Introduction to the Bayesian continual reassessment method (CRM) for phase one clinical trials

JoAnn Alvarez, MA
joann.alvarez@vanderbilt.edu

Department of Biostatistics
Center for Quantitative Sciences
Vanderbilt University School of Medicine

2013 October 11

Overview

VANDERBILT
UNIVERSITY

Phase 1 goal

Find the right dose

- The higher the dose,
  - the more efficacious
  - the higher the toxicity

VANDERBILT
UNIVERSITY

- The higher the dose,
  - the more efficacious
  - the higher the toxicity
- We balance efficacy and toxicity by choosing a targeted toxicity level (TTL) (probability of toxicity)

VANDERBILT
UNIVERSITY

- The higher the dose,
  - the more efficacious
  - the higher the toxicity
- We balance efficacy and toxicity by choosing a targeted toxicity level (TTL) (probability of toxicity)
- Then find dose with targeted toxicity level: MTD, maximum tolerated dose

VANDERBILT
UNIVERSITY

Phase 1 design

Goal is to

- get information about the best dose
- treat each patient ethically

VANDERBILT
UNIVERSITY

Phase 1 design

Goal is to

- get information about the best dose
- treat each patient ethically: with the treatment best supported by the current evidence

VANDERBILT
UNIVERSITY

Design of phase one study involves best way to acheive:

- get information about the best dose
- treat each patient with the treatment best supported by the current evidence

Main design issue is what dose to give each patient

Basic idea of CRM

- after each patient outcome is observed, dose-response relationship is re-estimated
- next patient is given the dose that is the current estimate of the MTD

VANDERBILT
UNIVERSITY

- patients arrive sequentially
- observation Y_j on each patient is whether they have a toxic response

Dose-response relationship

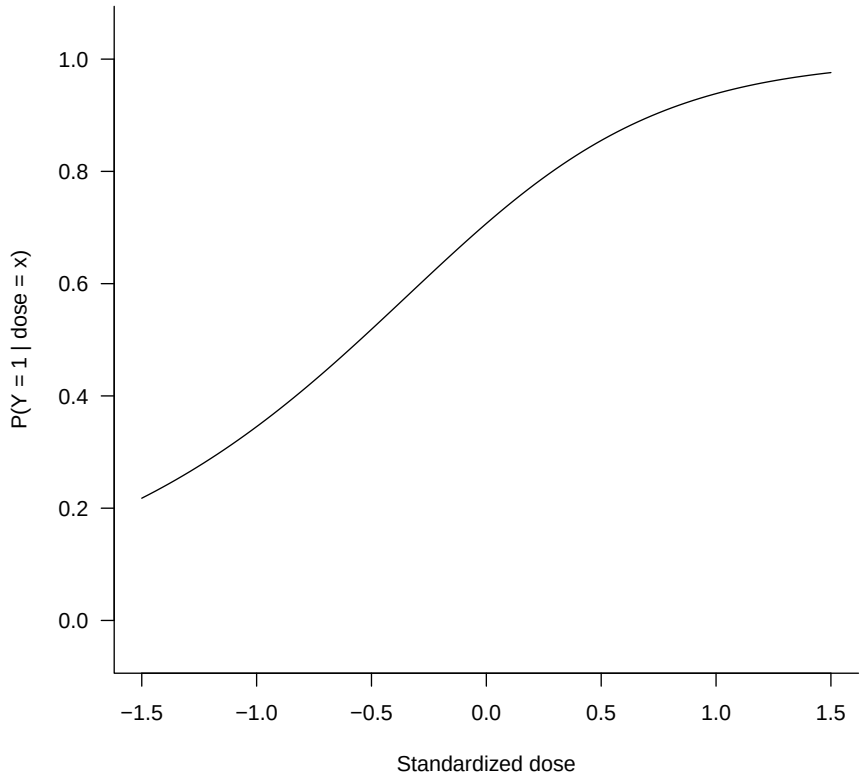

VANDERBILT
UNIVERSITY

- Let θ be the TTL.
- Objective is to find corresponding dose, x^* (MTD)

Dose-response model

- Need a model for the dose-response relationship.
- Choose any one-parameter function $\psi(x, a)$, monotonic in x and a.
- We assume that there exists an a_0: $\psi(x^*, a_0) = \theta$
- Can think of a_0 as a population parameter

VANDERBILT
UNIVERSITY

- since each patient gets the current best estimate of the MTD, more information is collected for values near the true MTD
- model not expected/required to be accurate at doses far from the MTD
- model expected to perform well near the MTD.
- only need one-parameter model

VANDERBILT
UNIVERSITY

Objective is to find the value of a that gives x^*

$$a \colon \psi^{-1}(\theta, a_0) = x^*$$

This value is a_0.

