The Stata Journal

The *Stata Journal* publishes reviewed papers together with shorter notes or comments, regular columns, book reviews, and other material of interest to Stata users. Examples of the types of papers include 1) expository papers that link the use of Stata commands or programs to associated principles, such as those that will serve as tutorials for users first encountering a new field of statistics or a major new technique; 2) papers that go "beyond the Stata manual" in explaining key features or uses of Stata that are of interest to intermediate or advanced users of Stata; 3) papers that discuss new commands or Stata programs of interest either to a wide spectrum of users (e.g., in data management or graphics) or to some large segment of Stata users (e.g., in survey statistics, survival analysis, panel analysis, or limited dependent variable modeling); 4) papers analyzing the statistical properties of new or existing estimators and tests in Stata; 5) papers that could be of interest or usefulness to researchers, especially in fields that are of practical importance but are not often included in texts or other journals, such as the use of Stata in managing datasets, especially large datasets, with advice from hard-won experience; and 6) papers of interest to those who teach, including Stata with topics such as extended examples of techniques and interpretation of results, simulations of statistical concepts, and overviews of subject areas.

The *Stata Journal* is indexed and abstracted by *CompuMath Citation Index*, *Current Contents/Social and Behavioral Sciences*, *RePEc: Research Papers in Economics*, *Science Citation Index Expanded* (also known as *SciSearch*), *Scopus*, and *Social Sciences Citation Index*.

For more information on the *Stata Journal*, including information for authors, see the webpage

http://www.stata-journal.com

Subscriptions are available from StataCorp, 4905 Lakeway Drive, College Station, Texas 77845, telephone 979-696-4600 or 800-STATA-PC, fax 979-696-4601, or online at

http://www.stata.com/bookstore/sj.html

Subscription rates listed below include both a printed and an electronic copy unless otherwise mentioned.

U.S. and Canada		Elsewhere	
Printed & electronic		**Printed & electronic**	
1-year subscription	$ 98	1-year subscription	$138
2-year subscription	$165	2-year subscription	$245
3-year subscription	$225	3-year subscription	$345
1-year student subscription	$ 75	1-year student subscription	$ 99
1-year institutional subscription	$245	1-year institutional subscription	$285
2-year institutional subscription	$445	2-year institutional subscription	$525
3-year institutional subscription	$645	3-year institutional subscription	$765
Electronic only		**Electronic only**	
1-year subscription	$ 75	1-year subscription	$ 75
2-year subscription	$125	2-year subscription	$125
3-year subscription	$165	3-year subscription	$165
1-year student subscription	$ 45	1-year student subscription	$ 45

Back issues of the *Stata Journal* may be ordered online at

http://www.stata.com/bookstore/sjj.html

Individual articles three or more years old may be accessed online without charge. More recent articles may be ordered online.

http://www.stata-journal.com/archives.html

The *Stata Journal* is published quarterly by the Stata Press, College Station, Texas, USA.

Address changes should be sent to the *Stata Journal*, StataCorp, 4905 Lakeway Drive, College Station, TX 77845, USA, or emailed to sj@stata.com.

The *Stata Journal* (ISSN 1536-867X) is a publication of Stata Press. Stata, Stata, Stata Press, Mata, Mata, and NetCourse are registered trademarks of StataCorp LP.

The Stata Journal (2014)
14, Number 2, pp. 407–417

Stata command for calculating adverse event and efficacy stopping boundaries for phase II single-arm trials

Bryan M. Fellman
MD Anderson Cancer Center
Houston, TX
bmfellman@mdanderson.org

Abstract. Many programs and functions in statistical packages focus on the final stage of clinical trials, that is, the data analysis. In this article, I aim to assist in the early stages of clinical trials, specifically, the design of phase II single-arm trials. I present the new command `stopbound`, which calculates stopping boundaries and operating characteristics based on monitoring an adverse event, efficacy, or an adverse event and efficacy.

