Krams, M., Lees, K.R. and Berry, D.A. (2005) *The past is the future: innovative designs in acute stroke therapy trials.*<u>Stroke</u>, 36 (6). pp. 1341-1347. ISSN 0039-2499

http://eprints.gla.ac.uk/21686/

Deposited on: 23 January 2012

The Past Is the Future : Innovative Designs in Acute Stroke Therapy Trials

Michael Krams, Kennedy R. Lees and Donald A. Berry

The online version of this article, along with updated information and services, is
located on the World Wide Web at:

http://stroke.ahajournals.org/content/36/6/1341

The Past Is the Future
Innovative Designs in Acute Stroke Therapy Trials

Michael Krams, MD; Kennedy R. Lees, MD; Donald A. Berry, PhD

More than 74 000[1] patients with acute ischemic stroke have been randomized into clinical trials over the past 35 years to investigate new therapies.[1] Only one treatment, thrombolysis with recombinant tissue plasminogen activator, has emerged from these investigations.[2] Efforts to establish acute neuroprotectant therapies have yet to succeed.[3,4]

Have we squandered our resources? Has methodological rigidity delayed development of a new treatment or prolonged investigation of an ineffective therapy? Here we present a flexible and more efficient approach to clinical trial design and analysis. We have the potential to improve the use of scarce patient resources and to accelerate development of promising agents.

In medical practice, we respond to a patient if a dosage seems inadequate by either changing the dosage or switching to another medication. We cautiously change treatment after reviewing new evidence: side effects, intractable symptoms, and poor adherence. We might express our estimate of how much the change may improve the patient's condition in terms of probability. We repeat this process every time we update the treatment plan in light of new important information. Why not take the same approach to clinical trials?

The proposed approach to the design and conduct of clinical trials uses Bayesian methods that make careful use of high-quality available past (prior) evidence to refine the inference from accumulating evidence in the ongoing clinical trial. This approach may: (1) enhance investigation of single agents or combination therapies; (2) make earlier and more reliable choices of dose for use in pivotal trials; (3) accelerate the progression from phase II into phase III trials all the way to a potentially seamless switch; and (4) treat trial participants more effectively by adaptively allocating more resources to therapies that are performing well while reducing support for less promising treatment arms.

Concepts

The basis of all statistical inference is probability. The frequentist approach to inference deals with probabilities of data for given hypotheses or particular values of unknown parameters. Bayesian probabilities apply as well to hypotheses and parameters themselves.[5] The difference is critical. The Bayesian approach is tailored to learning on the basis of evidence. Bayesian probabilities can be calculated at any time and can be updated continually as information becomes available.

A consequence of the Bayesian approach is the ability to calculate probabilities of the results of future observations given the current uncertainty in the parameters. For example, predictive probabilities allow for addressing whether and which observation to take next. This ability is fundamental in designing experiments.

The Bayesian approach is tailored to making decisions. Designing a clinical trial is a decision problem. Optimal designs are those that maximize gain. Gain or loss depends on the goals of the designer. For example, the goal may be to deliver a good medicine.

Bayesian designs can be arbitrarily complicated. However, with computer power available today, even very complicated designs can be simulated many thousands of times. This allows for evaluating the design's false-positive rate and other operating characteristics that are usually viewed as being frequentist measures. The design might be modified to have operating characteristics that are acceptable to regulatory or funding agencies. In a sense, this strategy is using the Bayesian approach as a tool for building a good frequentist design.

