Early Detection of Accrual Problems in Clinical Trials

If you are a researcher, a member of an Institutional Review Boards (IRB), or a member of a Data and Safety Monitoring Board (DSMB), you need good statistical tools for the initial planning and the ongoing monitoring of clinical trials. In particular, you need to carefully consider the accrual rate--how rapidly are patients being recruited into your clinical trial.

With good tools, you would be able to construct realistic targets for their sample sizes rather than promising a sample size that could not be delivered in a reasonable time frame and within a limited research budget. You would get an early warning when accrual rates are suffering. This would allow you to take appropriate corrective action before too much harm was done. This corrective action could include early termination of your trial if that what was needed to keep you from throwing good money after bad. Right now, there are no methods for planning and monitoring accrual rates in clinical trials other than subjective inferences and ad hoc rules. We want to fix this problem.

The process of evaluating methods for the initial planning and continual monitoring of clinical trials is a complex one that will require support from multiple research grants. We are seeking funding from the Kansas City Area Life Sciences Institute (KCALSI) to develop some preliminary data to support an NIH grant using the R03 or R21 mechanism. The KCALSI support will allow us to develop a simple proof of concept, whereas the R03/R21 will let us prepare a working software prototype.

Once this work is done, we want to turn our attention to drop-out rates and adverse event rates. These are conceptually just the flip side of the same coin. With accrual rates, you are worried about "too little/too late" but with drop outs and adverse events, you are worried about "too much/too soon." While both dropouts and adverse events represent problems considerably more complex than problems with accrual, the statistical tools that work for the planning and ongoing monitoring of accrual rates should extend nicely into planning and evaluation of drop out rates and adverse event rates.

This grant addresses the goal of the KCALSI to develop cross-cutting enabling technologies that improve the research capabilities in all areas of medicine. The long term goal of our research is to provide information technology solutions that allow researchers to conduct clinical trials more efficiently and more ethically.

Aim #1. Develop a specific framework for initial planning of accrual rates. We want to produce some tools that will help researchers to thoughtfully and carefully plan their clinical trials.

Aim #2. Build a set of graphical and analytical tools to examine accrual rates on an on-going basis. We want to provide information to help researchers and IRBs to review their clinical trials on a regular basis to see if their accrual targets are being met.

Aim #3. Set up a web server running prototype software developed under the first two aims. The web server would enhance our interactions with groups outside the Kansas City area because they could easily review our concepts in action. Any programs developed with the grant will be open source.

Background and significance

The most common reason why clinical trials fail is that they fall well below their goals for patient accrual. According to one source[1], more than 80% of all clinical trials fall short of their accrual goals. The net result is too many studies that appear with inadequate sample sizes and confidence intervals that are so wide that they are effectively uninterpretable.

The problem of inadequate accrual is also of concern to IRBs and DSMBs. If a study has an inadequate accrual rate, that changes the balance of risks and benefits. The patients that are enrolled in the study experience inconvenience and possibly pain, and may even suffer increased risk but do not provide information with sufficient precision to produce scientifically valid results.

It is well understood that IRBs have inadequate tools for gauging the performance of clinical trials[2]. The NIH has recognized this as well and has documented the need for better tools in an Program Announcement.[3] The latter calls in particular for

> *"Development of appropriate outcomes measures and quality indicators for the IRB review process for measurement of adequate protection of human subjects. Development and testing of a framework for assessing IRB review quality; determination of when variability in IRB outcomes would be acceptable and when such variation would indicate inconsistent quality."*

We plan to provide quantitative tools for the planning and on-going monitoring of accrual rates in clinical trials. We will do this through the use of Bayesian prior distributions, statistical process control charts, and Bayesian posterior predictive distributions. We will also develop tools on a web server running R and BUGS.

Bayesian prior distributions. Although Bayesian data analysis methods are increasing in popularity[4], a commonly cited complaint about this approach is the requirement that the data analyst must specify prior distributions. A prior distribution reflects current knowledge in the form of a statistical model. In our case the prior distribution reflects what is currently known about the anticipated accrual rate of a clinical trial. To obtain a credible prior distribution one needs knowledge about the projected accrual rate as well as the uncertainty of the accrual rate.

