CURRENT OPINION

Pharmacoeconomics 2005; 23 (4): 323-332
1170-7690/05/0004-0323/$34.95/0

Pharmacoeconomic Analyses Using Discrete Event Simulation

J. Jaime Caro

Caro Research Institute, Boston, Massachusetts, USA

Abstract

To date, decision trees and Markov models have been the most common methods used in pharmacoeconomic evaluations. Both of these techniques lack the flexibility required to appropriately represent clinical reality. In this paper an alternative, more natural, way to model clinical reality – discrete event simulation – is presented and its application is illustrated with a real world example.

A discrete event simulation represents the course of disease very naturally, with few restrictions. Neither mutually exclusive branches nor states are required, nor is a fixed cycle. All relevant aspects can be incorporated explicitly and efficiently. Flexibility in handling perspectives and carrying out sensitivity analyses, including structural variations, is incorporated and the entire model can be presented very transparently. The main limitations are imposed by lack of data to fit realistic models.

Discrete event simulation, though rarely employed in pharmacoeconomics today, should be strongly considered when carrying out economic evaluations, particularly those aimed at informing policy makers and at estimating the budget impact of a pharmaceutical intervention.

It is by now well accepted that pharmacoeconomic analysis[1,2] and the setting of health policy[3] require a model of the disease and its management.[4,5] In exceptional circumstances, a single study – a randomised clinical trial, for example – may provide all of the necessary information on costs and outcomes. Even then, a model is indispensable to address the economic impact of the intervention in actual practice[6] because, at a minimum, the experimental data must be applied to a real-world setting in order to reflect clinical reality.[7]

These economic models have typically been structured and analysed using decision trees.[8] This technique has been very successfully applied despite recognition a decade ago of the severe limitations of this approach when applied to medical problems.[9] In particular, decision trees impose a rigid structure based on mutually exclusive 'outcomes', they do not explicitly consider time, and they are very inefficient because every analysis requires computing all possible pathways (the 'branches'), sometimes multiple times.

Thus, an alternative approach – the Markov model – has been increasingly employed.[10] Instead of basing the model on mutually exclusive outcomes as the decision tree does, a Markov model represents the course of a disease in terms of mutually exclusive 'health states' and the transitions among them.[11] While this technique considers time more explicitly and can be analysed very efficiently, it retains some structural rigidity that can make appropriate representation of clinical reality difficult.

The requirement that all the aspects of the disease be denoted by a 'state' in a Markov model forces the

analyst to render features that are naturally continuous as discrete, such as weight gained (yes or no rather than the actual amount of weight gain), severity in specific categories, and so on. To begin to approximate their continuous nature, an unwieldy number of states have to be specified. For example, nearly 20 states are required just to reflect weight changes of ±40 pounds in increments of 5 pounds. This problem is compounded if the implications of the state change over time, as each instance then generates a new state. For example, the first month after a hospitalisation for psychiatric illness would be one state, the second month another, and so forth. A similar proliferation is imposed if the subsequent course of the disease depends on previous history, such as when the risk of adverse events depends on prior exposure to a given drug.

More critical still is the constraint in a Markov model that each patient can be in only one state at a time, leading to the requirement for multiple distinct states to represent all possible combinations (e.g. a minimum of four states are required to represent the combinations of 'depressed, yes/no' and 'psychotic, yes/no'). Needless to say, these conditions either force the analyst to invoke considerable simplifications (the option usually implemented today) or to accept immense complexity and inefficiency. Even with simplification, there remains no way to impose a hierarchy or sequence among the states within a cycle. For example, the analyst cannot specify that a visit to the doctor is to precede a dose change within the same cycle; these must be combined into a compound state where the temporal order is lost, or they have to be forced into separate cycles. Although, in practice, many 'work-arounds' have been developed to reduce some of these limitations – and many have even been incorporated into software – the description of a clinical problem only in terms of states and the transitions among them remains very restrictive.

In this paper, an alternative, more natural, way to simulate clinical reality is presented. Firstly, a typical example of a pharmacoeconomic problem is described. Then the essential elements of this approach – discrete event simulation (DES) – are defined and illustrated with the example. Next, the implementation of this type of model is delineated. Finally, the merits and disadvantages of this approach are discussed. Given the preponderance of Markov model and decision-tree techniques in pharmacoeconomics today, comparisons are made where appropriate, but the main purpose of this paper is to describe DES rather than contrast the methods or provide guidance on which to choose for which type of problem.

