# Adaptive designs in clinical trials

Francois Vandenhende, Ph.D.

Clinbay

francois@clinbay.com

UCL, March 6, 2009

# Outline

- Introduction to Bayesian adaptive dose-response trials
- Illustration of adaptive dose-response trials in
  - Phase I (safety)
  - Phase II (efficacy)
- Software
- Conclusions

# Definition: Adaptive Dose-Response Trials (AD-RT)

- Designs with prospectively defined dose-adapting plans.
- Controlled multi-arm clinical trials
- Planned interim analyses at selected times
- Outcomes at each interim:
  - Early stopping:
    - Futility
    - Efficacy
  - Adapt treatment allocation in next cohort

# Adaptive design process map

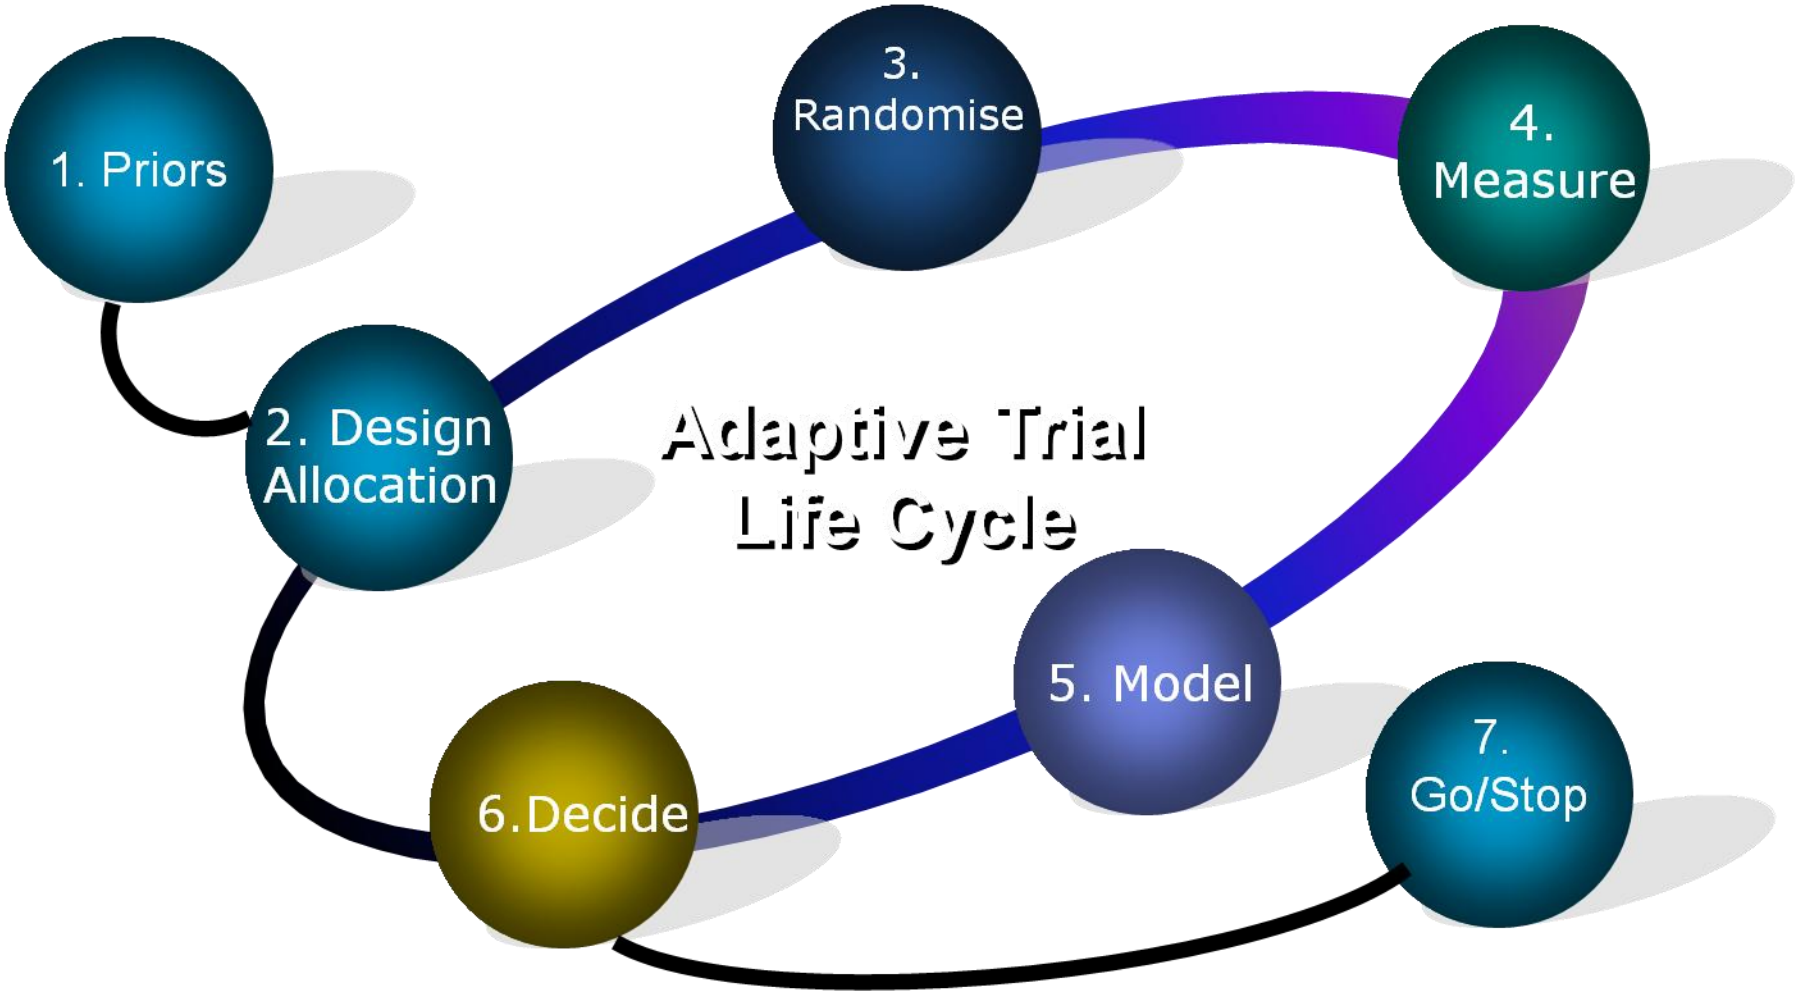

# Advantages and Challenges

## Advantages

- Test more treatments/ doses
  - Fewer retained at the end
- Early decisions
  - Ethical benefit
  - Accelerated development
- Limit risks of failed trials

## Challenges

- Statistical
  - Multiplicity/ Control of error rates
  - Predictions
- Trial Logistics
  - Access to data (EDC)
  - Drug Supply
- Feasibility
  - Recruitment rate slow relative to time to response.

