BAYESIAN ANALYSIS IN SMALL CLINICAL TRIALS

John Scott, Ph.D.
FDA/CBER/OBE/DB
The Science of Small Clinical Trials
November 27, 2012

Disclaimer

The findings and conclusions in this presentation have not been formally disseminated by the Food and Drug Administration and should not be construed to represent any Agency determination or policy.

Rare diseases at CBER

- CBER regulates vaccines, blood products and cell, tissue and gene therapies
- Therapeutic blood products for rare deficiency syndromes
  - Clotting disorders
  - Immune disorders
- Cell and gene therapies for single-gene defects and other rare diseases
  - Rare cancers
  - Hemoglobinopathies

Small clinical trials at CBER

- Outside of vaccines, more the rule than the exception
- Some recent approvals:
  - Ceprotin (protein C concentrate): open label historically-controlled study in 18 subjects
  - RiaSTAP (fibrinogen concentrate): accelerated approval based on clot firmness in 14 subjects
  - Corfact (FXIII concentrate): PK in 14 subjects
- Also cord blood, scorpion anti-venom, IGIV, ATryn, etc.

Bayesian analysis: The big picture

- Bayesian analysis provides a framework for:
  - Leveraging existing data
  - Synthesizing evidence of different types
  - Learning as we go
  - Estimating things we actually care about
    - e.g., the probability that a treatment works
- Bayesian analyses have formed the basis for a number of device approvals
- Only one biologics approval (to my knowledge)

Interpretations of probability

- The Frequency interpretation:
  - The probability of an event is the long-term frequency with which it occurs
  - Held by "Frequentists"
- The Subjectivist interpretation:
  - The probability of an event is the degree of belief a rational person has that the event will occur / has occurred
  - Held by (some) "Bayesians"
- Individual stances often malleable

Some probabilities

- The probability that a coin flip will be heads
- The probability that it will rain on Saturday
- The probability that JFK was killed by a lone gunman
- The probability that Drug A will work better than Drug B for your patient

The Bayesian approach

- All relevant information should be used in an analysis
- Acquired data can be used to continuously update our degree of belief
- Emphasis on making optimal decisions given available information, rather than testing hypotheses