1. A Typical Example of a Pharmacoeconomic Problem

Most pharmacoeconomic problems involve quantifying the economic implications of a healthcare intervention.[12] This requires:

- defining the target population, the environment, the aspects of the disease that are at issue, and the intervention;
- structuring the possible course of patients in a logical, realistic order over time;
- considering the events that may occur, together with their health and economic implications;
- and providing a computational means to derive the chosen measures of value.[13]

For example, consider the management of patients with bipolar mood disorder ('manic-depressive illness'). Recently, the use of newer ('atypical') antipsychotics has been considered as an alternative to the established mood stabilisers such as lithium or antiepileptic drugs.[14] The atypical antipsychotics may have the same ability to control manic symptoms as mood stabilisers, but with fewer adverse effects and no need for routine monitoring. The question naturally arises: will this new more costly intervention improve the health of these patients at a reasonable incremental cost?

To answer this question, the analyst needs to start by creating a simulated target population. This is done by specifying the relevant characteristics of the patients to be simulated. These characteristics – attributes such as age, experience with previous treatments, employment status – are relevant in the sense that they affect some aspect of the problem, such as the course of disease or the effects of treat-

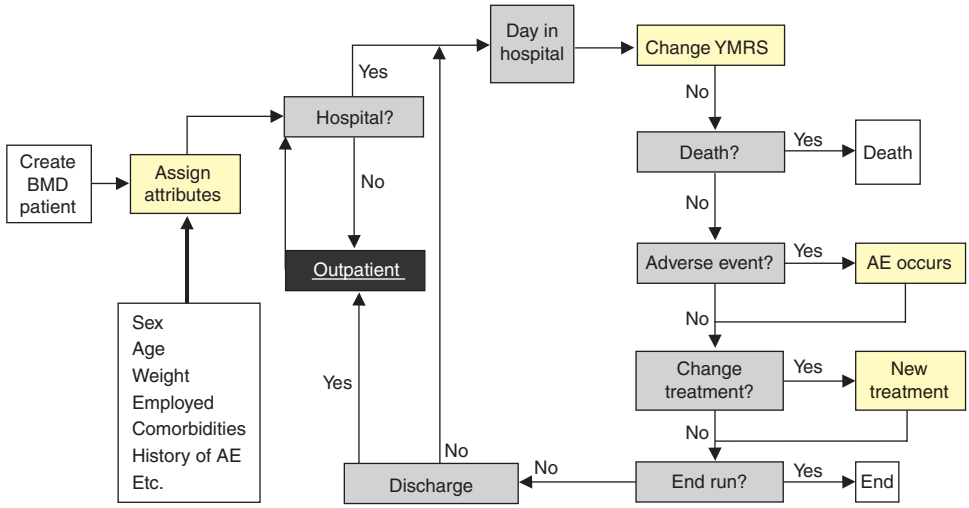

Fig. 1. Schematic representation of the initial course of patients with bipolar mood disease in acute mania. The outpatient submodel is very similar to the inpatient one, in that the YMRS is computed, the occurrence of death, adverse events and treatment changes are considered, and the time horizon may be reached. If the YMRS reaches the patient's personal threshold for readmission, then the patient moves back to the hospital. The attributes listed in the clear rectangle are specified because they affect other aspects of the model. For example, age affects the probability of death, employment status affects indirect costs, and so on. Yellow rectangles indicate that values are assigned; grey rectangles indicate points where the course may take one of several paths; the grey square indicates that resources are consumed; and the black box with underlined text points to additional detail (i.e. a submodel) not shown. **AE** = adverse event; **BMD** = bipolar mood disease; **YMRS** = young mania rating scale.

ments. Additional attributes can be specified to characterise the disease, for example its duration and severity, and to describe the intervention (dose, schedule, etc.). The environment in which this simulated treatment population exists also needs to be specified, including such things as the country or region, the calendar time and the currency. Thus, in one analysis we might want to consider a population of employed males and females aged between 30 and 50 years, with no previous treatment for bipolar disease who present with acute mania (defined by a score of >20 on the Young mania rating scale [YMRS][15]). We might also wish to assess the value of beginning treatment with an atypical antipsychotic in the US in 2004, relative to starting a mood stabiliser such as lithium. Although a real model of this situation would need to specify many more features, this gives an idea of what would need to be considered. In addition, we need to define various technical parameters such as the discount rate for costs and health benefits and the time horizon for the analysis.