Bayesian Applications

Bayesian methods and decision theory are widely used in medicine and industry. A number of medical devices have been approved by the Food and Drug Administration on the basis of Bayesian experimental designs and analyses. Many phase I and II oncology trials have been designed and conducted at M. D. Anderson Cancer Center from the Bayesian perspective.[6,7] An example is a trial that uses
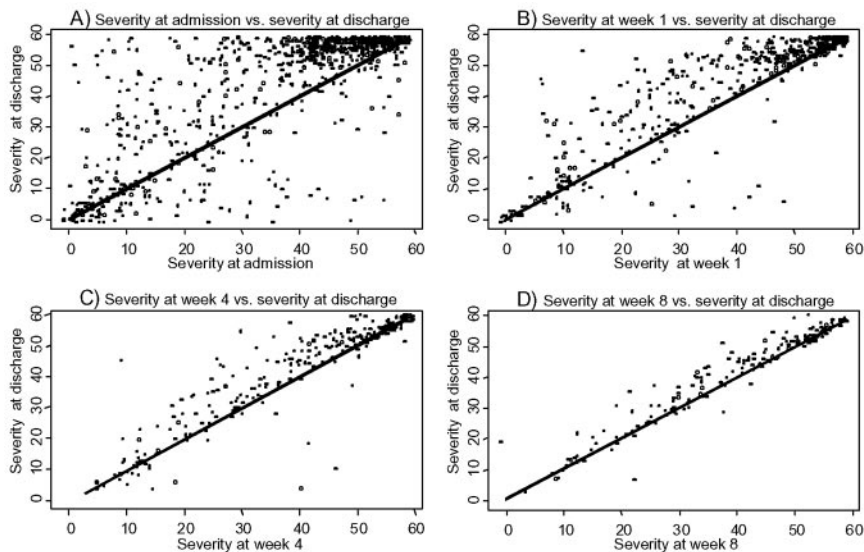
Figure 1. "Longitudinal model." Relationship between early response and final outcome in untreated acute stroke patients. Data from Copenhagen Stroke Study.[9] Stroke severity as measured on Scandinavian Stroke Scale (SSS) (0: comatose patient; 58: no gross neurological deficit). Each circle represents 1 patient. Y-axis: SSS at discharge from stroke unit. X-axis: SSS at earlier time points. Black line: linear regression through the data. A, SSS at admission does not correlate strongly with SSS at discharge. However, there is partial information even at baseline. Patients with mild strokes tend to do well (cloud in right upper corner), whereas patients with severe strokes tend to do poorly (left lower corner). B, SSS at week 1 vs discharge. C, Week 4 vs discharge. D, Week 8 vs discharge. Over time there is an increasingly strong relationship between the SSS score at an earlier time point and time of discharge.

adaptive randomization in which patients are more likely to be assigned treatments that are performing better in the trial.[8]

Historical Data

Our first step to designing a stroke study is to assess existing data and to model potential trial outcomes. Despite different treatment strategies prior information can be incorporated into a hierarchical model that will combine these existing data with the future results of the trial being planned.[6–9] The term "hierarchical" indicates how differences in design are expressed. The more similar designs tend to lie close to the top of hierarchical model. The model also contains terms that increase the variance and hence express the uncertainty among dissimilar studies.

The ASTIN trial provides a good example for the use of historical data.[10] To predict likely recovery profiles of stroke patients over a 90-day period, based on initial severity, under the assumption of "no experimental treatment," data from the Copenhagen Stroke Study[11] were used to model physiological recovery in untreated acute stroke patients (Figure 1).[12] Real data from the trial would gradually be introduced to update this "longitudinal model."

A wealth of high-quality data lies dormant that could be used to inform the design and conduct of future stroke trials. The Virtual International Stroke Trials Archive (VISTA) offers a mechanism for accessing valuable data sets and using them to benefit future patients: entire stroke trial data sets or records from placebo groups can be documented, securely stored, and, subject to approval by a committee of original investigators and sponsors, accessed for analysis. VISTA involves data from a wide range of countries, sites, and trials and reflects the natural history of patients recruited into stroke trials. Stroke trialists are invited to contribute to and utilize this resource (contact K.R.L., k.r.lees@clinmed.gla.ac.uk).

The Bayesian Approach

The Bayesian approach is one of continual learning. Instead of the current practice of leaving information accruing during a trial untouched in a sealed database as the trial progresses, we make immediate use of it. Our knowledge base is continuously updated, and aspects of the trial design such as allocation of patients to certain dosage levels are gradually modified. The accruing data changes our levels of uncertainty as expressed by probability estimates. A prior estimate of the probability of an uncertain event is updated to a posterior probability each time a new piece of information becomes available.[6,7,9] This continuous learning need not be transparent to investigators or anyone in a position of influence over the trial: the decision tree can be built in advance and "delegated" to a computer, which is closely supervised by an independent data monitoring committee, as in ASTIN.

Adaptive Designs

Insufficient understanding of the dose–response and inappropriate choice of doses taken into confirmatory studies plague drug development. Lubeluzole is one example of a neuroprotective drug development program that would have benefited from better understanding of the dose–response: plasma levels of lubeluzole achieved in the pivotal trials were lower than those necessary for neuroprotection in experimental models.[13] In retrospect, a different phase II trial design and incorporation of information from the experimental situation may have improved the choice of dose for pivotal trials or prevented considerable unnecessary expense and use of resources by terminating the program early.

