Adaptive Designs : Bayesian & Non-Bayesian Approaches

Andy Grieve

Statistical Research Consulting Centre

Pfizer Global R&D, Sandwich, UK

Andy_P_Grieve@Sandwich.Pfizer.com

Outline

- Classification of Experimental Designs
- Adaptive Interim Analyses
- Response Adaptive Designs
  - Randomised-Play-The-Winner (RPW)
    - Ethics
  - Up-and-Down
  - Continuous Reassessment Method (CRM)
  - General Approach
- Issues / Discussion

Palmer – Classification of designs

Statistical Methods in Medical Research 2002; 11: 381-402

- Parallel Group, Fixed Sample Size
  - Eg Bradford-Hill : streptomycin & treating pulmonary tuberculosis (*Br. Med. J*, 1948)
- Data-Dependendent Designs
  - Sequential (Abraham Wald , 1940's)
  - Group sequential (Armitage et al, 1960's)
  - Adaptive Interim Designs (Bauer et al, 1990's)
  - Response-adaptive designs
  - Bayesian decision theoretic designs

Use of Response Adaptive Designs in Pharmaceutical R&D

- O'Quigley, Pepe and Fisher (Biometrics 1990) CRM in phase I oncology studies
- Eli Lilly (UK) – early phase I studies dose-titration in control of diabetes
  – phase II depression
- SKB (UK) – adaptive FIM studies
- Astra-Zeneca (UK) – Dose Escalation Studies tolerability / efficacy
- Pfizer – phase II dose selection stroke, pain

Statistical Research & Consulting Center

Randomise Play the Winner (RPM)

RPW Design

- RPW designs described by Urn models
- At beginning of trial
  - Urn contains α balls of each of two colors (W&R) representing 2 treatments
  - When a patient is to be treated a ball is chosen at random

Statistical Research & Consulting Center

RPW Design

- RPW designs described by Urn models
- At beginning of trial
  - Urn contains α balls of each of two colors (W&R) representing 2 treatments
  - When a patient is to be treated a ball is chosen at random (with replacement)
  - When the response is known the urn content is updated as follows
- If the patient was allocated to treatment t (either W or R) and responds positively, β balls of colour t are added to the urn otherwise γ of colour s (the complement of t) are added.
- In time the urn will contain a higher proportion of colored balls associated with the more successful treatment
- RPW(α,β,γ) design

ECMO Bartlett et al (1985)

- Newborn infants with severe respiratory failure - Mortality
- Extra Corporeal Membrane Oxygenation vs Traditional Ventilator
- Phase I trials >50% survival on ECMO
- Optimal Therapy : survival < 20 %
- Chose Randomised Play-the-Winner (RPW)
  - speedy outcome - anticipated response diff -> small sample size - scientific/ethical dilemma

Randomised Play-the-Winner - Urn Model (ECMO)

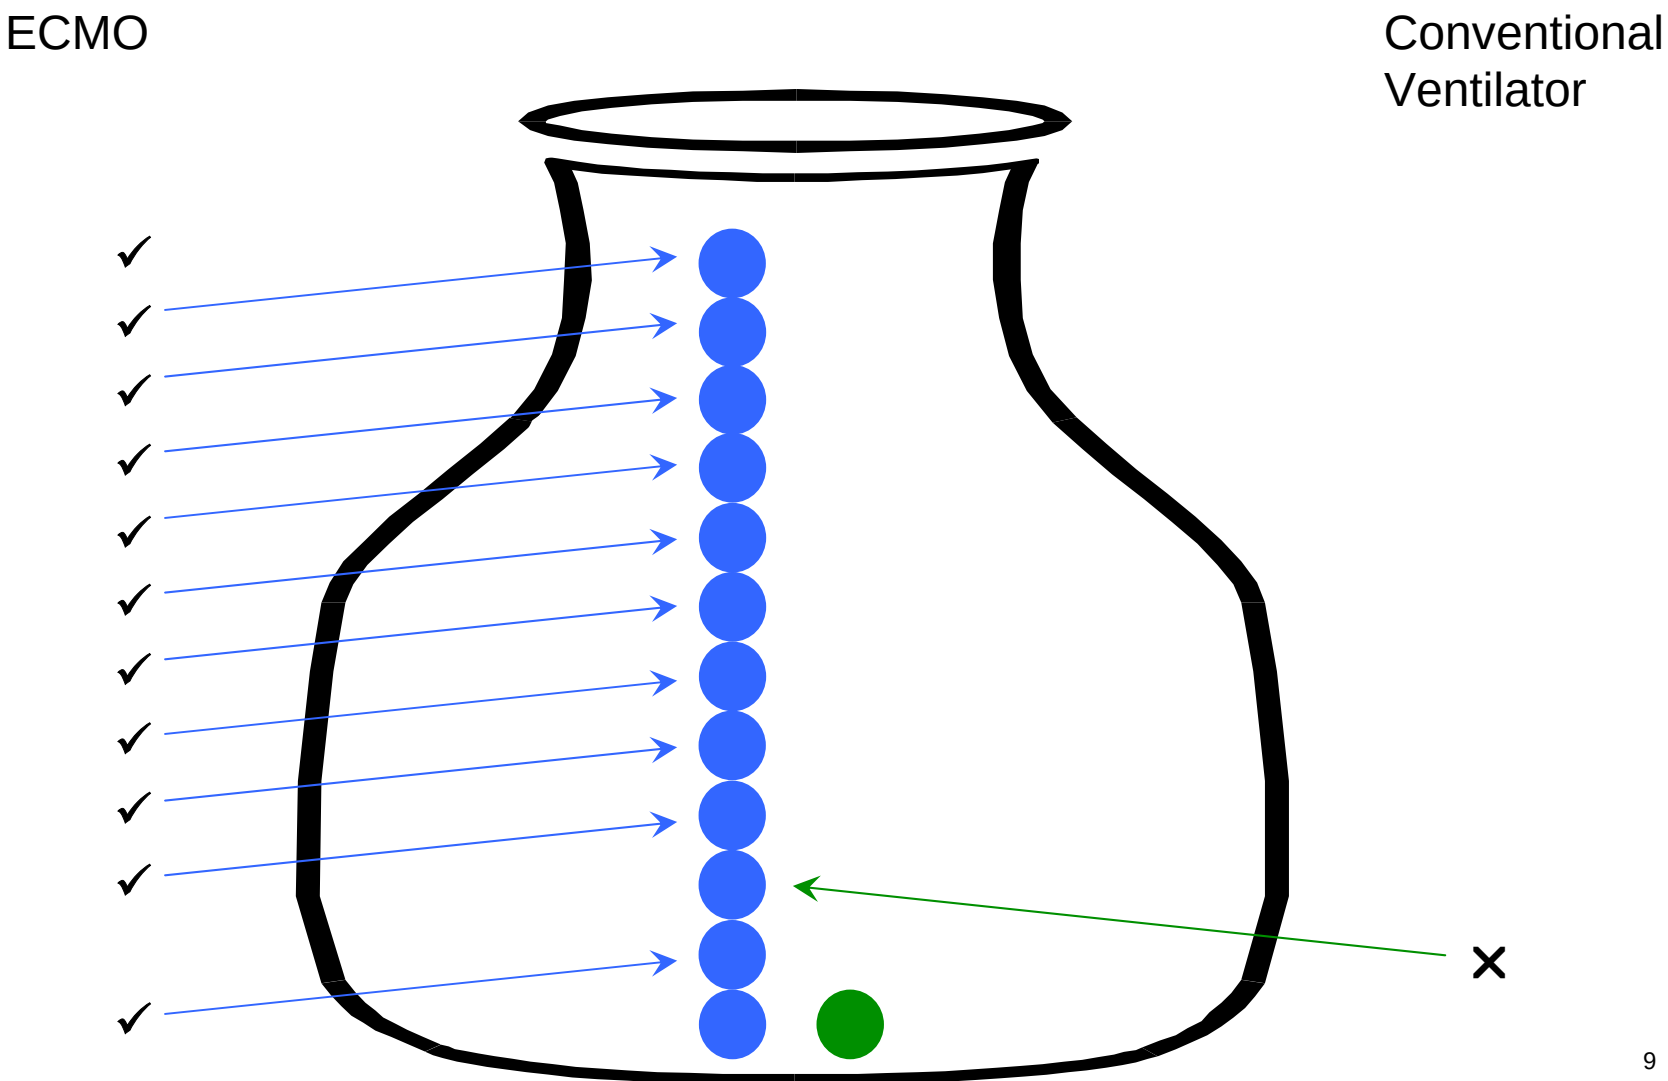

Randomised Play-the-Winner - Urn Model (ECMO) : Issues

- Was the urn model sensible ?
  - Other parameters

Randomised Play-the-Winner - Urn Model

ECMO

Conventional Ventilator

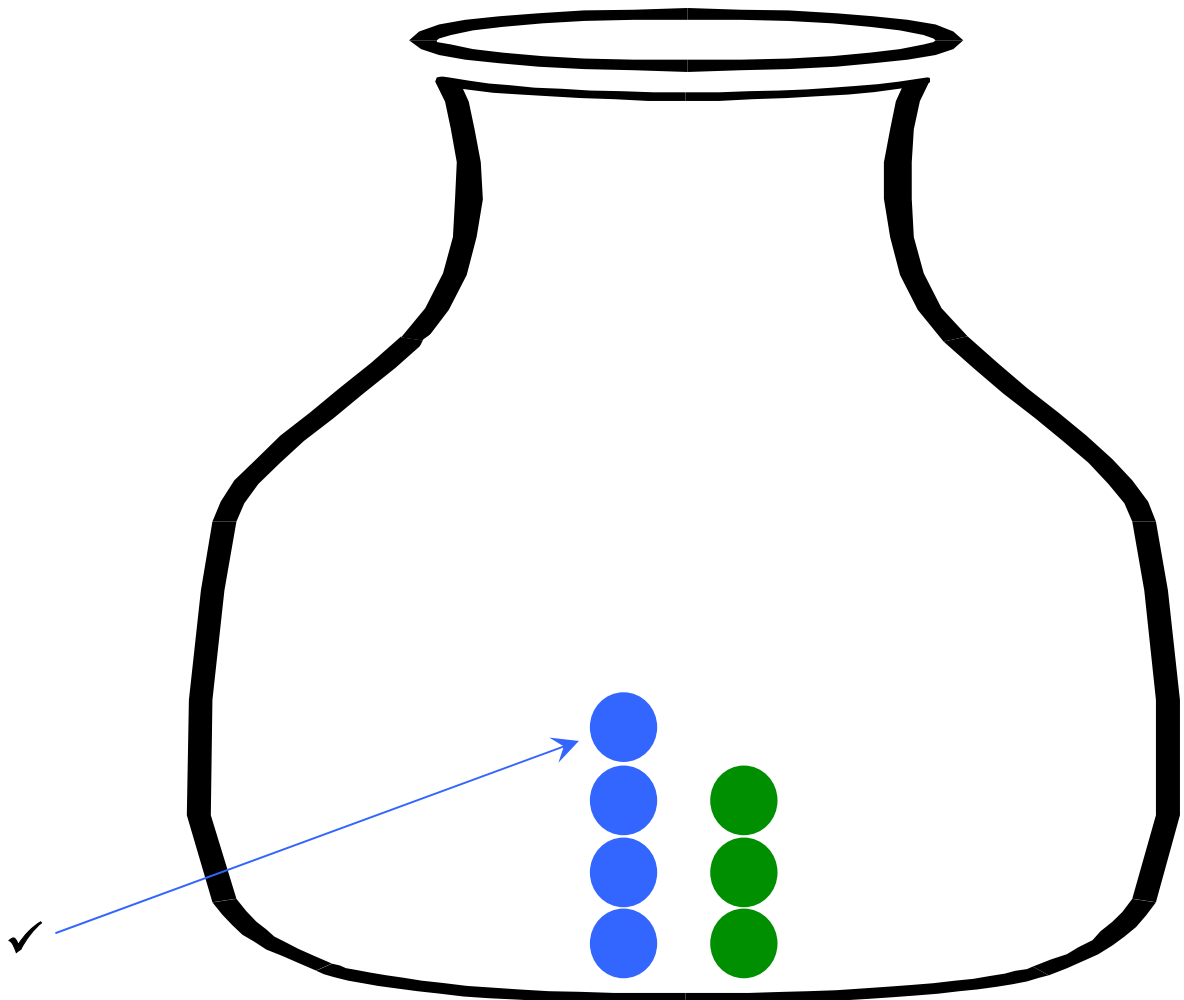

Randomised Play-the-Winner - Urn Model (ECMO) : Issues

- Was the urn model sensible ?
  - Other parameters
  - Begin with randomised block
- How reliable are the results - 11/11 vs 0/1 ?
  - Ranking and selection procedure
  - Minimum number of patients

Statistical Research & Consulting Center

Ethics of Adaptive Designs

Statistical Research & Consulting Center

Ethical Principles and Clinical Trials

- Individual ethics - doing what is best for subjects in current trial
- Collective ethics - doing what is best for future patients who stand to benefit from the results of current trial
- Tension - "Concern for the interests of the subject must always prevail over the interest of science and society" (Declaration of Helsinki)

Tension between individual / Collective Ethics

Individual Ethics	———	Collective Ethics
Adaptive Trials	———	Randomised Trials
?	———	Objective, unbiased evidence

Individual Ethics and Adaption (Clayton, BJCP,1982; Armitage, ISR, 1985)

- At start : ignorance (equipoise ?)
  - ➔ Randomisation
- Information accumulates
- Patients tend to be randomised to the “best” treatment

Individual Ethics and Adaption (Clayton, BJCP,1982; Armitage, ISR, 1985)

- Suppose 9:1 randomisation
  - Ethically can we randomise to the inferior treatment ?
  - How much information is enough ?

