Introduction to Bayesian adaptive study designs

JoAnn Alvarez

Department of Biostatistics
Vanderbilt University School of Medicine

November 17, 2011

Outline

VANDERBILT UNIVERSITY

Overview Bayesian adaptive designs

What are they?

adaptive some aspect of the study design may change during the study depending on observed values after trial has begun

- patient covariates or outcomes
- all changes prespecified (prospectively adaptive)

Bayesian adaptive decisions can be based on Bayesian quantities (a posterior probability or predictive probability of trial success) or use a Bayesian final analysis

VANDERBILT
UNIVERSITY

Types of adaptive features

- early stopping/ sample size re-estimation
- treatment group randomization weights (adaptive randomization)

VANDERBILT
UNIVERSITY

Basic process of an adaptive design

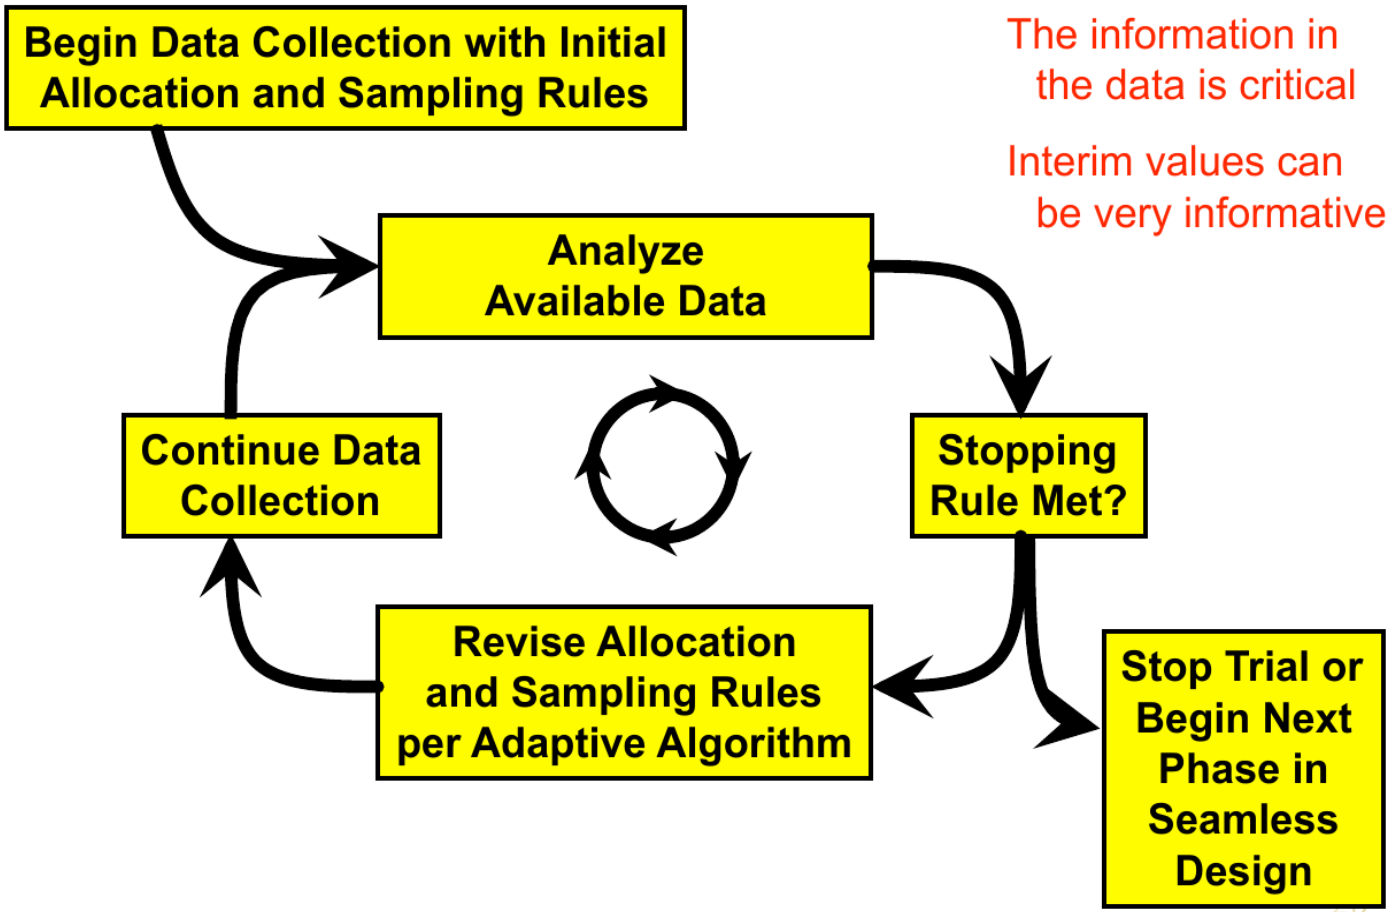

VANDERBILT UNIVERSITY

Image stolen from Jason Connor

Purpose of adaptive designs

- more ethical treatment of patients: decrease number of patients and patient exposure time to ineffective therapies
- more efficient drug development
- better use of resources