VANDERBILT
UNIVERSITY

The model

We assume

$$\forall \theta, \forall x^*, \exists! a_0 : \psi(x^*, a_0) = \theta.$$

VANDERBILT
UNIVERSITY

$$\forall\theta, \forall x^*, \exists! a_0 : \psi(x^*, a_0) = \theta.$$

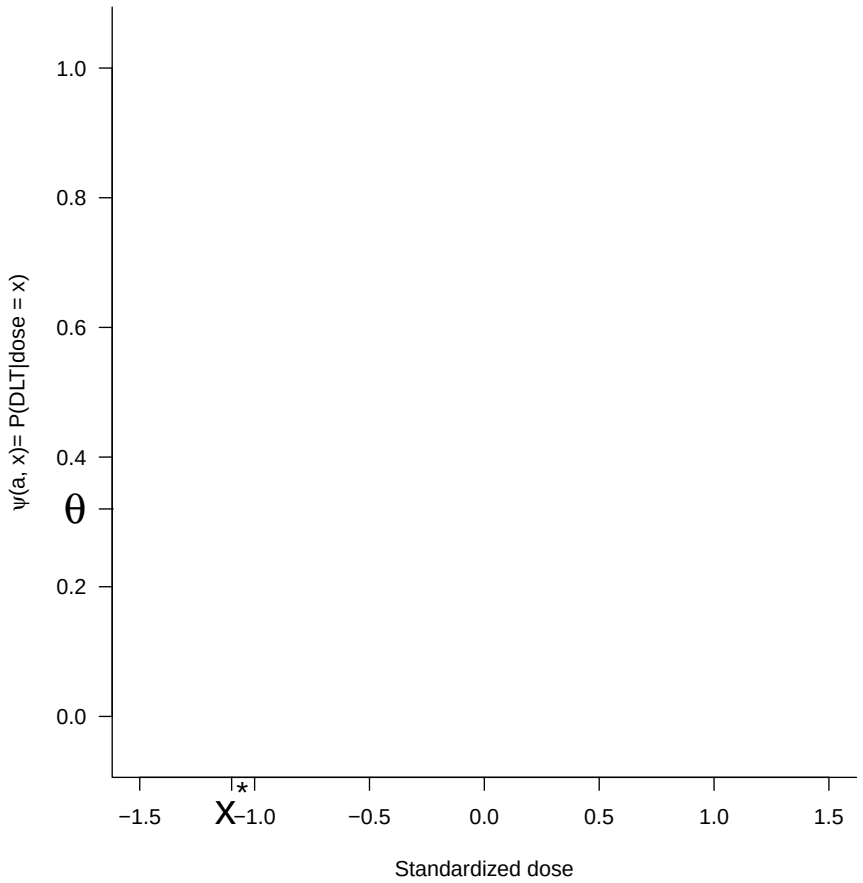

VANDERBILT
UNIVERSITY

$$\forall\theta, \forall x^*, \exists! a_0 : \psi(x^*, a_0) = \theta.$$

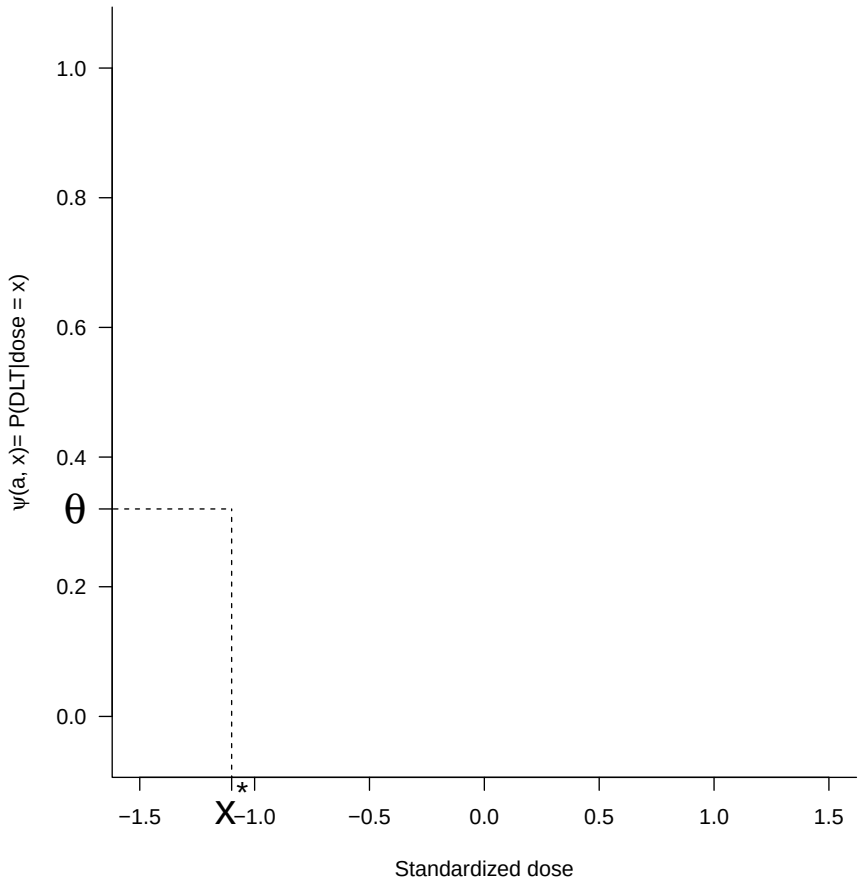

VANDERBILT
UNIVERSITY

$$\forall\theta, \forall x^*, \exists! a_0 : \psi(x^*, a_0) = \theta.$$

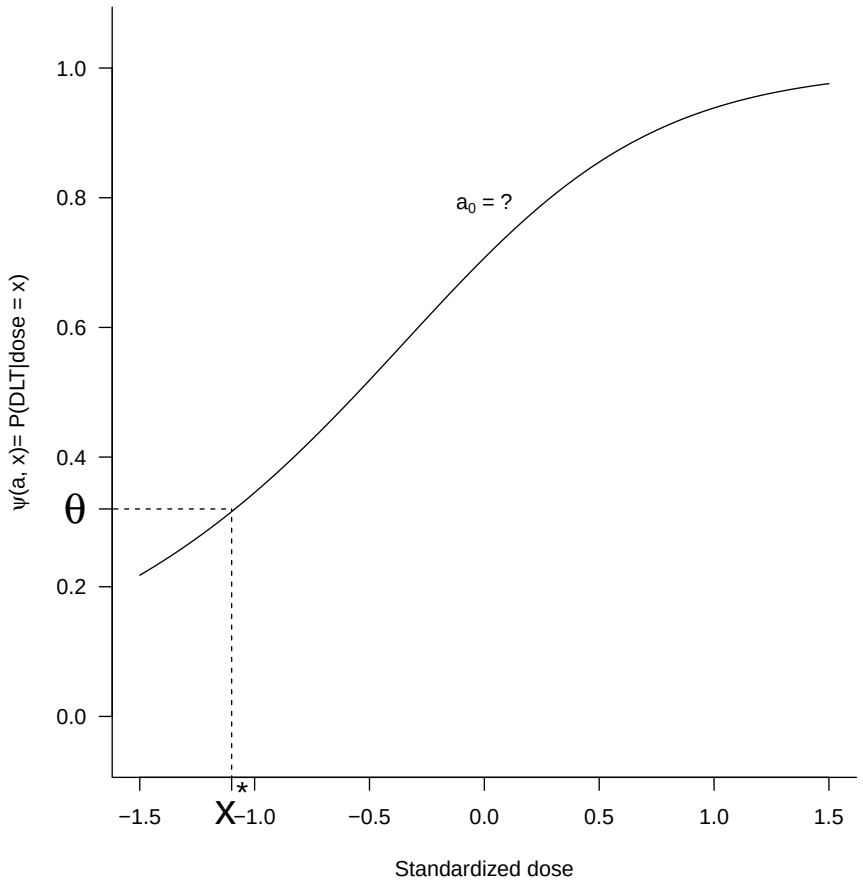

VANDERBILT
UNIVERSITY

Example dose-response curve family

$$\psi(x, a) = \frac{(\tanh(x) + 1)^a}{2^a}$$

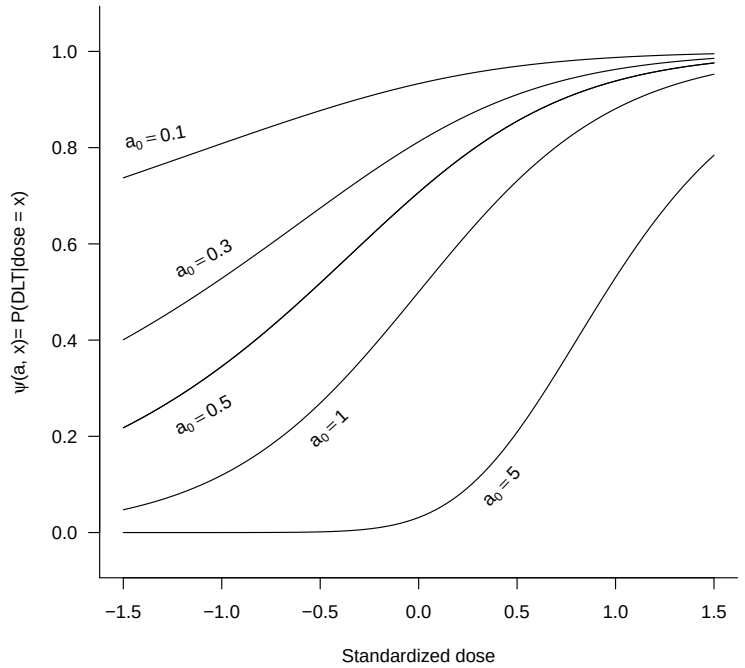

VANDERBILT
UNIVERSITY

Dose-response models