- What can we gain ?

Ethics - AZT : Mother/Newborn HIV Transmission

- Pregnant women randomised to placebo or AZT (A:239 - P:238)
- Endpoint : newborn HIV +ve (A:20 - P:60)
- Zelen & Wei - Randomised Play-the-Winner
  - # A:360 - P:117
  - HIV A:30 - P:30
  - CI Randomised : 11-23%
    RPW : 9-25% (efficiency ?)

Total Number of HIV +ve Babies from Simulated Trials

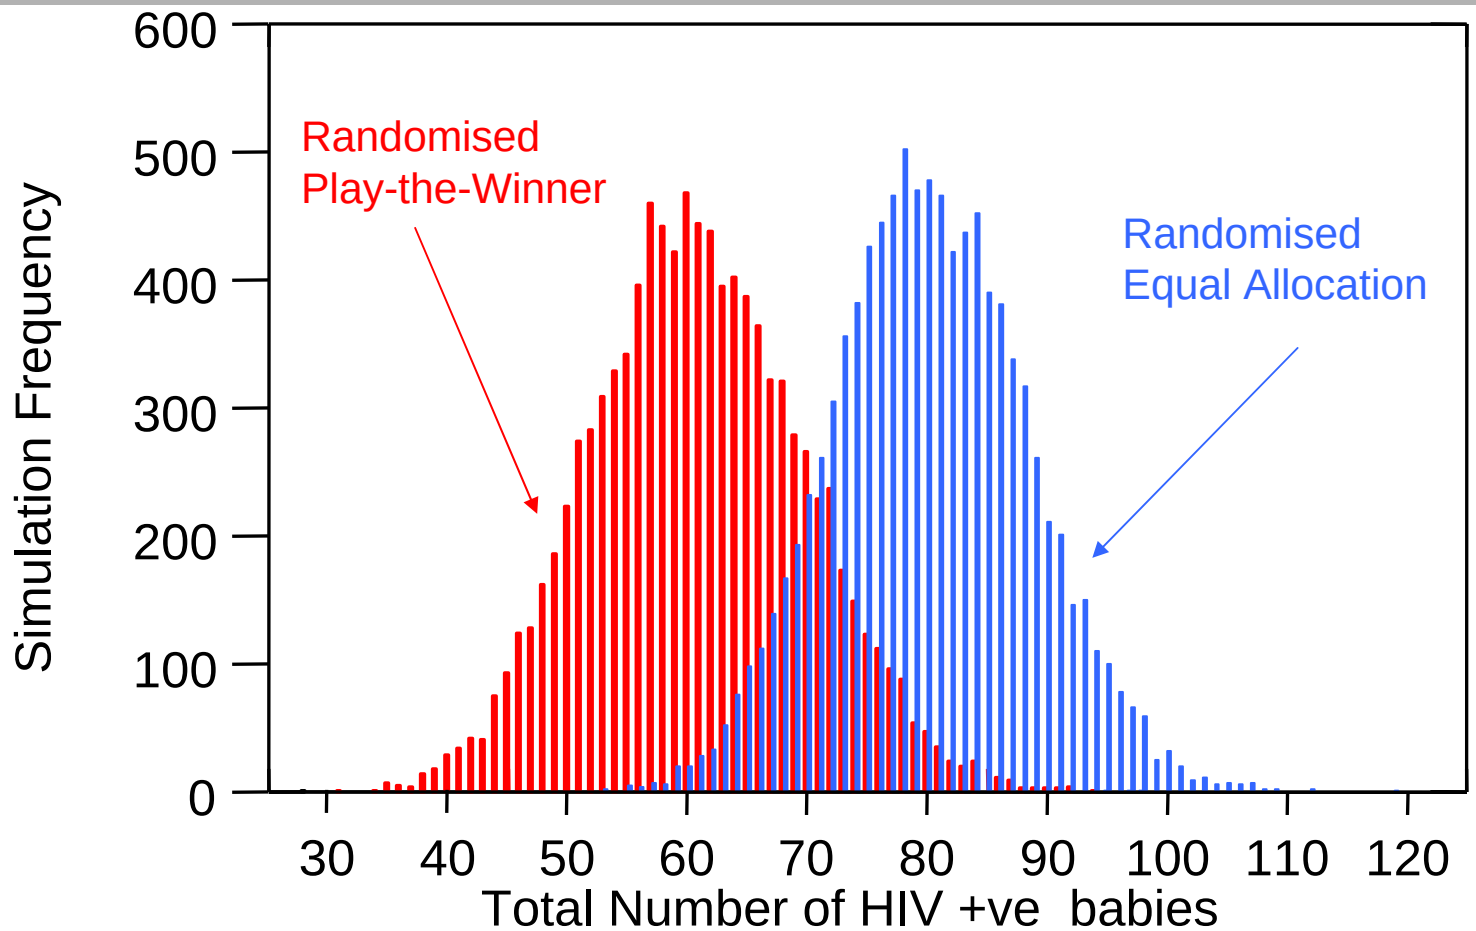

Up-and-Down Design

Background

- New compound - anti migraine
- Activity from 0.5 mcg
- Different mode of action from 500 mcg - more like elitriptan/sumitriptan
- Dose range is therefore 0.5 mcg - 500 mcg
- Need to reduce this range before conducting a dose response study
- Window of opportunity
- Placebo, 0.625, 3.125, 12.5, 62.5, 312.5 mcg - limited number because of dose form - intravenous : syringe sizes
- What is dose at which 50% of patients respond ? Seen as 20% > than placebo rate (30%)
- Response :
  - Change within 2 hours from severe or moderate headache to mild or no headache - Glaxo defn.
- Need enough patients around optimum dose to have confidence in estimate
- May not achieve this with standard parallel group (equal n) design

The method

- Allocates patients to dosing groups (usually unequally)
- Dose finding process
- Nth patient gets allocated to dose depending on response of (N-1)th patient
- First patient gets placebo or 12.5 mcg

Impact of Mechanical Head Trauma

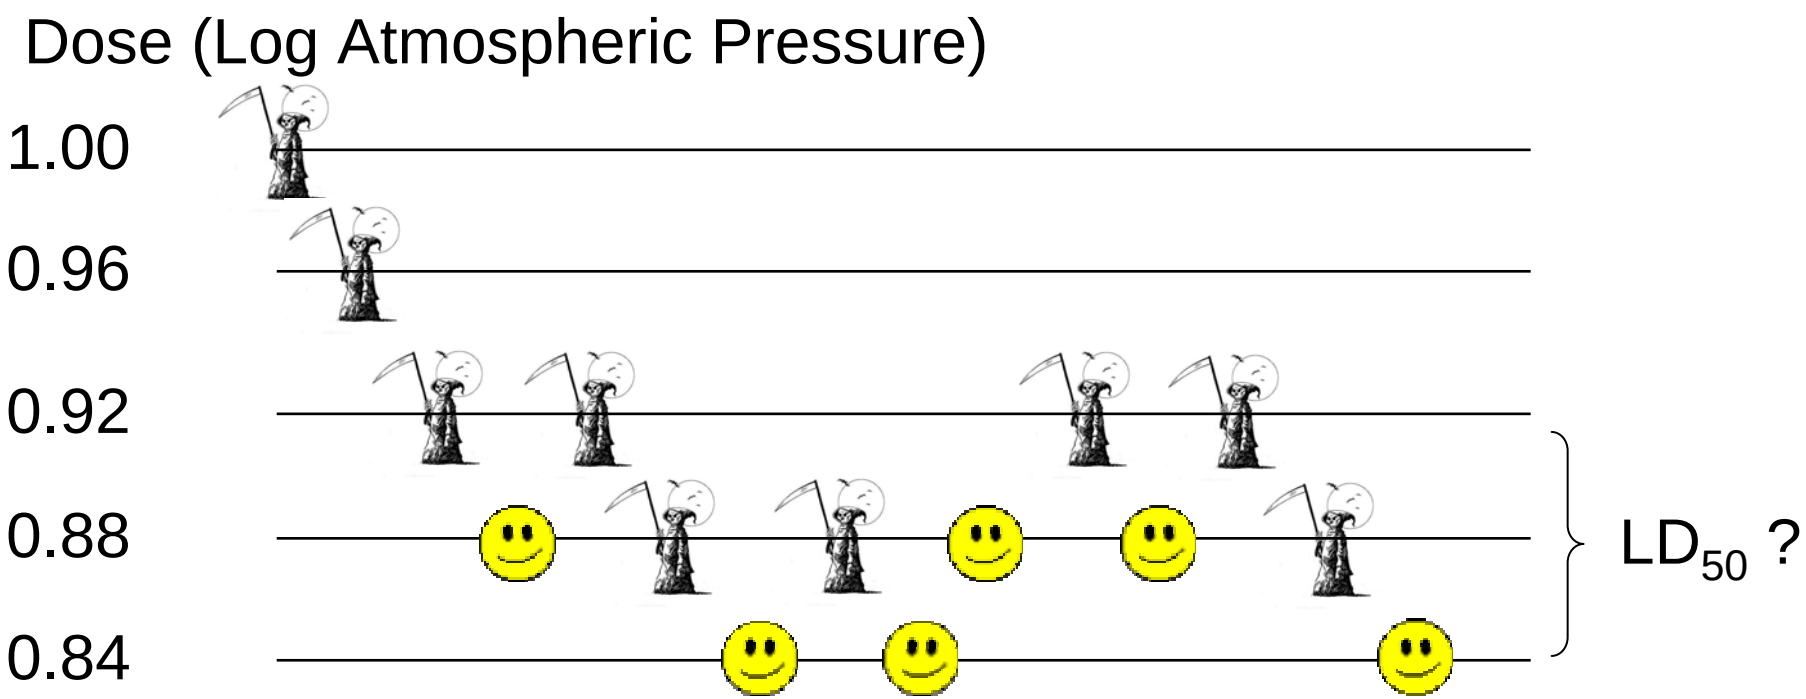

- Died

- Survived

Random Walk

LD_{50}=E(limiting distribution)