Critics of the Bayesian approach believe that researchers should take a disinterested and objective perspective on their research. While this point is certainly merits extensive debate, we come down on the side of the Bayesian statistics.

From the perspective of accrual rates, however, there should be no debate. If you do not have at least an inkling of how many patients you are likely to recruit during the course of your study, you are unqualified to perform the research.

The trick is not getting people to offer an opinion about accrual rates, but instead getting them to recognize the uncertainty inherent in the accrual rates and quantify it in terms of a prior distribution. Eliciting a prior distribution in general is a difficult task. Most researchers tend to understate the true amount of uncertainty when you ask them about prior distributions[6].

In the context of eliciting a prior distribution for accrual rates, there are at least two possible sources for achieving an informative prior distribution. Internal review boards might utilize past accrual

rates from previous clinical trials. This is a reasonable approach as long as the proposed trial has similar accrual properties as the historical trial. In the case where the current trial has different properties than the past trials we must elicit a prior distribution from expert opinion. A common statement in the elicitation literature is that psychometric tools are needed to design a valid and reliable tool for eliciting prior distributions from clinical researchers[6].

Statistical Process Control Charts. The ongoing monitoring of clinical trials is a difficult process. Entire books[7,8] have been devoted to the statistical complexities as well as expensive software suites.[9] The problem is that if you perform hypothesis tests and allow for the possibility of early stopping based on efficacy, your sample size at the end of the study is no longer fixed but is now a random variable. Failure to adjust the final hypothesis test and confidence interval can produce seriously biased results.

Much of the difficulty with ongoing monitoring of clinical trials disappears, however, when the focus of the interim analysis is on variables not directly related to efficacy. These endpoints can be examined at multiple time points without altering the final test of efficacy. Accrual rates in particular can and should be monitored on a regular basis without changing the final data analysis.

Control charts are an excellent tool for assessing the stability of the accrual process and evaluating whether the accrual process is consistent with the demands of the research. Nowhere is the strength of control charts better in Donald Wheeler's book, **Understanding Variation**.[10] Mr. Wheeler points out

> *While every data set contains noise, some data sets may contain signals. Therefore, before you can detect a signal within any given data set, you must first filter out the noise. (page 30).*

Control charts use empirically derived limits that have been proven in a wide range of setting to minimize the chances of interpreting noise as a signal while still maintaining a strong probability of detecting a signal when it is present.

While control charts used to be limited to industrial applications, today they are used much more widely.

> *Control charts are among the most important and widely used tools in statistics. Their applications have now moved far beyond manufacturing into engineering, environmental science, biology, genetics, epidemiology, medicine, finance, and even law enforcement and athletics.*[11]

Bayesian posterior predictive distributions.

Predictive distributions are the main tool we propose to predict the end of a trial. Predictive distributions reflect all of the levels of uncertainty in the accrual rate. Using posterior distributions, the Bayesian paradigm offers a flexible foundation for estimating a predictive distribution. The posterior predictive distribution is estimated from two sources of variation. The first source of variation comes from the uncertainty in the model. The second source of uncertainty comes from the posterior distribution of the model's parameters. Propagating the model's uncertainty with the parameters' uncertainty produces the posterior predictive distribution. Bayesian posterior predictive distributions are utilized for model adequacy[4] in addition to actual temporal prediction. Draper[12] and many other statistical scientists view posterior

predictive distributions as the cornerstone of statistics.

Before any actual accrual data are collected, the parameters'posterior distribution is really just a function of the prior distribution. As data comes in on the actual accrual rates in the trial, the posterior predictive distribution becomes a weighted average of information from the prior distribution and information from the distribution of the accrual times. This prevents a researcher from overreacting to a small bit of bad news early in the trial. As more and more data on accrual appears, the weight on the prior distribution decreases and the weight on the actual data increases. If the actual accrual times are very slow, a sufficient number of them will appropriately pull down even the most optimistic of initial projections.