Keywords: st0344, stopbound, Bayesian clinical trial design, adverse event, efficacy rate, futility, simulation, operating characteristic

1 Introduction

A good design is essential for the success of a clinical trial; however, many statistical packages lack support for the design phases of clinical trials. Thus in this article, I introduce a command for a widely used method in clinical trial design for calculating Bayesian decision criteria and stopping boundaries to warrant early termination of a clinical trial; see Thall, Simon, and Estey (1995), Thall, Simon, and Estey (1996), and the extension by Thall and Sung (1998). Stopping boundaries are usually implemented for monitoring a severe adverse event (AE) (for example, toxicity) or efficacy rates for a new experimental treatment compared with the standard-of-care historical rates. These rules are put in place to protect the safety of the patients entering the trial: one would like to stop a trial if the rate of AEs is much higher than anticipated to prevent further patients from exposure to the high risk of events. One may also stop a trial because of low efficacy rates to treatment, to protect future patients from receiving an ineffective therapy.

These trials are typically conducted to test whether an experimental therapy is safe and sufficiently efficacious to begin larger randomized clinical trials. The methods compare posterior probabilities of the experimental drug with the historical standard, where both probabilities are modeled by Beta(α, β) distributions, and the parameters α and β represent the number of successes and failures, respectively. Alpha and beta change; at first, they are from the prior, but as the data accumulate, they become the sum of the assumed prior successes or failures and the observed successes or failures.

In this article, I introduce a new command to Stata users. The command `stopbound` calculates AE or efficacy stopping boundaries and operating characteristics, or both, for phase II single-arm trials. I present basic methods, general use, syntax, options, and four examples to illustrate its use.

2 Methods and general use for the stopbound command

2.1 Models

The first design aspect the user must input are the parameters for modeling the probability of AE or efficacy or both for the standard and the experimental treatment. The probability for the standard treatment is modeled by a Beta(α_T, β_T) distribution, where α is the number of events observed and β is the number of nonevents. This prior on the standard treatment is based on historical knowledge and data. Instead of entering the α_T and β_T parameters for the historical standard, one can also enter the maximum (minimum) acceptable AE (efficacy) rate, which will be called the null rate. When entering the null rate, the user is essentially assuming a known mean of the historical standard treatment. This would be similar to having a large historical population on which to base the standard treatment prior, such as information on 10,000 patients. For AE (efficacy), the null rate is the maximum (minimum) allowable rate that the investigator is willing to accept to consider the therapy safe (efficacious). The difference between the null rate and the probabilistic definitions of the standard treatment is that the latter incorporates the sampling uncertainty in the comparator. Using the probabilistic approach requires simulation and takes a little longer for computation. An example of both approaches is given in section 4.1.

The prior probability for the experimental treatment is modeled by Beta(α_E, β_E). These values can be whatever the user wants; however, it is recommended that α_E and β_E are chosen such that their sum is small (as to not overwhelm the observed data) and such that the two priors have the same mean. The sum $\alpha_E + \beta_E$ performs as the sample size of the prior knowledge of the experimental treatment. A weakly informative or noninformative prior is used because little information is known about the new experimental agent; therefore, the prior for it should not be highly informative. A weakly informative or noninformative prior allows the data to speak. For example, if we have information on 100 patients on the standard treatment and 30 of them had an AE, the distribution of the standard treatment should be Beta$(30, 70)$. Knowing the standard distribution, the user could make the prior distribution on the experimental treatment $\sim$ Beta$(0.60, 1.40)$, which has a mean of 0.30, the same as a Beta$(30, 70)$. Again this is only a recommendation. If more information is known about the experimental treatment, the prior could be made as informative as the user would like.

2.2 Sample size

Other design aspects of the trial that must be entered are the maximum number of patients to be enrolled, cohort size, and the minimum number of patients accrued before applying stopping rules. For example, one might want to consider the latter if the investigator does not want to stop the trial early based on limited data, even if the stopping criteria suggest the trial should be terminated.

2.3 Stopping criteria

The next step is to generate the rules to monitor the trial so that stopping boundaries can be calculated and used to terminate the trial in case of an unacceptable number of AEs or a lack of efficacy. The stopping rule for AEs is $\Pr(\theta_E > \theta_S|\text{data}, n) > p_U$. This means the trial will stop if the posterior probability that the experimental treatment AE rate is greater than the standard-of-care AE rate is greater than the upper bound p_U, which typically ranges from 0.95–0.99. Generally speaking, after viewing the data, there is greater than p_U chance that the experimental treatment AE rate will be higher than the standard-of-care AE rate. The stopping rule for efficacy is $\Pr(\theta_E > \theta_S|\text{data}, n) < p_L$. This means the trial will stop if the posterior probability that the experimental treatment is more efficacious than the standard of care is less than the lower bound p_L, which is typically ≤ 0.10. Likewise, after viewing the data, there is less than p_L chance that the experimental treatment will have an event rate greater than the standard of care. This is known as a futility stopping rule.

