SPM Background and Documentation

Need for modeling
Stroke is the third leading cause of death in the US and the leading cause of disability among adults. In recent years several interventions have been demonstrated to be effective in preventing and treating stroke and a host of new therapies are in development. Decisions about the use of treatments for stroke prevention and treatment are difficult for several reasons. First, there are a large variety of options to consider. Second, the overall value of many options is not always apparent; in addition to any health benefits, many stroke-related treatments have risks and some are extremely costly. Third, decisions about the appropriate use of treatments for prevention and treatment of stroke are complicated by the intricate epidemiology of stroke, in which there are multiple competing risks to consider. Finally, decisions about stroke-related treatments involve a variety of stakeholders, each with their own set of information needs and priorities. Stakeholders include patients, physicians, other providers of stroke services, practice organizations, payers, regulatory authorities, and manufacturers.

Development of the PORT model
To facilitate the evaluation of stroke prevention and treatment strategies, the Stroke Prevention Patient Outcomes Research Team (Stroke PORT) developed a simulation model of stroke development and outcome. This model is written in S+ (the Stroke Policy Model, SPM-S), and is included (both code and documentation) in the Final Report of the Stroke PORT to the Agency for Health Care Policy and Research (now the Agency for Healthcare Research and Quality, AHRQ). The model has been peer-reviewed, and publications addressing the model's r ationale, basic structure, and application are listed among the references. The model has also passed a number of internal validity checks, and a manuscript describing our comprehensive approach to validation is in progress.
Developing the SPM-S served to satisfy a number of objectives. First, the model was successfully used as an organizing principle for the PORT. For example, our basic administrative goal was to give preference to projects and analyses which would serve to better elucidate clinical and policy issues involving the epidemiology, prevention, and treatment of stroke. Since the SPM-S summarizes much of this information, it was used to help operationalize our decisions about priorities -- in particular, analyses which provided, improved, or updated one or more of the parameters of the model were given priority over those which did not.
A second objective was to summarize what is currently known about the epidemiology, prevention, and treatment of stroke in a concise format that accommodates updating. We believe that the focus on updating addresses a concern of funders, namely that regardless of whether the work product is an evidence report, a practice guideline, or a summary of an entire topic area (as provided by a PORT), a considerable amount of time, money and expertise will have been invested, and it is in the everyone's interest to see that the product does not become quickly outdated. An advantage of a model is that it facilitates easy updating - in particular, as additional work is done in the topic area, these efforts generate new parameter estimates, which replace the parameter estimates used in the previous version of the model. Accomplishing the objective does, however, require that at least a minimal infrastructure remains in place to support this updating.

A third objective was to advance the state-of-the-science in decision and cost-effectiveness analysis of chronic disease. In particular, recent advances in computer technology now allow us to build highly complex models which can potentially represent much more of the detail involved with clinical decision-making than was previously possible. For example, models can now include more downstream outcomes, more health states, more intricate (e.g., multi-step) prevention and treatment strategies, be tailored to consider the implications of many more patient characteristics, and so forth. In addition, computation-intensive approaches such as bootstrapping and other resampling techniques can now facilitate a more extensive assessment of the precision of the conclusions that can be drawn from a model. However, most decision models do not make full use of current computing power, and in fact many obstacles - both technical and methodological - must be successfully surmounted before today's computing power can be effectively harnessed. In developing the SPM-S, our goals included identifying the barriers to successfully implementing a very detail-intensive simulation model, determining how these barriers can potentially be overcome, and sharing any general insights generated by this exercise with the decision modeling community as a whole.

A fourth objective, closely related to the third, was to advance the state-of-the-science for decision modeling as it specifically pertains to stroke. We believe that we have succeeded in doing so, and that the SPM-S is now a reference standard model in this field.
A final objective was to determine how best to present modeling to decision-makers. Another focus of funders involves how the results of the academic studies they support can be used to actually improve clinical practice. We have been particularly sensitized to these issues when presenting the published results of our analysis of acute stroke treatments. This analysis showed that when taking the traditional societal perspective that considers all direct medical and non-medical costs and following patients until death, acute stroke treatments which succeed in making only modest shifts in the pattern of disability after stroke would nevertheless have a large public health benefit, even if these treatments were quite costly. But hospitals, managed care organizations, and other decision makers are not typically responsible for all components of cost, and their frame of reference is much shorter than the lifetime of the patient. We quickly learned that in order to be credible to decision-makers (and thus have the opportunity to have an effect on actual practice), decision models must (1) allow tailoring of inputs to meet local needs; (2) allow tailoring of outputs (e.g., to consider some rather than all components of cost); and (3) must be meticulously documented, so that users can have confidence in the model's basic structure, algorithms, and input parameters. Our subsequent work in this regard is discussed below.

Subsequent modeling efforts

More recently, we have developed a new simulation model based on many of the principles and much of data from the SPM-S. This model is programmed in C++, and is denoted as the SPM-C. The main advantages of the SPM-C are (1) it is programmed in a more flexible and widely-used language than S+; (2) a much faster and thus more powerful simulation is possible using C++; (3) the structure of C++ makes it possible to develop an "intervention language" that allows a variety of interventions to be easily examined without significant reprogramming; (4) the SPM-C can easily support extensive on-line documentation; and (5) the SPM-C is suitable for use as a web-based application.

Current status

The SPM-S has been developed, validated for the analysis of the impact of acute stroke treatments, and comprehensively documented in the Final Report of the Stroke PORT. The SPM-C has been developed, and the current cycle of validation should be completed in May 1999. (The current cycle of validation checks the operation of the SPM-C's software and algorithms - the SPM-C uses essentially the same parameter estimates as the SPM-S and thus the validation of its parameters need not be repeated.) The SPM-C includes the prototype of a user-friendly interface.

