosis (negative angiography)		✓	✓
	Years older than 65	✓	✓	✓
	Gender	✓	✓	✓
	Systolic blood pressure (mm Hg)	✓	✓	✓
	History of myocardial infarction	✓	✓	✓
	Diabetes mellitus	✓	✓	✓
	Smoking	✓	✓	✓
	Obesity (body mass index over 30kg/m ²)	✓	✓	✓
	Family history of CHD	✓	✓	✓
	Lipid lowering drug use	✓	✓	✓
x''' :	Nitrate use	✓	✓	
	Total cholesterol (mmol/L)	✓	✓	
	Creatinine clearance below 80 (ml/min)	✓	✓	
	History of TIA, PVD or CVA	✓	✓	
	Angiotensin converting enzyme (ACE) inhibitor use	✓	✓	
	Calcium antagonist use	✓	✓	

Table 2: Baseline variables, their use in EUROPA and their availability in the CE-MARC and CECaT data sets. (TIA: transient ischaemic attack. PVD: peripheral vascular disease. CVA: cerebrovascular accident.)

Model selection

The strategy for model selection was to work up from the simplest model, examining data sources and assumptions until reaching a model adequately representing our situation. The model in Figure 3 could not be used because the CECaT data did not include some necessary covariates. The model in Figure 4 would require $x' \perp\!\!\!\perp x''' \mid (x'', \gamma)$, which is doubtful because x' is an indicator for symptomatic angina and x''' includes an indicator for use of nitrates, which are used to relieve and prevent angina.

We could have taken the approach of either Figure 5 or Figure 6, which would allow us to ignore x''' for CECaT patients. But these models would both require us to choose auxiliary variables z such that we could assume $x' \perp\!\!\!\perp x''' \mid (x'', z, \gamma)$. Alternatively, we could handle unobserved covariates in CECaT by specifying a joint conditional distribution for all variables in x''' , given x'' , and using the model in Figure 7.

We adapted the approach of Figure 7, avoiding the need to specify a joint distribution for x''' by choosing to impute only nitrate use for CECaT patients. Writing $x''' = (x^*, x^{**})$, where x^* denotes the binary covariate for nitrate use and x^{**} denotes the vector of covariates to be ignored, we used the model in Figure 8. The assumptions required were $x' \perp\!\!\!\perp x^{**} \mid (x'', x^*, \gamma)$ and $x^* \perp\!\!\!\perp x^{**} \mid (x'', \theta)$.

Figure 8: Selected model for updating β using the CE-MARC and CECaT data sets. Variables included are y (time to first CV event), x^* (nitrate use), x^{**} (total cholesterol, creatinine clearance below 80, TIA/PVD/CVA, ACE inhibitor use, calcium antagonist use), x' (symptomatic angina) and x'' (as in Table 2).

Implementation

Sampling methods provide a straightforward and efficient way to obtain an approximate posterior distribution for β . Specifically, we used MCMC methods in the freely available OpenBUGS software.¹⁴ An example program is provided in the online supplement to this paper.

Logistic regression models were assumed for the conditional distributions of the unobserved covariates, $p(x' \mid x'', x^*, \gamma)$ and $p(x^* \mid x'', \theta)$. Initial analysis using vague normal priors, $p(\gamma)$ and $p(\theta)$, led to very unstable Markov chains. To overcome such problems, Gelman et al.¹⁹ proposed a default method to assign weakly informative priors for logistic regression: shift each covariate by subtracting its mean, divide each continuous covariate by twice its standard deviation,²⁰ then assign independent Cauchy distributions to each parameter (centre 0, scale 2.5 for coefficients, scale 10 for intercept terms). These priors were chosen to give preference to reasonable values for the regression coefficients (around 50% prior probability of lying between -5 and 5).

We assigned priors $p(\gamma)$ and $p(\theta)$ using reasoning similar to that of Gelman et al. However, to avoid dependence of our priors on the CE-MARC and CECaT data, we shifted and scaled our covariates using summaries from EUROPA.

Also, since we would not anticipate large effect sizes in our context, and to avoid known limitations of Cauchy priors,²¹ we assigned normal priors with the recommended scale parameters. These are less conservative than Cauchy priors, giving each regression coefficient around 95% prior probability of lying between -5 and 5.

