15
Modeling Cost-Effectiveness

In this lecture we look at decision models for assessing alternative medical interventions in chronic disease. This is a good area to illustrate the basic concepts of cost-effectiveness analysis and to learn about modeling techniques that are useful in a number of other situations, when decision have complex consequences over a long period of time. Our discussion is a review of [12, 11]. These references, along with [8], also have more details about a real-life application to stroke prevention.

15.1 Elements of a decision model

Guidelines and recommendations for clinical practice can be developed by systematically evaluating and comparing alternative intervention strategies based on their predicted health outcomes and costs. Matchar (1994) discusses general decision analysis tools useful to this purpose. From a statistical point of view, the general framework for the comparison of alternative interventions is, to varying degrees of formalization, that of Bayesian decision theory: evaluation proceeds by deriving a joint predictive distribution of the outcomes under each strategy, and by using it to determine expectations of patient utilities, costs and so forth.

Determining the appropriate predictive distributions, utilities and costs, can sometimes be done directly, using a single data source. More often, though, one needs to combine several sources and explore hypothetical scenarios that have never been implemented before, such as a novel diagnostic technology, or a proposed legislation. Both these goals can be accomplished using a so-called decision model. The decision model serves as a central organizing focus, helping to choose data sets and determine which aspects and analyses are of primary interest. It is also a policy tool designed to assess the relative merits of various strategies. So the goals of a decision model are to integrate the evidence regarding alternative interventions and to provide predictions of expected clinical outcomes for different patient/intervention scenarios. Once the model is operational, the inputs are a patient profile and a specific intervention strategy, and the output is a joint predictive distribution for various outcomes of interest, such as patient's utilities, costs, survival times and so on. From the predictive distribution one can then derive the expectations that are typically used in practice for comparing interventions. For ease of exposition, we distinguish the following elements of a decision model:

- Health States and transitions;
- Duration Distributions;
- Intervention Effects;
- Health Outcomes and Costs.

We comment briefly on each of these elements.

15.1.1 Health States and Transitions

The core of the model is a representation of the progress of the disease, in the absence of clinical intervention. This begins with a list of relevant health states: an exhaustive and mutually exclusive classification of a patient's health at a given time, or age. Figure 15.1.1 illustrates a hypothetical elementary example: health states are represented by the ovals. Specifying the possible health states is a crucial element of modeling. Initially, it is based on informal incorporation of published results and expert opinion. Subsequently, the list and configuration of states may change as the result of estimation of other parts of the models. In particular, distinguishing between clinically different health states is important only if there also is a difference in terms of patient's prognosis, utilities or costs.

The health states provide a static representation of an individual's health. A model for the progress of the disease over time requires transitions between states. In Figure 15.1.1, these are represented by the arrows, and are mimicking a typical pattern for a chronic disease. For example, patients do not revert to the HEALTH state after entering the TIA state. It is useful to separate these individuals from the rest because they might be the target of of specific stroke prevention interventions, they may have higher

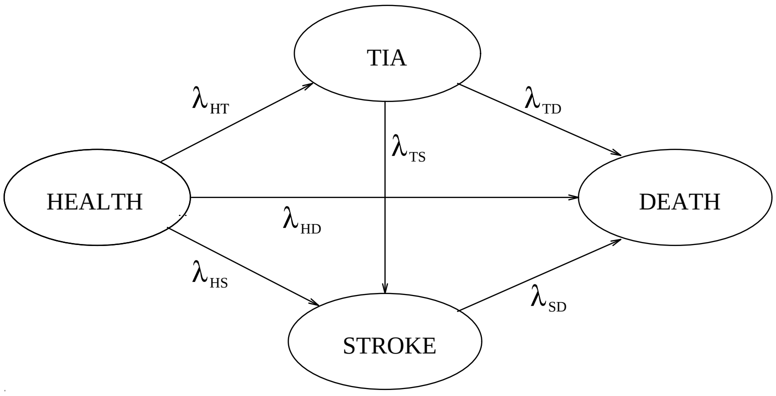

FIGURE 15.1.1. States and transitions of the disease process discussed in Section 2. The process is a highly stylized representation of typical stroke histories. Health states are an exhaustive and mutually exclusive classification of a patient's health at a given time, or age, and are represented by the ovals. TIA stands for Transient Ischemic Attach (a sort of mini-stroke). Transitions between states are represented by the arrows. Next to each arrow is the relevant transition rate.

medical bills and so forth. The resulting model is called a stochastic compartment model. Manton and Stallard (1988), give an in-depth discussion of stochastic compartment models.

15.1.2 Duration Distributions

The time at which transitions occur is usually modeled in terms of the patient's age, rather than chronological time. This in turn may be discrete or continuous. A convenient way of specifying the probabilistic structure of the model is to assign a duration distribution to each of the states.

Estimation of the duration distribution is typically based on epidemiologic studies. From a statistical viewpoint, there are several challenges. First, this is the stage in which it is necessary to model patient- and provider-specific features. This is typically implemented by generalized linear models, or by proportional hazard models, with covariates. Second, when there is a high number of duration distributions and covariates, estimation must often resort to information derived from different data bases or studies. In some cases, little is lost by focusing on one good study or data set properly addressing each of the questions of interest. Often, however, limitations of individual data sets make it necessary to formally combine information from different sources. Useful tools for addressing both problems can be developed within the framework of Bayesian hierarchical models.