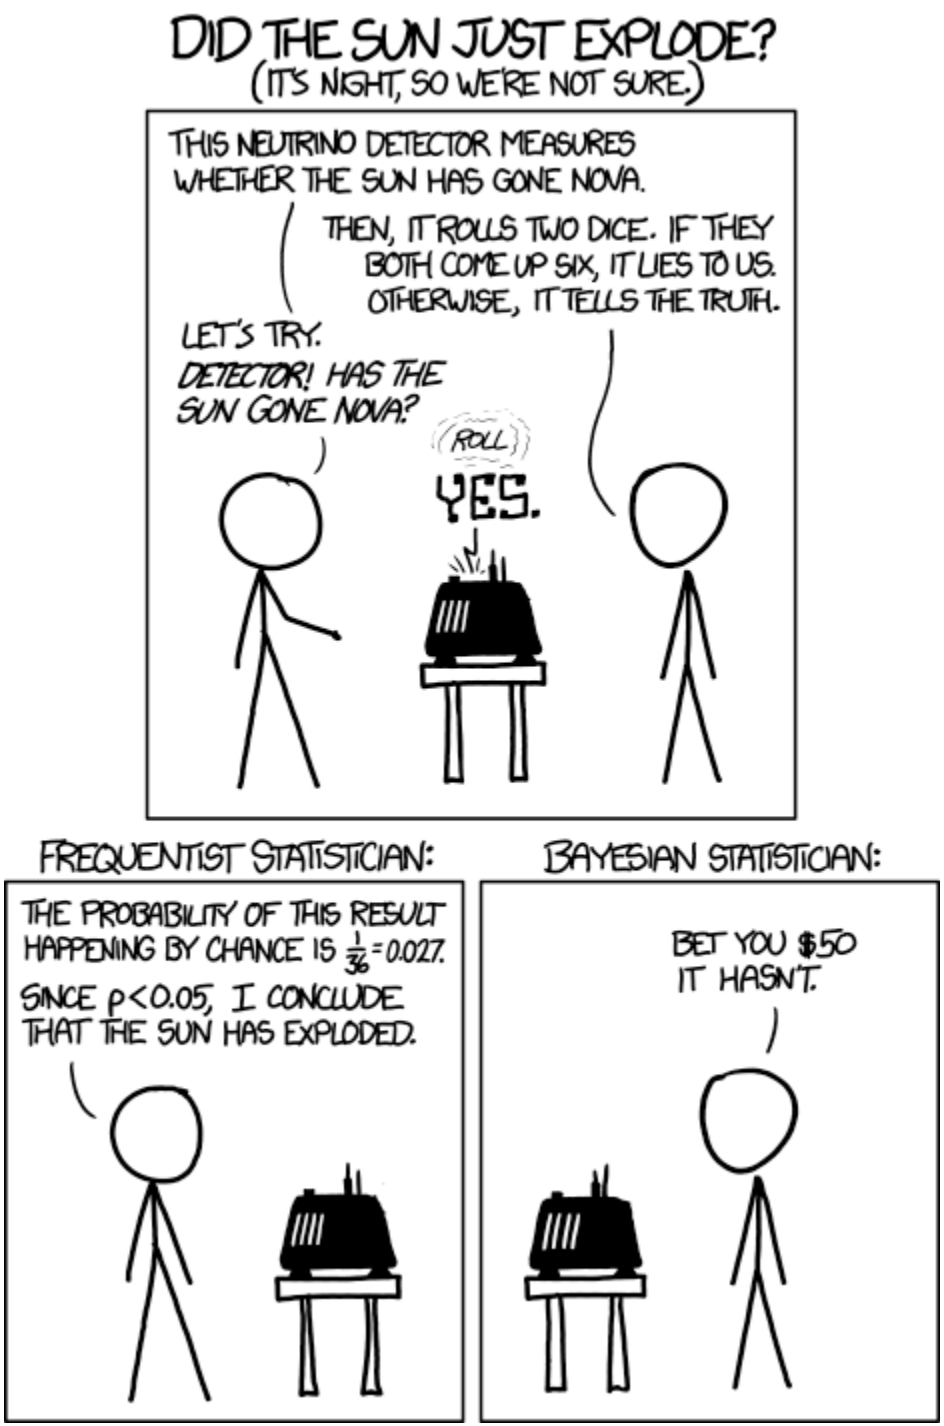

xkcd.com/1132

Steps in a Bayesian analysis

1. Summarize all relevant prior information with a *prior probability distribution*
2. Collect data
3. Combine the prior and the data into an updated *posterior probability distribution*
4. Learn from the posterior, possibly make a decision
5. Make the posterior into a new prior and repeat steps 2 – 4

Bayes' theorem

- Bayes' Theorem is a tool for calculating inverse probabilities:

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- Input:
  - A model for how a hypothesis generates data
  - Prior probability of the hypothesis
  - Something unimportant
- Output:
  - The probability that the hypothesis is true given the data

Bayes in small clinical trials

- From IOM *Small Clinical Trials* monograph:
  - Problem formulation
  - Sequential analysis
  - Meta-analysis
  - Prediction
  - Communication
- These boil down to two Bayesian strengths:
  - Evidence synthesis (overcome small trial lack of data)
  - Interpretability of conclusions

“Pivotal” Bayesian analyses

- When is a trial a ‘win’ for the product?
- In a Frequentist analysis, a win is $p < \alpha$
  - Preserves Type I error rate at α
- A Bayesian win is based on posterior tail probabilities
  - E.g. $Pr(S_T > S_C) > .95$
  - No direct relationship with Type I error rate
  - Depends on the prior
- Relatively few examples of this

Hemophilia A

- Hemophilia A is a rare bleeding disorder caused by a deficiency in the clotting protein, Factor VIII
- Standard therapy is FVIII replacement (plasma-derived or recombinant)
- Generally very effective as prophylaxis and on-demand treatment
- A typical FVIII product clinical development program:
  - Phase I PK / preliminary safety study (n = 10-20)
  - Phase III Safety / efficacy study (n = 80 – 104)

FVIII safety concerns

- The major safety concern with new FVIII products is neoantigenicity
  - Patients can develop neutralizing antibodies ("inhibitors") to the new FVIII protein
  - Can dramatically decrease hemostatic efficacy of FVIII
- Phase III trial size is driven by the need to demonstrate low inhibitor formation rate
  - Upper bound of 95% CI for inhibitor rate < 6.8%
  - Satisfied by ≤1/80 or ≤2/104 (events/subjects)