Some events were censored at known times $c_{(2)i}$ because of non-CV death, withdrawal or the end of follow-up. Assuming ignorable censoring, the censored event times $y_{(2)i}$ were sampled conditional on $y_{(2)i} > c_{(2)i}$. Implementing this in OpenBUGS is straightforward.²² Also, some CE-MARC patients were recorded at baseline to have suffered previous acute coronary syndrome, but it was not specified whether this had been myocardial infarction or unstable angina. Therefore the myocardial infarction covariate for these patients was sampled from a Bernoulli distribution, informed by the ratio of myocardial infarction to unstable angina observed in the ENACT study.²³

Model fit was checked graphically by comparing empirical data summaries to their posterior predictive distributions and by calculating Bayesian p-values.^{15, 24, 25} For example, potential underestimation or overestimation of event rates was investigated by assessing censoring of predicted event times. At each iteration of the Gibbs sampler, an event time was simulated for each CE-MARC patient conditional on the sampled value of β . For each iteration, we calculated the proportion π of predicted events that would be censored (at either $c_{(2)i}$ or at the final planned follow-up). We then compared a histogram of π (which had mean 90% and standard deviation 2%) to the actual proportion of censored times in the data (93%). This suggested some overestimation of event rates, with p-value $P(\pi > 93\%) = 0.1$. These checks were repeated in various subsets of the CE-MARC sample and under different scenarios for censoring, each producing similar results.

Results

The relevant input to the original cost-effectiveness model was a CV event rate λ for a ‘base case’ patient with a given set of baseline characteristics x . This rate was a deterministic function of β and x . For PSA, a rate λ was calculated from each sampled β .

In our analysis, we had two such rates as model inputs: one for a base case patient with significant stenosis (entering the model in either the FN or TP state), and another for an otherwise identical patient with CHD but only insignificant stenosis (entering the model in the TN state). Using our sampled values of β , the approximate posterior distributions of these rates are plotted

separately in Figure 9.

Figure 9: Histograms (left) showing posterior distributions of the rate of CV events (λ) for base case patients with significant or insignificant stenosis. Forest plot (right) showing the prior/posterior mean of each hazard ratio (e^{β_i}) with 90% highest prior/posterior density (HPD) intervals.

By construction, the prior distribution of λ in Figure 9 for a patient with significant stenosis was exactly the distribution obtained from EUROPA, as used in the original CE-MARC cost-effectiveness analysis. The posterior was very similar to this prior. This was because the update had little impact on the hazard ratios $e^{\beta_1}, \dots, e^{\beta_{16}}$, or on e^{β_0} which had prior mean 0.018 (highest prior density interval 0.008 to 0.028) and posterior mean 0.017 (highest posterior density interval 0.008 to 0.026).

However, the posterior of λ for a patient with insignificant stenosis was clearly different. On average, in the posterior, the event rate for a base case patient with insignificant stenosis was around half that of an otherwise identical patient with significant stenosis. Our posterior sample suggested around a 97% probability of the event rate being lower ($e^{\beta_{17}} < 1$) in patients with insignificant stenosis. The assumption made in the original analysis, that λ is the same in all three initial states, is therefore unlikely to be true.

Despite this, conclusions from the updated cost-effectiveness analysis, using the posterior distributions in Figure 9, were in agreement with those obtained using only the prior information from EUROPA. The base case scenario we considered was males, under 65, suffering from symptomatic angina and taking ACE inhibitors, who were otherwise average and without any other risk factors listed in Table 2. Analysis using the EUROPA priors gave 55% probability of the MRI-CA strategy being cost-effective at willingness to pay £30,000 per QALY. Expected costs and QALYs of MRI-CA were £13,920 and 10.73 respectively. The updated analysis produced the same probability and similar expected QALYs (10.76), but lower expected costs (£13,825). For all eight strategies compared, similar differences were seen between results of the two analyses. For example, SPECT-CA had expected costs and QALYs £13,943 and 10.73 before the update, then £13,848 and 10.75 after.

Our update only had a notable effect on event rates in the TN state. The proportion of patients entering TN would only vary between strategies through very small differences in mortality due to misdiagnosed patients undergoing unnecessary CA (Figure 1). Therefore comparisons between strategies would

be affected very little, potentially explaining the agreement in our two analyses. For example, although costs were lower in the updated analysis for all strategies, incremental costs between strategies were similar in both analyses. Had our update significantly affected event rates in the FN and TP states (for example, if estimates from EUROPA were less precise or were in clear disagreement with the observed data), then there might have been more impact on our conclusions because the proportion of patients entering these states depends more heavily on diagnostic accuracy of a strategy.