Once the population to be simulated has been specified, the possible course of the disease needs to be detailed. For example, patients with acute mania will most likely be hospitalised, but some may be managed as outpatients. Those who are hospitalised will then be treated with either an antipsychotic or a mood stabiliser (or both if that option is allowed) and during each day in the hospital various things may happen: an adverse effect to the medication may be experienced, the clinical condition (i.e. the mood) may improve or deteriorate, the patient may be discharged, and so on. In any case, resources are consumed and time passes, and both need to be properly counted and possibly adjusted according to some quality weights and discounting rates. When the chosen time horizon is reached, the simulation must end and then the various measures of value must be computed – total and net costs, average time in a manic state, etc. (see figure 1).

2. Discrete Event Simulation

A DES[16] is a natural way to model the initial course of patients with bipolar mood disease. Indeed, the structure of the simulation closely replicates the course outlined in figure 1, and the components of the model correspond to the required elements. The fundamental components of the technique are described below.

2.1 Entities

A central component of DES is the entity. In general, entities are the items that flow through the simulation – the protagonists of the events. In clinical simulations of disease, most of the entities will be patients, but there can be other types – caregivers, for example. In contrast to decision trees and Markov models, which do not specify the patient and instead focus exclusively on outcomes or states, the patient is an explicit element in a DES.

Patients have attributes (e.g. age, sex, duration of disease), with each individual having a specific value for each characteristic. These values are defined at the start of the simulation and may be updated as events unfold: age increases, disease severity decreases, drug dose is titrated up, and so on. These updates can happen at particular points in time (e.g. age could be updated every 6 months), when the simulated patient experiences events (e.g. starting a new treatment might reduce the level of mania) or even 'continuously' (e.g. increasing weight every day). The analyst specifies when and how these updates take place, depending on the requirements of the problem. e.g., in figure 1, the YMRS score is updated every day in hospital (yellow rectangle that follows 'Day in Hospital'), weight would be updated if an adverse event occurs, and so on. Thus, all the necessary features of each patient with bipolar disease – age, sex, weight, employment status, previous treatments, previous adverse effects, current level of mania (score on the YMRS), current treatment, location of care, etc. – can be directly specified.

A particularly important attribute in pharmacoeconomic evaluations is the quality of the patient's life. This is incorporated in a DES by defining the corresponding attribute(s). This can be a simple value that is used to adjust survival, or it can be a more complex set of values that carry the scores of various quality-of-life instruments. Either way, these can be updated over time as events happen and they can be reported on their own or included as weights in the QALY measure.

Other quantities that are important in describing the environment of the simulation (e.g. currency) and the technical details that will govern the analysis (e.g. time horizon, discount rate) are encoded in variables. Their values may change during the simulation. In our example, we would specify the variable 'calendar time' to be '2004' at the start of the simulation, 'setting' to be 'US', 'discount rate' to be '3%' for both costs and benefits, 'time horizon' to be '5' years and so on.

2.2 Events

The second major element of the simulation is the events that may occur. An event is very broadly defined as anything that can happen during the simulation. Thus, it can be the occurrence of an adverse event, admission to hospital, a change in dose, or even the failure to show up at work. This concept extends well beyond the 'transitions' in a Markov model, as the event need not imply a change in the patient's state, though it is still important to consider it in the simulation. Events in our bipolar disease simulation might be: start treatment, discharge from hospital, become non-compliant, gain weight, visit doctor, attempt suicide, lose job, return to hospital, and so forth until the end of the simulation occurs. These events can happen in any logical sequence and even simultaneously; they can recur if that happens in reality and they can change the course of a given patient's experience by influencing that patient's attributes and the occurrence of future events with no restriction on 'memory'. e.g., the amount of weight gained can depend on previous weight gain, age and sex, but can be altered by changing treatment or a visit to the dietician. The rates at which events occur can take any functional form supported by the data (or assumptions). They can be dependent

on any attributes or variables and these functions can change over time as appropriate.