A Bayesian proposal for FVIII

- Lee and Roth [*Haemophilia* 2005] proposed a Bayesian approach to inhibitor rate analysis
- They chose a "relatively non-informative" prior distribution, Beta(0.3, 3.9)
  - Median: 2%
  - 95% CI: (0%, 32%)
- The observed inhibitor rate in a single-arm trial would then be combined with the prior to form a posterior distribution on inhibitor rate

FVIII example

- This is the Lee and Roth prior:

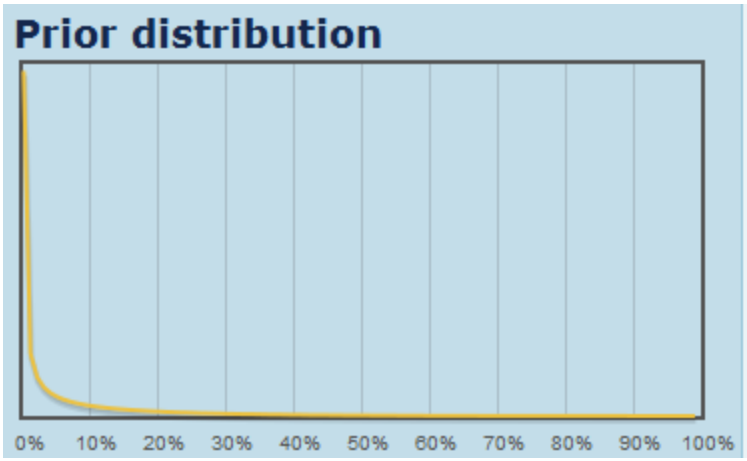

- Suppose we have 3/80 inhibitors in a study
- The likelihood looks like this:

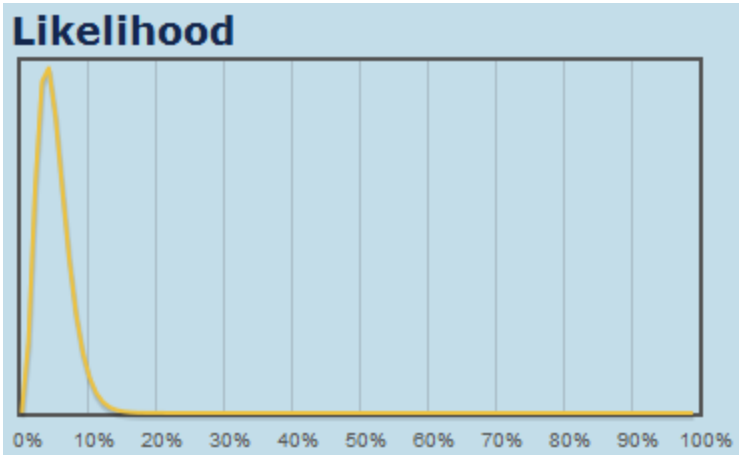

FVIII example (2)

- Using Bayes' Rule, we can combine the prior and the likelihood to form the posterior:
- But what do we do with the posterior?

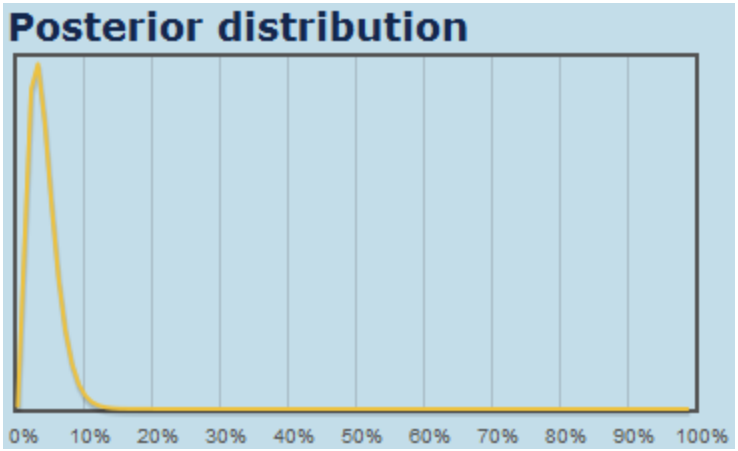

What is the posterior used for?