- logistic
- hyperbolic tangent
- power

VANDERBILT
UNIVERSITY

Dose-response models

Logistic model:

$$\psi(x, a) = \frac{e^{3+ax^*}}{1 + e^{3+ax^*}}$$

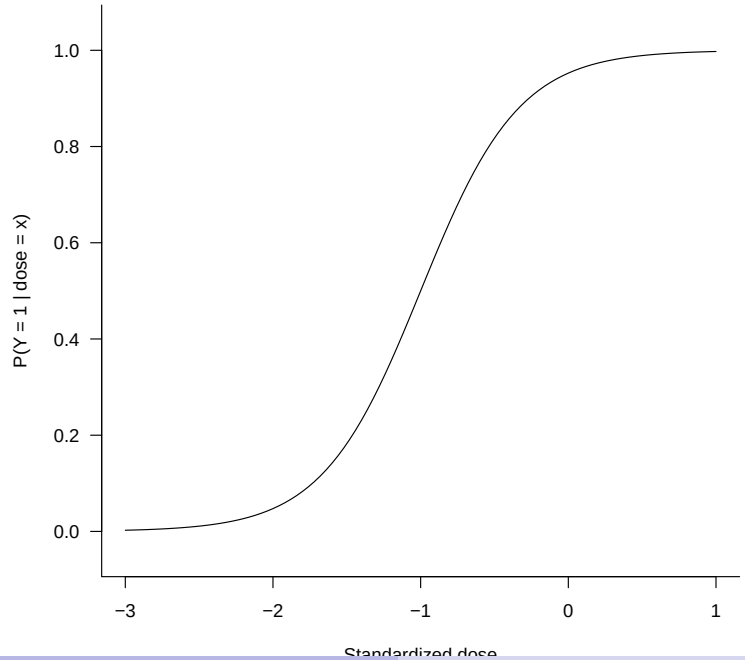

VANDERBILT
UNIVERSITY

- Suppose the j^{th} patient has enrolled and is ready to receive the treatment.
- The first $j-1$ patients have observed response data:

$$\begin{array}{cc} x(1) & Y_1 \\ x(2) & Y_2 \\ x(3) & Y_3 \\ \vdots & \vdots \\ x(j-1) & Y_{j-1} \end{array}$$

- Want to give patient j current best guess of MTD.

VANDERBILT
UNIVERSITY

Let the current posterior for a_0 be $f_{a_0}(a, \text{data})$

Ways to estimate P(DLT) for each dose

Estimate $\widehat{P(Y = 1|x = x_i)}$:

Plug-in estimator

Using the data from the first $j-1$ patients, calculate the posterior mean of a_0.
Plug it into ψ to get an updated dose-response curve.

Mean estimator

estimate the probability of toxicity it by its mean, integrating over all possible values of a_0 at each dose.

$$\widehat{P(Y = 1|x = x_i)} = \int_0^\infty \psi(x_i, a) f_{a_0}(a, \text{data})\, \mathrm{d}a$$

VANDERBILT
UNIVERSITY

Choose next dose

Now that we have updated $\widehat{P(Y = 1|x = x_i)}$, we have to choose the 'best' dose

Choose next dose

Patient j's dose, $x(j)$, will be

- dose that is closest to the the current estimate of the MTD: the dose that gives θ in the current estimate of ψ
- dose that gives the estimated P(DLT) closest to the TTL, θ