2.3 Time

The third fundamental component of a DES is time itself. An explicit simulation clock keeps track of the passage of time. This permits the analyst to clearly signal the start and end of the simulation and to create secondary clocks that track interim periods such as the length of stay in hospital, the time spent with symptoms, and, of course, the survival (quality-adjusted if appropriate) of patients. By making time explicit, a DES avoids one of the major problems of decision trees.[9] It also enables handling of time that is much more flexible than in Markov models, as there is no need to declare a cycle length. Time advances in 'discrete' jumps (hence the name) but the units can be minutes, days, months or whatever is convenient. Indeed, the simulation proceeds very efficiently because the clock is successively advanced to the time when the next event will occur, without wasting effort in unnecessary interim computations. Thus, when the patient with mania is in hospital, the simulation may advance in daily increments to allow for changes in the mood as treatment takes hold, and to permit decision rules about discharge based on response to treatment and other factors.

Once the patient is discharged the time period can lengthen until the next scheduled visit or a new adverse effect appears; this might be several weeks for one patient and months for another. In a very acute situation (management of a suicide attempt, for example) the time increments may shorten to hours or even minutes, if appropriate. The secondary clocks would track how long a patient has been on a treatment (which might influence adverse effects), the length of stay (which will affect cost), the duration of an episode (a determinant of QOL), and so forth.

2.4 Resources

An indispensable element for an economic evaluation is the explicit handling of resources. These are incorporated directly into a DES. An entity may consume a resource at any appropriate time. This consumption involves a defined number of units of the resource, used for a specific amount of time. Multiple resources may be consumed at the same time (e.g. our patient with mania may use on a particular day a hospital bed, a dose of an antipsychotic, a unit of doctor's time, and may lose a day of work).

While the patient is using a particular resource, it may be unavailable for use by others, depending on its capacity. This may lead to queues forming, something that is rarely considered in healthcare evaluations today because analysts typically assume an infinite capacity of each resource, despite the reality being clearly otherwise. A DES explicitly provides for queues,[17] even if they can be avoided in a simulation by setting the capacity of the resources to exceed the maximum number of simultaneous users. Limited availability of psychiatric beds may force a patient with mania to leave the queue and go to an outpatient facility; conversely, the opportunity cost implication of a stay in hospital shortened by treatment can be directly examined by simulating the smaller queue that results for other patients.

2.5 Means of Execution

The final crucial component is the means to execute the simulation, follow the desired logic and carry out all the calculations. This happens in an orderly fashion (figure 2) that ensures all the specifi-

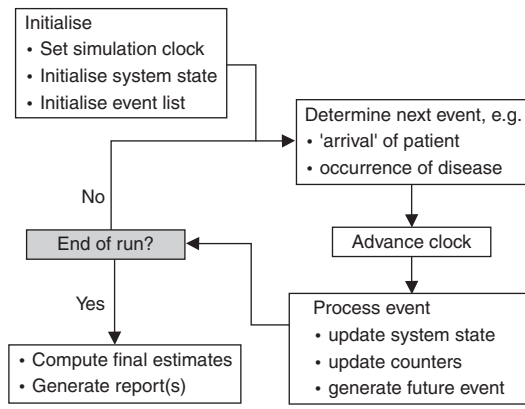

Fig. 2. Flow diagram of the computation process for a discrete event simulation.

cations detailed by the analyst are met. While very simple DESs can be carried out manually, most of those responding to real problems will require the use of a computer.[18] Available high-level simulators considerably simplify the process, facilitate creating and working with the models, and reduce computing times considerably.

3. Implementation

Many of the steps involved in a DES are common to all modelling approaches.[1] One must begin by formulating the problem in some detail, including the specifications for what the simulation is supposed to accomplish. This step also involves a description of the healthcare system that is to be modelled, in particular the details that pertain to the disease in question. From this solid foundation, the model can be designed conceptually (i.e. the logic of the simulation is laid out). It is very important to resist the temptation to allow the analyst's preferred technical approach (e.g. decision tree) to influence the conceptual design. Once the concept has been validated with help from relevant experts, data are sought to fit the model.

At this point, the development becomes specific to the technical approach. For a DES, the essential step is the means to process the events. This can be done with a spreadsheet,[19,20] particularly with liberal use of macros, but this is not very efficient, in part because it tends to lead to the use of fixed-time intervals – squandering one of the advantages of DES. The most efficient method is to use a general-purpose programming language, such as Fortran.[21-24] In this way, each component of the simulation – including the event processing but also the timing routine, entity definitions and statistical counters – is coded as a subroutine, and these are called only when necessary. While the complex logic of a realistic model can be represented very compactly by this means, specialised programming skills are required and the resulting program does not provide a very transparent view of the model.