R and BUGS software.

A description of the R language appears on the Comprehensive R Archive Network.[13]

> *R is a language and environment for statistical computing and graphics. It is a GNU project which is similar to the S language and environment which was developed at Bell Laboratories (formerly AT&T, now Lucent Technologies) by John Chambers and colleagues.*
>
> *R provides a wide variety of statistical (linear and nonlinear modelling, classical statistical tests, time-series analysis, classification, clustering, ...) and graphical techniques, and is highly extensible.*

There are numerous libraries available for R including a set of programs for statistical process control charts[14] and an R interface to BUGS (see below).

WinBUGS is the Windows implementation of a DOS program called BUGS.[15] BUGS stands for Bates Using Gibbs Sampling. BUGS provides a set of programming tools to perform a Bayesian data analysis using new simulation methods: Gibbs Sampling, of course, but also the Metropolis-Hastings algorithm and Markov Chain Monte Carlo (MCMC). The Gibbs sampler and MCMC[16] are extremely important tools for doing modern Bayesian data analysis, in fact, Bayesian data analysis for complex models might not be possible with out them. The simulation allows updates of the posterior distribution of model parameters conditional on the observed data and random draws from the model parameters.

Open source software.

The term "open source" needs to be defined carefully, because it can mean many things. The R language is available as free software under the GNU General Purpose License (GPL).[17] The GPL makes the program and the source code freely available with one important restriction: All code added to a GPL program must be distributed (if it is distributed at all) at no cost. Users of the GNU GPL and similar licenses refer to this as "copyleft" (in contrast, of course, to copyright). The philosophy of copyleft is that programmers remove restrictions on how software can be used by not allowing others to use modified versions for commercial purposes. This assurance encourages community development of programs because no one individual can profit from the free work donated by others.[18]

Although WinBUGS[19] is free, it is not covered by the GNU GPL. An open source version of BUGS, called OpenBUGS[20], is, however, is available under the GNU GPL. Another program, BRUGS,[21] combines OpenBUGS with an interace to R, and this

package is also covered under the GNU GPL.

We plan to distribute any code developed under this grant as open source using the GNU GPL. This will encourage greater use of our programs and will allow other collaborators to work on their own enhancements.

Developing our prototype as an open source project does not preclude the possibility of later developing a commercial product. You could, for example, show a programming team a prototype developed under the open source model and ask them to create a proprietary product that performs the same functions. At this point, however, we anticipate that this project's chances for success are maximized by utilizing an open source model for all our software development.

Preliminary studies.

We do not have a large number of preliminary studies. We plan to use the grant money to develop case studies from existing data sets and run some computer simulations using theoretical distributions for accrual rates. It is the prototype work developed by this grant that will serve as the material for a preliminary studies section of an NIH R03 or R21 grant.

We have, however, developed a few of our ideas on a hypothetical data set. The actual data values are real, but they come from a retrospective study rather than a prospective trial. Nevertheless, we can illustrate some of our general concepts by treating this data as if it came from a prospective trial.

Consider a hypothetical research study that you started in January 1997 with the intention to recruit 12 patients per year (one per month) over a ten year period, for a total sample size of 120 patients. By the end of June 2004, (roughly 7 1/2 years), the study has enrolled 42 patients (Table 1).

2/26/1997	4/ 4/1997	7/ 7/1997
7/25/1997	2/ 5/1998	2/15/1998
3/ 6/1998	7/ 3/1998	8/ 3/1998
2/ 8/1999	3/19/1999	4/20/1999
5/29/1999	6/21/1999	7/27/1999
9/ 6/1999	1/10/2000	1/11/2000
2/28/2000	3/ 3/2000	4/13/2000
5/30/2000	11/21/2000	12/18/2000
2/ 6/2001	4/30/2001	8/ 3/2001
1/20/2001	12/ 3/2001	12/ 7/2001
9/27/2002	10/ 1/2002	2/ 2/2003
3/ 3/2003	10/31/2003	11/ 4/2003
11/11/2003	1/ 5/2004	2/ 2/2004
4/15/2004	5/23/2004	6/ 2/2004

Table 1. List of accrual dates.