Discussion

We have presented a framework for updating decision model parameters based on regression equations when the available data do not include all of the necessary covariates. These methods are flexible and usually easy to implement in freely available sampling-based software. The graphical modelling approach encourages explicit statement of any independence assumptions. Using our motivating example, we demonstrated how these methods can be extended to more specific problems when necessary.

Our example had one unobserved covariate in each study. In other situations, it might be necessary to handle multiple unobserved covariates in each study. We would ideally specify multivariate distributions for these covariates. For example, for multiple unobserved binary covariates we might use a multinomial regression model. Also, for clarity, we demonstrated our methods using a complete case analysis. Limitations of such analyses are well known¹⁶ and Bayesian methods for dealing with intermittently missing data are described elsewhere.^{4,5} Use of these methods in combination with our proposed methods might require additional independence assumptions, or specification of multivariate distributions simultaneously describing unobserved and intermittently missing covariates.

We did not address the issue of how much weight should be given to the prior information. EUROPA was a large trial and produced very precise parameter estimates. However, the timing and setting of EUROPA (424 centres in EUROPE, 1997 to 2000) were very different to CE-MARC (a single centre in England, 2006 to 2009). Therefore the priors from EUROPA may have dominated to an unacceptable degree, which might explain why our model checks suggested some overestimation of event rates. A related issue is inconsistency between data sources. Trials have been criticised for their selection bias, with trial results often applied uncritically to a target population, despite differences in patient characteristics and disease severity.²⁶ Further work could investigate

how existing models for cost-effectiveness could be updated to target populations within this framework, with appropriate weighting of trial and target population data.

Imputation of unobserved covariates across multiple studies was inspired by methods from epidemiology.^{27–29} However, our methods have a different purpose. For example, Jackson et al.²⁷ fit one regression model using two data sets, each containing the same response variable but different subsets of the desired covariates. This was done to address confounding that might have arisen if the analysis was performed on a single data set, omitting unobserved covariates. In contrast, we considered a previously estimated model and introduced additional data to make an otherwise hopeless update possible, not to address confounding. We are unaware of any previous applications of similar methods in health economics.

In summary, we have presented methods to update an existing model using a data set with an incomplete set of covariates. However, these methods should not be applied blindly. Priors should be carefully chosen, independence assumptions should be justified, and discrepancies between model predictions and observed data should be investigated.

Acknowledgements

The authors thank Simon Walker and Professor Mark Sculpher (Centre for Health Economics, University of York) for providing the CE-MARC cost-effectiveness model, and Professor John Greenwood (Leeds Institute of Cardiovascular and Metabolic Medicine) for access to the trial data.

References

- ¹ Briggs A, Sculpher M. An introduction to Markov modelling for economic evaluation. *Pharmacoeconomics*. 1998 Apr;13(4):397-409.
- ² Briggs AH, Weinstein MC, Fenwick E, et al. Model parameter estimation and uncertainty analysis: A report of the ISPOR-SMDM modeling good research practices task force-6. *Value Health*. 2012;15:835-42.
- ³ Ades AE, Claxton K, Sculpher M. Evidence synthesis, parameter correlation and probabilistic sensitivity analysis. *Health Econ*. 2006 Apr;15(4):373-81.
- ⁴ Mason A, Richardson S, Plewis I, Best N. Strategy for modelling non-random missing data mechanisms in observational studies using Bayesian methods. *Journal of Official Statistics*. 2012 June;28(2):279-302.