- Point estimation: 3.6% inhibitor rate (median)
- Interval estimation: 95% *credible interval* (0.9%, 9.0%)
- Tail probabilities:
  - Pr($r < 2\%$) = 17.6%
  - **Pr($r < 6.8\%$) = 90.3%** (a win?)
- Run a new experiment using this posterior as the prior

Applying this approach

- One FVIII product (Xyntha) used a related approach to support licensure
- Critical difference:
  - The prior data came from previous studies of the same or predecessor products
- The historical data was down-weighted by 50% relative to the pivotal study
  - Why 50%?
  - Usually better to let the data guide the amount of borrowing

Where do priors come from?

- From your most recent posterior
- From the literature
- From case review
- From experts (elicitation)
- From nowhere
  - Default / non-informative priors
  - Priors for sensitivity analyses
    - Skepticism / optimism

Dealing with subjectivity

- Major criticism of Bayes: subjectivity
  - Different priors = different posteriors
- To deal with this:
  - Carefully justify priors
  - Explore posteriors under a variety of priors ("sensitivity analysis")
  - Use non-informative priors
    - Often equivalent to Frequentist methods

Sensitivity to the prior

- The posterior depends on the prior
  - But how much?
- The posterior is a *compromise* between the prior and the data
- Lots of data and/or lots of uncertainty in the prior:
  - The data will dominate the posterior
- Very little data and/or very little uncertainty in the prior:
  - The prior will dominate the posterior