Clearly you have problems. Your actual accrual rate is a meager 5.6 patients per year, and now it is probably too late to fix things. You would have to recruit at a rate more than 30 patients per year over the remainder of the study to overcome this shortfall.

Let's be honest--you knew it was a bit of a stretch to get 12 patients per year, and now you have to more than double that accrual rate. Wouldn't it be nicer if you had noticed the problem two years into the study rather than 7 1/2 years out?

So what have the accrual rates been telling us for the past 7 1/2 years? The logical way to evaluate this question is to plot the accrual rates, but you have some choices as far as what rates to plot.

You could plot the data using monthly rates (Figure 1). Notice that the data is a series of mostly zeros and ones. The accrual could also be summarized as the number of patients recruited per quarter (Figure 2) or as a yearly rate (Figure 3).

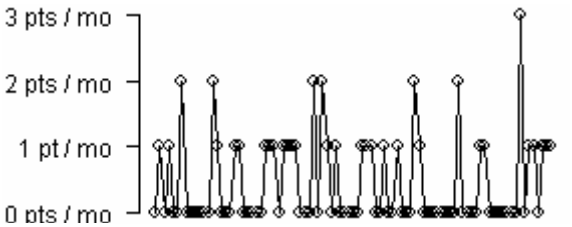

Figure 1. Plot of monthly accrual rates.

Figure 2. Plot of quarterly accrual rates.

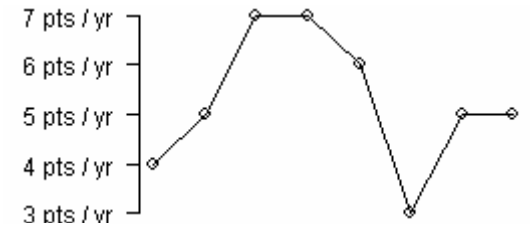

Figure 3. Plot of yearly accrual rates.

It is not clear what scale is most logical for this data. Should you analyze the data monthly? That allows you to respond quickly if a trend appears, but such a short time span leads to a lot of imprecision in any individual monthly value. A yearly trend provides a more stable estimate, but does not allow you to respond quickly to sudden shifts. Quarterly data offers the best (worst?) of both worlds.

The problem of which time interval to choose is a classic "Goldilocks" problem. Too large an interval leads to an inordinately long waiting time before you can accumulate enough data to take action. Too small an interval leads to highly granular data with little precision. We propose to avoid this problem entirely by focusing on the date gap.

The date gap is simply the amount of time that elapses between the recruitment of two successive patients. The hypothetical study started on January 1 and recruited the first two patients on February 26 and April 4. The gap between the start of the study and the first patient is 46 days and the gap between the first and second patients is 37 days. A plot of the date gaps is instructive (Figure 4):

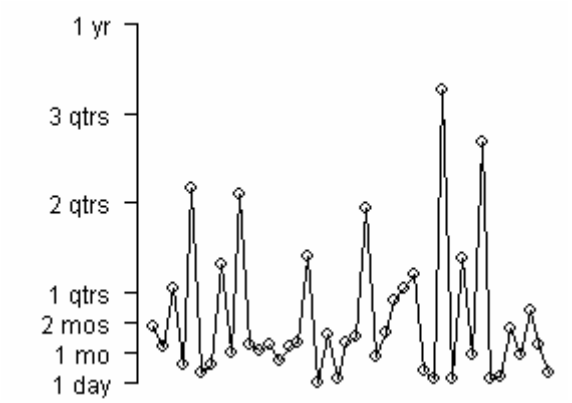

Figure 4. Plot of date gaps.