Results from an up-and-down Design in Migraine

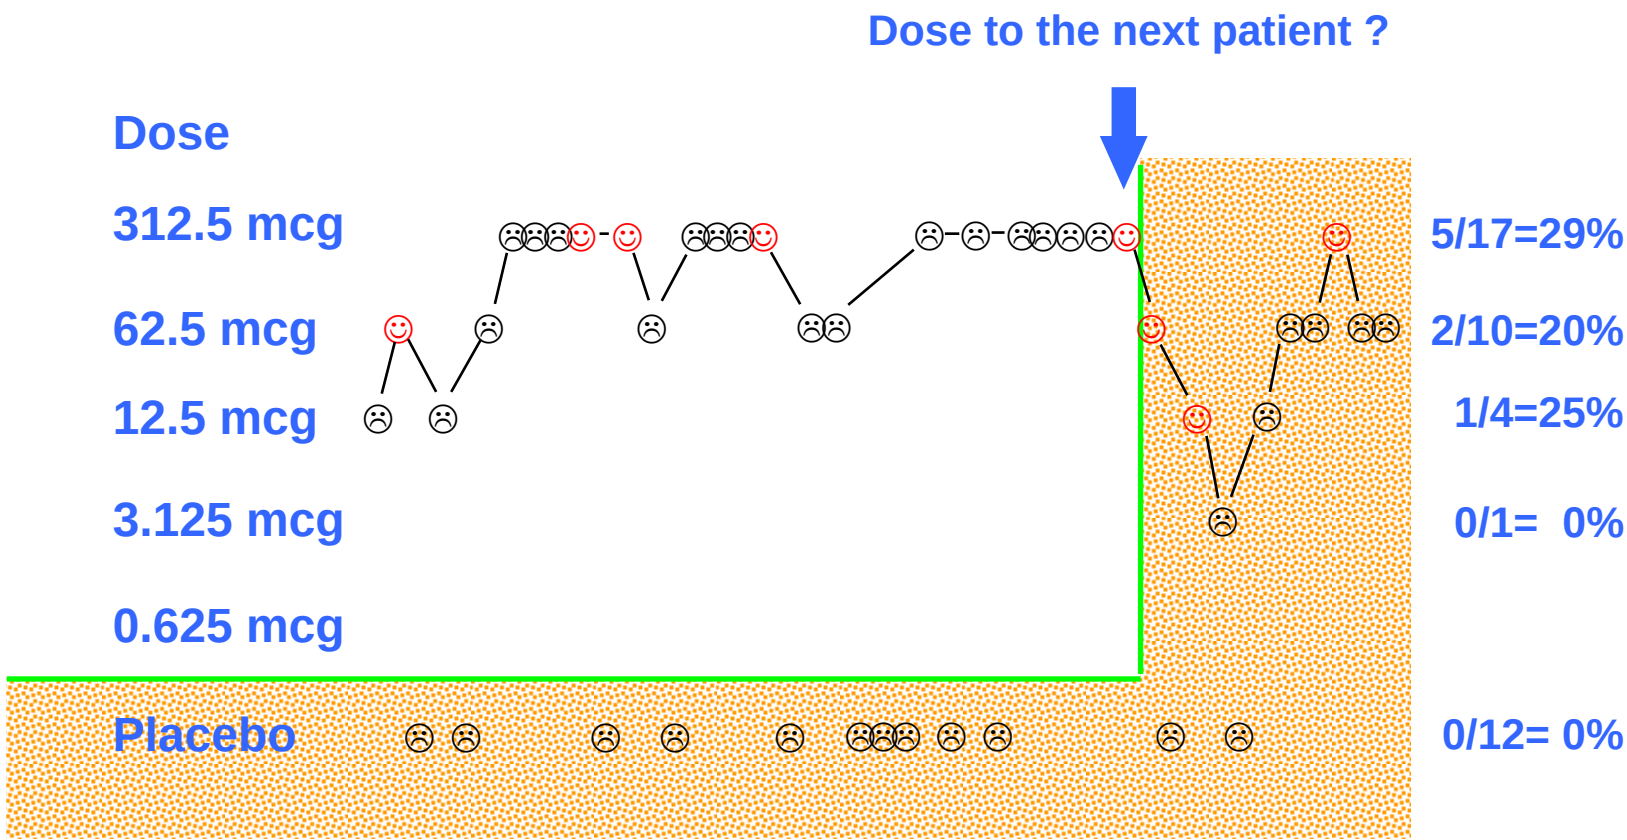

Statistical Research & Consulting Center

Posterior Distribution for the ED50

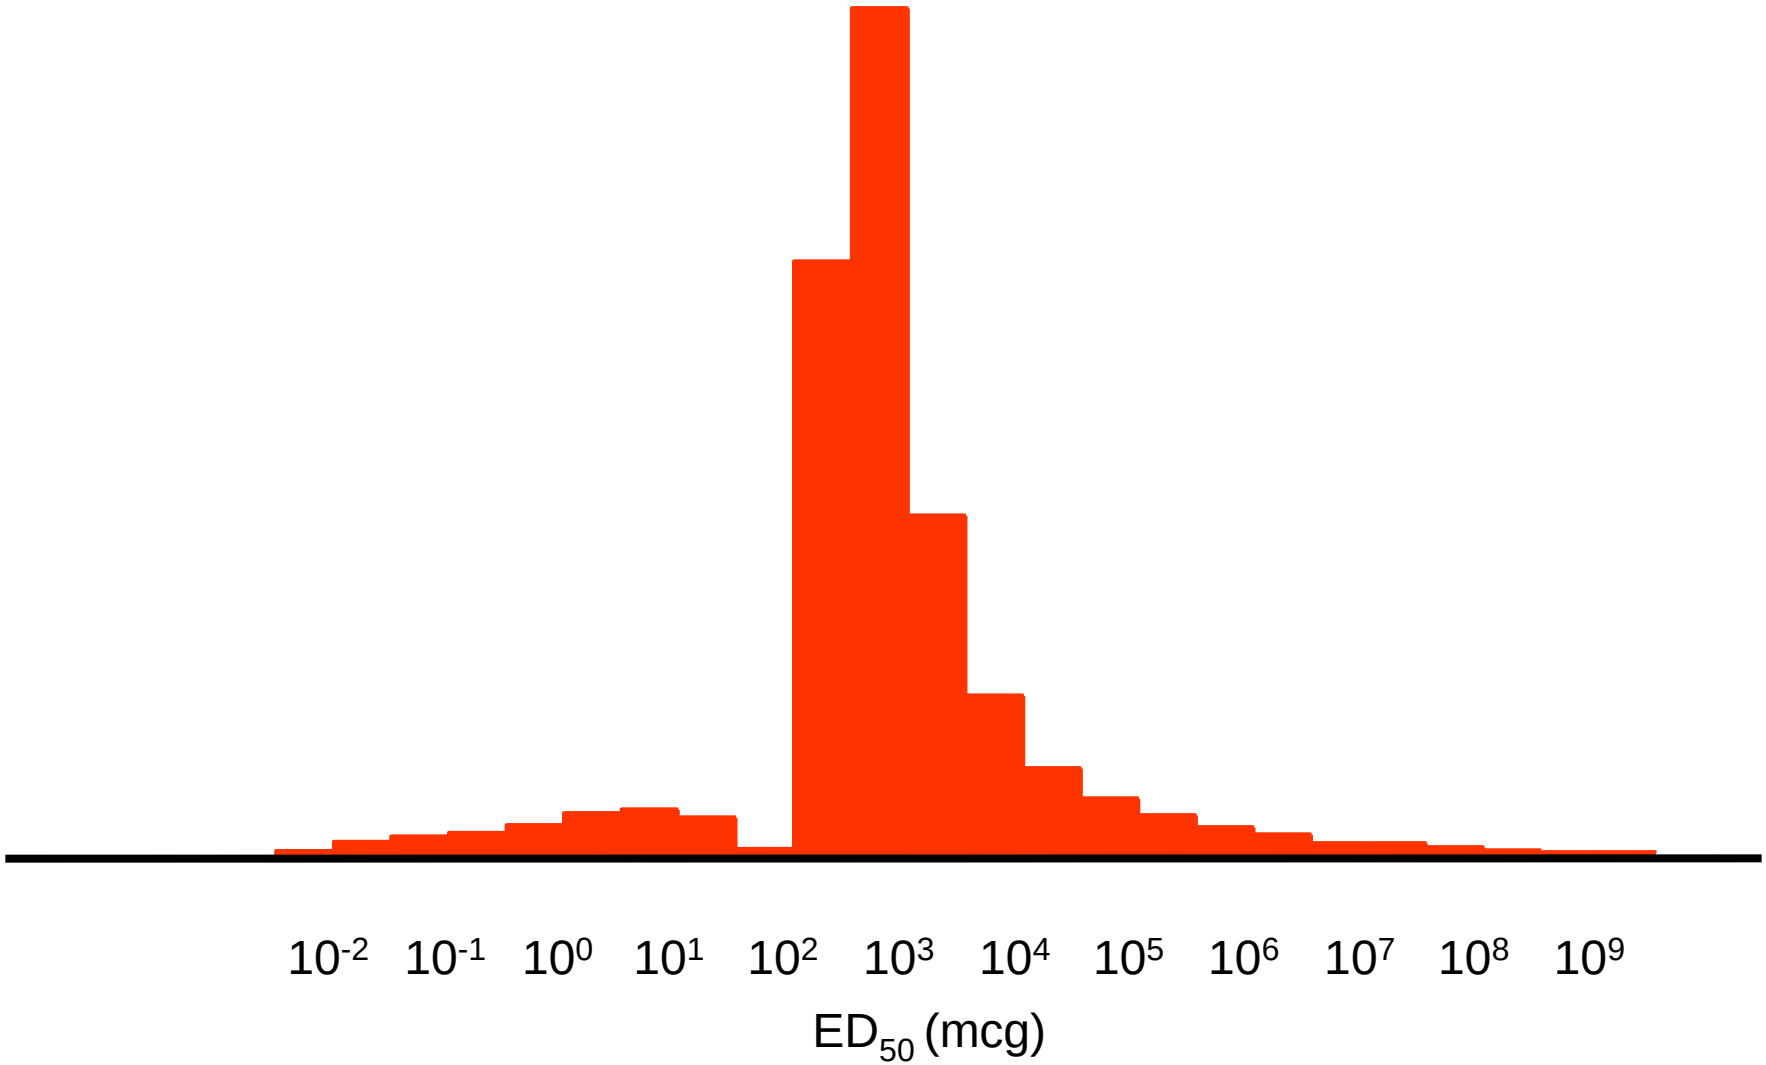

Statistical Research & Consulting Center

Knowledge about Dose response

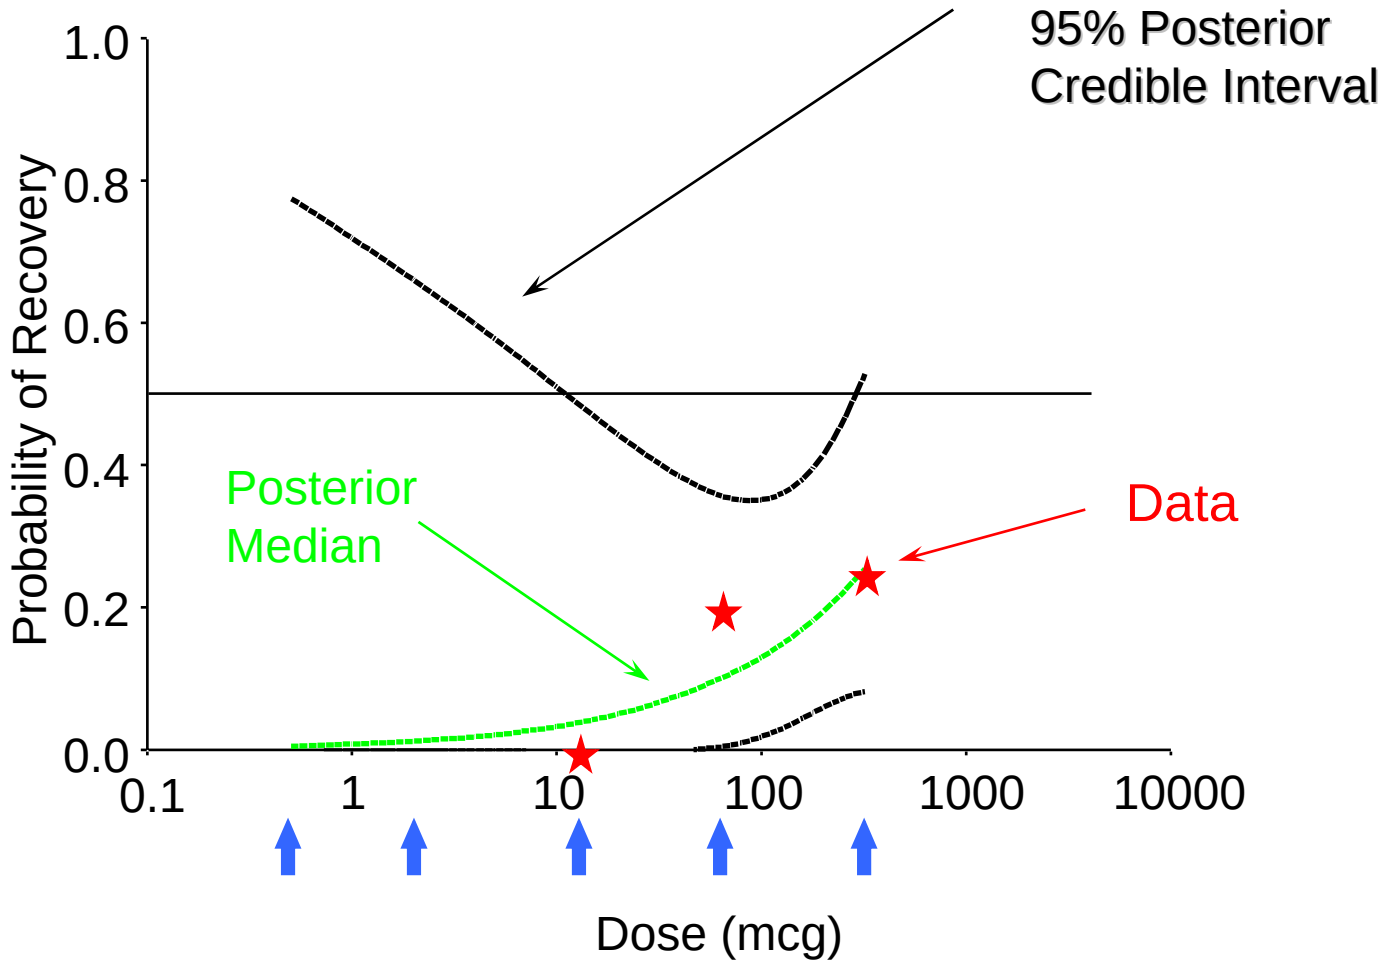

Statistical Research & Consulting Center

Extrapolation

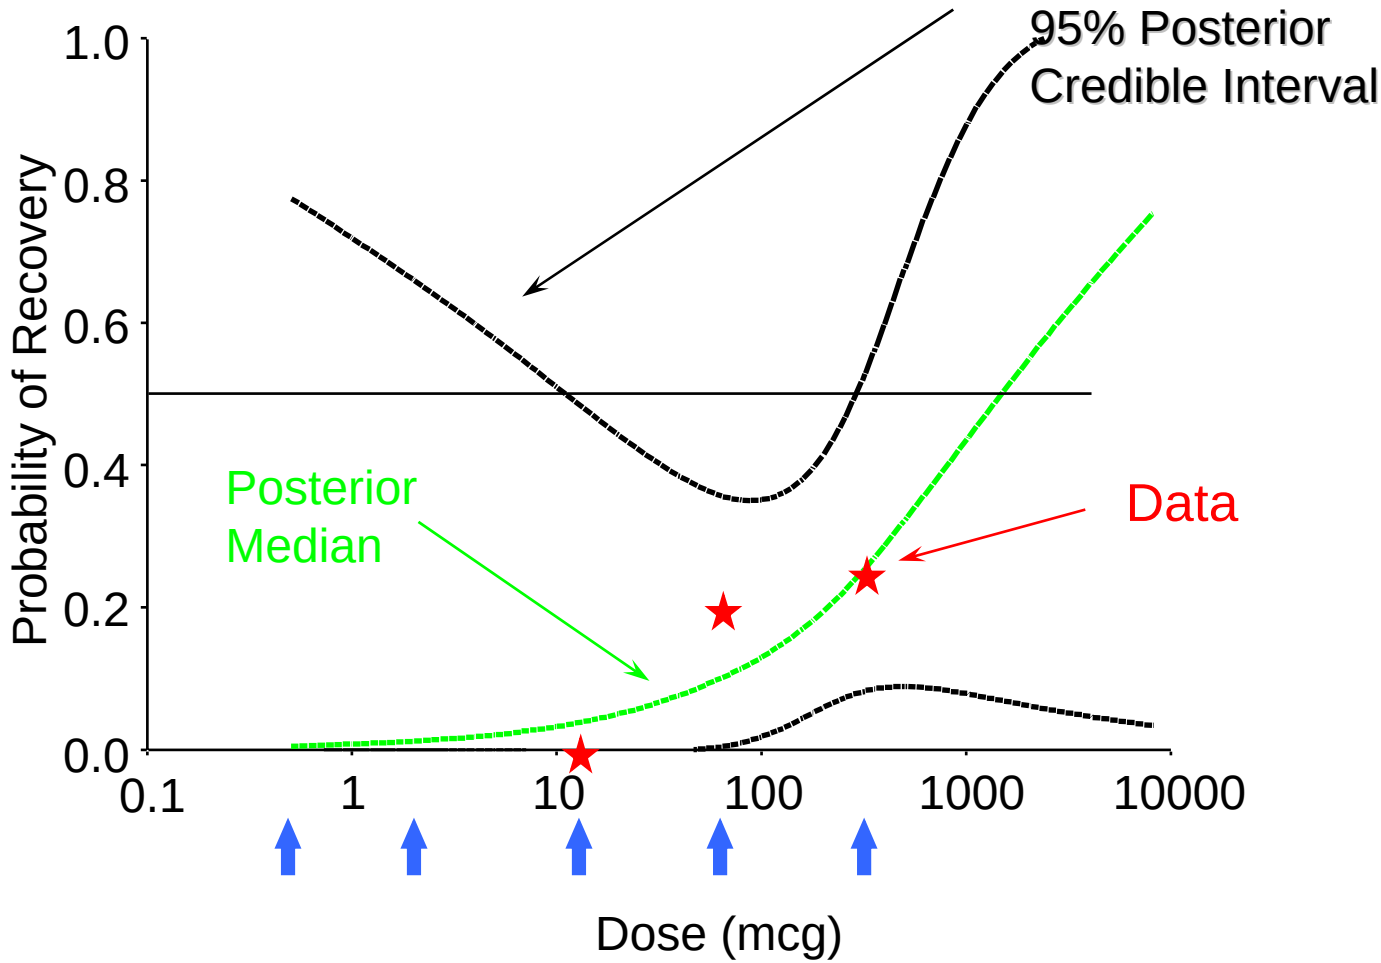

Statistical Research & Consulting Center