2.4 Operating characteristics

The next step is to supply a list of true or hypothetical adverse event (or efficacious) rates for the experimental treatment. This allows for operating characteristics to be calculated from simulation to estimate the probability of trial termination, the average number of events, and the average number of patients treated under each scenario. It is recommended that the probability of stopping the trial when the true rate is equal to the standard-of-care rate be kept as close to 0.10 as possible.

3 The stopbound command

3.1 Syntax

```
stopbound, type(#) nMAX(#) [dist nMIN(#) cohort(#) mcmc(#) sims(#)
   aT0(#) bT0(#) pT0(#) aT(#) bT(#) pU(#) aR0(#) bR0(#) pR0(#)
   aR(#) bR(#) pL(#) truepT(numlist) truepR(numlist) pNN(numlist)
   pNY(numlist) pYN(numlist) pYY(numlist)]
```

3.2 Options

Main options

`type(#)` specifies which stopping boundary to calculate, where `0` = AE; `1` = efficacy; `2` = independent AE and efficacy; and `3` = dependent AE and efficacy. `type()` is required because at least one boundary must be specified.

`nMAX(#)` specifies the maximum number of patients to accrue. `nMAX()` is required.

`dist` specifies that experimental treatment be compared with standard of care using distributions. Without this option, the experimental treatment is compared with a constant.

`nMIN(#)` specifies the minimum number of patients at first interim look. The default is `nMIN(1)`.

`cohort(#)` specifies the cohort size. The default is `cohort(1)`.

`mcmc(#)` specifies the number of Markov chain Monte Carlo replications to perform when calculating the posterior probability for comparing the experimental distribution with the standard distribution. The default is `mcmc(1000000)`.

`sims(#)` specifies the number of simulations to be performed for calculating operating characteristics. The default is `sims(10000)`.

AE boundary options

`aT0(#)` specifies alpha for prior beta distribution on AE for the standard therapy. This option is required when `dist` is specified and `type(0)`, `type(2)`, or `type(3)` is specified.

`bT0(#)` specifies beta for prior beta distribution on AE for the standard therapy. This option is required when `dist` is specified and `type(0)`, `type(2)`, or `type(3)` is specified.

`pT0(#)` specifies the maximum acceptable AE rate. This option is required if `type(0)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`aT(#)` specifies alpha for prior beta distribution on AE for the experimental therapy. This option is required if `type(0)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`bT(#)` specifies beta for prior beta distribution on AE for the experimental therapy. This option is required if `type(0)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`pU(#)` specifies stopping criteria for AE. This option is required if `type(0)`, `type(2)`, or `type(3)` is specified.

Efficacy boundary options

`aR0(#)` specifies alpha for prior beta distribution on efficacy for the standard therapy. This option is required when `dist` is specified and `type(1)`, `type(2)`, or `type(3)` is specified.

`bR0(#)` specifies beta for prior beta distribution on efficacy for the standard therapy. This option is required when `dist` is specified and `type(1)`, `type(2)`, or `type(3)` is specified.

`pR0(#)` specifies the minimal acceptable efficacy rate. `pR0()` is required if `type(1)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`aR(#)` specifies alpha for prior beta distribution on efficacy for the experimental therapy. This option is required if `type(1)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`bR(#)` specifies beta for prior beta distribution on efficacy for the experimental therapy. This option is required if `type(1)`, `type(2)`, or `type(3)` is specified without the `dist` option.

`pL(#)` specifies stopping criteria for efficacy. This option is required if `type(1)`, `type(2)`, or `type(3)` is specified.

Independent AE and efficacy options

`truepT(`*numlist*`)` specifies a list of hypothetical toxicity rates.

`truepR(`*numlist*`)` specifies a list of hypothetical response rates.

Dependent AE and efficacy options

`pNN(`*numlist*`)` specifies the probability of having no efficacy and no AE.

`pNY(`*numlist*`)` specifies the probability of having no efficacy and having AE.