15.1.3 Intervention Effects

The model components described so far can be used to evaluate the implications of no intervention (or standard care). To determine the implications of other available clinical strategies, an effective option is to modify the incidence rates in the natural history model, using intervention-specific risk ratios. These are obtained by reviewing and analytically synthesizing published research and other information on clinical strategies. Ideally, evaluation of risk ratios should be based on randomized clinical trials. In the example of Section 15.3 we will look at an intervention (action) that reduces the risk of stroke in TIA patients. Interventions could potentially be very complex. For example, only certain subsets opf patients may receive treatment, depending on their characteristics, previous medical history, previous interventions and so on.

In this phase there is clear opportunity to apply Bayesian meta-analysis. When meta-analysis is based on hierarchical modeling it can be cast in a unified framework with the estimation of the natural history model. An important advantage of Bayesian meta-analysis in this context is that it provides a formal assessment of the uncertainty in the evaluation of the risk ratio. Sources of uncertainty that can be incorporated include sampling variation within each study, between study variability and selection bias. Recent references to Bayesian methods in meta-analysis of clinical trials include Berry (1990), DuMouchel (1990), Eddy (1992), Eddy, Hasselblad and Shachter (1992), Morris and Normand (1992).

15.1.4 Health outcomes and costs

Consequences of clinical interventions are usually evaluated based on a number of different outcomes. Formally the problem is probably best captured by a multi-way utility function, whose dimensions may include survival times, patient's individual utilities, financial costs and so on. In the context of a stochastic compartment model, costs and individual utilities can be incorporated by attaching a value to each health state, and transition. The cost value represents the cost of remaining in the given health state for one unit of time, or the cost of making a certain transition. Similarly, individuals are assigned a utility value depending on their health state at each point in time. Intermediate values are interpreted accordingly. The quality-adjusted life year (QALY) measure for a health intervention is then determined by integrating this utility weight over an individual's life subsequent to the health intervention. Practical evaluation of individual utility is based on direct elicitation by interview of patients, physicians, or the general public. Cost estimation is typically based on administrative data sets.

15.1.5 Cost-effectiveness

This can in turn be used to perform various kinds of comparisons, such as cost-utility analysis (Kamlet 1991), cost-effectiveness analysis or one-dimensional utility maximizations by linear combination of components. Cost-effectiveness analyses is a fundamental tool in addressing this sort of of medical decision problems. The basic idea is that the two most important dimensions of the problem are the efficacy of the intervention, measured by life years, QALY's or other patient utilities, and cost. Two of the key quantities in CEA are an intervention's average cost and average effectiveness. These may be used to derive the marginal cost-effectiveness ratio (MCER) – for any two interventions the MCER is defined as the difference in average cost divided by the difference in average effectiveness. These averages are calculated over the unknown patient histories θ. Policy-makers use cost, effectiveness, and the MCER to help allocate scarce health care resources [14]. For example, "league tables" typically give point estimates of MCERs for various interventions.

15.2 Computing

15.2.1 Notation

The clinical interventions available are denoted by $\mathbf{a}$; x is the vector of covariates defining patient profiles; y are all the the data, including those used for estimating duration distribution, costs, and risk ratios; (λ, ξ) is the vector of unknown parameters of the natural history model, where λ is the subvector of the rates associated with the duration distributions, and ξ includes all the other parameters; κ_a is the vector of risk ratios for intervention $\mathbf{a}$; finally ω is the vector of outcome measures. The parameter vectors will typically be structured hierarchically, but we do not need to account for this explicitly in the notation of this section. In this setting the state of the world θ is a complete patient history over time. From the complte patient history, we can derive the quantities of interest, such as life lencth, cost and so on. So ω can be thought of as a function of θ.

The natural history model can be identified with the distribution:

$$p(\omega|x, \lambda, \xi), \tag{15.2.1}$$

defined implicitly by the assumptions on duration distributions, costs and utilities. Also, if you carried out a Bayesian analysis to to make inferences about the model parameters, you will have a posterior distribution, denoted by

$$p(\lambda, \xi|y). \tag{15.2.2}$$

Outcome predictions for each patient profile can then be obtained by the equation:

$$p(\omega|x, y) = \int p(\omega|x, \lambda, \xi)p(\lambda, \xi|y)d\lambda d\xi. \tag{15.2.3}$$

Similarly, intervention-specific predictions are derived by

$$p_a(\omega|x, y) = \int p(\omega|x, \lambda\kappa_a, \xi)p(\lambda, \kappa_a, \xi, |y)d\lambda d\xi d\kappa_a, \tag{15.2.4}$$

where $\lambda\kappa_a$ denotes point-wise multiplication of the elements in the vectors.

If $\omega = (\omega_1, \omega_2)$, where ω_1 is cost and ω_2 is effectiveness, the marginal cost-effectiveness ratio for intervention $\mathbf{a}'$, relative to intervention $\mathbf{a}$ (say standard care), for a patient with profile x, is

$$\text{MCER} = \frac{\int \omega_1 p_{a'}(\omega_1|x, y)d\omega_1 - \int \omega_1 p_a(\omega_1|x, y)d\omega_1}{\int \omega_2 p_{a'}(\omega_2|x, y)d\omega_2 - \int \omega_2 p_a(\omega_2|x, y)d\omega_2} \tag{15.2.5}$$

The Bayesian approach to prediction, based on averaging over unknown values of the parameters in (15.2.3) and (15.2.4), can be contrasted to the *ad hoc* approach of estimating the model by replacing point estimates for the parameters in (15.2.1). The latter is easier computationally, but it leads to predictive distributions that are under dispersed, as no accounting is taken of uncertainty about the parameter estimates. Analysis based on replacing point estimates in (15.2.1) are often accompanied by sensitivity analysis. This consists of repeating the analysis wvarying the assumed input, to evaluate their impact on the conclusions.