Fortunately, there are now several high-level simulation software packages available.[25-27] These allow the analyst to lay out the logic, and specify the attributes, variables and input values. The software takes care of the event processing, the simulation clock, the accumulation of outcomes, and all other 'housekeeping'. Depending on the simulation software, the layout and features may correspond very closely to the manner in which healthcare models are conceptualised and implemented. Indeed, the diagram in figure 1 is directly from one of these software packages (Arena®).[28] In that software, all the modules can be 'opened' (by double-clicking) to reveal all the equations or values that control their action.

Once the model is coded in software and debugged, the analysis in a DES proceeds by specifying the initial system conditions (i.e. starting values for all attributes and variables) and simulation settings (e.g. duration, time units, number of replications). The software then carries out the simulation by applying the logic to each entity (patient) using random numbers to obtain specific values from distributions and to determine whether probabilistic events occur at a given time to a given patient. Thus, a DES is, by definition, an individual patient, stochastic simulation. With currently available software and an up-to-date personal computer, full simulations covering tens of thousands of patients take only seconds. e.g., the simulation implementing the economic analysis of bipolar mood disorder briefly described here completes 50 replications, each consisting of 20 000 patients, in 2.3 seconds. Much more complex simulations (e.g. vaccination of all adolescents in the US over a century) can take several hours to run.

The last step in the implementation is to analyse the results of multiple runs of the model in order to infer the effect of whatever intervention was being studied, conditional on other inputs and the assumptions made. Although this step is very rudimentary in healthcare evaluations today – the results of a run are commonly reported along with some 'sensitivity analyses' – there are well-established methods for both terminating (scenarios where the model is run to a pre-specified time or condition)[29,30] and steady-state (when the model is run until the system stabilises with no further changes in major parame-

ters)[31] analyses of the simulation output. Indeed, some simulation software packages include an application for doing just that.[18]

4. Discussion

There is increasing recognition that modelling is essential to a pharmacoeconomic analysis.[32] Models allow more extensive understanding of the impact of a disease and its treatment than is possible with a detailed description of the experience of any particular cohort of patients. Although decision makers may remain somewhat sceptical of models,[2,33,34] there is no other way to provide realistic estimates of the consequences of a therapy in the actual circumstances encountered in a specific setting.[7] Perhaps because it was already well entrenched in medical decision-making, decision analysis was adopted by early pharmacoeconomists as the technique for structuring the economic models.[35,36] Despite recognition of the major limitations of this technique, even the move to Markov models has been kept largely within the same confines – many Markov models are still laid out using the familiar decision-tree structures, though this is completely unnecessary.[37] DES is a less limiting technique than Markov models and decision trees. Although it has a long history in operations research,[38] it has, to date, been very rarely employed to assess the value of healthcare interventions.[39]

The lack of uptake of DES is a shame, because this method imposes very few restrictions on the analyst. There is no need to force the course of disease into mutually exclusive states nor to cycle at fixed artificial intervals. Any aspect that conveys value or accrues cost can do so, even if it is a 'transition', and these quantities accumulate efficiently as they occur, rather than requiring 'folding forward and backward'. All of the necessary components are directly and explicitly part of the method. There are no key elements left unspecified; e.g., patient characteristics are not a definite part of either decision trees or Markov models. A vital function, the handling of time, is clear and as accurate as needed in DES.

Although largely unpublished, many 'workarounds' have been developed to mitigate the problems with decision trees and Markov models. Thus, software programs now commonly provide tables or other structures for specifying patient characteristics, and enable the use of Monte Carlo techniques to allow drawing from distributions rather than using single values and the modelling of individuals instead of cohorts. However, these improvements do not remove the fundamental limitations of Markov models and decision trees, and even at the most practical level, why use an inadequate technique when a better, more complete one is available? It's like choosing the kitchen knife to tighten a screw instead of the screwdriver.