Choice of Dose - Method 1

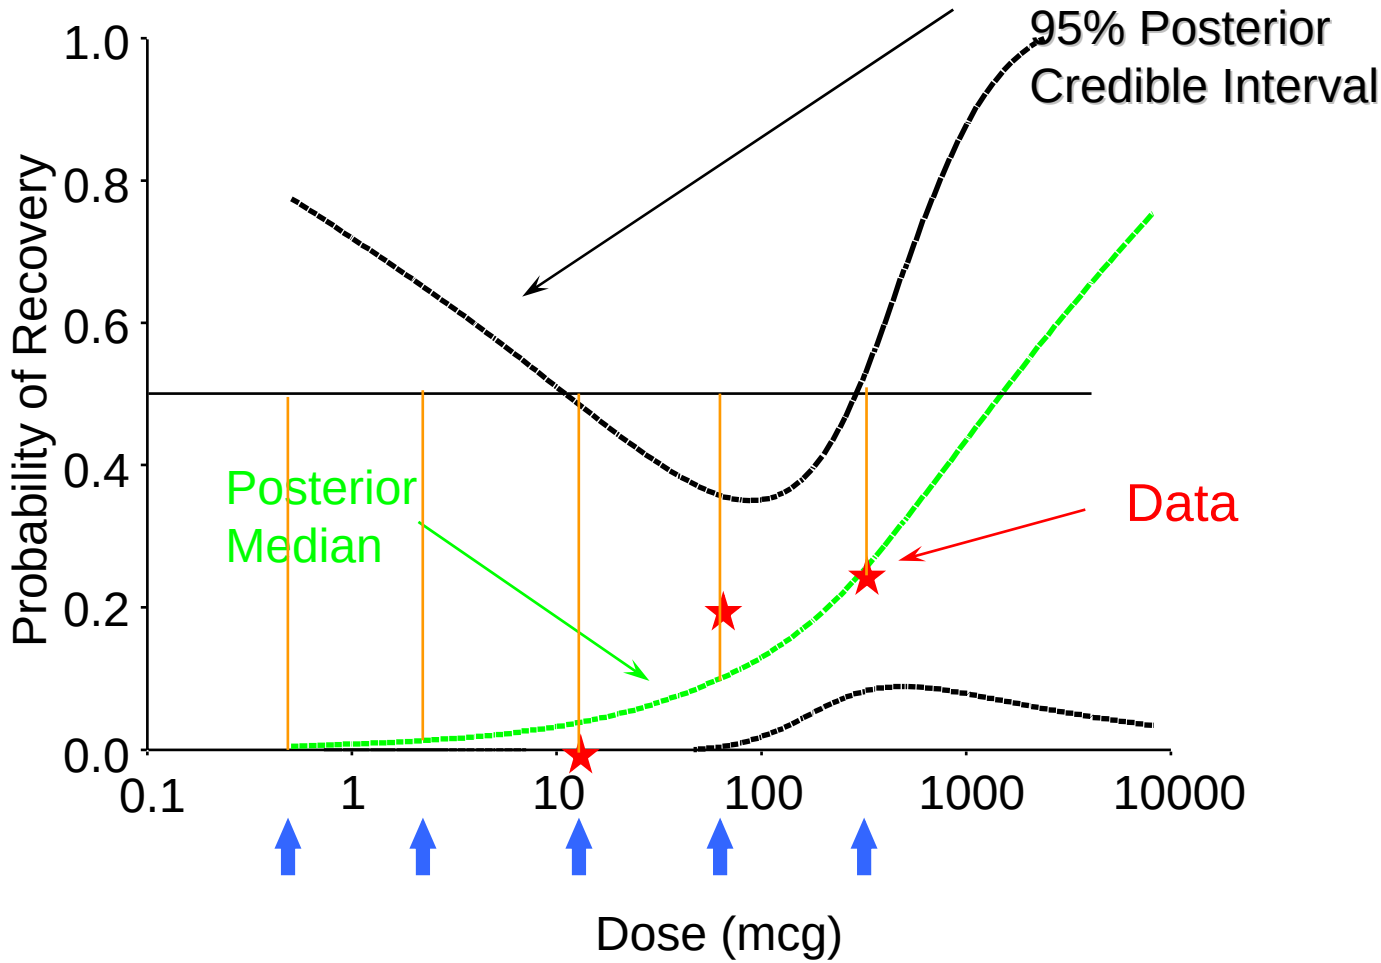

Other Methods - First Step

Specify what you are interested in

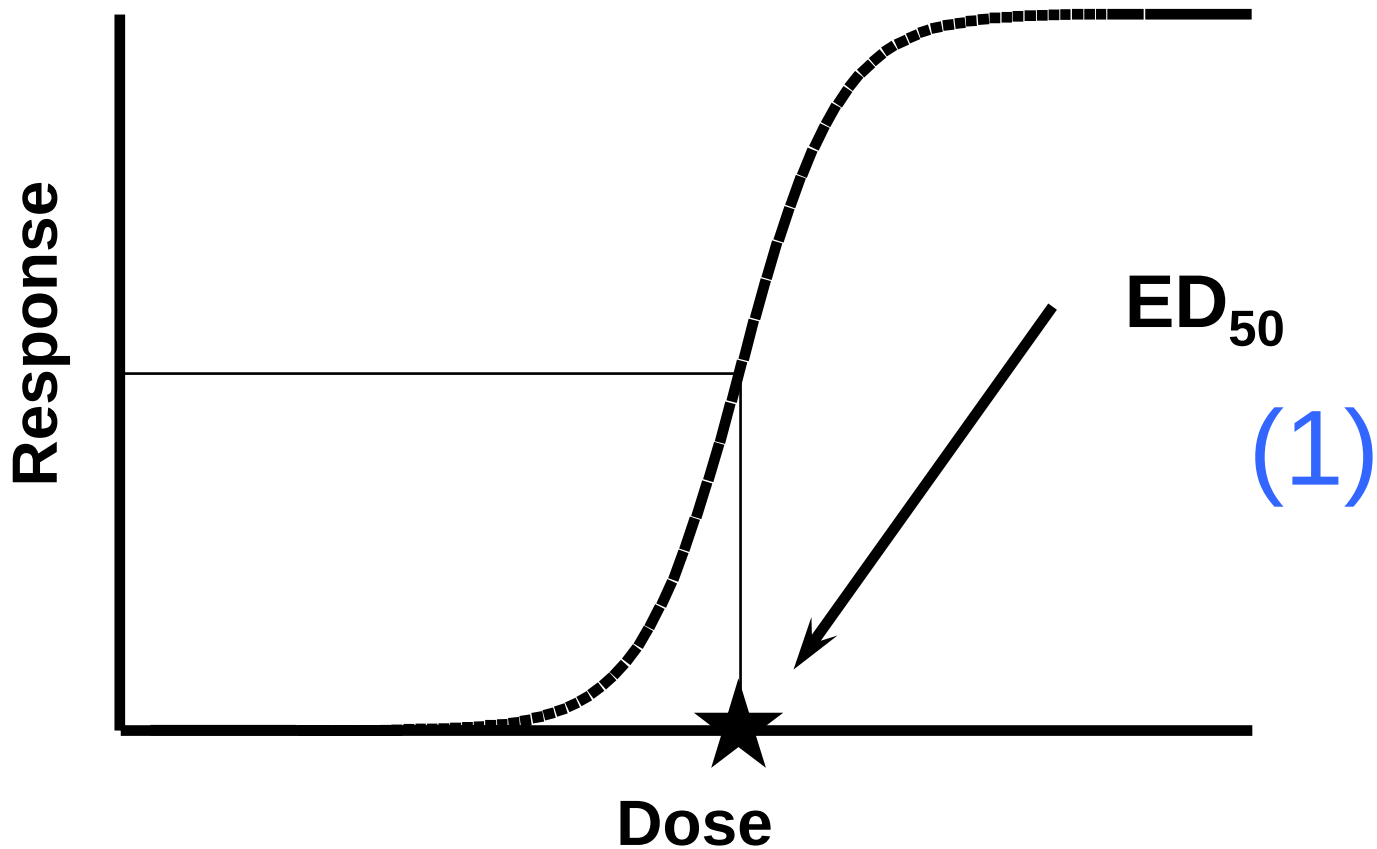

Other Methods - First Step

Specify what you are interested in

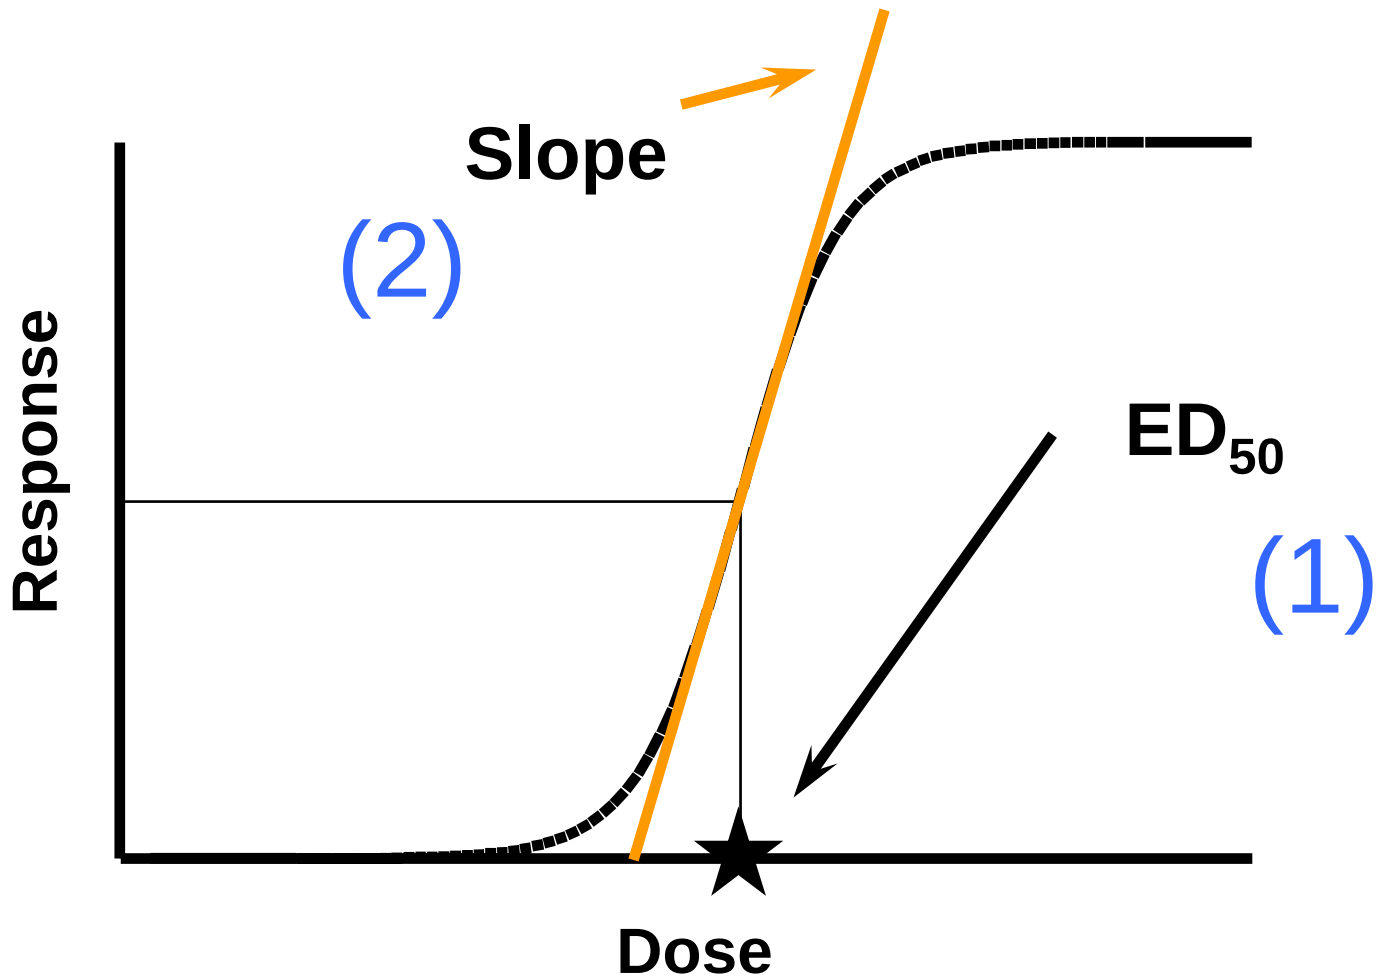

Other Methods - First Step

Specify what you are interested in

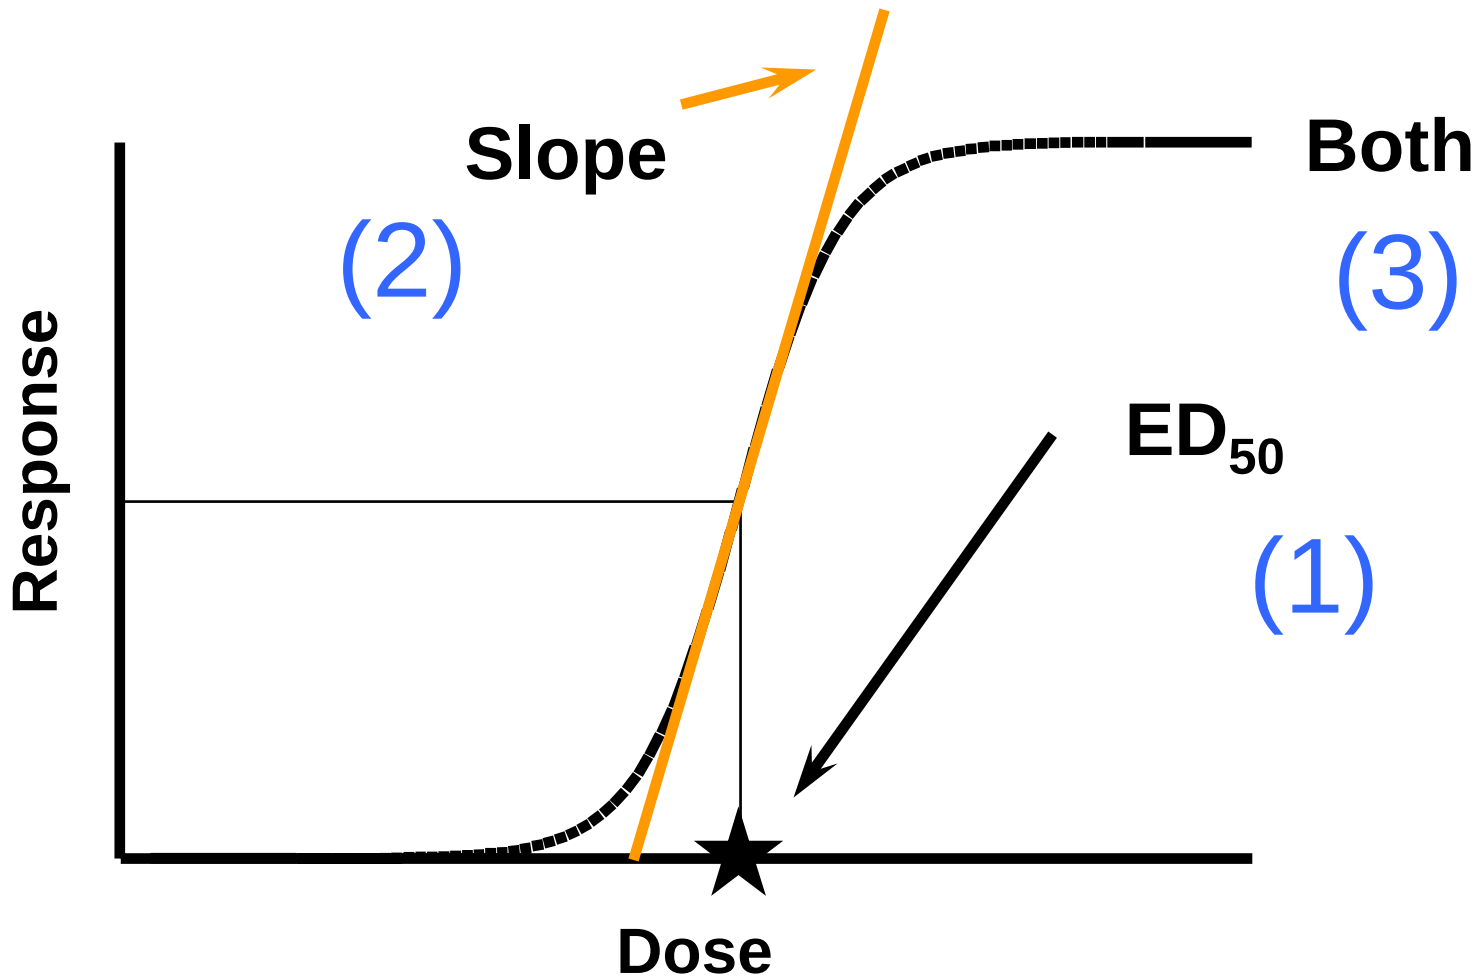

Second Step - Choose a Measure of Performance