We illustrate the benefits of modeling and using a Bayesian adaptive design for efficient learning about the dose–response[12] with ASTIN[10] as an example. Parallel group designs often test only a small number of treatment arms, comparing them against placebo. Suppose the objective is to identify the minimal dose yielding near-maximal efficacy (ED95). The appropriate dose can never be found more accurately than the distance between the doses studied. It helps to increase the number of doses: in ASTIN there were 16. However, a traditionally powered design with 16 arms would be enormous. Adaptive treatment allocation in a sequential design is more efficient: outcome data accrue in
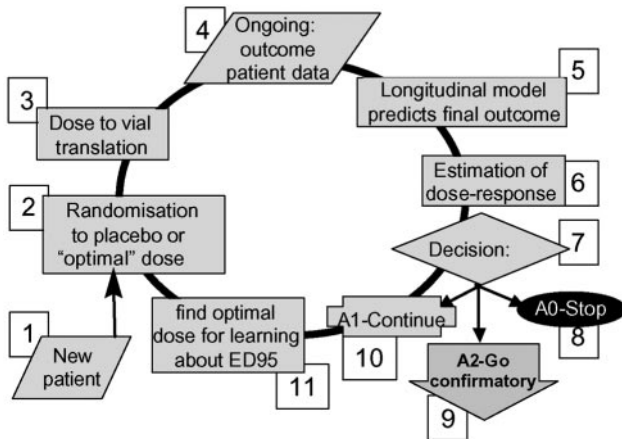
Figure 2. Design for response-adaptive learning about the dose–response with dynamic termination rule. Data on the primary end-point is captured in real time. The accruing data are continuously reassessed to help with 2 decision problems. (1) To which treatment should the next patient be allocated to optimize learning about the research question (dose–response, ED95)? (2) Should the study continue to accrue information or is there sufficient data to stop the trial either for futility or efficacy? The diagram follows the passage of 1 patient through the system. In reality, several patients will be progressing through the system at any one time and the database will contain data on commencing, ongoing, and complete patients. [1] that the patient presents for randomization (baseline data are wired into system); [2] Patient is randomized in blinded fashion to placebo or "optimal" dose to learn about research question; [3] Dose assigned is converted to particular vial numbers, allowing for blinded administration of study drug; [4] Patient's response data are entered into the system as they progress through the study; [5] Patient's final outcome is predicted using a longitudinal model (see Figure 1) (the prediction is substituted by the final response, as soon as it becomes available); [6] Based on the currently available data, the system updates the "estimate" of the dose response curve and its uncertainty; [7] Each day the algorithm implements a decision rule and recommends to either: [8] A0: stop the study because of futility (based on the posterior probability that the treatment has an effect smaller than a minimum clinically relevant size) or [9] A2: stop dose finding and moves to a large confirmatory study (based on the posterior probability that the treatment has an effect larger than some clinically relevant size); [10] A1: continue dose finding study (the recommendation of the system is reviewed by the IDMC, which incorporates clinical judgment and factors in safety issues); [11] The dose allocator chooses a dose from a list of possible doses that will optimize learning about the ED95 or some aspect of the dose–response curve. The database used to determine the dose is continually updated as outcome data from patients are gathered.

real time, the data are modeled to estimate the dose–response, and our decision as to which treatment to allocate to the next patient is conditional on the latest updated estimate of the dose–response (Figure 2). Patients will preferentially be allocated to informative treatment arms. The goal is to close in on the appropriate dosage level and then efficiently minimize the variance about a parameter of interest. In ASTIN we chose to minimize the variance around the point estimate of the treatment effect at the estimated ED95. In other words, we concentrated our effort around doses that seemed to produce near-maximal efficacy. We can explore a wide range of possible doses at the start of the trial without having to waste patients on treatment arms with low information value later in the trial.

Thanks to high-speed computing, adaptive treatment allocation is not limited to one-dimensional problems. Figure 3

illustrates simulated examples of learning about the dose–response (surface) for a single investigational drug and a combination of 2 drugs.