# Bayesian adaptive design

- Relies Bayes theorem to summary treatment effects at any time:

$$p(\theta \mid y) \propto p(\theta)\, p(y \mid \theta)$$

- Mix of study data and prior information
  - Weight of likelihood increases with sample size
- Decisions based on :
  - Posterior probability of success/failure
    - intuitive and interpretable risk estimators
  - Posterior predictive distributions
    - E.g., Predictive power at final analysis.
  - Utility functions

# Decisions

- Posterior distribution:

$$p[g(\theta) \mid y]$$

  - Examples: g(θ) = drug effect vs pbo
    - Efficacy decision if

      Pr[g(θ)>ε|Y] is large (eg, >95%)

    - Futility decision if

      Pr[g(θ)>ε|Y] is low (eg, <5%)

# Decisions

- Posterior predictive distribution
  - Decisions based on predictions
  - Distribution of future responses:
    - Given current data, and
    - Unconditionally to any fixed parameter value.

$$p(y_{new} \mid y) = \int p(y_{new} \mid \theta, y)\, p(\theta \mid y)\, d\theta$$

  - Example: proportion of future patients with a response high enough should be large

$$\Pr[y_{new} > \varepsilon \mid y] > \tau$$

# Decisions

- Predictive power for a test at study completion
  - Interim data = Y
  - Sample size at completion = N
  - α is the type I error rate for the test
  - $\beta(\theta_0)$ is the type II error rate for the test at $\theta=\theta_0$
  - Then, predictive power is

$$1-\int \beta(N \mid \alpha, \theta, y)\, p(\theta \mid y)\, d\theta$$

# Adaptive allocator

- Utility-based
- Enrolment in cohorts
  - Flexible treatment arms : N patients among x arms
  - Fixed treatments (N/arm)
- Utility of flexible arms computed after each interim analysis (larger is better).

$$E(d) = \int U(d,\theta)\, p(\theta \mid y) d\theta$$

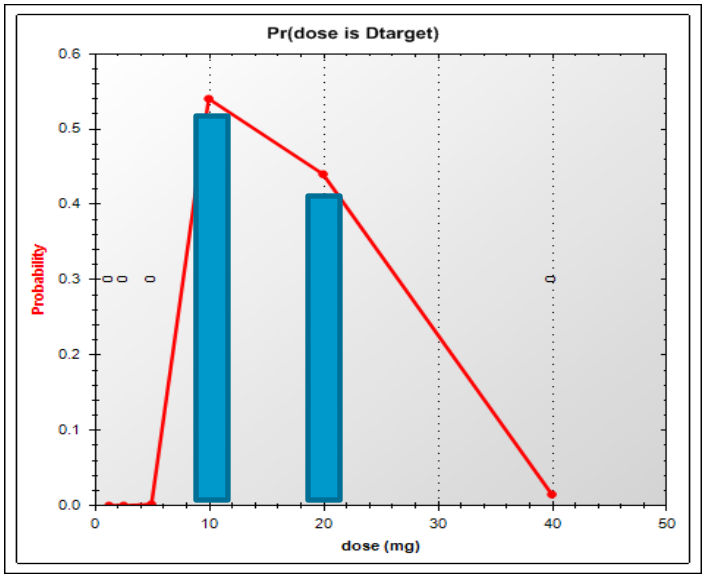


  - Examples
    - Variance of a parameter
    - Cost/Benefit ratio for dose
- Randomise proportionally to utility values.

# ILLUSTRATIONS IN:

- PHASE I (SAFETY)
- PHASE II (EFFICACY)

# ADAPTIVE DESIGN IN PHASE I

Dose-escalation methods:

- CRM : O'Quigley et al. (1990), Chevret (1993), Faries (1994), Goodman et al. (1995)…

- EWOC: Babb et al. (1998), Zacks et al (1998), Shih et al. (1999), Tighiouart et al (2005),…

# First Human Dose

- First-in-man, single dose escalation trial of a cancer product
- Low dose tested before escalating up
  - Cohort of 3 patients/dose
  - Interim safety review drives dose escalation & trial termination
- Dose-limiting toxicity (DLT)
  - Overall summary of subject's tolerability evaluation
  - Binary response (No DLT/ DLT) per subject.
- Goal of study:
  - To identify the maximum tolerated dose (MTD)
  - Dose at which DLT rate is not too large: 20 to 30%.

# Dose-response model

Logistic regression:

- Response = Yes/No

$$Y_{ij} \sim Bern(p_j)$$

$$\log\left(\frac{p_j}{1-p_j}\right) = \alpha + \beta * dose_j$$

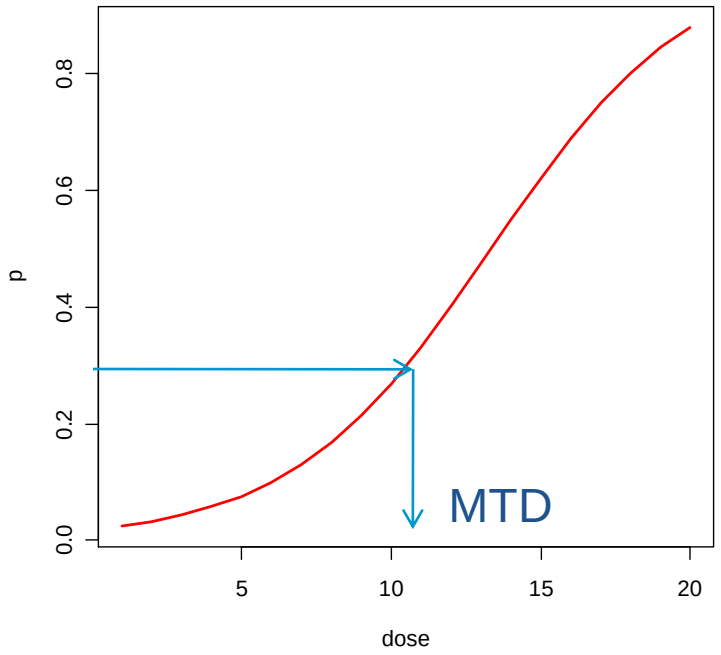


2 parameters:

- α : logit score under placebo (α=0 ⇔ p=50%)
- β : log-odds ratio for a change with dose. Monotonic increase with dose when β>0.

# Bayesian inference - Priors

Prior distribution of model parameters:

- Before study starts, elicitation of
  - α: Non-tolerability rate under placebo

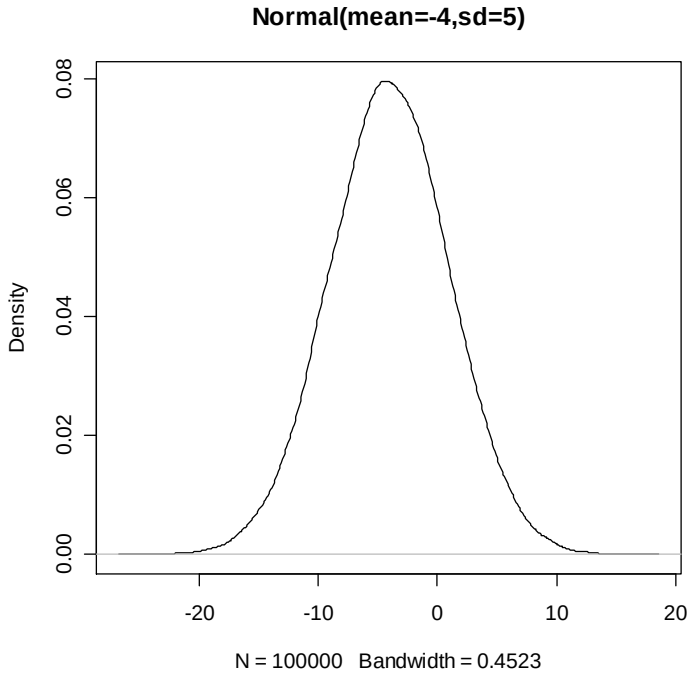


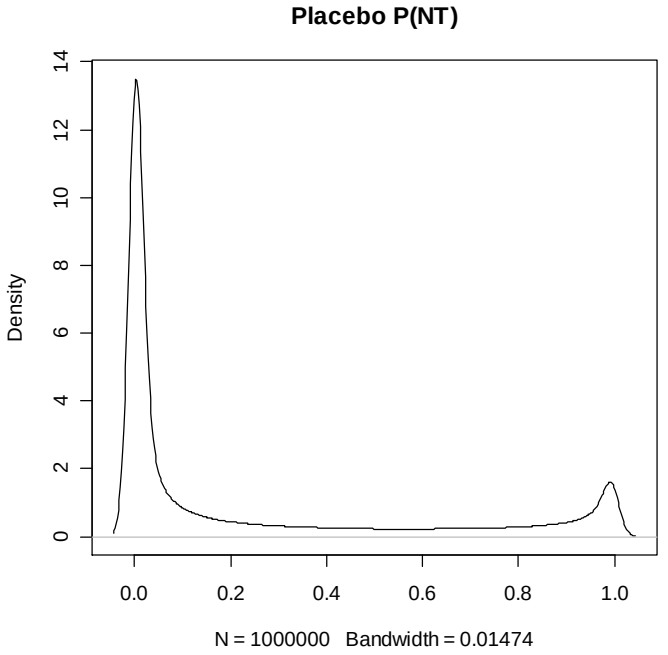

# Bayesian inference - Priors

– Exp(β): change in NT odds ratio when increasing dose by 1 unit.