Bayesian predictions produce answers that incorporate both the variability in the unknown patient histories θ and the unknown model parameters (λ, ξ, κ). Sometimes it is interesting to separate the two sources of variability in presenting the results of the analysis. For example, there is a growing recognition that reporting estimates of MCERs is necessary but not sufficient, and that describing their uncertainty with respect to model parameters is critical to decision making [6, 3, 2, 7]. In the example of Section 15.3 we will see how this can be implemented. Also, the joint distribution of the outcomes can be used to assess the strength of the dependence between outcomes, as a function of clinical interventions.

O'Brien [10] distinguished between models based upon "prospective stochastic analysis" (such as randomized trials where the investigators have access to information about cost and effectiveness for individual patients), and "secondary analysis of retrospective data" (where the investigators do not typically have access to data on individual patients, and information about cost and effectiveness may be derived from different sources). They then discuss alternative approaches to estimating the precision of MCERs obtained from the first type of analysis. Among these approaches are univariate and multivariate sensitivity analyses, asymptotic approximations, and resampling. Resampling is recommended as a particularly promising approach, as it supports a formal assessment of precision based solely on the information obtained from the sample. In a similar vein, Gardiner *et al.* [3] discuss large sample approximate confidence intervals for CEA where the primary outcome is survival. General reviews of statistical issues arising in CEA are provided by Mullahy and Manning [9], Manning *et al.* [7] and Siegel *et al.* [13].

Describing the precision of MCERs obtained from secondary analysis of retrospective data is more complex. A general framework for quantifying uncertainty about MCERs obtained through stochastic decision models (e.g., decision models implemented using simulation). The input to stochastic decision models can include prospective stochastic analyses, secondary analyses of retrospective data, expert opinion, patient utility data, and possibly more. Methods based on Bayesian inference are more computationally intensive than large-sample approximation, but they are also more flexible in handling complex distributional assumptions as well as a variety of outcome measures of interest [12].

15.2.2 Simulation

Practical implementation of decision models can require a substantial computing effort. In realistic decision models, distributions (15.2.1), (15.2.2) and (15.2.3) have very high dimensionality or are intractable analytically and must be determined numerically. An efficient strategy is to set up a general simulation-based framework for the computations. This can address jointly several steps that are often thought of as separate. These are 1. Estimating model parameters, perhaps by combining multiple sources of information; 2. Predicting outcomes and 3. Performing a probabilistic sensitivity analysis.

Estimation of the model distribution (15.2.1) usually requires Monte Carlo methods. The prevalent strategy for computing (15.2.2) in hierarchical models are simulation methods based on Markov chain Monte Carlo [4]. Briefly, one generates a sample of parameter values from (15.2.2) by iteratively updating a subgroup of parameters each time. Each iteration provides a value of (λ, ξ, κ), where κ denotes the whole vector of risk efficacy parameters. The collection of iterations gives a sample from the desired posterior distribution. The advantage is that one breaks down a complex problem into many simpler problems.

In decision modeling, the most important distribution is the predictive distribution (15.2.3); we can regard the model equations (15.2.1) and the posterior distribution (15.2.2) as intermediate steps, and (λ, ξ, κ) as nuisance parameters. An efficient strategy for generating simulated samples of outcomes measures, and incorporating the uncertainty about the model parameters, is to set up a single simulation scheme to perform estimation, combination of information and prediction. The basic idea is to sample from the joint distribution of ω and (λ, ξ, κ), given x and y, and then use only the simulated values for the outcome, which constitute a draw from (15.2.3). In practice this can be implemented by nesting a prediction step within each Markov Chain Monte Carlo iteration. Each iteration generates (λ, ξ, κ); given these one can use a Monte Carlo step to simulate one draw from (15.2.1).

Using the resulting sample of outcomes for decision making is then a simple matter of summarization, which can be done using standard graphical tools such as boxplots or histograms. Important summaries such as cost-effectiveness ratios can be evaluated by Monte Carlo averages. Predictions incorporate variability of model parameters (λ, ξ, κ). Probabilistic sensitivity analysis with respect to other quantities relevant to the decision model can be performed by simulating a value from the postulated distribution of such quantities at each step of the simulation scheme.

15.3 Example

We now illustrate these general concepts with a specific hypothetical example. Consider a decision model whose goal is comparing an intervention to usual care. The core of the model is the stochastic process of Figure 15.1.1, describing the progress of the disease under usual care. It is a highly stylized representation of typical stroke histories. Health states are an exhaustive and mutually exclusive classification of a patient's health at a given time, or age, and are represented by the ovals. Transitions between states are represented by the arrows. The time at which transitions occur is random and is guided by duration distributions. To make things simple, we are going to assume that the probability of making a transition during the next unit of time is independent of the time already spent in the current state and in the previous states. The result is a so-called Markov model, with exponentially distributed waiting times in each states.

We will indicate by λ_{HT} the transition rate from HEALTH to TIA (Transient Ischemic Stroke). A similar notation is used for the other transition rates. Also, we will use the notation

$$\begin{aligned}\lambda_H &= \lambda_{HT} + \lambda_{HS} + \lambda_{HD} \\ \lambda_T &= \lambda_{TS} + \lambda_{TD}\end{aligned}$$

Under standard care, each will spend an exponential amount of time in the TIA state, with mean $1/\lambda_T$, and will move to the STROKE state with probability λ_{TS}/λ_T. Similar considerations apply to the other states. If you are not familiar with these concepts, you can look at a chapter on finite-state continuous-time Markov processes in a stochastic processes book such as [1].