Note that the vertical axis uses a mixture of units of measurements (days, months, and years). This emphasizes the fact that a date gap is self scaling. If events occur frequently, the data will fit in the portion of the vertical axis where units are measured in days or weeks. If the events occur rarely, the data will fit in the portion of the graph where units are measured in months, quarters, or even years.

Another very important feature of the date gap is the fact that each time a patient is recruited into the trial, another point appears on the chart. This is in contrast to the three previous charts where you have to wait until the end of a month, quarter, or year before you can plot an additional data point. By recasting accrual rates in terms of the number of days between successive patients, we have liberated this problem from arbitrary calendar boundaries.

An important unsolved question is how to set control limits properly for a chart with

highly skewed data. The control limits is traditionally divided into zones, but for this data set, some of the zones are in the negative territory. This suggests that perhaps the control limits should be set using an asymmetric rule, that the traditional use of zones in a control chart should be modified for skewed data, or possibly both. Another intriguing possibility is to transform the data prior to computing control limits.

Although a prominent statistician has recommended the use of transformation in control charts (see page 265 of [22]), in practice, this is rarely done.

Another intriguing prospect is to apply a CUSUM chart to this technique. An accrual rate of one per month implies an average date gap of 30 days. If you plot the cumulative sum of the deviations of each individual date gap from the target of 30 days you get the following chart (Figure 5).

The data has shown a clear and consistent problem from the very first date gap. After 10 patients, the study is more than a year behind schedule. While there are a few points in time where the accrual seems to be

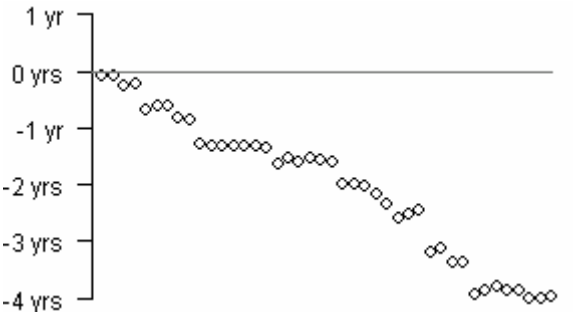

Figure 5. CUSUM plot of date gaps.

making up for lost time, more often than not, the study is falling further and further behind. The decision rules for CUSUM charts are not very well defined (see Chapter 13 of [23]) so additional research is needed.

You can extrapolate the CUSUM chart to the target sample size to get an estimate of the completion date (Figure 6). The diagonal line with a downward slope indicates the

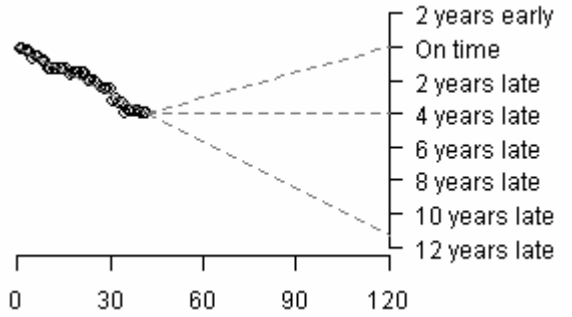

Figure 6. CUSUM plot with projections.

final fate of the study (11 years late) if the accrual patterns remain unchanged. The flat line represents the final fate of the study (4 years late) if you were able to fix things and get your accrual rate back up to the original target of one new patient per month. The line with an upward slope represents the rate you would have to recruit at if you still wanted to finish this study on time. The slope of this line is 19. This tells you that for the remaining 78 patients, your average date gap would have to be 19 days faster than the target of 30. In other words, you have to find a new patient every 11 days for the rest of the study if you wanted to finish on time.