A “noninformative” prior

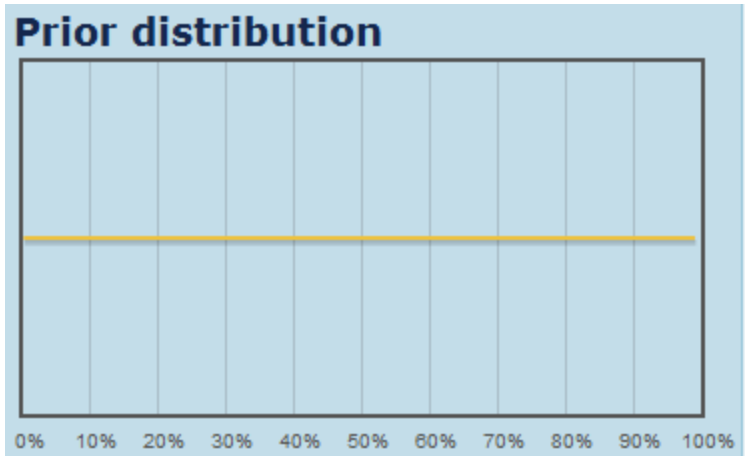

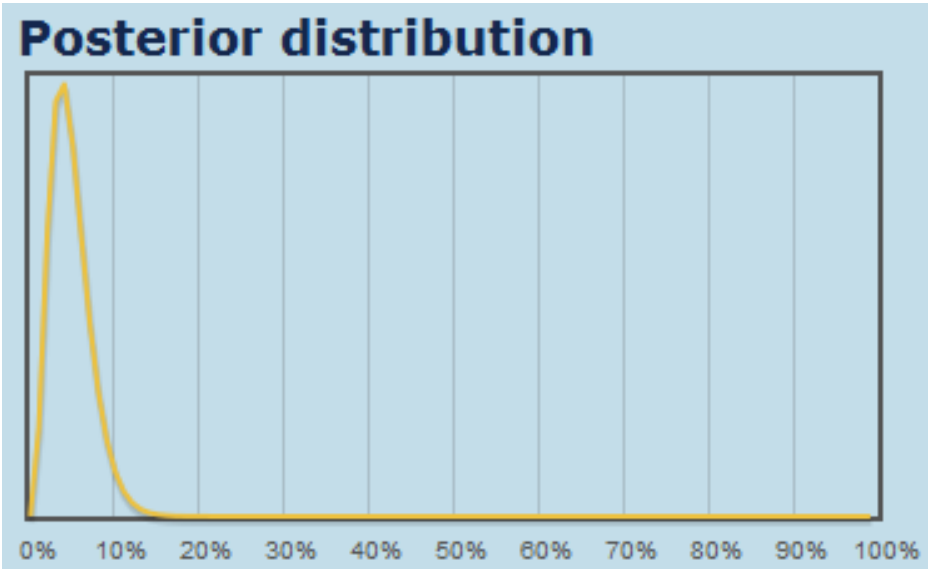

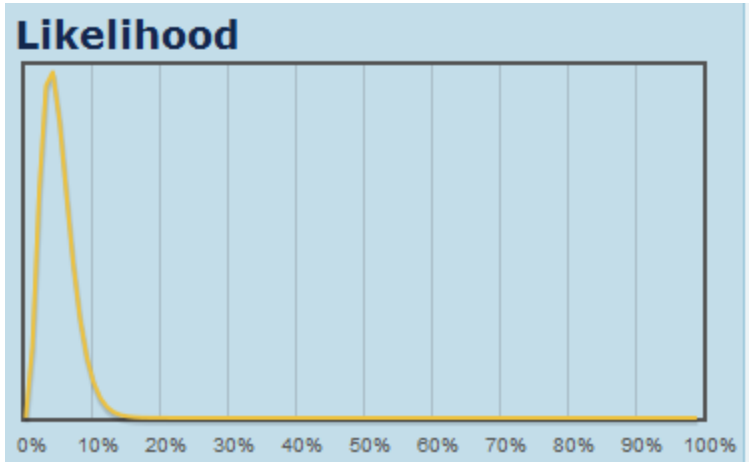

An incorrect prior

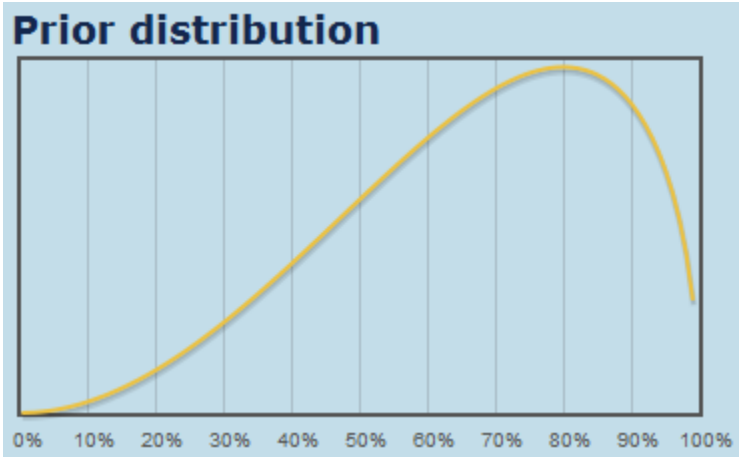

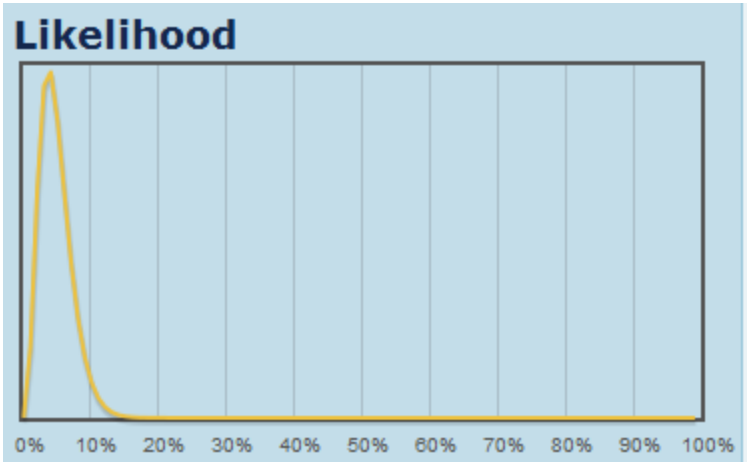

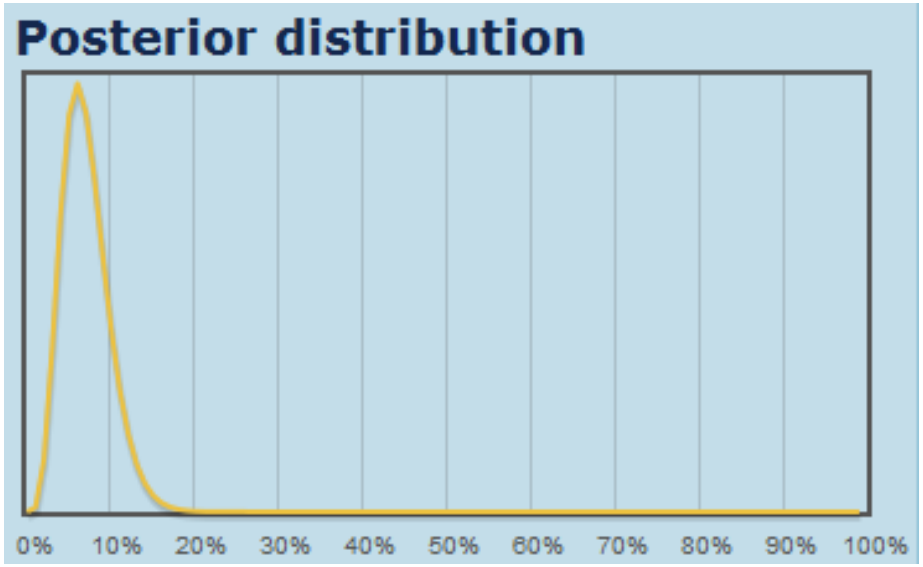

An extremely incorrect prior

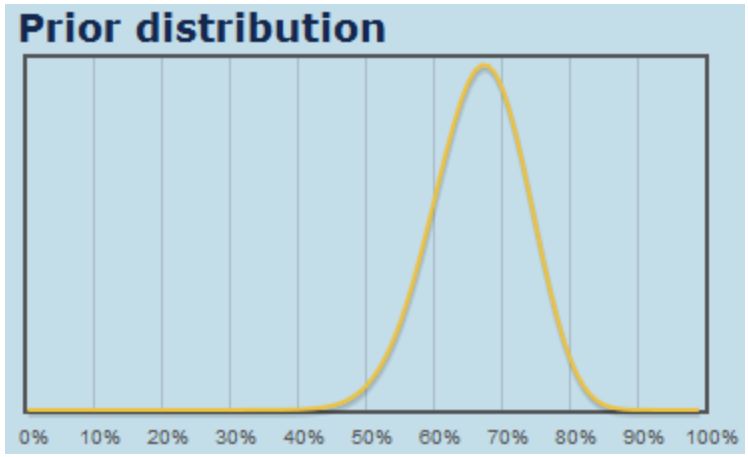

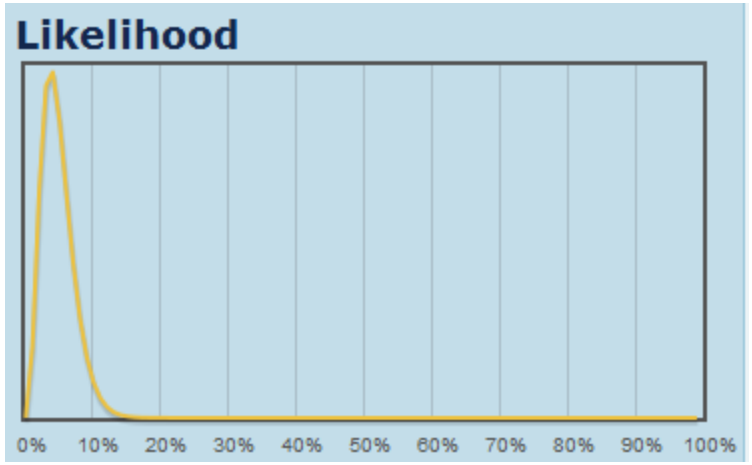

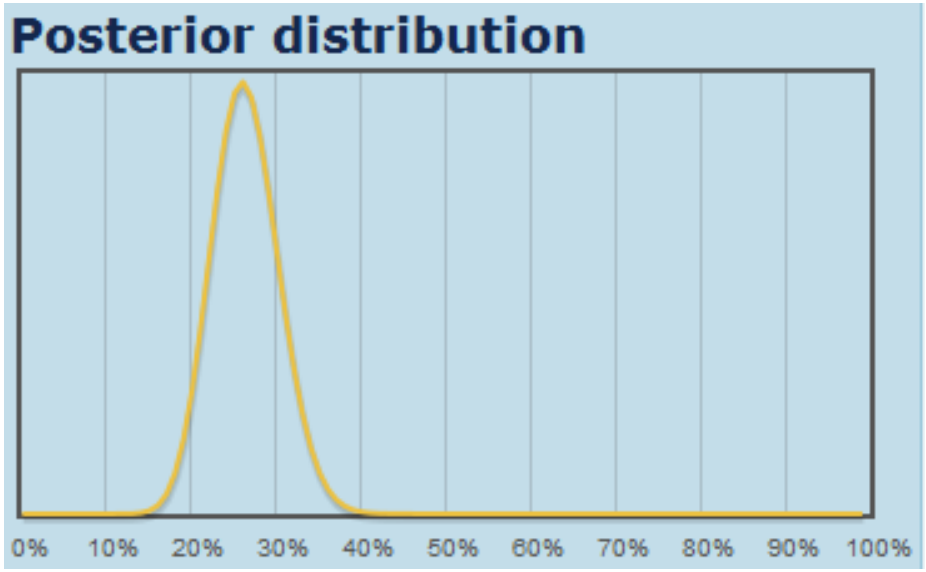

Exploratory Bayes in small clinical trials

- Bayesian methods can also be used for exploratory analyses:
  - Previous data can be leveraged to help understand what's in front of us
  - Bayesian interpretations are very nice for things like understanding safety signals
- BayesWeb.com is a software tool for non-statisticians to explore these ideas

Interpreting safety signals

- Suppose an AE of concern has occurred in a small clinical trial
- Investigators and regulatory medical officers need to make a decision based on potential risk to future subjects and patients
- Generally an informal application of expert opinion
- A rough quantitative understanding of risk might help