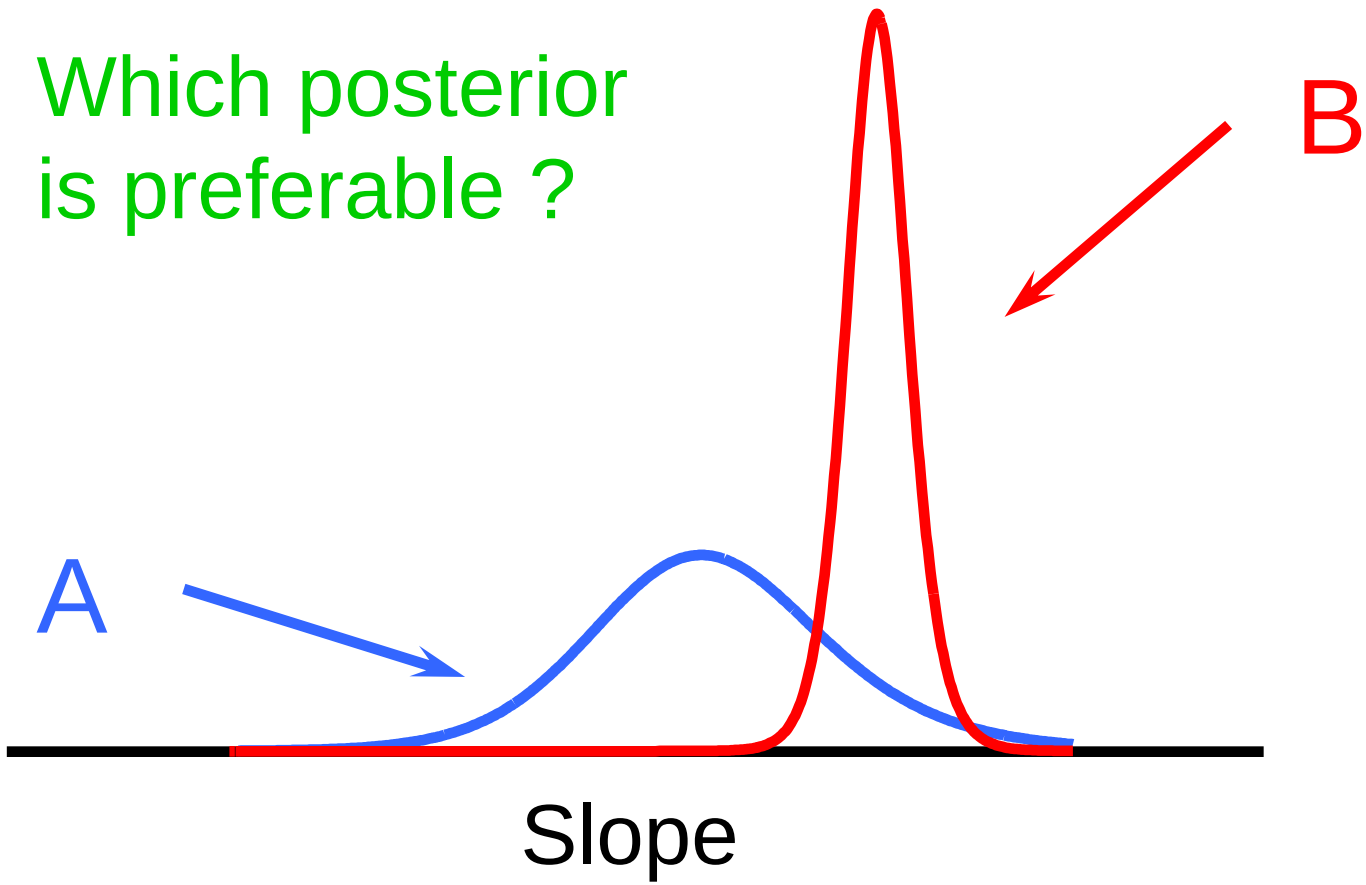

SRCC
Statistical Research & Consulting Center

Predictive Calculations

Dose (mcg)	Potential Data	Predictive Probability	Posterior Std. Dev.	Expected Std. Dev.
312.5	1	0.2608	0.8553	0.8187
	0	0.7392	0.8058	
62.5	1	0.1250	0.5673	0.9017
	0	0.8750	0.9495	
12.5	1	0.0096	0.4457	**0.6508**
	0	0.9904	0.6528	
3.125	1	0.0092	0.3842	0.8328
	0	0.0008	0.8370	
0.625	1	0.0094	0.3389	0.7824
	0	0.9906	0.7924	