Termination Rules

In drug development, most projects fail. It would be preferable to stop failing clinical trials as soon as possible. When successful in finding a dose that provides clinically meaningful benefit, we would prefer a rapid transition from dose–response exploration to a confirmatory study. Although traditional designs permit a few interim analyses for futility, the Bayesian approach as applied in ASTIN feature continuous reassessment of the data, with a computer algorithm advising an Independent Data Monitoring Committee (IDMC) on whether to continue or stop the trial. At the start of the trial, there is great uncertainty around our estimates of dose–response and ED95, but as trial data accrue, this uncertainty shrinks. The stopping rule in ASTIN continuously asked the following questions: (1) Does our estimate of the dose–response suggest that there is $<10\%$ chance of success for any dose (success was defined as a >3-point recovery over and above placebo as measured by a stroke scale)? If so, then stop for futility. (2) For the best dose, is the response good enough to conclude that there is $>90\%$ chance of success? If so, then stop for efficacy and switch to a confirmatory trial, comparing the "best dose" against placebo.

We considered, but did not use, a more sophisticated decision–theoretic approach to stopping the trial.[12] Clinical practice and clinical research involve making decisions, eg, choosing sample size. It is impossible to precisely predict the consequence of a particular decision. But it is possible to associate a predictive probability to each possible result and its consequence. A numerical assignment to a consequence indicating the overall value of a consequence is called a "utility." Economists distinguish utility from dollar value because realistic values also depend on the usefulness of the consequence. The utility of any particular consequence of a clinical trial design should reflect its consequent impact on patients with the disease, including patients inside and outside the trial.[14] Say that we could define the value of a successful treatment to any one stroke patient to whom it would be deployed. A decision–theoretic stopping rule would ask: Where can each individual patient contribute maximal value, in the trial learning about the dose–response or in a confirmatory trial? Clearly, the traditional p-value seldom reflects utility, but rather serves only to provide a common standard. Ethically, the utility approach is appealing. Its focus is the overall set of patients with stroke, trying to maximize the value of each patient entering clinical research programs to optimize treatment for the overall population and the individual patients. For a more detailed discussion, see Lewis et al[15] and Cheng et al.[16]

Simulation-Guided Trial Design

Computer simulation of clinical trials can help to improve the final design and learn about its characteristics. A Bayesian approach requires an initial alignment of assumptions and agreement on which models to use. Early interactions with
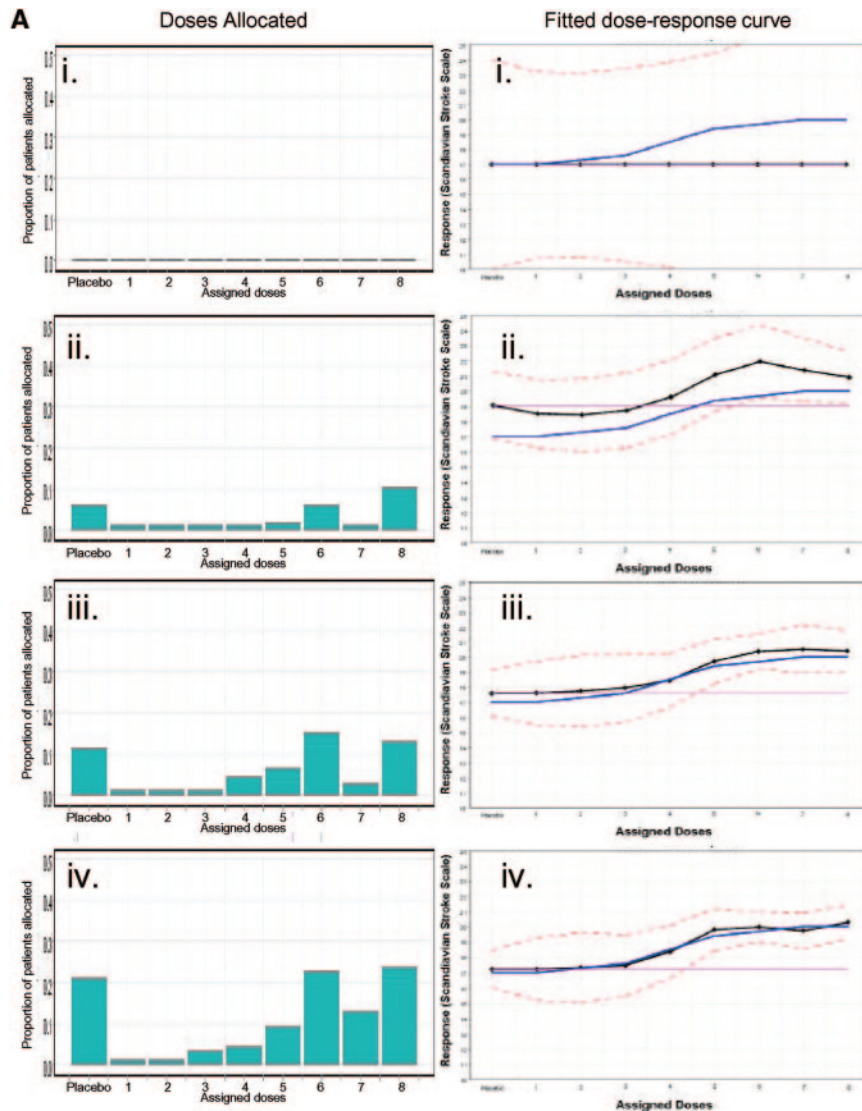
Figure 3. Response-adaptive treatment allocation—learning about the dose–response as the trial progresses. Two scenarios: (A) single drug in 8 doses; (B) two drugs in 3 and 4 doses, respectively. The 4 panels provide snapshots over time (i.: at start; iv.: at end of trial). Bar charts to the left illustrate adaptive treatment allocation to different treatment arms over time. Figures to the right illustrate estimated dose–response (surfaces) and learning over time. Response is a continuous variable, eg, change from baseline to day 90 on a stroke scale (higher value indicates improvement). A, Single drug scenario: we start with no assumptions, ie, with a flat prior for the estimate of the dose–response (black line). The blue line represents the true dose–response, which at the start of the trial is unknown to us. With data coming in from the trial, we will be able to more accurately estimate the true dose–response. Red lines: 95% credible intervals around estimated dose–response. B, Now we extend the methodology to two drugs. Think of "drug one" as a neuroprotectant with little information about its dose–response (flat prior, identifiable as a dotted black line parallel to the "drug one" dose axis) and-"drug two" as a thrombolytic with information about its dose–response (informative prior, identifiable as a curved dose—response, black solid line, above the dose axis for "drug two" on the right edge of the dose–response surface). Collectively, the combination of these assumed dose–responses form the prior dose–response surface (in color). In gray, the "true," but initially unknown, dose–response surface, which we want to identify. As we acquire more data, our estimate of the dose–response surface (in color) gets updated and approaches the true dose–response surface with increasing accuracy. Vertical bars at each treatment cell indicate 95% credible intervals around the mean estimated treatment effect.