`pYN(`*numlist*`)` specifies the probability of having efficacy and having no AE.

`pYY(`*numlist*`)` specifies the probability of having efficacy and AE.

4 Examples

4.1 Monitoring AE only

Suppose a single-arm phase II trial is being designed to test a new therapeutic agent T to fight a type of cancer, and we want to make sure it is less toxic when compared with standard treatment S. We want to monitor toxicities for the new agent, so the purpose of the following design is to ensure that treatment T is not more toxic than treatment S. The method will compare the posterior probabilities of toxicity for patients in the

trial on treatment T with historical data on treatment S. Suppose we have searched historical data and found that 100 patients similar to those entering this trial have received standard treatment. Of those 100 patients, 25 experienced toxicity; therefore, the prior for θ_S should be Beta$(25, 75)$. Recall the example in section 2.1; the prior for T will be Beta$(0.5, 1.5)$. We decide to stop the trial if, based on the data, the posterior probability of the toxicity rate for T is going to be more than S is greater than 0.95, that is, $\Pr(\theta_T > \theta_S|\text{data}, n) > 0.95$. Additionally, the trial is limited to a maximum of 36 patients, and patients will be entered in cohorts of size 3.

Now that we have all the information on the trial design, the stopping boundaries are calculated as follows:

```
. stopbound, type(0) dist nMIN(1) cohort(3) nMAX(36) aT0(25) bT0(75) aT(0.5)
> bT(1.5) pU(0.95) truepT(0.05 0.15 0.25 0.35 0.45)
Toxicity Stopping Boundaries:
-----------------------------
The following are greater-than-or-equal boundaries:
T/N means to stop if the number of toxicities after
treating N patients is greater than or equal to T.
STOP          post prob
-----------------------
3/3           .981757
4/6           .963045
5/9           .951541
7/12          .982698
8/15          .978032
9/18          .974453
10/21         .971091
11/24         .968677
12/27         .966731
13/30         .965193
14/33         .963867
15/36         .962854
P(true)  P(stop)   p10  p25  p50  p75  p90  Avg # pts    Avg # toxicities
------------------------------------------------------------------------
.05      .0002     36   36   36   36   36   35.994       1.8119
.15      .0111     36   36   36   36   36   35.6667      5.3227
.25      .0962     36   36   36   36   36   33.5868      8.3966
.35      .386      6    18   36   36   36   27.5874      9.6624
.45      .78       6    6    18   33   36   18.7224      8.4273
```

So for any given number of patients N entered into the trial, if there are T toxicities or more at that point, the trial will terminate. For instance, if we have enrolled 18 patients and have observed 9 or more toxicities, the trial would terminate because of excessive toxicity. The chances of stopping the trial for different true rates of toxicity for T are also presented. If the true rate of toxicity is 0.45, there is a 78% chance we would stop the trial early. If the true rate was equal to the 0.25 rate (mean of the standard treatment), then the probability of stopping the trial early is 0.0962, which approximates to a 10% chance; see section 2.4.

Below is an example of the same study design as above except the null rate approach is used. Here the maximum acceptable toxicity rate is set at 0.25, which is the mean of the prior for θ_S in the probabilistic approach.

```
. stopbound, type(0) nMIN(1) cohort(3) nMAX(36) pT0(0.25) aT(0.5) bT(1.5)
> pU(0.95) truepT(0.05 0.15 0.25 0.35 0.45)
Toxicity Stopping Boundaries:
-----------------------------
The following are greater-than-or-equal boundaries:
T/N means to stop if the number of toxicities after
treating N patients is greater than or equal to T.
STOP          post prob
-----------------------
3/3           .98368078
4/6           .96792853
5/9           .95906788
6/12          .95442069
7/15          .95230486
8/18          .95174547
9/21          .95216136
10/24         .95319298
11/27         .95461051
12/30         .95626317
13/33         .95805007
14/36         .95990275
P(true)  P(stop)   p10  p25  p50  p75  p90  Avg # pts    Avg # toxicities
------------------------------------------------------------------------
.05      .0003     36   36   36   36   36   35.991       1.778
.15      .0142     36   36   36   36   36   35.6229      5.3439
.25      .1559     15   36   36   36   36   32.5827      8.1833
.35      .5186     6    12   33   36   36   25.0395      8.7635
.45      .8714     6    6    12   24   36   15.9651      7.1962
```

The stopping boundaries are slightly different because of the sampling uncertainty used when comparing with the standard prior for θ_S in the first approach.