PATIENT	ORIGINAL STATE	YEARS IN STATE	RESULTING STATE
1	HEALTH	0.84	TIA
1	TIA	1.46	STROKE
1	STROKE	0.30	DEATH
2	HEALTH	9.75	DEATH
3	HEALTH	6.99	TIA
3	TIA	2.42	STROKE
3	STROKE	0.53	DEATH
4	HEALTH	0.02	STROKE
4	STROKE	0.45	DEATH
...			
...			
998	HEALTH	11.41	TIA
998	TIA	0.37	DEATH
999	HEALTH	0.43	TIA
999	TIA	0.20	STROKE
999	STROKE	0.10	DEATH
1000	HEALTH	4.68	TIA
1000	TIA	0.30	STROKE
1000	STROKE	0.07	DEATH

TABLE 15.3.1. Structure of the hypothetical cohort study of Section 15.3, from [12]. Each row corresponds to a transition.

the most appropriate type of data for estimating the transition rates is a cohort study. For illustration, we consider a hypothetical cohort study with 1000 patients, all with the same profile of relevant covariates, and all followed until death. A patient history consists of the states visited by that patient and the corresponding sojourn times. Table 15.3.1 presents the data. Each row represents a transition, so there can be multiple rows for each patient. Under our simple assumptions, transition rate between any two states can be estimated by the total number of transitions observed in the sample, divided by the time at risk in each state.

We are going to imagine that the intervention is designed to reduce the chance that an individual who had a TIA will experience a stroke. A published clinical trial indicates that, as a result of the intervention, the hazard of making a transition from the TIA state to the STROKE state is reduced by 20%, and that such reduction is independent of the amount of time spent in the TIA state. The risk multiplier κ, or efficacy or the intervention, is 80%. We will carry out the illustrative analysis under two scenarios. First the value of κ will be assumed known, then unknown with a probability distribution based on the uncertainty assessment reported in the original study. There are known costs $c_t = .1$ and $c_s = 2$ for each unit of time spent in the TIA state and in the STROKE state, respectively.

An additional cost $c_b = 3$ will be incurred if the intervention takes place. These are, say, in thousands of dollars. For simplicity, costs are assumed to be known and not discounted. Extensions are straightforward.

A standard approach for computing point estimates of the MCER is:

i) Compute point estimates of the model parameters assuming usual care. In our example, the model parameters are the six transition rates denoted by the λ's in Figure 1, and their point estimates are the number of transitions divided by the time at risk.

ii) Compute point estimates of the model parameters assuming intervention. In our example, this amounts to reducing the transition rate from TIA to STROKE by 20%.

iii) Compute the expected effectiveness and cost under both usual care and the alternative intervention. In general this requires simulating a large cohort of individuals under both scenarios and computing summary measures (Monte Carlo method). Here, however, the model is simple enough that all of these quantities can be computed in closed form as a function of the parameters, as shown below.

iv) Compute the MCER of the intervention, in comparison with usual care, by taking the ratio of the estimated differences in effectiveness and costs.

We now need some additional notation and the expressions for the expected cost and length of life mentioned in iii). We denote by $Q_{\mathbf{a}}$ and $Q_{\mathbf{a}'}$ the effectiveness measure under strategies A (usual care) and B. Here we take Q to be years of life, while in the next section we will move to quality-adjusted life-years (QALY). Similarly, we denote by $C_{\mathbf{a}}$ and $C_{\mathbf{a}'}$ the cost under strategies A (usual care) and B. Finally $\Delta_Q = Q_{\mathbf{a}'} - Q_{\mathbf{a}}$ and $\Delta_C = C_{\mathbf{a}'} - C_{\mathbf{a}}$. All of these quantities are population parameters. Point estimates will be indicated by a "hat". These parameters are related to the transition rates λ's by the expressions:

$$\begin{aligned} Q_{\mathbf{a}} &= \frac{1}{\lambda_H} + \frac{\lambda_{HT}}{\lambda_H \lambda_T} + \left[\frac{\lambda_{HT}\lambda_{TS}}{\lambda_H \lambda_T} + \frac{\lambda_{HS}}{\lambda_H}\right] \frac{1}{\lambda_{SD}}, \qquad (15.3.1) \\ C_{\mathbf{a}} &= c_t \frac{\lambda_{HT}}{\lambda_H \lambda_T} + c_s \left[\frac{\lambda_{HT}\lambda_{TS}}{\lambda_H \lambda_T} + \frac{\lambda_{HS}}{\lambda_H}\right] \frac{1}{\lambda_{SD}}. \end{aligned}$$

The logic behind the derivation of $Q_{\mathbf{a}}$ is this: all individuals remain in HEALTH for an average $1/\lambda_H$ years; A fraction λ_{HT}/λ_H of individual move to TIA and spend that an additional $1/\lambda_T$ years , on average; finally a fraction given by the expression in square brackets reaches STROKE and spends there $1/\lambda_T$ years, on average. The other expressions are derived similarly.