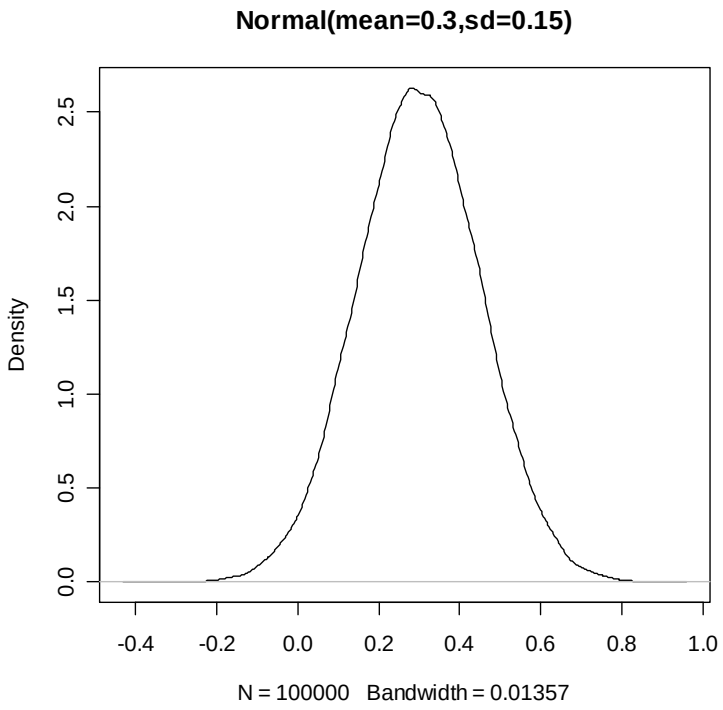


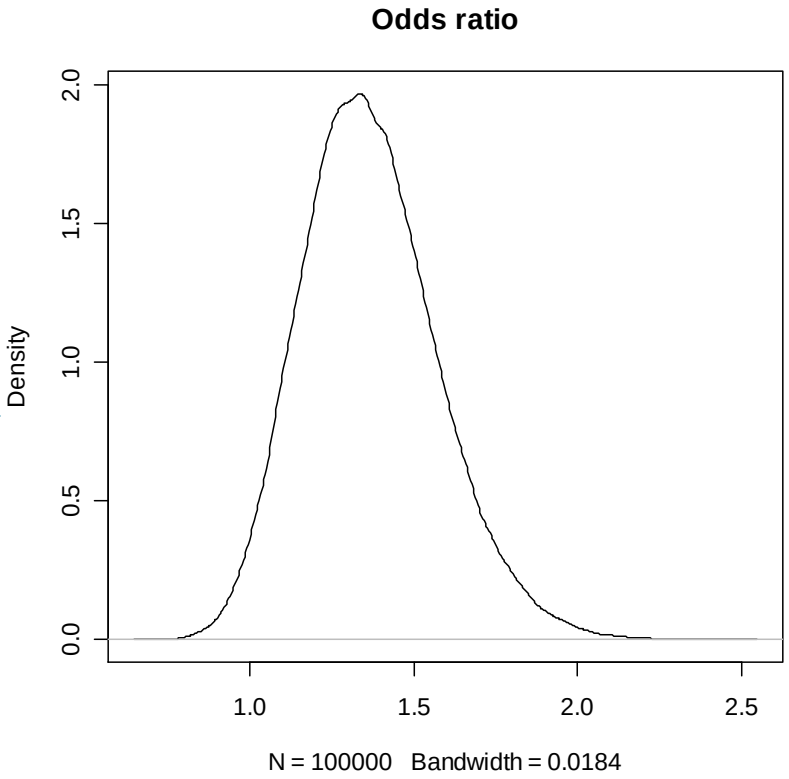

# Convenient reparameterization

- Dose in [Xmin, Xmax]
  - ρ = Pr[DLT|dose=Xmin]
  - γ = MTD
- Then,

$$\alpha = \frac{1}{\gamma - X\,\min}[\gamma \log it(\rho) - X\,\min \log it(0.3)]$$

$$\beta = \frac{1}{\gamma - X\,\min}[\log it(0.3) - \log it(\rho)]$$

# Bayesian Posterior Update

- Starts out of full conditionals of model parameters:

$$f(\beta \mid \alpha, Y) \text{ and } f(\alpha \mid \beta, Y).$$

- If not analytical form available, use known distribution & accept/reject samples (e.g., Metropolis/Hastings algorithm).
- Then, iterate on the following sampling scheme:

$$\beta_k \sim f(\beta \mid \alpha_{k-1}, Y)$$

$$\alpha_k \sim f(\alpha \mid \beta_k, Y)$$

- The chain converges when k is large to a random sample

$$\{\alpha_k, \beta_k\} \sim f(\beta, \alpha \mid Y)$$

# Model implementation in Winbugs

```
model{
# loop across subjects
for(i in 1:N) {
   logit(p[i]) <- alpha + beta*d[i]
   y[i] ~ dbern(p[i], n[i])
        }
   alpha ~ dnorm(-4,0.25)
   beta ~ dnorm(0.3,44)
}
```

# Trial Objective: MTD

- The objective of the trial is:
  - To estimate precisely the MTD

$$\hat{MTD} = \max d_j : \Pr(p_j <= 30\% \mid Y).$$

$$\hat{MTD} = \frac{\log(0.3/0.7) - \hat{\alpha}}{\hat{\beta}}$$

  - Stop enrolment when:

CV(MTD)<20%

# Prediction

- Predicting future DLT rate in next cohort of 3 patients at a dose:

$$p(y_{new} \mid y) = \int Bin(p\ ,3)\, f(p\ \mid y)dp$$

- MCMC estimation

$$Y_{k,new} \sim Bin(\ p_k,3).$$

$$\bar{Y}_{new}$$

# Adaptive Allocator

- Enroll subjects in cohorts of size=3
- Start with lowest dose possible
- Pre-specify list of admissible doses:
  - Xmin, X2, X3, …,Xmax.
- Adaptive allocator = posterior probability that dose xi is the MTD.
- Randomizer = All 3 subjects to dose with highest probability.

# Overdose control (EWOC)

- Ethically, one cannot assign unsafe doses.
- Do not test too high doses if lower safe doses have not been administered beforehand.
- Admissible doses $d_j$ :
  - $d_j$ < low quantile of MTD distribution
  - Or so that $Pr[p_j>30\%|Y]$ is low
- Practical limit:
  - No more than doubling the maximum administered dose.