A DES is very flexible. Multiple perspectives can be accommodated simultaneously with no special procedures. The costs, accrued as entities, pass through the simulation, can be accumulated separately according to the type of resources consumed and then aggregated as appropriate for each perspective. Although the information requirements are the same, there is no need to develop separate, perspective-specific cost estimates for each Markov state or portion of a branch in a decision tree. The flexibility of a DES is particularly evident when it comes to addressing specific, real settings. Though this level of realism might be viewed as unnecessary for cost-effectiveness analysis, it is essential for proper budget-impact analysis.[2] To address the impact of a new treatment on a particular budget, the model structure must be adapted to reflect the actual management processes in the setting at issue, and the real arrival and attrition of patients must be incorporated in the assessment. This is readily done in a DES but very difficult with the other methods.

Given the many large sources of uncertainty in evaluations of healthcare interventions, sensitivity analysis is an indispensable step. All modelling techniques accommodate sensitivity analysis, but the ease with which it is done may determine how extensively it is used in practice. Recommendations not withstanding,[1,2] today's sensitivity analyses are largely restricted to changing a few input values through limited ranges.[40] This may not be due to

the analysts' shortcomings but rather to the barriers imposed by the methods chosen. Computational inefficiency and partitioning of the inputs, so that some are in the structure (e.g. probabilities) and others are not (e.g. patient characteristics), makes proper extensive sensitivity analyses very arduous.

In a DES, full probabilistic analysis happens as a matter of course. All inputs can be defined using distributions of any shape. This allows the analyst to incorporate both real variation (e.g. in the distribution of age) and uncertainty (e.g. in the effect of treatment). Moreover, by running the simulation many times, the results can incorporate additional variation due to the sampling process and the input distributions can, themselves, be allowed to vary (e.g. the cost of a hospitalisation might change from region to region). This allows the analyst to summarise variability in the results using confidence intervals (if appropriate) and to provide cost-effectiveness acceptability curves. With modern computing and software, hundreds of simulations can be carried out in a reasonable time to fully incorporate variability.

The ultimate goal of also varying the assumptions that define the model structure itself is even further out of reach for Markov models and decision-trees. The ability to do this structural sensitivity analysis is part of a DES;[41] scenarios reflecting different structural assumptions can be directly analysed, avoiding the false sense of security that is provided by an incomplete assessment. This is done by incorporating alternative structural paths in the simulation and controlling these with 'gatekeeper' variables. Although running time increases somewhat if many scenarios are analysed, the times typically remain in the range of minutes.

One of the most frequent complaints voiced by those asked to accept the results of a model is that it is a 'black box'.[33] As inputs are swallowed in an impenetrable thicket of calculations, the reviewer is understandably wary of the outputs. Thus, a major desirable feature of a modelling method is the clarity with which it can be presented to others.[42] While this is said to be one of the advantages of decision trees and Markov models, this clarity is often evident only in the extremely simplified examples chosen for didactic purposes – most real models have hundreds, if not thousands, of branches or multiple states. Although the software code for running a DES requires specialised skill, the high-level simulators such as Arena® present the model very transparently. The method for diagramming the structure is a straightforward flowchart that lays out the possible course of care (another method, event graphs,[43] is common in operations research). These flowcharts use the same elements already familiar to clinicians and other healthcare personnel for delineating disease processes. There are no components (e.g. 'tunnel states' or 'bindings on recurrent trees') introduced purely for the sake of accommodating the limitations of the modelling method. Each component reflects a real part of the problem and its purpose is apparent, with all its functions obvious and comprehensively at hand. In a DES, all the components – both structural and computational – are included directly in the model. The flow of patients is represented using well-established flow symbols and the calculations are displayed at the point they occur (the format depends on the software but most link the calculations directly to the place in the flow where they operate). This is not to say that DESs are necessarily simple – complex problems may require intricate model structures – but they can be transparently presented.

The main disadvantage of DES stems directly from its advantages; it facilitates realistic modelling to such an extent that it may promote more in-depth depiction of the problem than is warranted,[44] particularly given the available data. Detailing the use of every needle and bandage consumed in a hospital stay may be just as inappropriate as overly superficial models. Thus, the analyst must guard against incorporating details of the patient's management simply because it can be done. This can increase the information needs beyond what is reasonably available. The decision about what to model and at what depth must be independent of the method used and should reflect the needs of the problem. Various papers that provide guidance on choosing a modelling technique have been published.[45,46]

Concerns have been raised[47] about the possibility that opting for a DES will force the analyst to relinquish implementation of the model, and thus control, to a technical specialist versed in the technique's intricacies. Fortunately, this is no longer the case. Existing software makes it possible for a health economic analyst to carry out all the modelling steps with no need for a simulation specialist. The skills already acquired for decision tree and Markov modelling are readily transferred to DES.