relevant experts, including statisticians, clinicians, and regulators, can establish credible models and simulation can sort through possible scenarios to find the potentially best design. In developing and optimizing the design for ASTIN, hundreds of thousands of stroke trials were simulated. We confirmed that the design would perform to specification (eg, correctly adapt treatment allocation according to dose–response, learn about dose–response efficiently, choose the correct ED95, and stop early for the right reasons). The design was tested under extreme circumstances, including scenarios in which the true dose–response curve was flat, sigmoid, or up–down; different patient recruitment speeds; and different thresholds for futility and efficacy. Frequentist statistical characteristics such as type I and II error rates were adjusted and confirmed using Monte Carlo simulation. Even during the conduct of ASTIN, the IDMC undertook additional simulations to satisfy themselves that some of the responses they were seeing would be correctly handled by the computer algorithm.
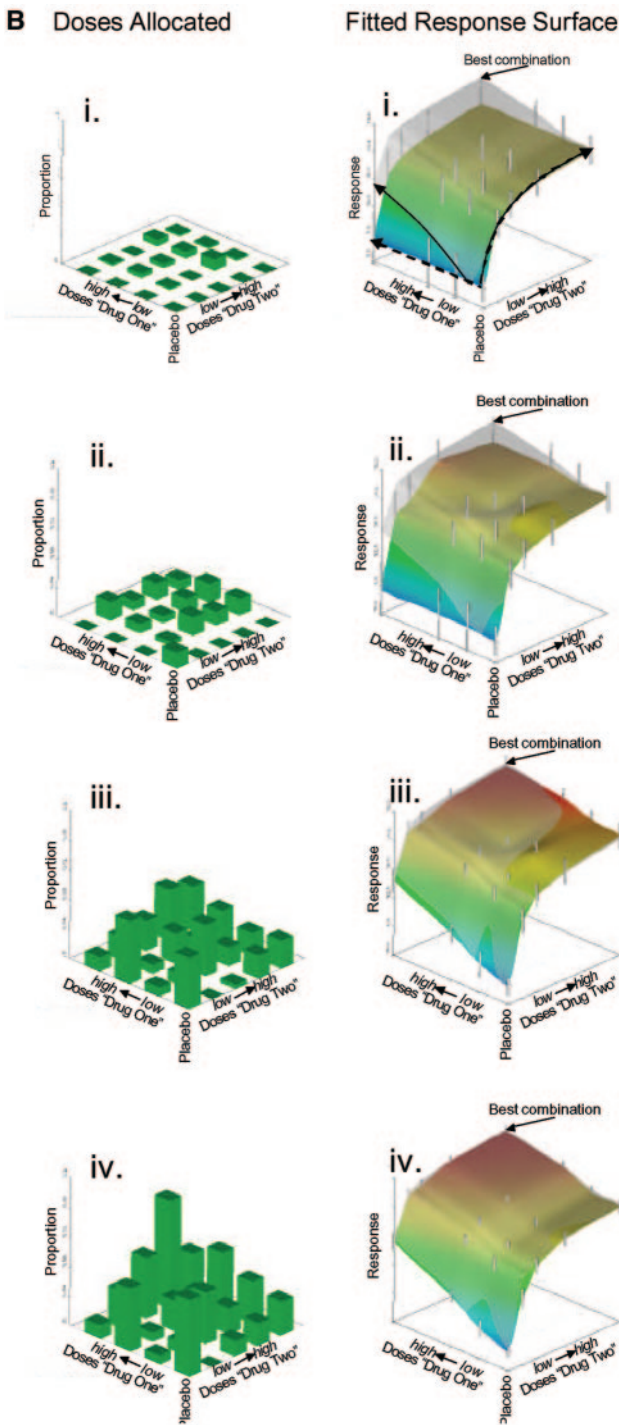
Figure 3. Continued