4.2 Monitoring efficacy only

Suppose we are designing a similar trial but now we want to apply a stopping rule for the case of a nonresponsive agent T. In this example, we will compare T to a constant. We want the response rate to be better than 0.40. With the null rate at 0.40, the prior on T could be Beta$(0.8, 1.2)$. We decide to stop the trial if, based on the data, the posterior probability of the response rate for T is going to be higher than 0.40 is less than 0.10, that is, $\Pr(\theta_T > 0.40|\text{data}, n) < 0.10$. The following Stata code would be used:

```
. stopbound, type(1) nMIN(1) cohort(3) nMAX(36) pR0(0.40) aR(0.8) bR(1.2)
> pL(0.10) truepR(0.20 0.30 0.40 0.50 0.60)
Response Stopping Boundaries:
-----------------------------
The following are less-than-or-equal boundaries:
R/N means to stop if the number of responses after
treating N patients is less than or equal to R.
STOP          post prob
----------------------
0/3           .08485245
1/9           .03328467
2/12          .04452425
3/15          .05205239
4/18          .05702922
5/21          .06023106
6/24          .06217455
7/27          .06321036
8/30          .06358214
9/33          .06346308
10/36         .06297874
P(true)  P(stop)   p10  p25  p50  p75  p90  Avg # pts    Avg # responses
-------------------------------------------------------------------------
.2       .9474     3    3    3    15   27   10.4283      2.0774
.3       .6747     3    3    18   36   36   19.0146      5.7229
.4       .345      3    12   36   36   36   26.5134      10.5992
.5       .1554     3    36   36   36   36   31.17        15.5528
.6       .0648     36   36   36   36   36   33.909       20.3172
```

Thus for any given number of patients N entered into the trial, if there are R responses or fewer at that point, the trial will terminate. If we have enrolled 18 patients and have seen only 4 or fewer responses, the trial would terminate for futility. This design would stop 35% of the time if the true response rate was 0.40. This would be extremely high in most trials, so changing parameters such as `pL()` might lower the probability of stopping early if the true rate was equal to the null rate and yield better operating characteristics for the trial design.

4.3 Monitoring AE and efficacy

Independent case

Now suppose we design a trial that has both a stopping rule for toxicity and a stopping rule for response, and we assume the two are independent of one another. We will use a Beta$(25, 75)$ on $\theta_{S,T}$ and a Beta$(80, 120)$ on $\theta_{S,R}$, so our priors on $\theta_{T,T}$ and $\theta_{T,R}$ could be Beta$(0.5, 1.5)$ and Beta$(0.8, 1.2)$, respectively. We will use an upper bound of 0.95 for monitoring toxicity and a lower bound of 0.05 for monitoring response. The following Stata code would be used:

```
. stopbound, type(2) dist nMIN(1) cohort(3) nMAX(36) aT0(25) bT0(75) aT(0.5)
> bT(1.5) pU(0.95) truepT(0.05 0.15 0.25 0.35 0.45) aR0(80) bR0(120) aR(0.8)
> bR(1.2) pL(0.05) truepR(0.60 0.30 0.40 0.30 0.20)

Toxicity Stopping Boundaries:
-----------------------------
The following are greater-than-or-equal boundaries:
T/N means to stop if the number of toxicities after
treating N patients is greater than or equal to T.

STOP          post prob
----------------------
3/3           .981583
4/6           .96311
5/9           .951417
7/12          .98248
8/15          .977937
9/18          .974392
10/21         .97146
11/24         .969048
12/27         .966809
13/30         .965265
14/33         .963777
15/36         .962792

Response Stopping Boundaries:
-----------------------------
The following are less-than-or-equal boundaries:
R/N means to stop if the number of responses after
treating N patients is less than or equal to R.

STOP          post prob
----------------------
0/6           .018148
1/9           .036685
2/12          .049564
3/18          .021876
4/21          .026949
5/24          .031057
6/27          .034286
7/30          .037125
8/33          .039044
9/36          .041023

P(true)
 T  R  P(stop) p10 p25 p50 p75 p90 Avg # pts  Avg # toxicities  Avg # responses
------------------------------------------------------------------------------
.05 .6  .009   36  36  36  36  36  35.7522     1.815              21.4191
.15 .3  .4764  6   12  36  36  36  25.7277     3.8608             7.719
.25 .4  .2343  9   36  36  36  36  30.2988     7.5811             12.1412
.35 .3  .6784  6   9   18  36  36  20.5539     7.2164             6.2027
.45 .2  .9718  6   6   9   12  24  11.0037     4.9521             2.2043
```