Posterior Inference
Prob. Response at 312.5 mg

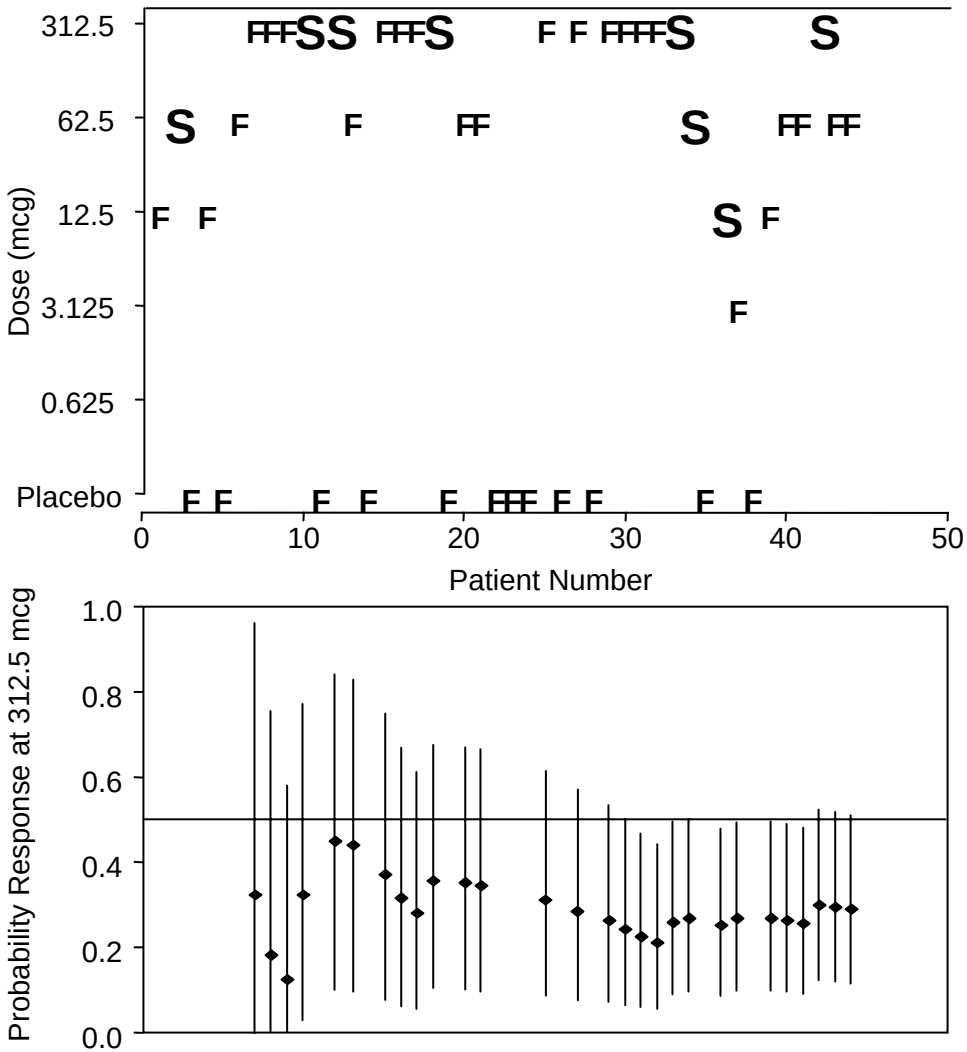

Continuous Reassessment Method (CRM)

Phase 1 Studies : The Continuous Reassessment Method (CRM

- Background in cancer studies
  - not on healthy volunteers
- low doses useless, high doses very toxic
- need to balance risks against potential benefits
- some prior information
- definition of MTD as dose at which a critical proportion π of patients suffer unacceptable toxicity
- adaptive dosing schedule to target in on a specified π
- local model of dose-response around π

CRM – Original Form

- Need set of doses d_i and prior estimates p_i of toxicity at each dose
- Re-label dose as: $p_i = \tanh(x_i + 1)/2$
- Giving the first estimate of dose response curve as :

$$x_i = \tanh^{-1}(2p_i - 1)$$

- Assume a local dose response curve :

$$\Pr(Y = 1 \mid x_i, \theta) = \left[\tanh(x_i + 1)/2\right]^{\theta}$$

where Y=1 if toxicity occurs

Statistical Research & Consulting Center

CRM – Original Form

- A “vague” prior $g(\theta)=\exp(-\theta)$ with mean 1 is assumed
- Suppose that you have a sequence of dose, response pairs (x_i,y_i) i=1, …, N are observed
- Posterior distribution for θ is

$$p(\theta \mid x,y) \propto \prod_{i}^{N} p(y_i \mid x_i,\theta)e^{-\theta}d\theta$$

Statistical Research & Consulting Center

CRM – Original Form

- Mean of the distribution is available to give information about θ
- Predictive probabilities

$$p_i = \int_a p(Y = 1 \mid x_i, \theta) p(\theta \mid \underline{x}, \underline{y}) d\theta$$

- Choose as next dose the one which gives p_i closest to the target π
- Continue until a pre-specified number of patients - final dose is the estimate

Issues with CRM

- simulations shown good performance
- designed for cancer trials but seems widely applicable
- needs more inputs (prior, defined MTD)
- critics of design have suggested stepped increments, repeated increments, starting from minimum possible dose
- critics have suggested logistic curves, non-parametric curves
- Can use cohorts, predefined stopping rules, eg if 6 patients treated with same dose stop.

CRM Dose Selection : Sedation of Infants / Cardiac Catheterisation

- Fabre et al (1998) - Br J Clin Pharm
- Aim : Find ED90 (90% sedated)
- Bayesian approach
- One parameter (α) logistic dose response
- Choose dose to "optimise" gain (utility) function
  - predictive probabilities

$$\pi_i = \int_a p(Y = 1 \mid x_i, \alpha) p(\alpha \mid \underline{x}, \underline{y}) d\alpha$$

- Choose as next dose the one which gives π_i closest to the target π (ED90)

CRM Design Infant Sedation

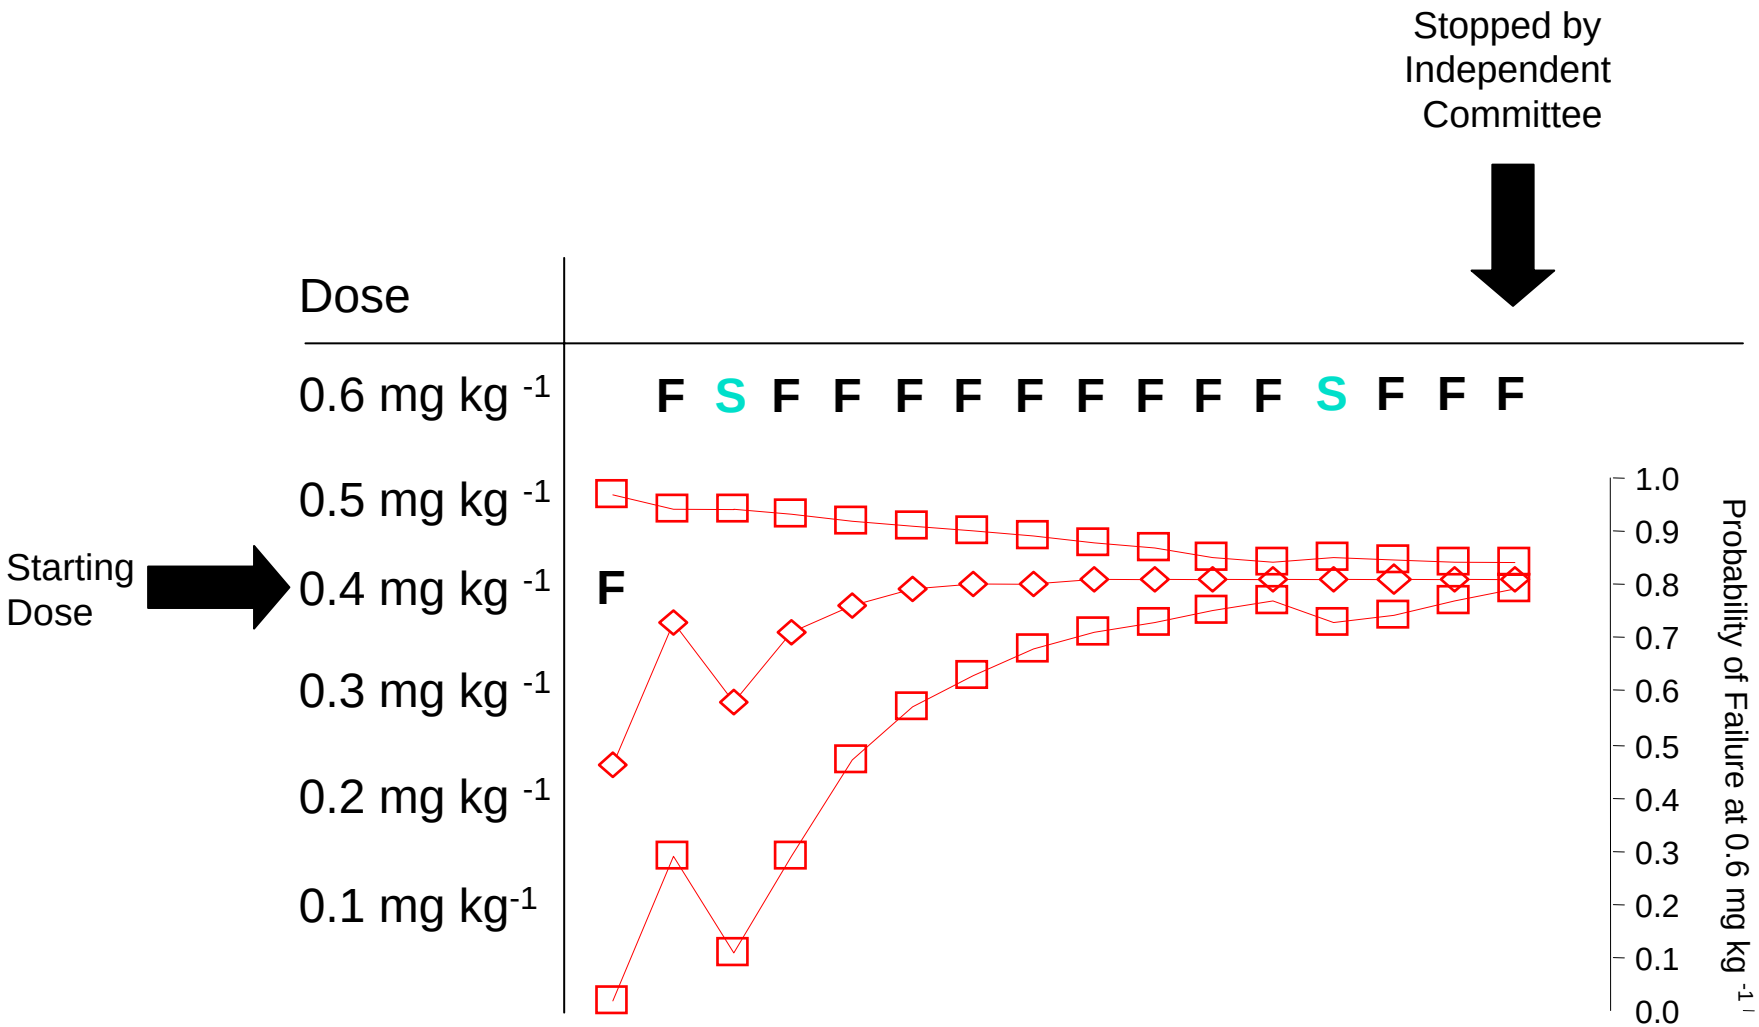

Fabre et al (1998) - Br J Clin Pharm

General Design

An Old Design Problem

- Non-linear response function
  - Optimal design available if we know the function
  - We don't know the function
- Solution :
  - Do some experiments
  - Learn a bit
  - Optimise
  - Learn a bit more
  - Optimise
  - etc

Dose Selection Standard design

- Placebo + 4 doses available where do I put them ?