TRANSITION RATES: USUAL CARE	TRANSITION RATES: TREATMENT
$\hat{\lambda}_{HT} = 0.111$	$\hat{\lambda}_{HT} = 0.111$
$\hat{\lambda}_{HS} = 0.050$	$\hat{\lambda}_{HS} = 0.050$
$\hat{\lambda}_{HD} = 0.051$	$\hat{\lambda}_{HD} = 0.051$
$\hat{\lambda}_{TS} = 0.914$	$\hat{\lambda}^*_{TS} = 0.914$ x $0.8 = 0.731$
$\hat{\lambda}_{TD} = 0.082$	$\hat{\lambda}_{TD} = 0.082$
$\hat{\lambda}_{SD} = 2.156$	$\hat{\lambda}_{SD} = 2.156$

$\hat{Q}_{\mathbf{a}} =$	5.569	$\hat{Q}_{\mathbf{a}'} =$	5.683
$\hat{C}_{\mathbf{a}} =$	.717	$\hat{C}_{\mathbf{a}'} =$	3.720

$\hat{\Delta}_Q = \hat{Q}_{\mathbf{a}'} - \hat{Q}_{\mathbf{a}} =$	.114
$\hat{\Delta}_C = \hat{C}_{\mathbf{a}'} - \hat{C}_{\mathbf{a}} =$	3.003
MCER $= \hat{\Delta}_C/\hat{\Delta}_Q=$	26.342

TABLE 15.3.2. Point estimate of MCER for the example of Section 2. First, we obtain point estimates of the transition rates under usual care, by computing the ratio of the number of transitions to the total time at risk in the cohort study. Second we compute estimates of the transition rates under the treatment, by applying the hazard ratios from the literature. Only one transition is affected in this case. Finally, using expressions (15.3.1) and (15.3.2) we determine the estimate of the MCER.

Also, if $\lambda^*_{TS} = \kappa\lambda_{TS}$ and $\lambda^*_T = \lambda^*_{TS} + \lambda_{TD}$, the corresponding expression under the intervention are:

$$\begin{aligned} Q_{\mathbf{a}'} &= \frac{1}{\lambda_H} + \frac{\lambda_{HT}}{\lambda_H\lambda^*_T} + \left[\frac{\lambda_{HT}\lambda^*_{TS}}{\lambda_H\lambda^*_T} + \frac{\lambda_{HS}}{\lambda_H}\right]\frac{1}{\lambda_{SD}}, \\ C_{\mathbf{a}'} &= c_b + c_t\frac{\lambda_{HT}}{\lambda_H\lambda^*_T} + c_s\left[\frac{\lambda_{HT}\lambda^*_{TS}}{\lambda_H\lambda^*_T} + \frac{\lambda_{HS}}{\lambda_H}\right]\frac{1}{\lambda_{SD}}. \end{aligned} \tag{15.3.2}$$

The estimated MCER is $26,342 per year of life saved. Table 15.3.2 summarizes the calculations. This figure might be entered into league tables and compared with other interventions that are competing for the same limited resources. Although one can rank point estimates in a league table, it is more interesting to study the entire distribution of cost, effectiveness and the MCER. Such distributions can then be used to determine: point estimates that account for the complex nonlinear relationship between inputs and outputs; probability intervals; and probability statements relevant for policy-making (e.g. regarding whether the intervention is more effective than usual care, whether the cost-effectiveness ratio is lower than a given value of interest, and so forth). Further support for considering the entire distribution of the MCER comes from the specific nature of comparisons between controversial health interventions. In many cases, there is the potential for large advantages and yet a substantial probability of a small advantage. This leads to a pronouncedly skewed distribution for the MCER, which is difficult to summarize appropriately using simply the mean MCER and confidence interval.

15.3.1 Inference

To derive a probability distribution for the cost, effectiveness and MCERm we need two elements: i) the joint conditional probability distribution of the unknown transition rates given the data ("a posteriori" distribution); ii) the probability distribution of the risk reduction due to the intervention. These distributions together imply a distribution over the cost-effectiveness of the treatment *via* expressions (15.3.1) and (15.3.2). In this way, estimating the inputs and assessing the precision of the results can be done jointly, using a common and coherent assessment of uncertainty. This is closely related to probabilistic sensitivity analysis. In fact one can think of it as a systematic, data- and expert-based process for providing a joint probability distribution on all unknowns, as well as a convenient computational framework for carrying out the analysis.

In the Markov model used in this section, it is simple to derive the a posteriori probability density function of the each of transition rates. Consider the derivation of $p(\lambda_{HT}|\text{Data})$. If N_{HT} is the number of transitions from HEALTH to TIA in the sample and t_H is the total time at risk in the HEALTH state, the likelihood function is $L(\lambda_{HT}) = (\lambda_{HT})^{N_{HT}} e^{-\lambda_{HT} t_H}$, because waiting times are independent and exponentially distributed. Assuming vague initial knowledge about the unknown transition rate, modeled via a so-called noninformative prior of the form $p(\lambda_{HT}) \propto 1/\lambda_{HT}$, and using Bayes rule, $p(\lambda_{HT}|\text{Data}) \propto p(\lambda_{HT})L(\lambda_{HT})$, the a posteriori probability density function of the rate λ_{HT} is:

$$p(\lambda_{HT}|\text{Data}) = \frac{t_H}{(N_{HT}-1)!}(t_H \lambda_{HT})^{-1+N_{HT}} e^{-\lambda_{HT} * t_H}, \tag{15.3.3}$$

a gamma distribution. Similar expressions hold for all the other rates. Alternative "a priori" assumptions can be used to incorporate expert opinion or information from published studies investigating the same transition rate. These are not pursued here, but can be accommodated explicitly within this framework. We assume that the probability distribution of the efficacy κ is a normal with mean .8 and standard deviation .05.

We are now in a position to sample from the distributions of the transition rates, by generating Gamma deviates, and to sample from the distribution of efficacy, by generating normal deviates. We generated 500 samples in both the known and unknown efficacy case. We used Splus, but most other statistical packages provide these deviates. For each sampled set of rates and efficacy, expressions (15.3.1) and (15.3.2) lead to easily computed values of the life expectancies, costs and MCER.