# How does the CRM allocator work?

**Posterior logistic model**
**(N=3 @ dose=1; No DLT).**

**Pr[Dose=MTD] & Allocator**

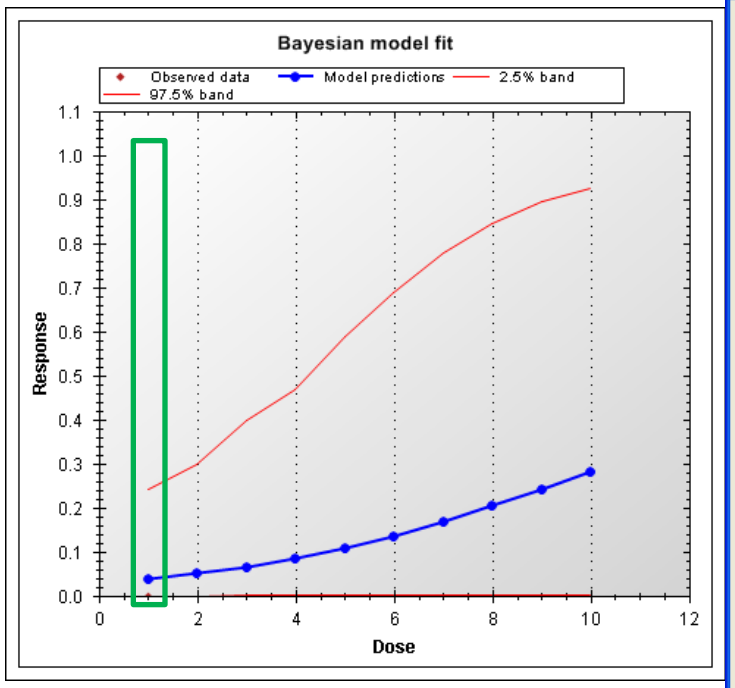


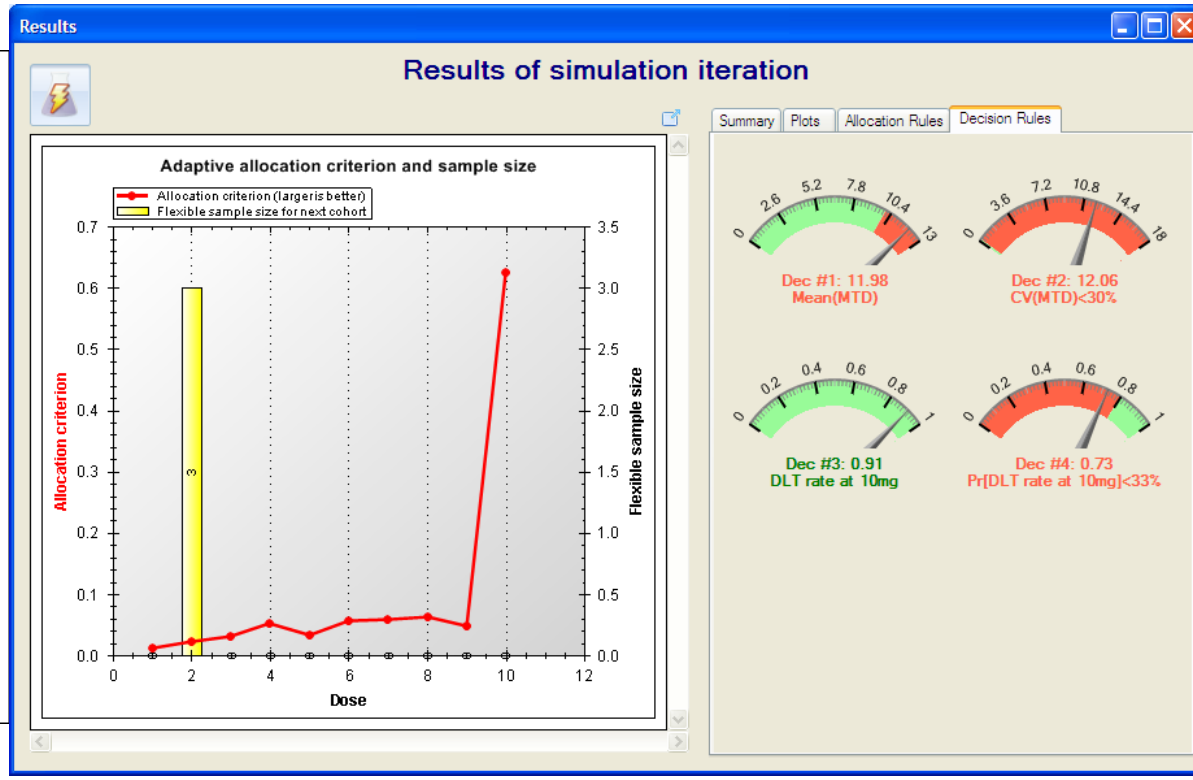

# Trial Simulations

- To assess operational characteristics of an adaptive design
- Example: 10 dose levels – DLT rate:

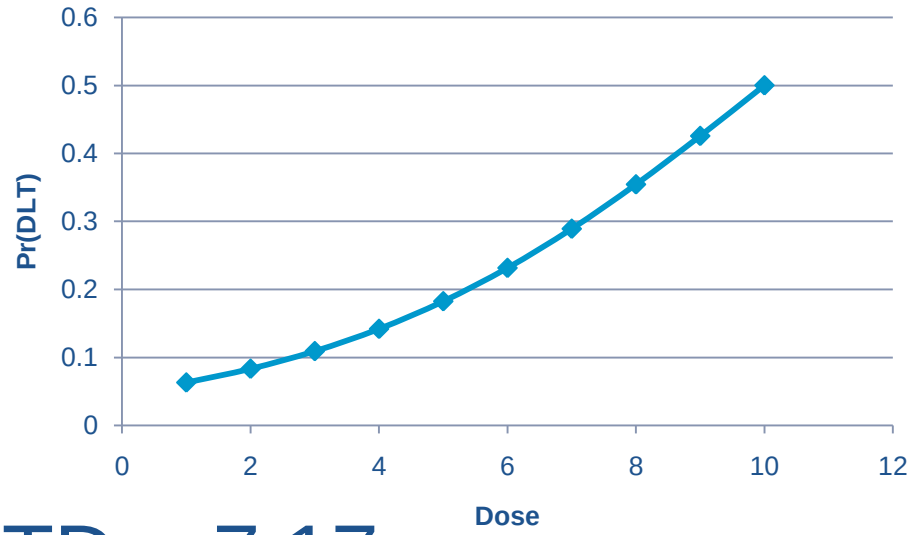


- True MTD = 7.17mg
- N=3/dose – Max size =60 (20 cohorts).