These projections are somewhat simplistic, and we hope to produce better projections through the use of posterior predictive distributions. In order to obtain a Bayesian posterior predictive distribution of when the trial will end, we must first define a prior distribution for the accrual rate. In this example, we used a single parameter exponential distribution. Then, utilizing Bayes theorem, you would update this prior distribution, using the observed 42 patients, to obtain a posterior distribution of the model parameter. The posterior predictive distribution of the unobserved time gaps (in this case 78 more patients) is estimated

using two sources of variation. The first source of variation comes from the uncertainty in the exponential model. The second source of uncertainty comes from the posterior distribution of the model parameter. Propagating the model uncertainty with the parameter uncertainty produces the posterior predictive distribution.

A careful Bayesian analysis should consider several prior distributions for the analysis.[15] In this example, two sets of informative prior distributions are considered, an optimistic prior distribution and a pessimistic prior distribution. The optimist prior would reflect an expert's opinion, before any accrual, that the trial will end early. The pessimist reflects an expert's opinion that the trial will end late. We also examined a flat prior distribution that made no strong assumptions about the accrual rates.

The Bayesian posterior predictive distribution for time of trial using optimistic prior information versus pessimistic prior information appears in Figure 7. The key feature of this analysis is not only data and opinion being combined, but both posterior predictions reflect an uncertainty in the predictions. For the optimist, a 95% interval for the completion time is 3.8 years late to 7.1 years late. For the pessimist, the 95% interval for the posterior predictive distribution ranges from 11.1 years late to 19.1 years late. Clearly, there is cause for concern, no matter what your initial expectations.

In addition to the informative priors, a review committee could choose a "non-informative" prior which relies mostly on the data - the opinion carries virtually no weight on the posterior distribution. This posterior predictive 95% interval for a non-informative prior is 8.6 to 14.8 years late.

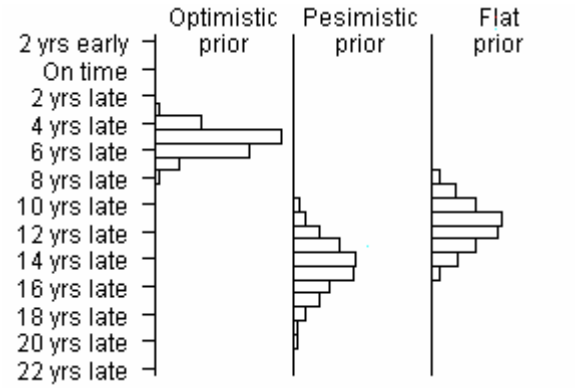

Figure 7. Posterior predictive distribution under three different prior distributions.

In contrast, what most review committees will do is to take the observed data and simply extrapolate to the finish with an average multiplied by 120 subjects. With this data would they would obtain a point estimate finish of 21.2 years. This analysis fails to acknowledge that there is always some level of uncertainty in the estimate of remaining accrual time. It is not the estimate of 21.2 years that is so troubling, it is the estimate that even allowing for uncertainty of accrual rates, you are very certain that the trial will finish many years behind schedule.

The uncertainty reflected in the posterior predictive distributions reinforces how important it is to carefully elicit all sources of uncertainty in accrual at the start of the study and set up a prior distribution for the accrual rate. If you believe that there is uncertainty about when the trial will finish midway through the study, why would you not have some level of uncertainty about when the trial will finish before the trial even starts?

Research design and methods

Since the goal of this grant is to develop some preliminary data for an NIH grant, there is no formal research design and no specific hypotheses to test. Instead, we plan to achieve three goals.

Our first goal is to develop a specific framework for initial planning of accrual rates using **Bayesian prior distributions**. We will develop a series of questions based on sound psychometric principles to carefully elicit all the sources of uncertainty associated with accrual rates. These questions might be supplemented with a graphical display of the researcher's assumptions.

Our second goal is to examine accrual rates on an on-going basis using **Statistical Process Control Charts** and **Bayesian predictive distributions**. We have already identified several sources for accrual data from actual clinical trials, and we plan to develop three good case studies using these clinical trials. We will also run some simple computer simulations examining the behavior of these approaches under an idealized setting.