VANDERBILT
UNIVERSITY

When to use adaptive designs

- patient response is quickly observed relative to the patient accrual rate
- large cost associated with each patient or with duration of study
- uncertainty about the minimum clinically significant difference (or any uncertainty in computing power)

VANDERBILT
UNIVERSITY

Challenges of adaptive designs

- require specialized statistical expertise and large time investment
- need more time in planning and development phase
- can require very prompt or real time data entry
- can require very good coordination between study sites
- regulatory approval
- can yield biased estimates

VANDERBILT
UNIVERSITY

How to compare different designs

operating characteristics

- power
- probability of type one error
- expected sample size (E(N)), maximum sample size, variablility of E(N)
- probability of early termination (PET) for futility, for efficacy under different values of θ
- practicality (including number of looks)

VANDERBILT
UNIVERSITY

How to find operating characteristics of a particular design

- Calculate analytically
- Simulate

Simulations can help identify the effects of different:

- priors
- accrual rates
- decision rules/ cut offs
- effect sizes
- nominal alpha levels (if doing a frequentist final analysis)
- maximum sample sizes

VANDERBILT
UNIVERSITY

How to find operating characteristics of a particular design via simulation

To find . . .

P(type I error):

- Simulate data under $\theta = \theta_0$.
- Find percent of trials that declare success/reject H_0.

power under different values of θ_1:

- Simulate data under $\theta = \theta_1$.
- Find percent of trials that declare success/reject H_0.

VANDERBILT UNIVERSITY

Bayesian quantities

prior $p(\theta)$

sampling distribution/data likelihood $p(\underline{X}|\theta)$

posterior distribution $p(\theta|\underline{X})$

posterior predictive distribution $p(\tilde{\underline{X}}|\underline{X})$

predicted probability of success

$P(\text{success at final analysis}|\underline{X})$

VANDERBILT
UNIVERSITY

Scenario

Suppose . . .

- We have a treatment and want to evaluate its effect on a binary outcome.
- On the current standard therapy, 15% of patients have the response.
- We hope that with the new treatment, the response rate will be higher, like 35% or 45%.
- We want to test this drug in one-arm experiment to see if it warrants further study.

VANDERBILT
UNIVERSITY

Specific design for this scenario

- Final analysis based on Bayesian posterior probability: $P(p > p_0|X)$.
- If $P(p > p_0|X) >$ cut off (success), we will do a randomized study with this treatment and the standard treatment.
- Will use a maximum sample size of 19.

VANDERBILT
UNIVERSITY

Specific design: prior distribution

- Recall we have a binary outcome.
- Can assume each patient's outcome has a binomial distribution. (This is our sampling distribution.)
- The Beta distribution is conjugate with binomial sampling distributions. $\rightarrow$ The posterior distribution will also have a Beta distribution

VANDERBILT
UNIVERSITY

Specific design: prior distribution

- $p \sim Beta(0.15, 0.85)$
- $E(p) = \frac{\alpha}{\alpha+\beta} = \frac{0.15}{0.15+0.85} = 0.15 = p_0$
- Can interpret the info in a Beta prior as having observed a prior sample of $\alpha + \beta$ (One in this case)

VANDERBILT
UNIVERSITY

Specific design: posterior distribution

- With this prior, the posterior distribution of p given the data is

$$p|X \sim Beta(\alpha + x, \beta + N - x)$$

- So the final analysis will just be a probability of a beta distribution:

$$P(p > p_0|X) > \text{cut off}$$

VANDERBILT UNIVERSITY

Specific design for this scenario

- Opportunity to stop the trial early for futility or efficacy at fixed interim point(s).
- Time of interim looks are based on number of observations.

VANDERBILT
UNIVERSITY

Specific design: efficacy stopping rule

- Criteria for stopping early efficacy based on **posterior probability** of superiority.
- Will stop and declare success if

$$P(p > p_0|X) > \text{cut off}$$

at any interim look.

VANDERBILT UNIVERSITY

Specific design: futility stopping rule

- Criteria for stopping early for futility based on Bayesian **predicted probability of success**, based on the number of successes observed so far.
- Will stop and draw no conclusion if

$$P(\text{success at final analysis}|X) < \text{futility bound}$$

at any interim look.