- ⁵ Mason A (2009). Bayesian methods for modelling non-random missing data mechanisms in longitudinal studies (doctoral thesis). Imperial College London. Retrieved from <http://www.bias-project.org.uk/research.htm>
- ⁶ Baio G. Bayesian methods in health economics. Boca Raton, Florida: Chapman and Hall/CRC; 2012.
- ⁷ Claxton K, Sculpher M, McCabe C, Briggs A, Akehurst R, Buxton M, et al. Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. *Health Economics*. 2005 Feb;14(4):339-347.
- ⁸ Greenwood JP, Maredia N, Radjenovic A, et al. Clinical evaluation of magnetic resonance imaging in coronary heart disease: the CE-MARC study. *Trials*. 2009 Jul 29;10:62.
- ⁹ Greenwood JP, Maredia N, Younger JF, et al. Cardiovascular magnetic resonance and single-photon emission computed tomography for diagnosis of coronary heart disease (CE-MARC): a prospective trial. *Lancet*. 2012 Feb 4;379(9814):453-60.
- ¹⁰ Walker S, Girardin F, McKenna C, et al. Cost-effectiveness of cardiovascular magnetic resonance in the diagnosis of coronary heart disease: an economic evaluation using data from the CE-MARC study. *Heart*. 2013 Jun;99(12):873-81.
- ¹¹ Fox KM, et al. Efficacy of perindopril in reduction of cardiovascular events among patients with stable coronary artery disease: randomised, double-blind, placebo-controlled, multicentre trial (the EUROPA study). *Lancet*. 2003 Sep 6;362(9386):782-8.
- ¹² Briggs A, Mihaylova B, Sculpher M, et al. Cost effectiveness of perindopril in reducing cardiovascular events in patients with stable coronary artery disease using data from the EUROPA study. *Heart*. 2007;Sep;93(9):1081-6.
- ¹³ Greenwood JP, Herzog BA, Brown JM, et al. Prognostic value of cardiovascular magnetic resonance and single-photon emission computed tomography in suspected coronary heart disease: long-term follow-up of a prospective, diagnostic accuracy cohort study. *Ann Intern Med*. 2016;165(1):1-9.
- ¹⁴ Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS - A Bayesian Modelling Framework: Concepts, Structure, and Extensibility. *Statistics and Computing*. 2000;10:325-37.
- ¹⁵ Gelman A, Carlin JB, Stern HS, et al. Bayesian data analysis, third edition. Boca Raton, Florida: Chapman and Hall/CRC; 2013.

- ¹⁶ Faria R, Gomes M, Epstein D, White IR. A guide to handling missing data in cost-effectiveness analysis conducted within randomised controlled trials. *Pharmacoeconomics*. 2014;32:1157-70.
- ¹⁷ Genders TS, Meijboom WB, Meijs MF, et al. CT coronary angiography in patients suspected of having coronary artery disease: decision making from various perspectives in the face of uncertainty. *Radiology*. 2009 Dec;253(3):734-44.
- ¹⁸ Sharples L, Hughes V, Crean A, et al. Cost-effectiveness of functional cardiac testing in the diagnosis and management of coronary artery disease: a randomised controlled trial. The CECaT trial. *Health Technol Assess*. 2007;Dec;11(49):iii-iv,ix-115.
- ¹⁹ Gelman A, Jakulin A, Pittau MG, Su YS. A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat.* 2008;2(4):1360-83.
- ²⁰ Gelman A. Scaling regression inputs by dividing by two standard deviations. *Statist. Med.* 2008;27:2865-73.
- ²¹ Ghosh J, Li Y, Mitra R. On the use of Cauchy prior distributions for Bayesian logistic regression. *Bayesian Anal.* advance publication. 7 March 2017. doi: 10.1214/17-BA1051.
- ²² Lunn DJ, Jackson C, Best N, et al. *The BUGS book: a practical introduction to Bayesian analysis*. Boca Raton, Florida: Chapman and Hall/CRC; 2012.
- ²³ Fox KA, Cokkinos DV, Deckers J, et al. The ENACT study: a pan-European survey of acute coronary syndromes. European Network for Acute Coronary Treatment. *Eur Heart J*. 2000 Sep;21(17):1440-9.
- ²⁴ Gelman A, Meng X, Stern H. Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*. 1996;6:733-807.
- ²⁵ Gelman A. Exploratory Data Analysis for Complex Models. *J. Comp. Graph. Stat.* 13(4):755-87.
- ²⁶ Hartman E, Grieve R, Ramsahai R, Sekhon JS. From sample average treatment effect to population average treatment effect on the treated: combining experimental with observational studies to estimate population treatment effects. *J. R. Statist. Soc. A*. 2015;178(3):757-78.
- ²⁷ Jackson C, Best NG, Richardson S. Bayesian graphical models for regression on multiple data sets with different variables. *Biostatistics*. 2009;10(2):335-351.

²⁸ Jackson CH, Best NG, Richardson S. Hierarchical related regression for combining aggregate and individual data in studies of socio-economic disease risk factors. *J. R. Stat. Soc. A.* 2008;171(1):159-78.

²⁹ Molitor NT, Best N, Jackson C, Richardson S. Using Bayesian graphical models to model biases in observational studies and to combine multiple sources of data: application to low birth weight and water disinfection by-products. *J. R. Stat. Soc. A.* 2009;172(3):615-37.