Seamless Designs

Traditionally, the time gap between a phase II dose–response finding study and a confirmatory phase III trial can be >1 year. However, in the absence of major issues raised during phase II, the transition from a learning phase to a confirmatory phase can occur seamlessly, with no pause in accrual. Investigators may not even appreciate the change; they would continue to get blinded dosing instructions.[17] There could be substantial savings in site set-up effort and opportunity costs.

With the creation of stroke trial networks, we may speculate to extend this idea further. Say that we could agree on the most suitable primary endpoint and other key characteristics for acute stroke trials. We then envisage conducting an ongoing experiment with no clear beginning or end. New therapeutic options are introduced as they mature from safety testing, and patients are allocated to whatever treatment promises maximal benefit to the overall stroke population. This might sound futuristic, but in a rudimentary form such designs are being implemented in oncology at M. D. Anderson Cancer Center.

Endpoints

Some clinical endpoints used in acute stroke trials (neurological stroke scales, measures of functional outcome such as modified Rankin score or Barthel Index) suffer from considerable observer dependent variability. Any trial design would benefit from endpoints with less variability. More meaningful endpoints may evaluate improvement on an individual patient basis rather than seeking a population response. The "responder" analysis used by the AbESTT investigators is an example. Here, the threshold for favorable outcome is adjusted according to initial stroke severity. Patients with National Institutes of Health Stroke Scale scores of 4 to 7 must achieve a final mRS of 0; patients with initial severity of 7 to 14 on National Institutes of Health Stroke Scale need to achieve an mRS of 0 to 1; and more severely affected patients are considered to have a favorable outcome with an mRS score of 0 to 2.[18] More sophisticated modeling of prognosis, for example by using data from VISTA, is conceivable. Gain functions integrating informative biomarkers (including imaging biochemical markers such as S100, etc) with centrally assessed clinical endpoints may enhance trial efficiency.[19] If there are concerns regarding regulatory acceptability of novel endpoints, we could discuss a strategy of using novel endpoints in the learning phase to inform adaptive treatment allocation but switch to more conventional endpoints for the final pivotal analysis.

Issues to Consider

Bayesian methods impose extra work. We need to assess and quantify available information and plan for extensive modeling and simulation. Clinical trialists, biometricians, and regulators must agree to take this approach. These experts need to assess the scientific credibility of the models and prior data.

The sequential design discussed relies on having at least some degree of exchangeability among patients after taking into consideration observed patient covariates. Bayesian methods can deal with some lack of homogeneity such as strong region and center effects, ie, a patient from a rural clinic in India may differ from a patient in New York. There may be time trends: nonpharmacological stroke therapy is improving, with wider introduction of acute stroke units and better management of risk factors. There may be accrual bias, with physicians becoming partially unblinded to trial results, eg, if a trial continues beyond the maximal sample size for the exploratory phase in a seamless design, it could be inferred by participating trialists that the project has not been stopped for

futility. This is not necessarily a disadvantage, provided that investigators cannot bias the treatment effect estimate, guaranteed through randomization and masking, there is no reason to conceal accumulating evidence of potential worth of the treatment undergoing study.