Our third goal is to set up a **web server running R and WinBUGS**. There are already several working examples that we can refer to, including a system developed here at Children's Mercy Hospital in collaboration with HeartLab. The cost for setting up the server is minimal. The programs we plan to use are free, and we will initially piggyback on an existing server (bipad.cmh.edu). Any programs for this server will be developed under and open source license. The open source license will help promote wide distribution of our initial work and will encourage collaboration from other experts in the area.

Project timeline.

Month 1.

- Receive funding from KCALSI.
- Submit paperwork for IRB approval of work on identifying possible case studies.
- Set up web server: accrual.cmu.edu.

Month 2.

- Outline all sources of uncertainty in accrual and identify appropriate probability distributions.
- Identify at least three research projects that might produce promising case study results.
- Install R and BRugs on server.

Months 3-6.

- Develop valid tools for eliciting prior distributions based on sound psychometric principles.
- Obtain actual accrual dates for case studies and prepare analyses.
- Set up and test some simple analyses using R and BRugs.

Months 7-9.

- Obtain feedback from selected local experts on the tools for eliciting prior distributions.
- Develop some simulation models. Realistic distributions for the simulation will be developed based on the distributional patterns of the case studies.
- Advertise R and BRugs analyses of accrual rates to selected local experts and ask for comments.

Months 10-12.

- Document case studies and simulation results in a form suitable for publication.
- Adapt server programs based on initial feedback. Start work on a user friendly interface.
- Develop R03 or R21 grant for submission to NIH.

Unique skills and experiences of the research team.

The first principal investigator (Dr. Simon) has substantial experience with control charts. At his previous job at the National Institute for Occupational Safety and Health, he helped set up a program in quality control from scratch. He has been invited repeatedly to the American Society of Andrology to lecture on the use of control charts to improve laboratory quality.

Dr. Simon's use of date gaps and control charts forms the foundation of a patent that Children's Mercy Hospital has applied for. The patent, A System and Method for Monitoring and Analyzing Data Trends of Interest Within an Organization, is currently under review by the U.S. Patent Office.

The inventors and Children's Mercy Hospital are currently negotiating with a major health care company for distributing software using date gap methods in control charts to monitor patient safety events. This grant represents an effort to expand the patent to work in a totally new arena.

Dr. Simon has also consulted on a software system developed by HeartLab to evaluate the growth patterns in echocardiogram measurements of the heart. A major component of this project is an interface between the database of clinical measurements and a program in R to use smoothing splines.

Finally, Dr. Simon is the author of the StATS (www.cmh.edu/stats), a web site with over one thousand pages that cover broad areas of research methodology and statistics. These pages including major sections on quality control and the ethical conduct of research.

The second principal investigator (Dr. Gajewski) has extensive experience in Bayesian data analysis. Over six years, as an Assistant Professor, he has researched and taught two years at St. Cloud State University and then four years at University of Kansas Medical Center. Over these years, Gajewski has developed Bayesian peer-reviewed statistical methodology motivated from problems that arise in the areas of transportation, sports, and health sciences. These peer-reviewed statistical papers have contributed to the development of Bayesian smoothing splines, Bayesian piecewise linear models, Bayesian non-normal probability mixed models, Bayesian latent variable modeling, and Bayesian measurement error .

Recently, Gajewski has published in the area of Bayesian clinical trial design specifically for determining sample sizes for Phase II trials (Mayo & Gajewski; Gajewski & Mayo). The specific methodology developed relies on modeling experts' opinions and incorporating an informative prior distribution to be utilized in sample size calculations. It was shown that the operating characteristics vary depending on the types of priors elicited from the clinicians. These types of priors include optimistic, pessimistic, and a mixture of opinions.

In addition to pure Bayesian developments, Gajewski has been involved in developing and validating behavioral instruments (questionnaires) from his work with nursing researchers. This requires skill in the area of psychometric development. In addition to substantive development, Gajewski has recently submitted a statistical methodological publication to the journal *Psychometrika*. In this paper, Gajewski investigates the feasibility of calculating credible intervals for Cronbach's alpha from items measured on the ordinal scale.