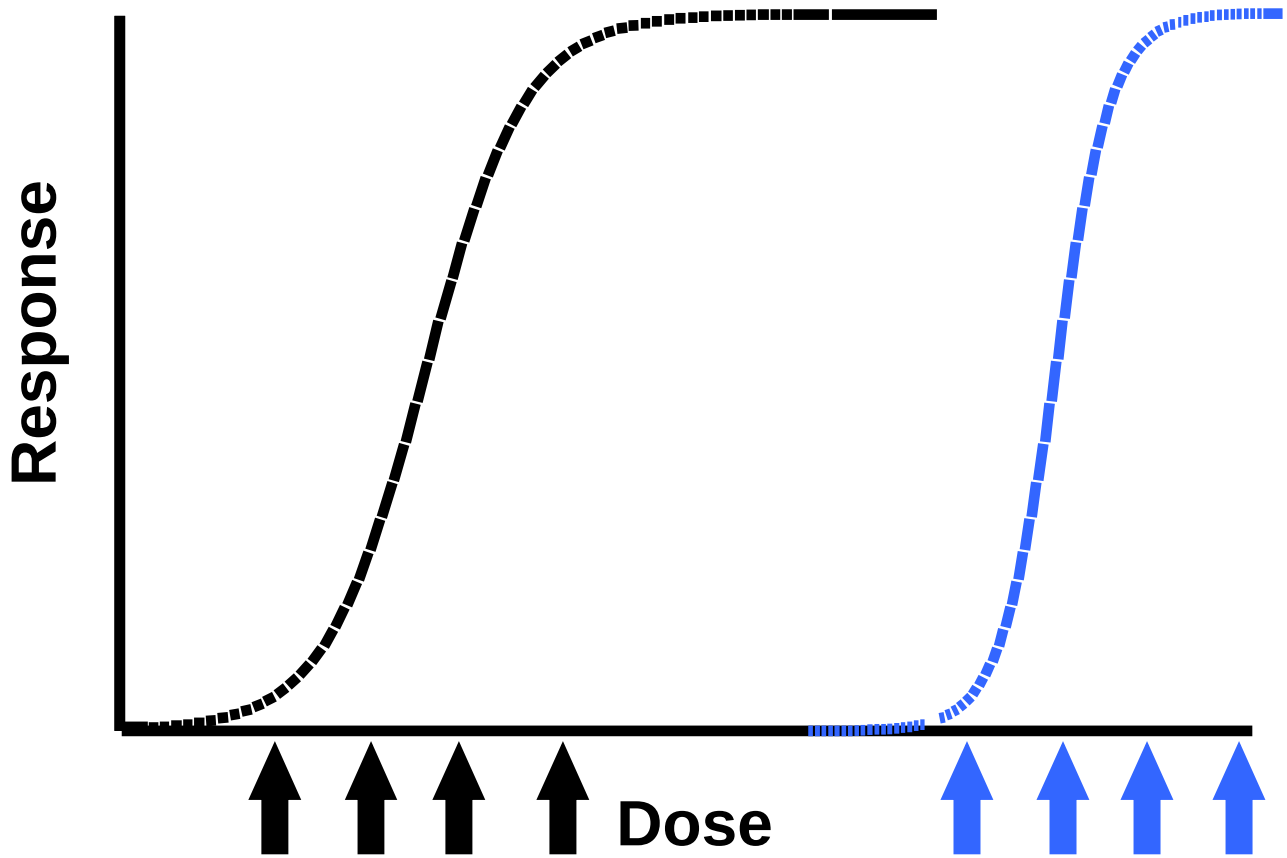

Issues in Dose Selection
Increase Number of Doses & Adapt

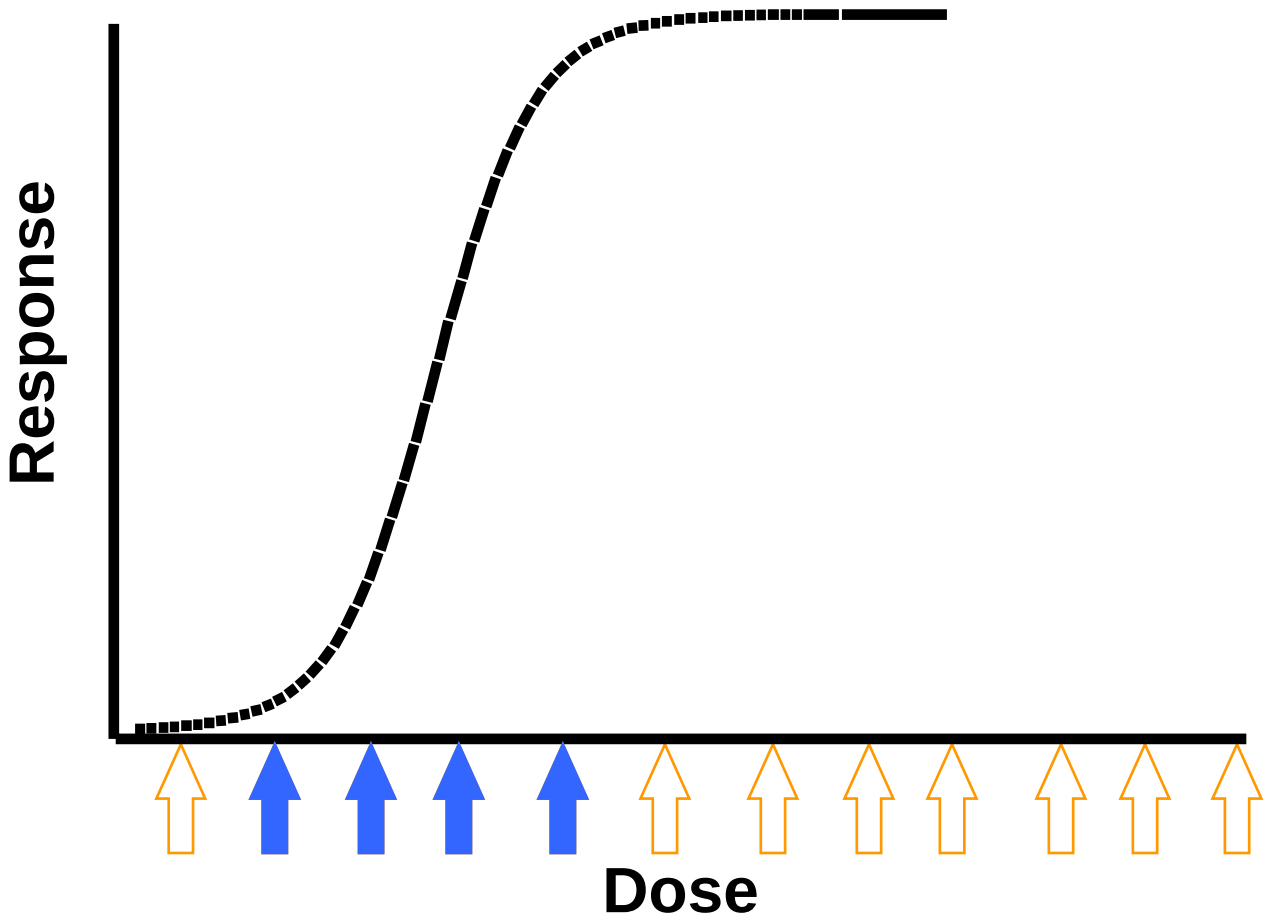

Improvements to Standard Design

- Increase number of doses (placebo + a large number - 15)
- Learn about dose-response (ED95) and adapt
- Prevent allocating patients to ineffective doses (ETHICAL)
- Model dose response
- Futility analysis / early decision making

- Trial 1
  - Dose-finding : is there a dose with sufficient efficacy to take into a confirmatory trial ? (ED95)
- Trial 2
  - Confirmatory : placebo controlled based on a single dose chosen in Trial 1 and sample sized based on learning about the size of effect and variability
- Independent or seamless

Design Process

Berry DA, Müller P., Grieve AP, Smith MK, Parke T, Blazek R, Mitchard N, Krams M (2002). Adaptive Bayesian designs for dose-ranging trials. Case Studies in Bayesian Statistics V, Springer, 99-181.