Other important topics pertinent to DES are not covered in this introductory presentation, but the literature from other fields has extensively addressed them. Subjects such as variance reduction,[48] steady-state analysis[31] and incorporation of continuous instead of discrete time elements,[25] hold the potential to greatly improve the application of modelling to pharmacoeconomic problems.

5. Conclusion

DES provides a means to represent healthcare processes realistically and, in turn, to better address questions posed in economic evaluations without forcing the analyst to accept unnecessary compromises. It is hoped that as analysts become more familiar with DES, its advantages will become evident and lead to widespread adoption.

Acknowledgements

This work was supported, in part, by a grant from Astra Zeneca AB. The author has no conflicts of interest directly relevant to the contents of this manuscript.

References

1. Weinstein MC, O'Brien B, Hornberger J, et al. Principles of good practice for decision analytic modeling in health care evaluation: report of the ISPOR Task Force on good research practices: modeling studies. Value Health 2003; 6: 9-17
2. Academy of Managed Care Pharmacy. Format for formulary submissions: version 2.0. AMCP, 2002 [online]. Available from URL: http://www.amcp.org/data/nav_content/formatv 20%2Epdf [Accessed 2002 Oct]
3. Greenberg PE, Arcelus A, Birnbaum HG, et al. Pharmacoeconomics and health policy. Current applications and prospects for the future. Pharmacoeconomics 1999; 16: 425-32
4. Chang K Nash D. The role of pharmacoeconomic evaluations in disease management. Pharmacoeconomics 1998; 14: 11-7
5. Weinstein MC, Toy EL, Sandberg EA, et al. Modeling for health care and other policy decisions: uses, roles and validity. Value Health 2001; 4: 348-61
6. Baltussen R, Leidl R, Ament A. Real world designs in economic evaluation: bridging the gap between clinical research and policy-making. Pharmacoeconomics 1999; 16: 449-58
7. Caro JJ. Disease simulation models and health care decisions. CMAJ 2000; 162: 1001-2
8. Akehurst R, Anderson P, Brazier J, et al. Decision analytic modeling in the economic evaluation of health technologies. Pharmacoeconomics 2000; 17: 443-4
9. Sonnenberg FA, Beck JR. Markov models in medical decision making. Med Decis Making 1993; 13: 322-38
10. Briggs A, Sculpher M. An introduction to Markov modelling for economic evaluation. Pharmacoeconomics 1998; 13: 397-409
11. Beck JR, Pauker SG. The Markov process in medical prognosis. Med Decis Making 1983; 3: 419-58
12. Eisenberg JM. Why a journal of pharmacoeconomics? Pharmacoeconomics 1992; 1: 2-4
13. Freund DA, Dittus RS. Principles of pharmacoeconomic analysis of drug therapy. Pharmacoeconomics 1992; 1: 20-31
14. Strakowski SM, DelBello MP, Adler CM. Comparative efficacy and tolerability of drug treatments for bipolar disorder. CNS Drugs 2001; 15: 701-18
15. Young RC, Biggs JT, Ziegler VE, et al. A rating scale for mania: reliability, validity, and sensitivity. Br J Psychiatry 1978; 133: 429-35
16. Banks J, Carson JS, Nelson BL. discrete event system simulation. Englewood Cliffs: Prentice-Hall, 1996
17. Jonasson O. Waiting in line: should selected patients ever be moved up? Transpl Proc 1989; 21: 3390-4
18. Law AM, Kelton WD. Simulation modeling and analysis. Boston (MA): McGraw-Hill, 2000
19. Caro JJ, Huybrechts KF, Klittich WS, et al. for the CORE Study Group. Allocating funds to prevention of cardiovascular disease in light of the NCEP ATPIII guidelines. Am J Managed Care 2003; 9: 477-89
20. Caro JJ, O JA, Klittich WS, et al. The economic impact of warfarin prophylaxis in non-valvular atrial fibrillation. Dis Manag Clin Outcomes 1997; 1: 1-7
21. Caro JJ, Salas M, O'Brien JA, et al. Modeling the efficiency of reaching a target intermediate endpoint: a case study in type 2 diabetes in the US. Value Health 2004; 7: 13-21
22. Caro JJ, Huybrechts K. Stroke Treatment Economic Model (STEM): predicting long-term costs from functional status. Stroke 1999; 30: 2574-9
23. Caro JJ, Ward A, O'Brien J. Lifetime costs of complications resulting from type 2 diabetes in the US. Diabetes Care 2002; 25: 476-81
24. Caro JJ, O'Brien JA, Migliaccio-Walle K, et al. Economic analysis of initial HIV treatment: efavirenz versus indinavir. Pharmacoeconomics 2001; 19: 95-104
25. Pegden CD, Shannon RE, Sadowski RP. Introduction to simulation using siman. Boston (MA): McGraw-Hill, 1990
26. Davies HTO, Davies R. Simulating health systems: modeling problems and software solutions. Eur J Oper Res 1995; 87: 35-44
27. Bowden RO. The spectrum of simulation software. IIE Solutions 1998; 30: 44-6
28. Kelton WD, Sadowski RP, Sadowski DA. Simulation with ARENA. Boston (MA): McGraw-Hill, 1998
29. Law AM. Statistical analysis of the output data from terminating simulations. Naval Res Logist Quart 1980; 27: 131-43
30. Alexopoulos C, Seila AF. Output data analysis. In: Banks J, editor. Handbook of simulation. New York: John Wiley, 1998