# Simulation Results (100 sims)- Size

**Pr[Stop before cap]=66%** **E[size]=40**

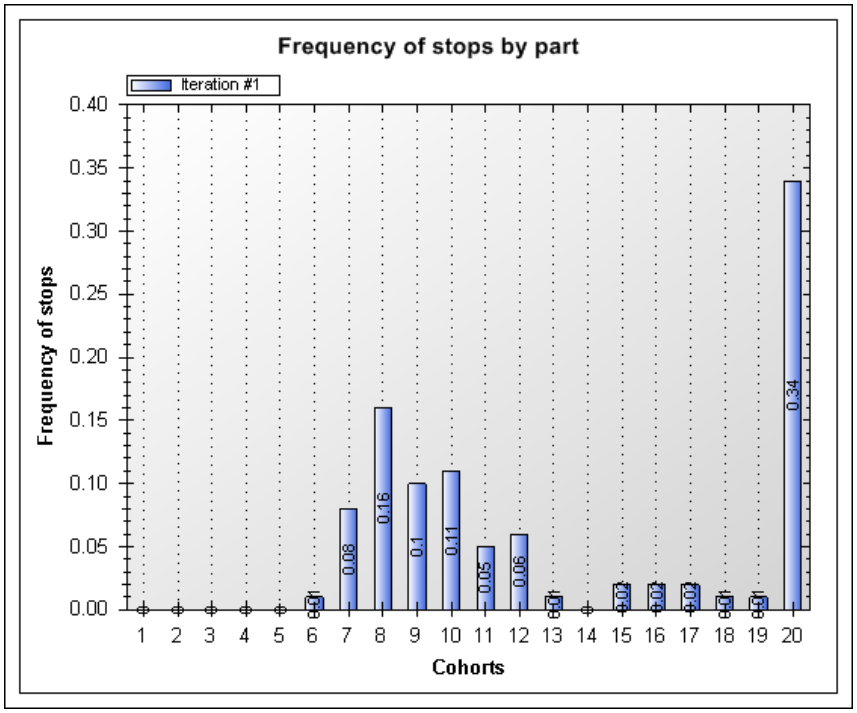


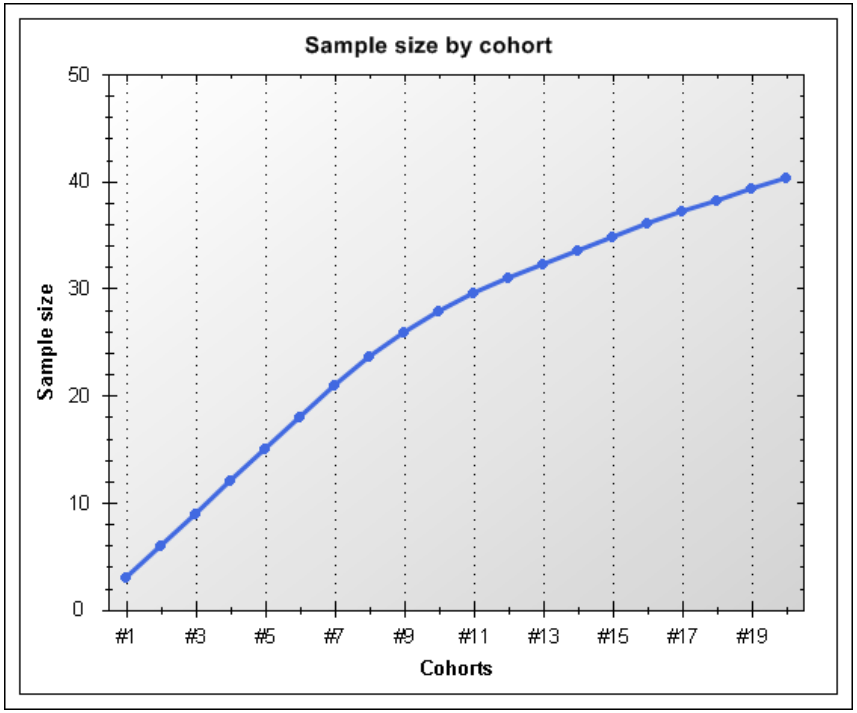

# Simulation Results - MTD

## MTD estimate => 7.13 mg

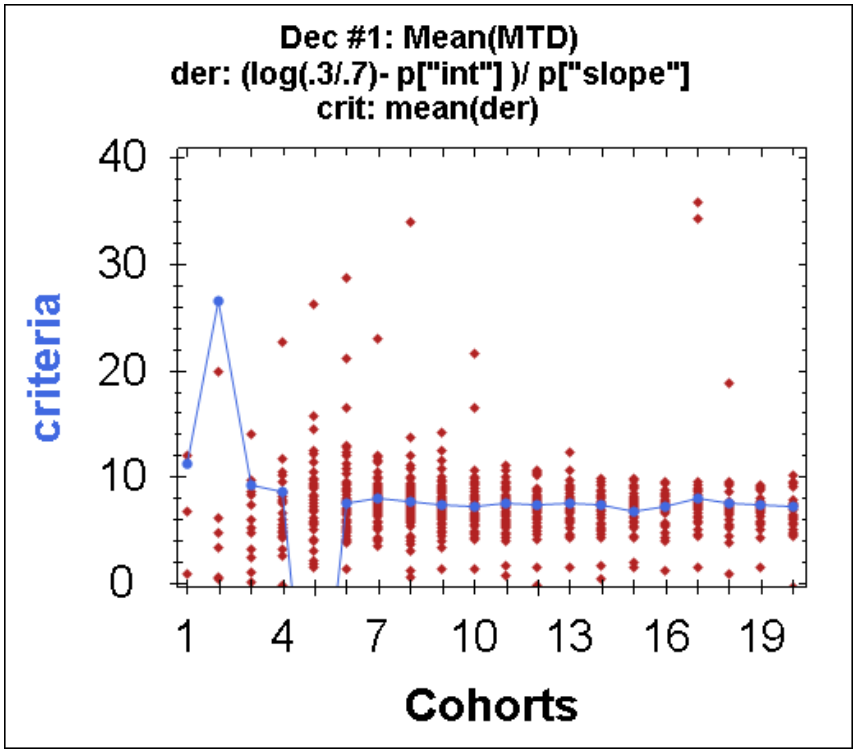


## Pr[DLT] vs dose

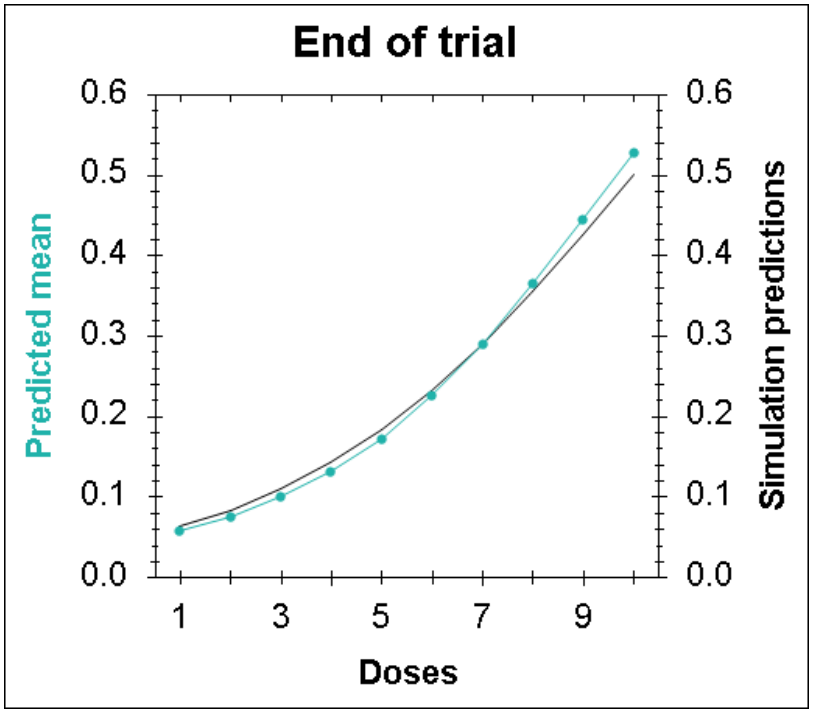

# Simulation results - Doses

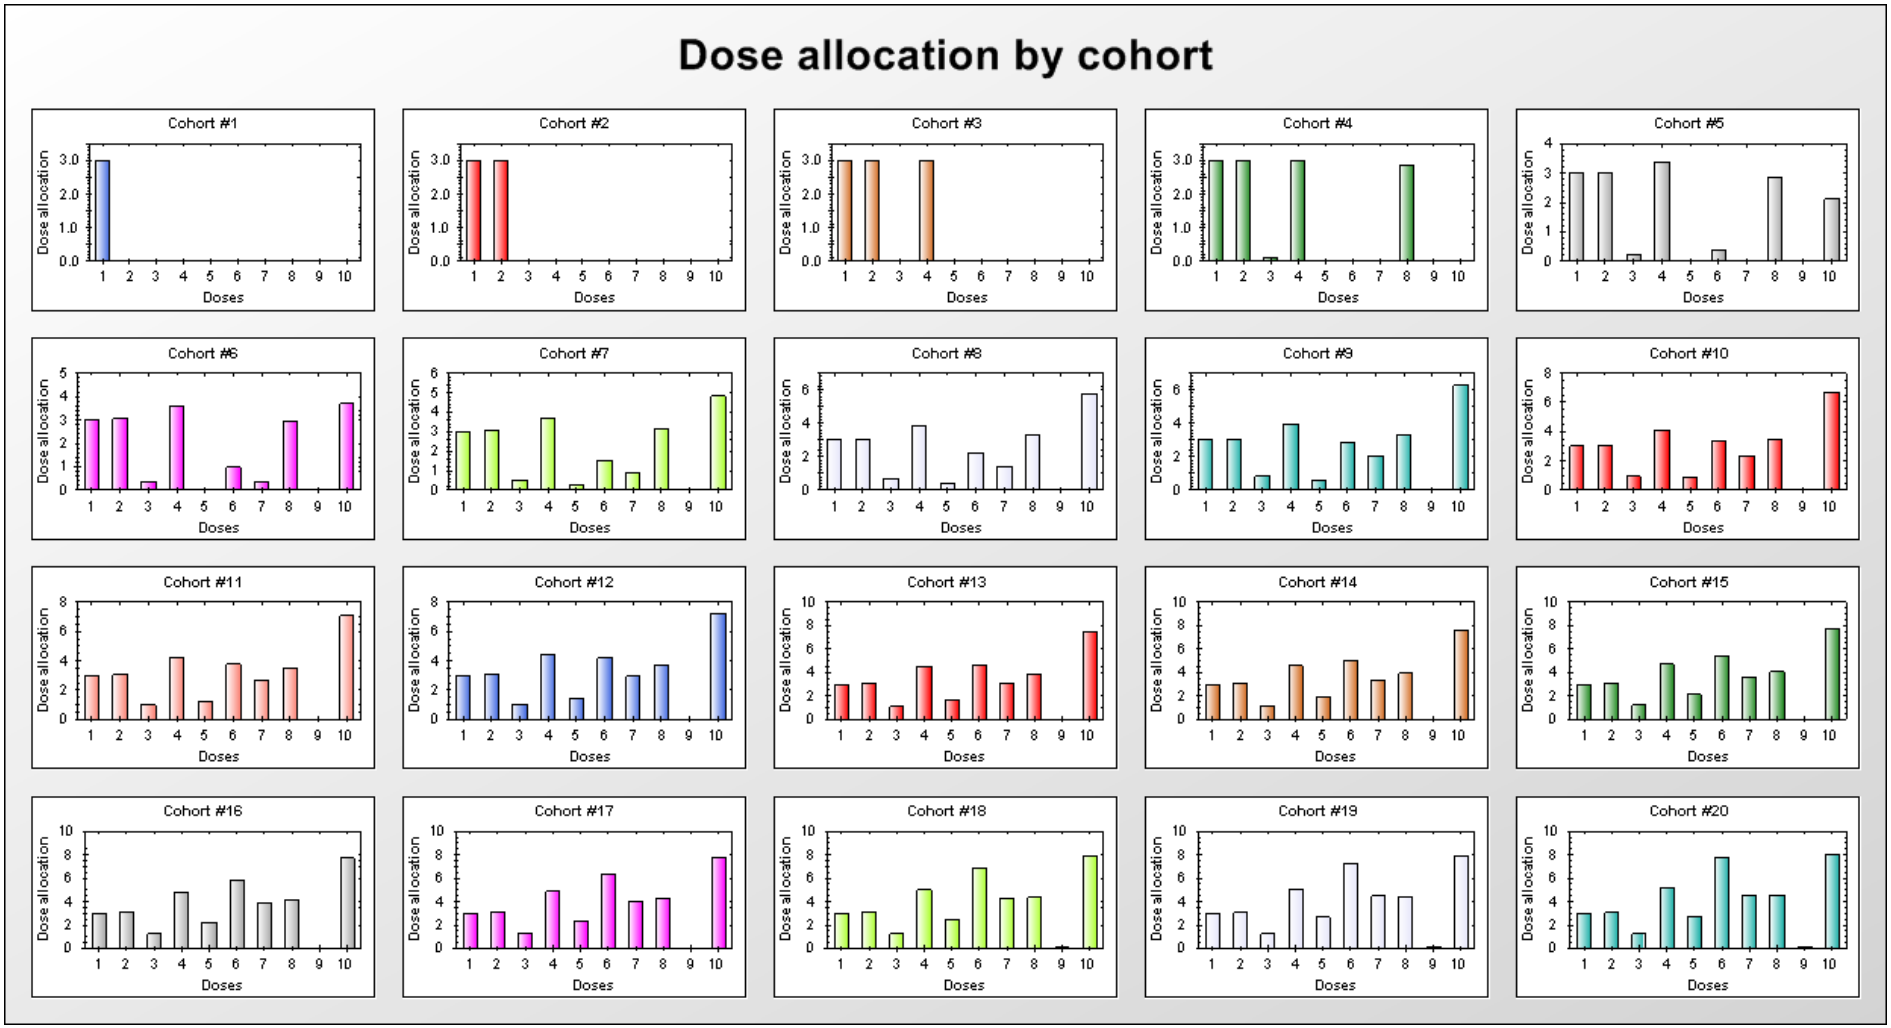

# Summary: Adaptive design in Phase I

- Goal is to estimate the MTD
- Bayesian method provides a probability measure that each dose is the MTD.
- Ethical benefit:
  - Stop when precision is sufficient
- Mix of statistical methodology, expert input and practical constraints
- Trial simulations help to predict future trial performance
- Several methods (CRM, EWOC,…) around a similar concept
- Multivariate/Mixed model extensions.
  - Several endpoints (efficacy/safety)
  - Several patient populations

Hybrid Bayesian/Frequentist Analysis

Two part efficacy study:

1. Proof of concept : MTD vs placebo
2. Dose-ranging

# ADAPTIVE DESIGN IN PHASE II

# Phase II Clinical Trial

- Neuropathic Pain Compound
- Change in VAS after 12 weeks.
- Doses of 0, 14, 28, 42, 70, 98, 140mg
- 2 Parts study:
  - POC :
    - N=20/arm – Pbo vs 140mg
    - T-test at ½ study part (interim) and at completion
  - Dose-ranging if POC successful
    - N=12 subjects/cohort; 1 pbo & 11 active
    - Maximum of 10 cohort in total (including POC): N<=136.
    - Goal: Find ED50 = dose producing 50% of the maximum effect at 140mg.