A badly formulated prior estimate can hinder a trial, because the experiment then needs to overcome the weight of incorrect data and assumptions. However, traditional trials suffer in the same way from poor models or poor estimates of standard error.

Well-chosen prior data and models for Bayesian designs tend to lead to small trials, but trials can be large, precisely when a large trial is necessary. In contrast, conventional trials without continuous scrutiny of the data may come to their predetermined end with an ambiguous conclusion.

Bayesian designs can make use of incoming data to inform future decisions and thereby reduce potential delay. The sooner a clinical endpoint reads out, the earlier it can impact future decisions. When the final assessment of treatment benefit occurs with some delay, such as in acute stroke trials, in which it is traditionally assessed 3 months after treatment start, longitudinal models can be help to predict final outcome using earlier readouts or biomarkers (Figure 1).

Mechanics

With the advances in computer technology, large-scale clinical trial simulations have become possible using sophisticated algorithms. However, innovation has a cost. It requires hard work, involving considerable upfront investment in establishing the software, running simulations, and fine-tuning the system for an optimal design to fulfill the user requirements of the study.

In performing the trial, it is key to have an IDMC of clinicians and statisticians knowledgeable in the specifics of the design and able to overlook the performance of the system as well as the usual concerns of IDMCs. Decisions regarding the treatment allocation must happen in real time. The IDMC reviews the performance of the system against predefined user requirements. The decision regarding stopping the trial requires IDMC endorsement: once the algorithm recommends stopping, the IDMC will review the relevant information. They may endorse the decision, but they may choose to override the algorithm's recommendation when there are strong grounds for doing so.

The mechanics of real-time data captured through fax, telephone, or Internet, have all been developed and require integration into the infrastructure of data management and trial logistics pertinent to large clinical trials.

Although it is particularly easy to administer a large number of different doses with intravenous compounds in a blinded fashion, it is also feasible to apply the principle to oral compounds, for instance, combining 2 tablets that are available at strengths of 0, $1\times$, $3\times$, and $4\times$ allows 9 equidistant doses. There are similar schemes to cover a dose range of 0 to 243 on a semi-logarithmic scale. Even where drug supplies are limited by cost or manufacturing logistics, it is possible to organize packaging in an efficient manner.

Early interactions with investigators, regulatory agencies, and other parties involved in the conduct of trials have been key to ensure the success of applying the Bayesian approach in ASTIN. Our interactions with regulatory agencies have been particularly rewarding. The Food and Drug Administration has recently staged workshops on Bayesian applications (http://www.prous.com/bayesian2004/), supporting learning about the approach and future applications where appropriate.

The Future

We have much to learn from the field of oncology. Bayesian methods promise a seamless research process in which preclinical data feed into clinical studies and in which phases I, II, and III blend together, bringing treatment advances to stroke patients as quickly and efficiently as possible. Although we are accustomed to considering the concepts of false-positives and false-negatives, we usually ignore the more common problem of a "false neutral": all the important research questions that have not been investigated, simply because a lack of efficiency of the scientific process. So rather than consuming yet another 74 000 patients with only marginal benefit to the overall population of stroke patients, we envisage an ongoing definitive stroke trial, which learns in real time about new treatments (or combination of treatments) with adaptive shifting of resources toward the most convincing therapeutic approach, within (dose–response finding) and across compounds. ASTIN has been a first step in this direction. Bolder approaches looking at combination therapies are currently being implemented and Figure 3B gives a glimpse on what an adaptive design for establishing the best combination of, for example, a fibrinolytic and a neuroprotectant might look like. Future designs may also include the option of modifying the trial design itself. Possible modifications include stopping early, changing entry criteria, expanding to additional sites, extending accrual beyond the trial's original sample size, or dropping or adding treatment arms.

The greatest need for innovation and the greatest room for improving drug development is effectively dealing with the enormous numbers of potential drugs that are available for development. The notion of developing drugs one at a time is part of the pharmaceutical culture, but this will change. Companies able to screen many drugs simultaneously and do so effectively will survive, and others will not. Drugs that are apparently more promising will move faster through the preclinical setting. Drugs that give disappointing data will languish.