  - This sounds like a Bayesian exercise

A tool for exploratory Bayesian safety analysis

- Bayesweb.com puts simple Bayesian machinery in the hands of physicians
  - Easy to access and use
  - Sufficiently general to handle a wide variety of potential events and prior states of belief
  - Doesn't require statistical feedback
  - Unintimidating (cf. BUGS, R, SAS)
- Didactic goal: Wider understanding and acceptance of Bayesian methods

BayesWeb.com

- Web application to maximize accessibility and (hopefully) minimize intimidation
- Users can "self-elicit" prior distributions and explore their choices numerically and graphically
- Experimental data are summarized with familiar Frequentist statistics
- Posterior can then be explored numerically and graphically to help understand risk
- Users can explore simple prior sensitivity analyses

The model

- BayesWeb assumes events follow a binomial distribution:

$$f(x \mid \theta) = \binom{n}{x} \theta^x (1-\theta)^{n-x}$$

- Prior is taken from beta family:

$$f(\theta) = \frac{1}{\mathrm{B}(a,b)} \theta^{a-1} (1-\theta)^{b-1}$$

- Posterior is then also in the beta family:

$$f(\theta \mid x) = \frac{1}{\mathrm{B}(a+x, b+n-x)} \theta^{a+x-1} (1-\theta)^{b+n-x-1}$$

Prior rationale

- Conjugate prior limits flexibility, but:
  - Computationally feasible with client-side scripting (MCMC would not be)
  - Admits to wide range of simple elicitation techniques
- The website provides seven options for eliciting priors
- Elicitation methods need to be:
  - Simple enough to be easily understood
  - General enough to apply to any binomial parameter estimation problem

Prior elicitation 1: By moments

- Users specify center of distribution by mean or mode
- Spread entered by variance, s.d. or credible interval
  - C.I. can be arbitrary probability
  - 1- or 2-sided (equal tails)
  - S.D. is estimated from C.I. by normal approximation

- Prior fit by moments: $a = \frac{\mu^2 - \mu^3}{\sigma^2} - \mu$, and $b = \frac{a}{\mu} - a$,

Prior elicitation 2: Interval probabilities

- Unit interval split into 5, 10 or 20 equally-spaced intervals
- Users provide prior belief that parameter lies in each interval
- Normalized if necessary
- Beta prior fit by moment matching with the induced discrete distribution on the midpoints of the intervals