# Normal Dynamic Linear Model (NDLM)

- Semi-parametric regression for normal responses

The model is as follows:

$$Y_{ij}|x_j \sim N(\theta_j, \tau),$$

where $\tau$ is the residual precision (i.e., the inverse variance), and for j>1:

$$\theta_j = \theta_{j-1} + (x_j - x_{j-1})\delta_{j-1} + \omega_j, \quad \omega_j \sim N(0, \tau/W_\theta), \qquad (1)$$

and

$$\delta_j = \delta_{j-1} + \gamma_j, \quad \gamma_j \sim N(0, \tau/W_\delta).$$

- W is the variance inflation factor. It determines the extend of smoothing in D-R curve:
  - W => 0: rigid fit => linear regression
  - W => ANOVA

# Quantile variance allocator

- Quantile q=50% $g(d_{(q)}) = target(q).$

$$Var[g(d_{(q)})] = \sum_{k=0}^{K} Var[g(d_k)]\, \mathrm{p_k}(d_{(q)}),$$

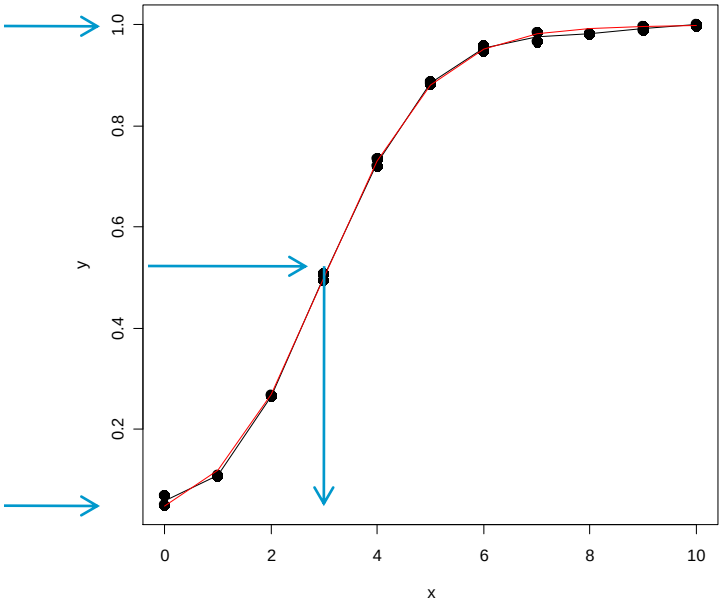


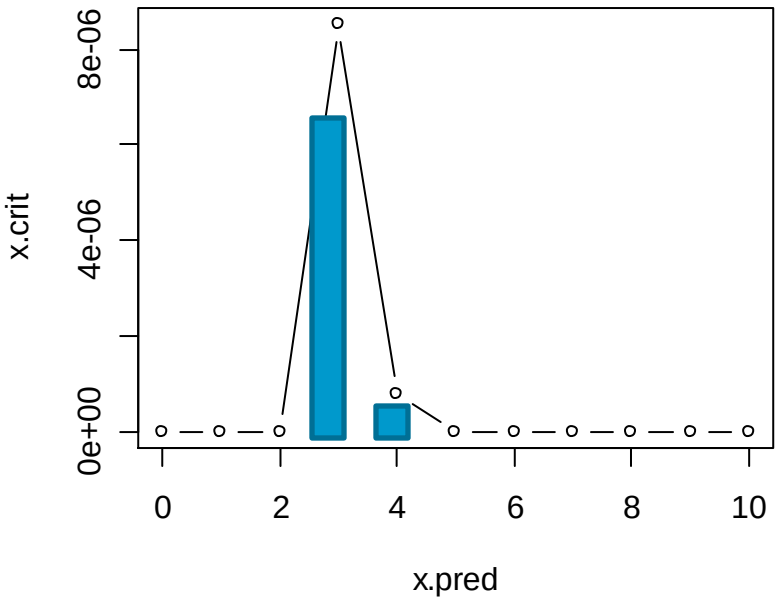


- Randomizer : Biased-coin proportional to utility value.

# Stopping rules

- POC : Frequentist test $p<0.05$ one-sided
- Dose response:
  - Efficacy: N=30 patients at any single dose
  - Trial failure : 10 cohorts enrolled.

# Decimaker Software

- Developed by ClinBAY
- Adaptive design & Bayesian analysis software
  - Trial simulations
  - Inference
- Main Features
  - CRM method
  - NDLM
  - Models for normal & binary data
  - Frequentist & Bayesian tests
  - D- & C-optimal designs
- Interactive & batch-mode execution.

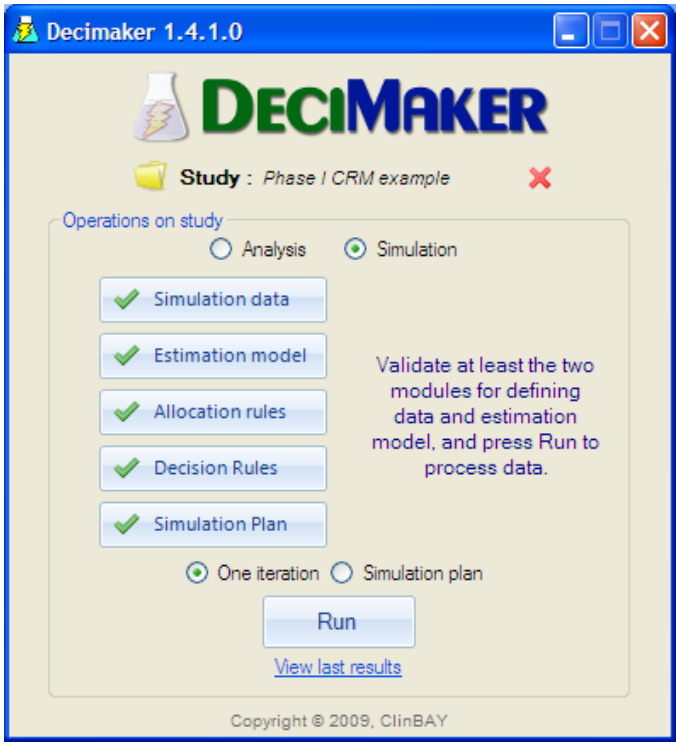

# Decimaker Architecture

- Uses R and Winbugs for computation
  - R/Dcom server link to GUI
  - Versatile
  - Powerful
- Graphical interface in .NET
  - Microsoft GUI
  - Nicer, dynamic graphics
  - Inter-operational with MS products
- Clinical trial oriented
  - Support, validation, audit trails,..