There are two things we can do today. One is to share our work experience and raw data to allow model-based approaches to clinical drug development (through VISTA). The other is to be open-minded and willing to experiment with innovations available today and be willing to embrace those shown to be useful.[17]

References

1. Kidwell CS, Liebeskind DS, Starkman S, Saver JL. Trends in acute ischemic stroke trials through the 20th century. *Stroke*. 2001;32: 1349–1359.
2. National Institute of Neurological Disorders and Stroke rt-PA Stroke Study Group. Tissue plasminogen activator for acute ischemic stroke. *N Engl J Med.* 1995;333:1581–1587.
3. Grotta J. Neuroprotection is unlikely to be effective in humans using current trial designs. *Stroke*. 2002;33:306–307.
4. Lees KR. Neuroprotection is unlikely to be effective in humans using current trial designs: an opposing view. *Stroke*. 2002;33:308–309.
5. Berry DA. Clinical trials: is the Bayesian approach ready for primetime? Yes! *Stroke*. 2005; In press.
6. Berry DA. Statistical innovations in cancer research. In Holland J, Frei T, et al, eds. *Cancer Medicine,* 6th ed. London: BC Decker; 2003:465–478.
7. Berry DA. Bayesian statistics and the efficiency and ethics of clinical trials. *Stat Sci.* 2004;19:175–187.
8. Giles FJ, Kantarjian HM, Cortes JE, Garcia-Manero G, Verstovsek S, Faderl S, Thomas DA, Ferrajoli A, O'Brien S, Wathen JK, Xiao L-C, Berry DA, Estey EH. Adaptive randomized study of idarubicin and cytarabine versus troxacitabine and cytarabine versus troxacitabine and idarubicin in untreated patients 50 years or older with adverse karyotype acute myeloid leukemia. *J Clin Oncol.* 2003;21:1722–1727.
9. Berry DA. *Statistics: a Bayesian perspective.* Belmont, CA: Duxbury Press; 1996.
10. Krams M, Lees KR, Hacke W, Grieve AP, Orgogozo JM, Ford GA; for the ASTIN Study Investigators. ASTIN: an adaptive dose-response study of UK-279,276 in acute ischemic stroke. *Stroke.* 2003;34:2543–2548.
11. Jorgensen HS, Nakayama H, Raaschou HO, Vive-Larsen J, Stoier M, Olsen TS. Outcome and time course of recovery in stroke. the Copenhagen Stroke Study. *Arch Phys Med Rehabil*. 1995;76:399–412.
12. Berry DA, Mueller P, Grieve AP, Smith MK, Parke T, Krams M. Bayesian Designs for dose-ranging drug trials. In: Gatsonis C, Kass RE, Carlin B, Carriquiry A, Gelman A, Verdinelli I, West M, eds. *Case Studies in Bayesian Statistics,* vol 5. New York: Springer-Verlag; 2002: 99–181.
13. Diener HC, Cortens M, Ford G, Grotta J, Hacke W, Kaste M, Koudstaal PJ, Wessel T. Lubeluzole in acute ischemic stroke treatment: A double-blind study with an 8-hour inclusion window comparing a 10-mg daily dose of lubeluzole with placebo. *Stroke.* 2000;31:2543–2551.
14. Berry DA. Decision analysis and Bayesian methods in clinical trials. In: Thall PF, ed. *Recent Advances In Clinical Trial Design and Analysis*. New York: Kluwer Press; 1995:125–154.
15. Lewis RJ, Berry DA. Decision theory. In: Armitage P, Colton T, eds. *Encyclopedia of Biostatistics*, vol 2. New York: John Wiley & Sons; 1998:1109–1118.
16. Cheng Y, Su F, Berry DA. Choosing sample size for a clinical trial using decision analysis. *Biometrika*. 2003;90:923–936.
17. Inoue LYT, Thall P, Berry DA. Seamlessly expanding a randomized phase II trial to phase III. *Biometrics*. 2002;58:264–272.
18. Adams HP, Leclerc JR, Bluhmki E, Clarke W, Hansen MD, Hacke W. Measuring outcomes as a function of baseline severity of ischemic stroke. *Cerebrovasc Dis*. 2004;18:124–129.
19. Stroke Therapy Academic Industry Roundtable 4. Recommendations for advancing development of acute stroke therapies. *Stroke.* 2005; In press.