Figure 15.3.1 shows a sample from the distribution of the transition rates from health to the other states in the model. Similar samples obtain for the other transition rates. Figure 15.3.2 shows the resulting joint distribution of the increase in costs and increase in effectiveness. This is done under two scenarios: the top panel is derived incorporating the uncertainty about the λ's and ignoring the uncertainty about κ. The bottom panel incorporates uncertainty about both the λ's and κ. From this distribution one can evaluate the probability that the intervention will lead to improved life expectancy, which in this case appears to be close to one. In this formulation, the treatment is costly in two ways: the initial amount and the prolonged sojourn in the TIA state, which is random. Therefore parameter values leading to a higher efficacy also lead to higher cost. This determines the dependence that is observed in the joint distribution.

Finally, Figure 15.3.3 summarizes the distributions of the MCER. Again, the top row refers to the Bayesian inference results, presented under two scenarios: risk reduction known to be 20% and risk reduction assumed to be normally distributed with standard deviation .05 around the 20% value. Both distributions naturally incorporate the correlation between marginal cost and marginal effectiveness, due to the fact that both are functions of the same set of underlying model parameters, via (15.3.1) and (15.3.2). The more dispersed distribution on the right corresponds to the unknown rate case. The distribution of the MCER is sensitive to variability in the reduction rate, as should be expected.

The joint distribution of marginal costs and marginal QALY's emerging from the simulation contains all the information needed for decision making and can be thought of as the "goal" of an analysis of uncertainty in cost-effectiveness. Summarization of this joint distribution requires care. A standard approach, as we have seen, is to compute the mean marginal effectiveness $\hat{\Delta}_Q$ and the mean marginal cost is $\hat{\Delta}_C$. This is interesting but should be complemented by additional information. If the MCER cutoff in a league table is given, one can quickly calculate the probability that the intervention passes the CE test.

Determining the distribution of the MCER can lead to misleading analysis when there are draws where the intervention reduces costs, or it results in a decrement in QALY. In the first case the intervention dominates standard care, while in the second it is dominated. A more appropriate summarization is to estimate the probabilities of dominance and, separately the MCER conditional on lack of dominance. Also, in some draws, the effectiveness of the two strategies is very close. Thus, the denominator of the MCER is very small, implying that the MCER is large in absolute value. The distribution of the MCER contains a small proportion of very large values, and is skewed. The mean of the distribution, that is the mean of the ratios, can be a poor summary of the distribution of the MCER.

One can also conceive of circumstances where both the numerator and denominator are nearly zero. The fact that the MCER distribution has a very long right tail is valuable information for policy making and should not be discounted. On the other hand, standard strategies for reporting the results, such as point estimates —especially the mean— can be misleading, and should be replaced by probability intervals whenever feasible. It is important to recognize that outlying sampled values should be acknowledged but shouldn't dominate the reported MCER. Each outlier generated by a large numerator has a clear interpretation—the strategy with the smaller costs is clearly to be preferred. In the rare circumstance where both marginal cost and marginal effectiveness are near zero, the choice is a toss-up, as the strategies are nearly identical in terms of costs and benefits. Inspection of a two-dimensional scatterplot such as Figure 15.3.2 would reveal this information. Accordingly, if the outliers are located in, say, the top 2% of data, it is still acceptable to have a 95% probability interval from the 2.5-th to the 97.5-th percentile. Since effectiveness is usually the primary culprit, the problem could be addressed by defining marginal effectiveness/cost ratios rather than cost/effectiveness ratios. This is intuitively appealing, as effectiveness-cost ratios describe "bang for the buck".

15.4 Exercises

1. There is almost no connection between this problem and the material in this lecture, but it's a fun problem and and it doesn't fit anywhere else either. It's probably due to DeGroot.

Two firms A and B are to make competing secret bids on the price at which they will supply a computer system to a government agency. The firm with the lower bid will get the job. Firm A knows that its actual cost of supplying the system will be c dollars, and it has a prior distribution on the price that firm B will bid. Assume there is no cost in preparing and submitting the bid, and that the firm has the option not to bid. What price should firm A bid to maximize the profit if:

a. The prior is exponential with mean μ.

b. The prior is an arbitrary density f.

c. Answer parts a. and b. if there is a fixed cost of k dollars to firm A for preparing and submitting a bid.

2. A machine can be either functioning (F), in repair (R) or dead (D). Under the normal course of operations, the probability of making a transition between any of the three states, in a day, is given by this table.

		TO		
		F	R	D
	F	.92	.06	.02
FROM	R	.45	.45	.10
	D	.00	.00	1.0

These proabalilities only depend on where the machine is today, and not on its past history of repairs, age, etc (this is obviously an unreasonable assumption, but I want you to be done with this homework before the end of time...).

When the machine functions, the daily income is \$1000. When it is in repair, the daily cost is \$150. There is no income after the machine is dead. Suppose you can put in place a regimen that gives a lower daily income (\$900), but also decreses the probability that the machine needs repair or dies. Specifically, the transition table under the new regimen is

		TO		
		F	R	D
	F	.96	.03	.01
FROM	R	.45	.45	.10
	D	.00	.00	1.0

New regimen or old? Assume the decision maker is risk neutral and that the machine is functioning today. You can use an analytic approach or a simulation.

3. Reincarnation of the problem 2. A diabetes patient can be either in his/her normal state (F), hospitalized (H), or deceased (D). On each day the probability of making a transition between any of the three states, is given by this table.