31. Pawlikowski K. Steady-state estimation of queuing processes: a survey of problems and solutions. Commun Assoc Comput Mach 1990; 22: 123-70

32. Siegel JE, Torrance GW, Russell LB, et al. Guidelines for pharmacoeconomic studies: recommendations from the panel on cost effectiveness in health and medicine: panel on cost effectiveness in health and medicine. Pharmacoeconomics 1997; 11: 159-68

33. Hatoum HT, Kong SX. How much faith can we have in pharmacoeconomic analyses? Pharmacoeconomics 1994; 6: 584-6

34. Sheldon TA. Problems of using modeling in the economic evaluation of health care. Health Econ 1996; 5: 1-11

35. Weinstein MC, Stason WB. Foundations of cost-effectiveness analysis for health and medical practices. N Engl J Med 1977; 296: 716-21

36. Glick H, Kinosian B, Shulman K. Decision analytic modeling: some uses in the evaluation of new pharmaceuticals. Drug Inf J 1994; 28: 691-707

37. Hazen G. Stochastic trees: a new technique for temporal medical decision modeling. Med Decis Making 1992; 12: 163-78

38. Schmidt JW, Taylor RE. Simulation and analysis of industrial systems. Homewood: Richard D Irwin, 1970

39. Jun JB, Jacobson SH, Swisher JR. Application of discrete event simulation in health care and clinics: a survey. J Oper Res Soc 1999; 50: 109-23

40. Agro KE, Bradley CA, Mittmann N, et al. Sensitivity analysis in health economic and pharmacoeconomic studies: an appraisal of the literature. Pharmacoeconomics 1997; 11: 75-88

41. Fu MC. Optimization via simulation: a review. Ann Operations Res 1994; 53: 199-247

42. Drummond MF, Jefferson TO. Guidelines for authors and peer reviewers of economic submissions to the BMJ. BMJ 1996; 313: 275-83

43. Schruben LW. Simulation modeling with event graphs. Commun Assoc Comput Mach 1983; 26: 957-63

44. Law AM. Simulation model's level of detail determines effectiveness. Ind Eng 1991; 23: 16-8

45. Karnon J. Alternative decision modelling techniques for the evaluation of health care technologies: Markov processes versus discrete event simulation. Health Econ 2003; 12 (10): 837-48

46. Barton P, Bryan S, Robinson S. Modelling in the economic evaluation of health care: selecting the appropriate approach. J Health Serv Res Policy 2004; 9 (2): 110-8

47. Karnon J, Brown J. Selecting a decision model for economic evaluation: a case study and review. Health Care Manag Sci 1998; 1: 133-40

48. Avramidis AN, Wilson JR. Integrated variance reduction strategies for simulation. Oper Res 1996; 44: 327-46

Correspondence and offprints: Dr *J. Jaime Caro*, Caro Research Institute, 336 Baker Avenue, Concord, MA 01742, USA.
E-mail: jcaro@caroresearch.com