If both the true toxicity and response rates are at the mean of the standard, then the trial would stop 23% of the time. If there are very good response and toxicity rates, then the trial would stop only 1% of time. On the other hand, if there are very poor rates for both, then the trial would stop 97% of the time.

Dependent case

Suppose that we do not assume independence of toxicity and response. In this case, there are four possible probabilities—λ_1, λ_2, λ_3, and λ_4—which can be presented in the 2×2 table below, where λ_1, λ_2, λ_3, and λ_4 must sum to 1:

	Toxicity	
Response	No	Yes
No	λ_1	λ_2
Yes	λ_3	λ_4

Therefore, the probability of toxicity is $\pi_T = \lambda_2 + \lambda_4$, and the probability of response is $\pi_R = \lambda_3 + \lambda_4$.

Now suppose we design a trial that includes a stopping rule for toxicity and a rule for response; however, now we assume the two are dependent of each other. The same design parameters are used as in the independent example. The table below includes five possible scenarios of interest. For example, scenario 3 would be the best scenario, resulting in a high response rate ($0.55 + 0.05 = 0.60$) and low toxicity rate ($0.05 + 0.05 = 0.10$). Scenario 1 would be the most undesirable case with high rates of toxicity and low rates of response, while scenario 2 has toxicity and response rates equal to the standard rates.

	Scenario 1 $\downarrow R \uparrow T$		Scenario 2 Null		Scenario 3 $\uparrow R \downarrow T$		Scenario 4 $\uparrow R \uparrow T$		Scenario 5 $\downarrow R \downarrow T$	
					Toxicity					
	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
Response										
No	0.55	0.25	0.55	0.05	0.35	0.05	0.40	0.10	0.75	0.05
Yes	0.05	0.15	0.20	0.20	0.55	0.05	0.20	0.30	0.15	0.05

The following Stata code would be used to produce the operating characteristics for this design:

```
. stopbound, type(3) dist nMIN(1) cohort(3) nMAX(36)
> aT0(25) bT0(75) aT(0.5) bT(1.5) pU(0.95) aR0(80)
> bR0(120) aR(0.8) bR(1.2) pL(0.10) pNN(0.55 0.55 0.35 0.40 0.75)
> pNY(0.25 0.05 0.05 0.10 0.05) pYN(0.05 0.20 0.55 0.20 0.15)
> pYY(0.15 0.20 0.05 0.30 0.05)
```

(*output omitted*)

P(true)										Average		
(-R,-T)	(-R,+T)	(+R,-T)	(+R,+T)	P(stop)	p10	p25	p50	p75	p90	# pts	# tox´s	# resp´s
.55	.25	.05	.15	.9793	3	3	3	12	21	8.3106	3.3252	1.6541
.55	.05	.2	.2	.4046	3	6	36	36	36	24.9606	6.2398	9.9475
.35	.05	.55	.05	.0662	36	36	36	36	36	33.846	3.4011	20.331
.4	.1	.2	.3	.6555	3	6	21	36	36	20.4144	8.139	10.2412
.75	.05	.15	.05	.9482	3	3	3	15	27	10.5459	1.0674	2.1133

5 References

Thall, P. F., R. M. Simon, and E. H. Estey. 1995. Bayesian sequential monitoring designs for single-arm clinical trials with multiple outcomes. *Statistics in Medicine* 14: 357–379.

———. 1996. New statistical strategy for monitoring safety and efficacy in single-arm clinical trials. *Journal of Clinical Oncology* 14: 296–303.

Thall, P. F., and H.-G. Sung. 1998. Some extensions and applications of a Bayesian strategy for monitoring multiple outcomes in clinical trials. *Statistics in Medicine* 17: 1563–1580.

About the author

Bryan M. Fellman is a research statistical analyst in the Department of Biostatistics at the University of Texas MD Anderson Cancer Center.