Infrastructure -- basic

The SPM is being developed as a public use application, with administration provided by the Stroke Policy Program at the Center for Clinical Health Policy Research at Duke University. Infrastructure resources (e.g., project coordinator, librarian, programmer, website manager) are housed within the Center (or subcontracted for through the Center). The Center is currently soliciting funds from government, foundations, and industry in order to support this infrastructure. Researchers, teachers and public health professionals will be encouraged to use the SPM for academic and other non-commercial applications.

Top of Page

Next steps: model development

This section briefly summarizes the overall plan for model development. Additional details for accomplishing some steps are presented separately.

Improve input parameters

Transition functions (i.e., time-dependent event rates) are the core of the information describing the epidemiology of stroke. Our current transition estimates are the result of making a request to the investigators of the Framingham Study to reanalyze their data in a way which is consistent with the SPM (i.e., the main difference between these analyses and those previously published by the Framingham in vestigators is that we censored patients at the first model-relevant event). These analyses were quite useful, but with access to the raw data we could have obtained even more information (in particular, we could have used bootstrapping to obtain state-of-the-science estimates of the precision with which the transition functions were estimated). The Framingham Study has a mechanism by which investigators can submit a grant to obtain use of archived versions of their raw data (i.e., these data are not current, but are complete to within a few years of the present, and would be sufficient for these purposes). We propose to write such a grant, and re-estimate the relevant transition functions. In the process, we also propose to address various statistical issues pertinent to estimating transition functions in the presence of competing risks.

Another area in which the model's input parameters can be improved significantly is in its cost data, which are primarily based upon utilization observed in United States Medicare files during 1991-1993. Some salient areas for additional data collection pertain to (1) updating the Medicare utilization data; (2) collecting additional data on the relationship between stroke disability and daily costs (this relationship being both crucial to the acute stroke analysis and relatively little studied in the literature); and (3) obtaining international utilization data (i.e., thus allowing the model to be extended to non-US applications).

Improve functionality

There are four natural areas for improving the functionality of the model. First, in large part because the SPM was originally developed for analysis of stroke prevention strategies, the model contains relatively little detail about events near the time of the stroke. For example, the model simply counts costs during the initial hospitalization, during the first 30 days, and so forth, where these costs are based upon aggregate patterns of utilization (among Medicare beneficiaries during 1991-1993). The model does not currently support comparison of different rehabilitation strategies, different models of care for acute stroke management, and so forth. However, these improvements have been anticipated in the current C++ code, and thus adding this functionality, while detail-intensive, should not require a significant change in the SPM-C's basic algorithms.

A second area for improvement is in accommodating time-dependent covariates. For example, in actual practice patients develop hypertension over time, the level of carotid stenosis often increases over time (thus suggesting that screening should not necessarily take place once and for all, but might instead be performed at regular intervals), and so forth, and it would be helpful for the model to be able to reflect these underlying stochastic processes. These improvements have also been anticipated in the current C++ code.

A third area for improvement involves replacing analysis of a single patient-type (e.g., 65 year old men with atrial fibrillation and a history of transient ischemic attack) with analysis of entire patient populations - the goal being to answer the public health question of what would be the impact of implementing an intervention on a specific population. The programming changes required to accomplish this task should be straightforward, and should essentially amount to adding a patient sampling loop at the beginning of each iteration of the simulation as well as a mechanism for specifying the characteristics of the population under consideration.

A final area for improvement involves the development of an intervention language - that is, a mechanism for specifying the components of the str oke prevention / treatment strategies under consideration in such a way which does not require additional programming for each intervention. This application requires sophisticated software development, and any general algorithms developed would have great potential to advance the field of decision modeling in general.

Improve documentation

The basic insight is that in order to be credible models must be well documented, and that on-line documentation (whether as part of a stand-alone program or through a website) allows the content matter to be very extensive yet presented at a level of detail, which is tailored to the needs of users. We will determine what form of on-line documentation (e.g., data elements, level of detail, format such as pull-down menus) is desired by users, and then to implement that documentation. Our efforts would involve focus-group and other interviews of decision-makers to determine what information is required, and then reprogramming (performed by subcontractors) of our existing user interface to implement the recommendations.

Improve interface / develop website

As a closely related application, we want to improve the general user-friendliness of the

SPM interface. The application will eventually be placed on the Internet in two formats. The first format would primarily be used as a test of the interface, and the goals would be to communicate public health messages regarding stroke (e.g., its epidemiology, prevention, and treatment) in as efficient and user-friendly fashion as possible. In particular, we would make a large but finite number of model runs, each based upon a clinical- or policy-relevant scenario, and store the results of these runs within a table. The user would not actually run the SPM, but would be directed to select the analysis of interest, with the results immediately extracted from the table. The website would also note that a more extensive version of the SPM is available, and would contain a hypertext link to its address.
The second format would be intended to serve the needs of more experienced modelers. Here, users would have considerably more flexibility regarding the specification of inputs and outputs, and the model would be run in real time. As discussed above, this version of the model would be made available without charge for research, teaching and public health, and could be licensed for commercial purposes.

Continue validation
Any efforts at improving the code (e.g., intervention language) will require ongoing efforts at validation, in order to ensure that the code modifications operate as intended.
Assess benefits to users
This includes tracking the users and uses of the model. We would perform various surveys, focus group interviews and so forth in order to assess current patterns of use of the model and how the model might be improved to better meet user needs. A substantial part of this assessment can also be built into the user interface.