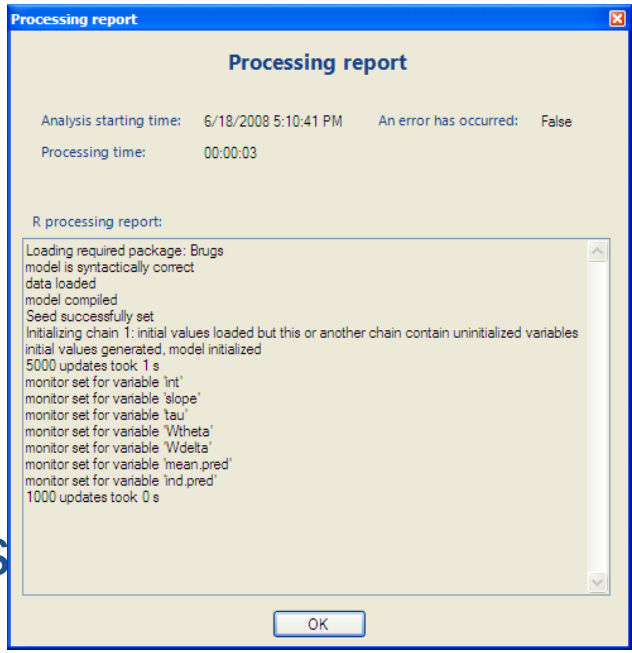

# Software components

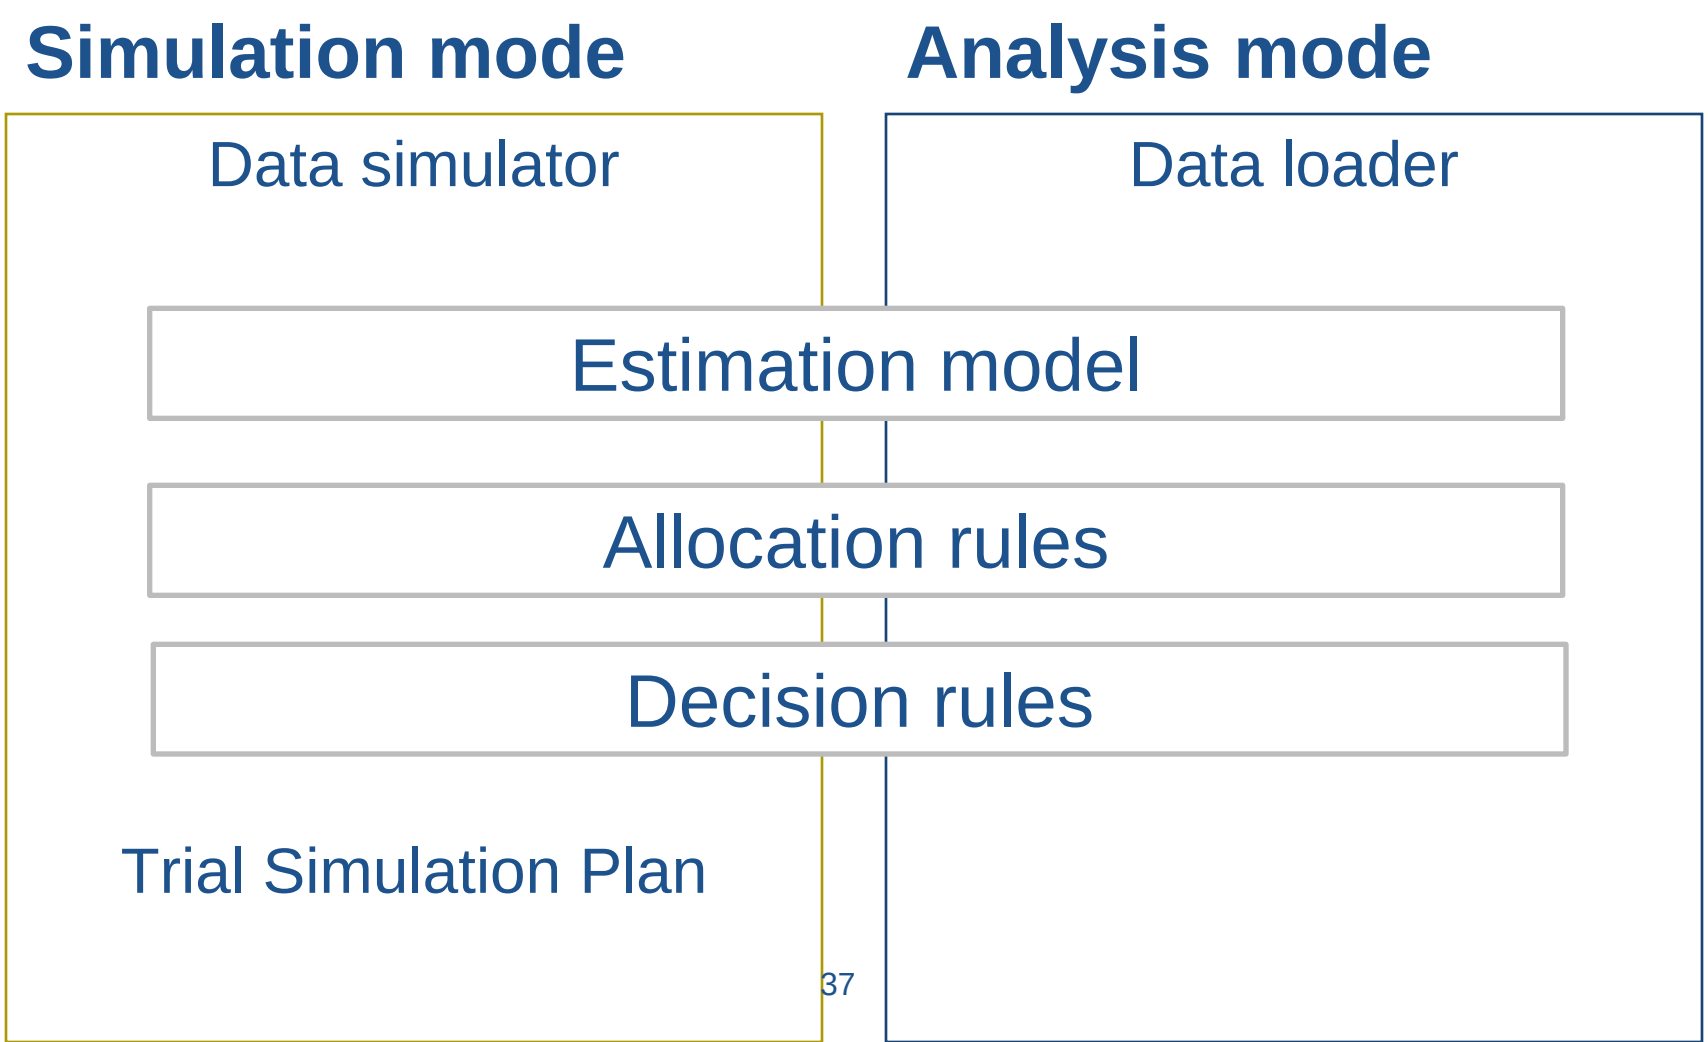

# Demo

- Phase II – 2 Part Study
  - Simulation results

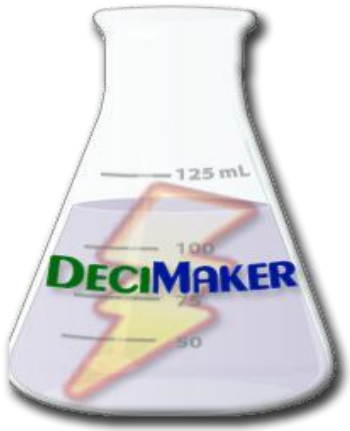

# Summary: Adaptive design in phase II

- Mixed Bayesian/Frequentist Inference possible
  - Extensions: Group-sequential methods, predictive power
- Semi-parametric regression for dose-response:
  - Weaker assumptions than model-based
  - Slight loss of efficiency
  - Worst case scenario reverts to ANOVA
- Multi-Part/Seamless Phase trials:
  - « Keep the ball rolling »
  - Pre-planning of resources is more demanding.

# Conclusions

- Bayesian adaptive designs permit to:
  - Design studies based on quantitative measures of risks/benefits
  - Modify designs in real-time to optimize these measures.
- This requires specialized skills and software:

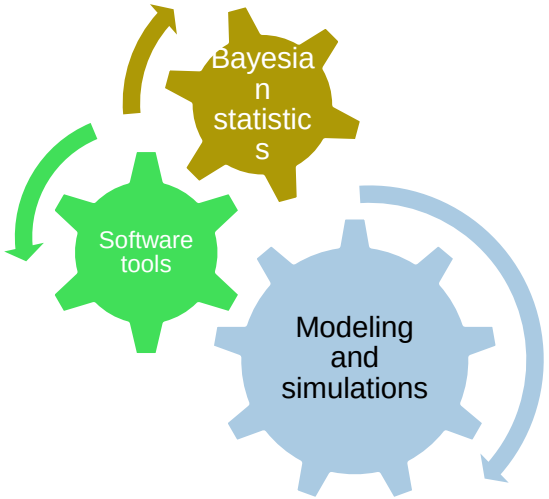

Thank you!

# ANY QUESTION?

A job with us? See www.clinbay.com/-Career-.html