		TO		
		F	H	D
	F	.92	.06	.02
FROM	H	.45	.45	.10
	D	.00	.00	1.0

These proabalilities only depend on where the patient is today, and not on his/her past history of hospitalization, age, etc (this is obviously an even more unreasonable assumption, but again I want you to be done with this homework before the end of time...).

When the patient is in F the quality of life is .99 compared to full health and the cost is \$5. When in the hospital, the the quality of life is .8 and the daily cost is \$250. Suppose you can take a drug that costs \$20 per day, but also decreses the probability of hospitalization ad death. Specifically, the transition table under the new regimen is

		TO		
		F	H	D
	F	.96	.03	.01
FROM	H	.45	.45	.10
	D	.00	.00	1.0

What is the marginal cost effectiveness ratio of the drug? Measure effectiveness in QALYs. Suppose that the patient is normal today. Again, you can use an analytic approach or a simulation.

4. Describe how you would approach the problem of estimating the cost effectiveness of regulating emission of pollutants into the environments. Pick a real or hypothetical pollutant/regulation combination, and briefly describe data collection, modeling and utility elicitation.

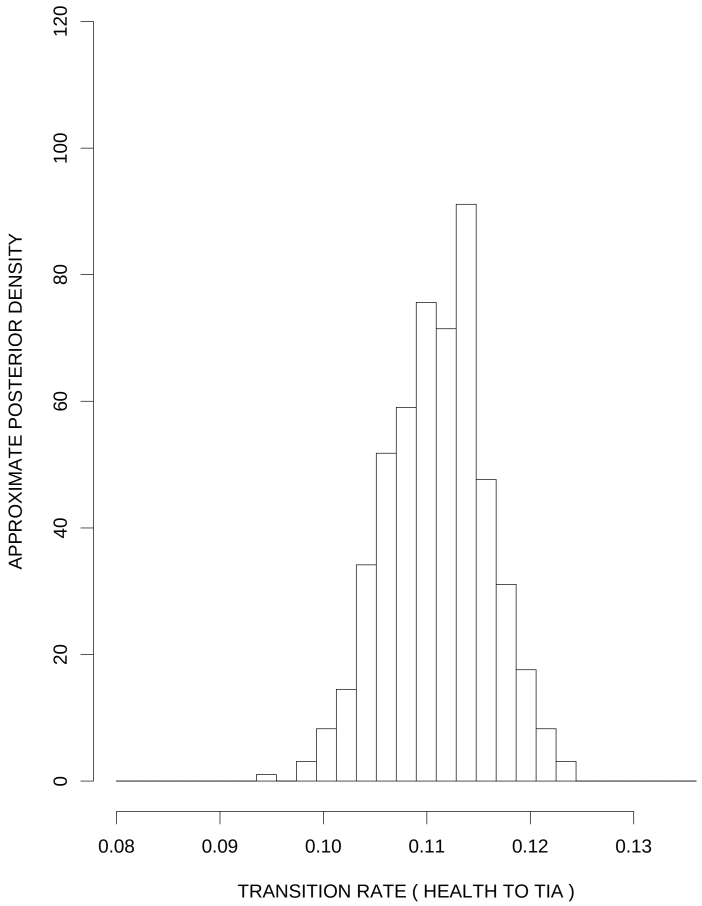

FIGURE 15.3.1. Comparison of distributions of transition rate from HEALTH to TIA.

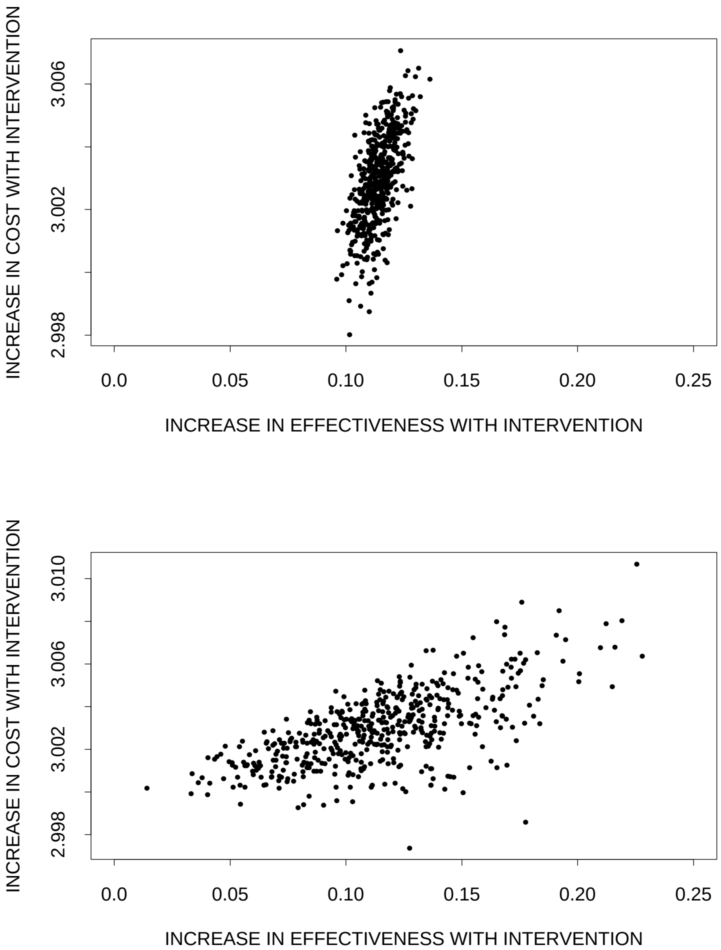

FIGURE 15.3.2. Comparison of joint distribution of the increase in costs and increase in effectiveness from adopting the intervention. The top panel is derived incorporating the uncertainty about the λ's and ignoring the uncertainty about κ. The bottom panel incorporates uncertainty about both the λ's and κ.