Prior elicitation 3: Historical data

- Users provide:
  - Number of events (x_h)
  - Number of subjects (n_h)
  - Optional down-weighting factor (w)

- Beta prior fit as $a = wx_h$, $b = w(n_h - x_h)$

Prior elicitation 4: Hypothetical subjects

- Following Berry & Stangl, users are asked to give:
  - Probability of an event for 1st hypothetical subject; interpreted as prior mean, $\mu = \frac{a}{a+b}$
  - Conditional probability for 2nd hypothetical subject interpreted as mean of an updated prior, $\mu^+ = \frac{a+1}{a+b+1}$
- Beta prior reconstructed as:

$$a = \frac{\mu(1-\mu^+)}{\mu^+ - \mu}, \quad b = \frac{(1-\mu)(1-\mu^+)}{\mu^+ - \mu}$$

Prior elicitation 5: Prior mode & ESS

- Following Sambucini, users are asked for
  - Their "best guess," interpreted as prior mode, *m*
  - "How many subjects worth of data" the guess is based on, n_0
- Prior is fit as:

$$a = n_0 m + 1, \quad b = n_0 (1 - m) + 1$$

Other prior options

- More than one of these methods can be used simultaneously
  - Prior is formed as compromise
- Users can also request a non-informative prior
  - Uniform or Jeffreys
- Beta parameters can be input directly

Technical details

- BayesWeb runs by client-side scripting
  - Responsive
  - Can be run offline (theoretically)
  - Computationally limited
- Written entirely in HTML, CSS and JavaScript
- JavaScript libraries used:
  - jQuery / jQuery UI
  - Flot
  - PragMath

Example

- Suppose one inhibitor is observed in a 20 subject PK study of a new FVIII product
- Use exploratory analyses to answer:
  - How much risk does this represent to future subjects?
    - $Pr(\theta > .068)$
  - Should the clinical development be paused / stopped?
    - $Pr(\theta > .04)$
- Potential priors:
  - Beta(0.3, 3.9) [Lee & Roth, *Haemophilia* 2005]
  - 2% mode, 90% CI (0%, 20%)

Mid-trial Bayesian analysis

- In a Phase III trial, if inhibitors occur early, less likely trial will be a success
- Mid-trial analysis:
  - Start with a prior on inhibitor rate
  - When an inhibitor is observed, calculate a posterior beta distribution of the inhibitor rate
  - This distribution can be used to calculate:
    - Tail probabilities of clinical interest
    - Probability of win (using beta-binomial distribution)

Fixed sample example

- In a 104 subject design (2 inhibitors allowed)
- Suppose first inhibitor is observed after subject #40 enrolled
- With Beta(0.3, 2.9) prior, we have Beta(1.3, 41.9) posterior
  - $\Pr(\theta > .068) = 9\%$
  - $\Pr(\theta > .02) = 56\%$
  - $\Pr(\text{win}) = 54\%$
- If first inhibitor is observed after subject #10 enrolled:
  - $\Pr(\theta > .068) = 56\%$
  - $\Pr(\text{win}) = 13\%$

Futility stopping example

- Start with a Beta(0.5, 0.5) prior
- Set futility threshold at win probability < 10% at time of first observed inhibitor
  - For $\theta = 10\%$, Pr(stopping) = 72%
  - For $\theta = 6.8\%$, Pr(stopping) = 57%
  - For $\theta = 2\%$, Pr(stopping in error) = 13%
  - For $\theta = 1\%$, Pr(stopping in error) = 9%
- Stop is usually quite early
- Oversimplification – can it be improved by incorporating observation time?

Limitations

- Simple parametric prior
  - Mixtures better? How to elicit?
- Simple model
  - Incorporate temporal association, dose-response, etc.?
- Single-arm
- Retrospective elicitation
- Not user-friendly enough
- Too user-friendly?

Acknowledgements

- Austin Hand
- Layla Sian

- cf. Scott, Hand & Sian (2011), *JBS* 21, 1030-1041

Challenge questions

1. The process of getting an expert to express their knowledge as a prior distribution for Bayesian analysis is called A. solicitation, B. elicitation, C. recitation
2. Bayesian analyses are well-suited for small trials because A. they don't work well for large trials, B. small trials are usually pointless, so we might as well try something new, C. data deficits of small trials can be mitigated by incorporating past information