Literature cited

1. Barnes K (2006) "Pharma giants risk reputation through clinical trial cost-cutting." Viewed on July 24, 2006. www.in-pharmatechnologist.com/news/ng.asp?n=68150-chiltern-india-cost-clinical-trial-regulatory

2. Emanuel EJ, Wood A, Fleischman A, Bowen A, Getz KA, Grady C, Levine C, Hammerschmidt DE, Faden R, Eckenwiler L, Muse CT, Sugarman J. (2004) "Oversight of human participants research: identifying problems to evaluate reform proposals." Annals of Internal Medicine (Aug 17); 141(4): 282-291.

3. National Institutes of Health (2006) "Research On Ethical Issues In Human Subjects Research." Viewd on July 24, 2006. grants.nih.gov/grants/guide/pa-files/PA-06-367.html

4. Gelman A, Carlin JB, Stern HS & Rubin, DB (2000), **Bayesian Data Analysis**, Chapman and Hall/CRC: Washington DC.

5. [[Text book by LANL person]]

6. Garthwaite PH, Kadane JB, & O'Hagan A (2005), "Statistical Methods for Eliciting Probability Distributions," *Journal of the American Statistical Association*, 100 (470), 680-701.

7. Jennison C, Turnbull JW (1999). **Group Sequential Methods with Applications to Clinical Trials**. New York, NY: Chapman & Hall.

8. Whitehead J (1997**). The Design and Analysis of Sequential Clinical Trials**. New york NY: John Wiley and Sons, Inc.

9. Insightful Corporation (2006). "S+SeqTrial." Viewed on July 24, 2006, www.insightful.com/products/seqtrial/default.asp

10. Wheeler, D (1999). Understanding Variation. The Key to Managing Chaos (2nd Edition). Knoxville, TN: SPC Press, Inc.

11. Stoumbos ZG, Reynolds, Jr. MR, Ryan TP, Woodall WH (2001) "The State of Statistical Process Control as We Proceed into the 21st Century." Journal of the American Statistical Association. 95: 992-998.

12. Draper D, Gittoes M (2004), "Statistical Analysis of Performance Indicators in UK Higher Education," Journal of the Royal Statistical Society A, 167(3), 449-474.

13. Leisch F (2006) "The Comprehensive R Archive Network." Viewed on July 24, 2006. cran.r-project.org

14. Scrucca L (2004) "qcc: Quality Control Charts." Viewed on July 24, 2006. cran.r-project.org/src/contrib/Descriptions/qcc.html

15. Spiegelhalter DJ, Abrams KR, Myles JP (2004) **Bayesian Approaches to Clinical Trials and Health-Care Evaluation**. New York NT: John Wiley & Sons.

16. Gilks WR, Richardson S & Spiegelhalter DJ (1996), **Markov Chain Monte Carlo in Practice**, London, UK: Chapman and Hall/CRC.

17. The GNU Project. "GNU General Public License." Viewed on July 24, 2006. www.gnu.org/copyleft/gpl.html

18. Stallman R. "Copyleft: Pragmatic Idealism." Viewed on July 24, 2006. www.gnu.org/philosophy/pragmatic.html

19. Stevens A. (2004) "The BUGS Project - WinBUGS." Viewed on July 24, 2006.

www.mrc-bsu.cam.ac.uk/bugs/winbugs/contents.shtml

20. Thomas A, O'Hara B. (2004) "OpenBugs: the Homepage." Viewed on July 24, 2006. mathstat.helsinki.fi/openbugs/Home.html

21. Thomas A, O'Hara B (2006) "BRugs: OpenBUGS and its R interface BRugs." Viewed on July 24, 2006. cran.r-project.org/src/contrib/Descriptions/BRugs.html

22. Joiner B (1994) **Fourth Generation Management: The New Business consciousness**. New York NY: McGraw-Hill.

23. Wheeler DJ (2004) **Advanced Topics in Statistical Process control: The Power of Shewhart's Charts**. Knoxville TN: SPC Press, Inc.