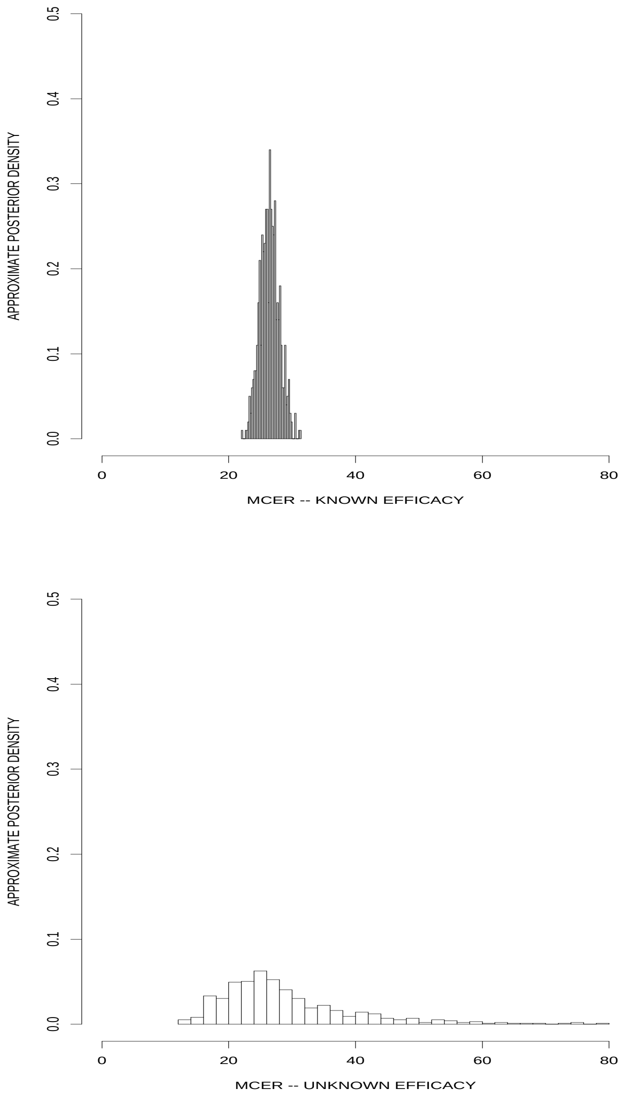

FIGURE 15.3.3. Comparison of distributions of the MCER under two scenarios: The top panel is derived incorporating the uncertainty about the λ's and ignoring the uncertainty about κ. The bottom panel incorporates uncertainty about both the λ's and κ. The uncertainty about the MCER is very sensitive to the inclusion of uncertainty about κ.

References

[1] U. N. Bhat. *Elements of Applied Stochastic Processes (Second Edition).* John Wiley & Sons, 1984.

[2] J. R. Bult, M. G. M. Hunink, J. A. de Vries, and M. C. Weinstein. A fully stochastic cost-effectiveness analysis. *Medical Decision Making*, 15:415, 1995.

[3] J. Gardiner, A. Hogan, M. Holmes-Rovner, D. Rovner, L. Griffith, and J. Kupersmith. Confidence intervals for cost-effectiveness ratios. *Medical Decision Making*, 15:254–263, 1995.

[4] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov Chain Monte Carlo in Practice.* Chapman and Hall, London, 1996.

[5] M. R. Gold, J. E. Siegel, L. B. Russell, and M. C. Weinstein, editors. *Cost-effectiveness in health and medicine.* Oxford University Press, Oxford, 1996.

[6] J. S. Hodges. Uncertainty, policy analysis and statistics (c/r: P276-291). *Statistical Science*, 2:259–275, 1987.

[7] W. G. Manning, D. G. Fryback, and M. C. Weinstein. Reflecting uncertainty in cost-effectiveness analyses. In Gold et al. [5].

[8] D. Matchar, G. P. Samsa, J. R. Matthews, M. Ancukiewicz, G. Parmigiani, V. Hasselblad, and J. Lipscomb. The stroke prevention policy model (sppm): Linking evidence and clinical decisions. *Annals of Internal Medicine*, 8:704–711, 1997.

[9] J. Mullahi and W. G. Manning. Statistical issues in cost-effectiveness analyses. In Gold et al. [5].

[10] B. J. O'Brien, M. F. Drummond, R. J. Labelle, and A. Willan. In search of power and significance: issues in the design and analysis of stochastic cost-effectiveness studies in health care. *Medical Care*, 32:150–63, 1994.

[11] G. Parmigiani, M. Ancukiewicz, and D. B. Matchar. Decision models in clinical recommendations development: the stroke prevention policy model. In D. A. Berry and D. K. Stangl, editors, *Bayesian Biostatistics*, volume 151 of *Statistics: Textbooks and Monographs*, pages 207–233, New York, NY, USA, 1996. Marcel Dekker.

[12] G. Parmigiani, G. Samsa, M. Ancukiewicz, J. Lipscomb, V. Hasselblad, and D. Matchar. Assessing uncertainty in cost-effectiveness analyses: Application to a complex decision model. *Medical Decision Making*, 17:390–401, 1997.

[13] C. Siegel, E. Laska, and M. Meisner. Statistical methods for cost-effectiveness analyses. *Controlled Clinical trials*, 17:387–406, 1996.

[14] F. Sloan. *Valuing health care*. Cambridge University Press, Cambridge, 1995.