VANDERBILT
UNIVERSITY

Specific design

- Will do futility and efficacy analysis after observing 10, 13, and 16 outcomes.
- Need to specify cut off for efficacy and the futility bound before trial begins
- Choose these to get good operating characteristics

VANDERBILT
UNIVERSITY

How the design performs

efficacy cut off = 0.95, futility bound = 0.05

p_{true}	Pr(win)	E(N)	sd(N)	PET Futility	PET Efficacy
0.15	0.123	12.258	3.181	0.816	0.063
0.35	0.806	12.765	3.658	0.179	0.614
0.45	0.950	11.469	2.893	0.050	0.856

efficacy cut off = 0.97, futility bound = 0.05

p_{true}	Pr(win)	E(N)	sd(N)	PET Futility	PET Efficacy
0.15	0.054	12.198	2.760	0.924	0.039
0.35	0.678	13.453	3.082	0.307	0.566
0.45	0.904	12.264	2.824	0.093	0.830

efficacy cut off = 0.95, futility bound = 0.15

p_{true}	Pr(win)	E(N)	sd(N)	PET Futility	PET Efficacy
0.15	0.096	10.989	2.617	0.867	0.060
0.35	0.708	11.652	3.152	0.282	0.593
0.45	0.890	10.927	2.340	0.109	0.837

How the actual trial would work

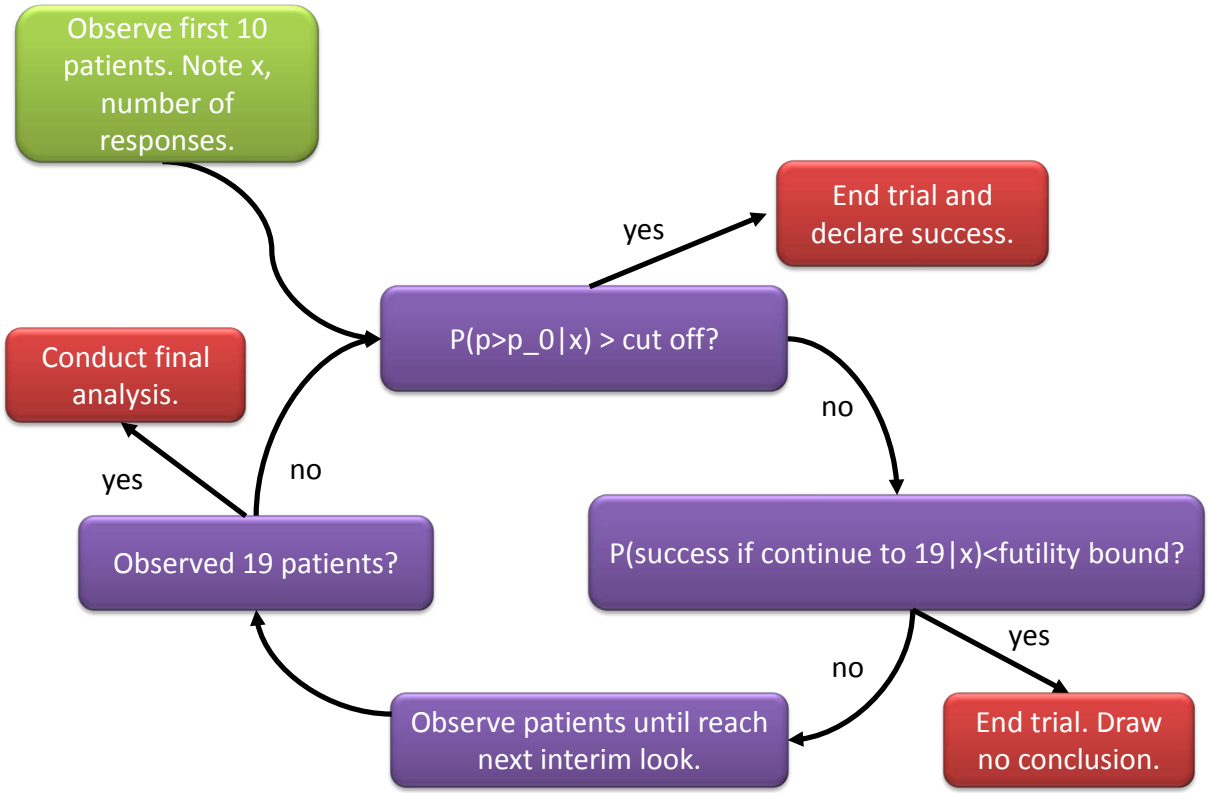

VANDERBILT UNIVERSITY

Basic simulation recipe

1. make assumptions about distributions, parameters, etc.
2. generate data set under those assumptions
3. record test statistic
4. repeat one billion times
5. look at distribution of the billion test statistics

VANDERBILT
UNIVERSITY

Tips for simulations in R

- make everything into functions
- pre-allocate rather than growing incrementally
- work in vectors rather than data frames
- be sure your number of simulations is adequate
- when you have loops, make sure everything that can be done outside the loop **is** outside the loop

VANDERBILT
UNIVERSITY