VANDERBILT
UNIVERSITY

Choose next dose

Patient j's dose, $x(j)$, will be

- dose that is closest to the the current estimate of the MTD: the dose that gives θ in the current estimate of ψ
- dose that gives the estimated P(DLT) closest to the TTL, θ

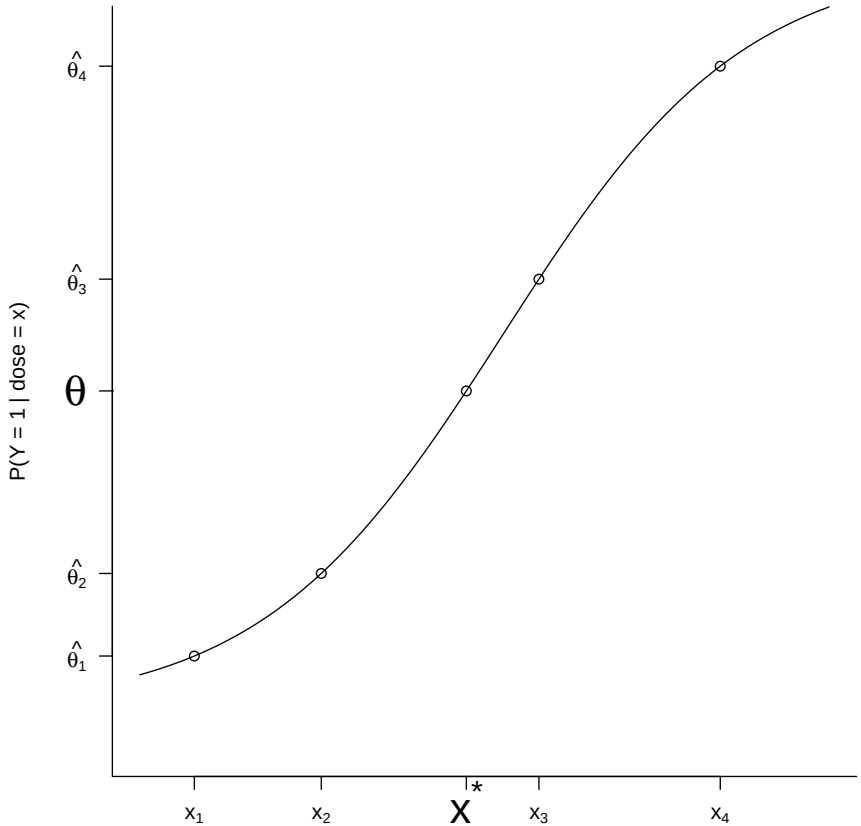

VANDERBILT UNIVERSITY

- Observe patient j's response, Y_j: either toxicity or no toxicity.
- Patient j's datum will be used to update the posterior distribution of a_0, $f_{a_0}(a, \text{data})$

VANDERBILT
UNIVERSITY

Likelihood function

For one patient:

$$P(Y_j = y_j) = (\psi(x_j, a))^{y_j}(1 - \psi(x_j, a))^{1-y_j}.$$

Based on the bernoulli pmf

VANDERBILT
UNIVERSITY

Bayes

posterior after observing j^{th} patient:

$$f_{a_0}(a, \text{data}_j) = \frac{(\psi(x_j, a))^{y_j}(1 - \psi(x_j, a))^{1-y_j} f_{a_0}(a, \text{data}_{j-1})}{\int_0^\infty (\psi(x_j, u))^{y_j}(1 - \psi(x_j, u))^{1-y_j} f_{a_0}(u, \text{data}_{j-1})\, \mathrm{d}u}$$

VANDERBILT
UNIVERSITY

Choice of prior

Common priors for a_0

- gamma
- uniform
- lognormal

Final estimate of MTD

- Can be estimated in same way as the dose for each patient is determined, since each patient is treated with the current estimate
- Determined after last patient is observed, if prespecified maximum n

VANDERBILT
UNIVERSITY

Variations

- modified CRM
- EWOC escalation with overdose control (Babb et al., 1998)
- Intervals of toxicity (Neuenschwander et al., 2008)

VANDERBILT
UNIVERSITY

Reference

O'Quigley, J., Pepe, M., and Fisher, L. (1990).
Continual reassessment method: A practical design for phase 1 clinical trials in cancer.
Biometrics, 46:43–48.

VANDERBILT
UNIVERSITY

Standardized doses

Calculated based on

- the probabilities of toxicity that the docs think the set of study doses have
- the assumed dose-response model
- an initial estimate of a_0

This info would be based on previous studies.

VANDERBILT
UNIVERSITY