Statistical Research & Consulting Center

Design Process

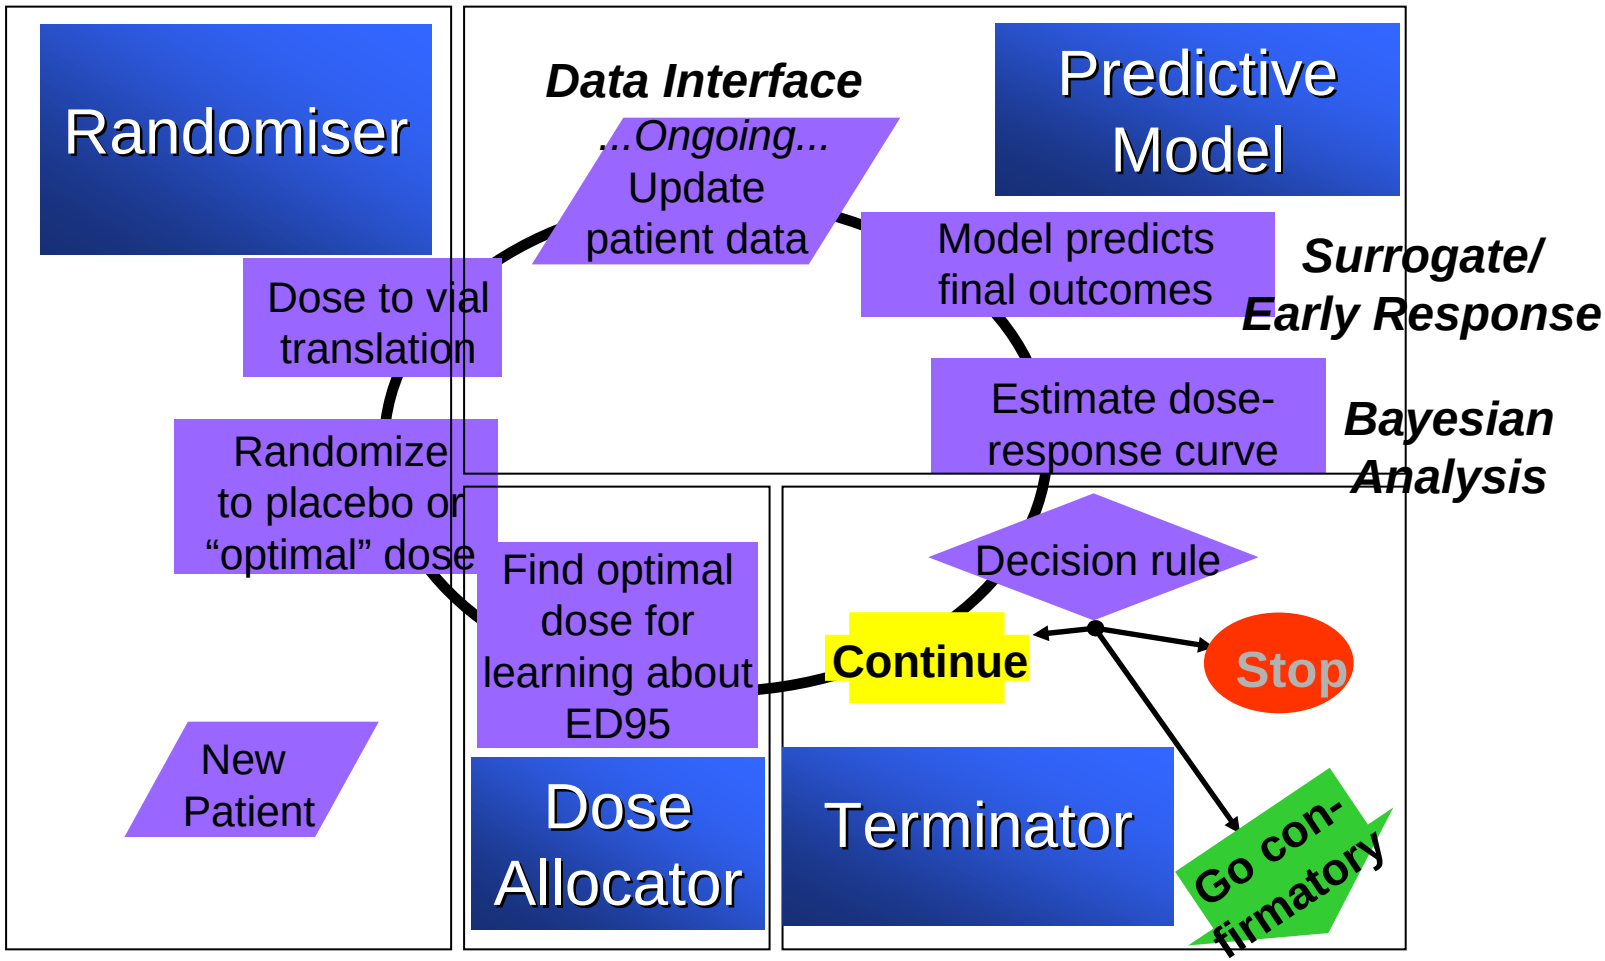

Building Predictive Models

Data from Copenhagen Stroke Data-Base

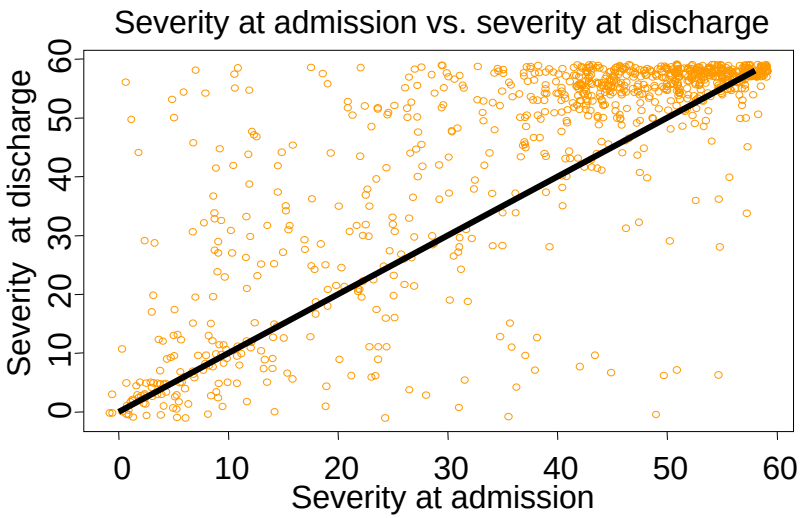

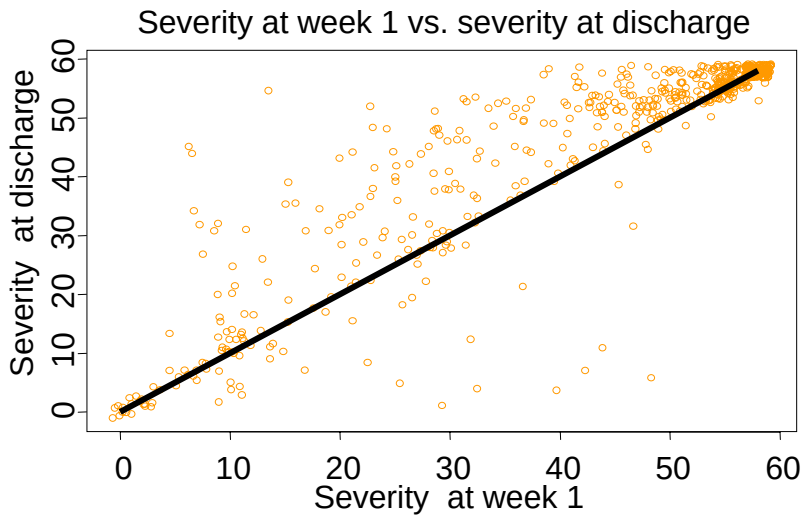

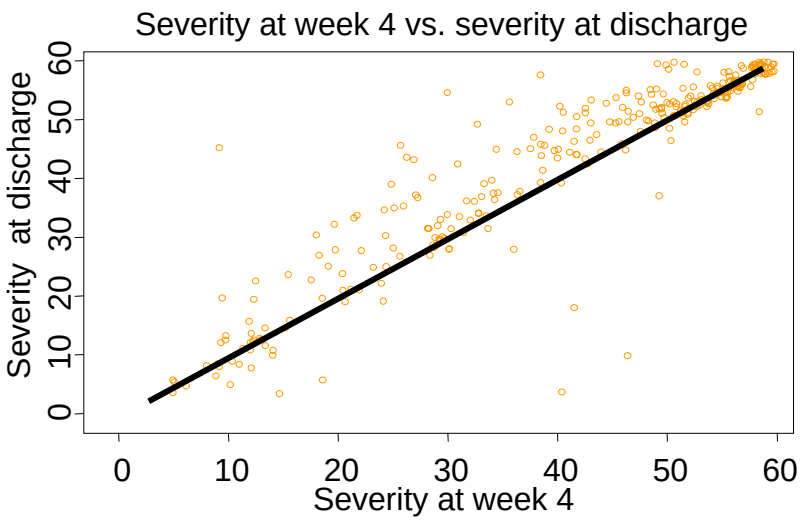

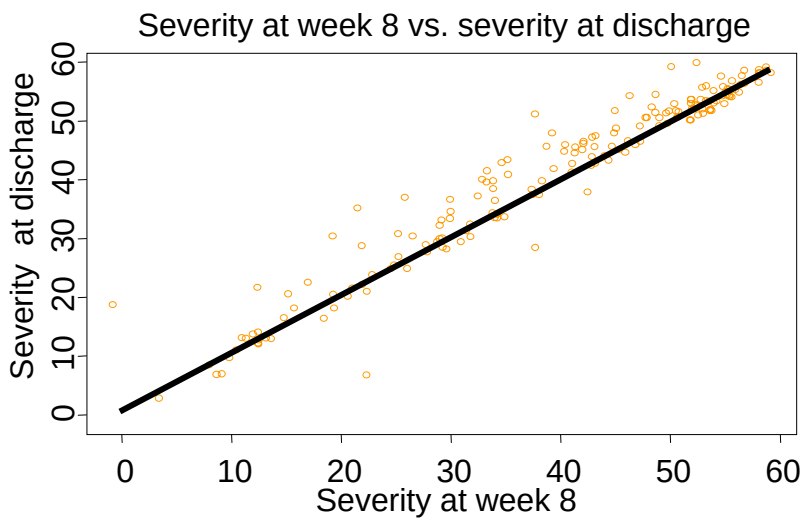

Issues

- How do we predict ?
  - Longitudinal model based on CSD
- How and what should we update ?
  - Dose response
  - Longitudinal model
- How do we model response ?
- Decisions
  - How do we choose a dose ?
  - How do we stop ?

The Dose-Response Curve : $\theta_j = f(z_j, \theta)$

Mean Response

Dose

Parameters

- Requirements
  - To model f (z ,θ) we need :

1. a flexible model, allowing non-monotone curves. and allocator)
2. analytical posterior updating (simulation required for terminator and allocator)
3. efficient (analytic) computation of expected utilities

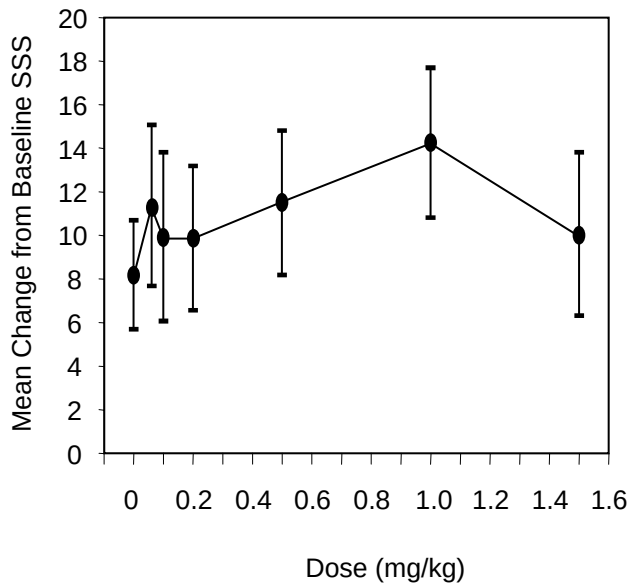

- Possibilities
  1. Splines
  2. Kernel Regression
  3. Normal Dynamic Linear Model

Normal Dynamic Linear Model

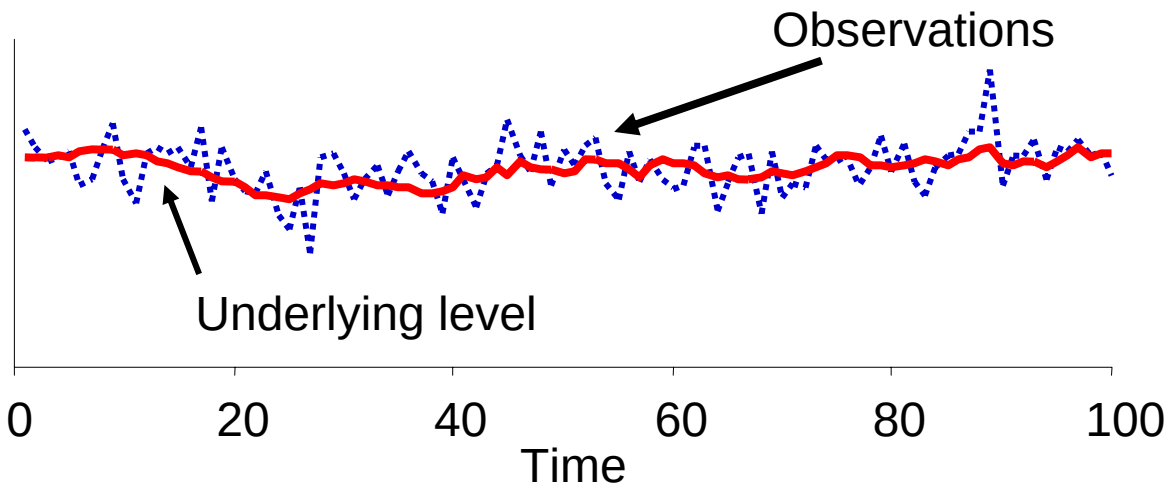

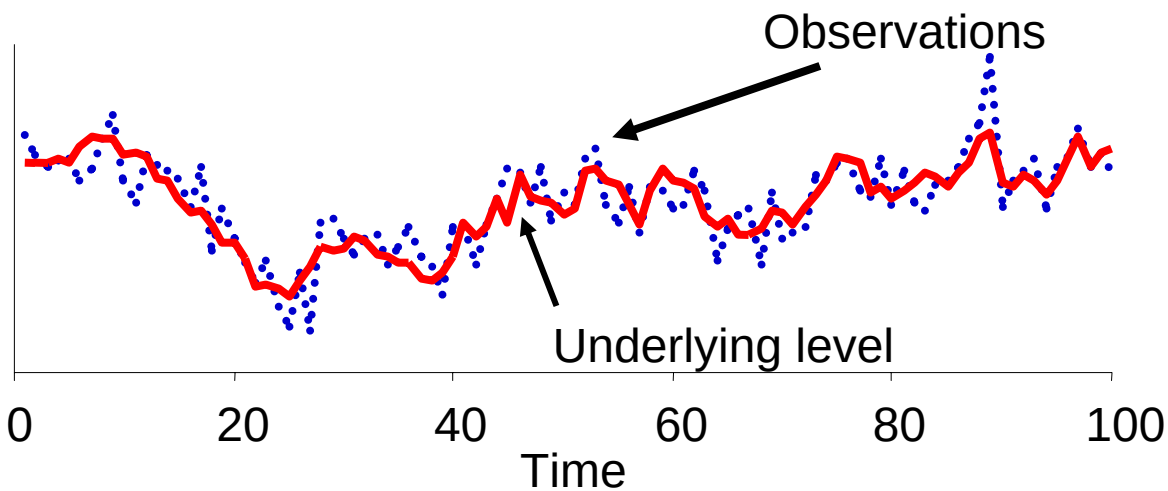

Normal Dynamic Linear Model

- Simplest NDLM – 1st Order polynomial
- Most of the important concepts and features
- Characteristics of NDLM's
- Observation Equation : $Y_t = \mu_t + \nu_t \ , \ \nu_t \sim N(0,V_t)$

  Observation Error
- System Equation : $\mu_t = \mu_{t-1} + \omega_t \ , \ \omega_t \sim N(0,W_t)$

  Random Walk

  Evolution Error
- Forecast Function : $E(Y_{t+k} \mid D_t) = E(\mu_t \mid D_t) = m_t$

  Posterior Mean

Modelling Dose Response

- We model f (z , θ) as a 2nd order polynomial NDLM (West and Harrison 1997):

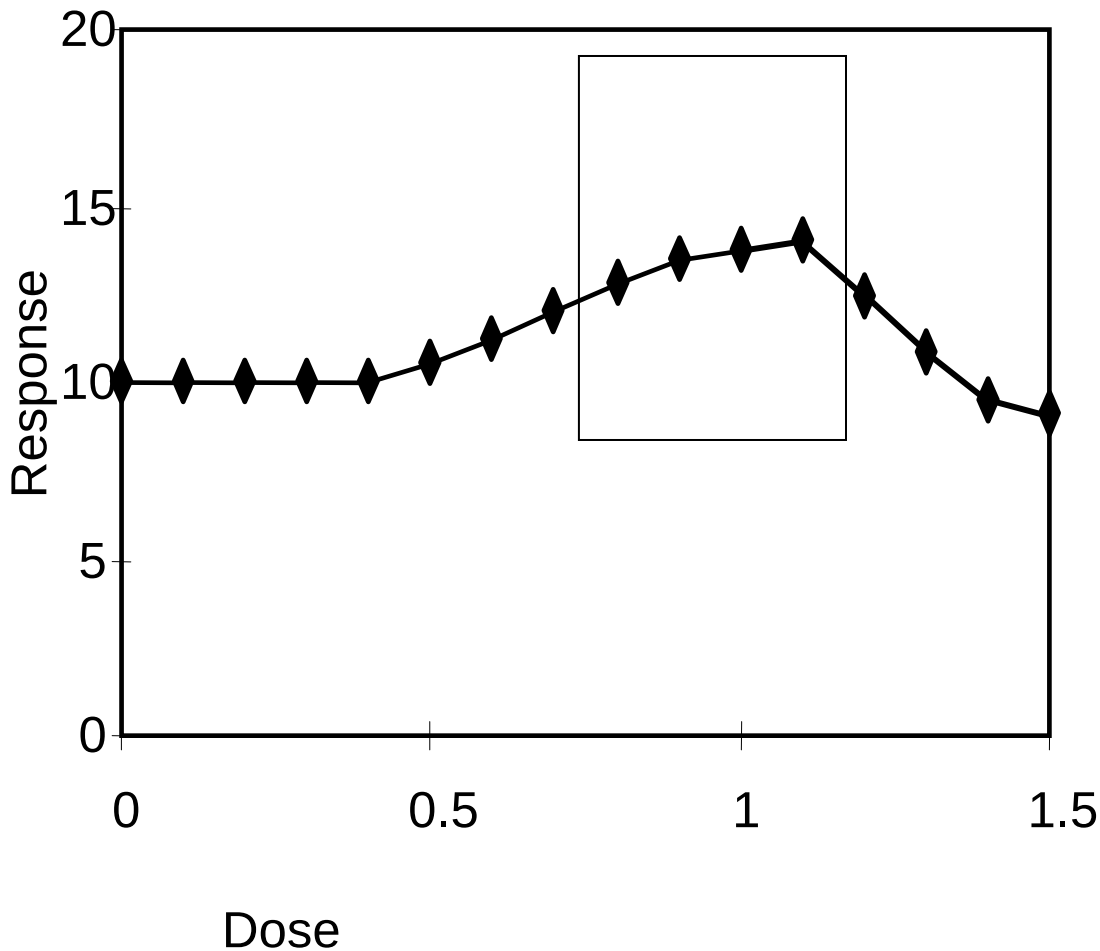

2nd Order Polynomial NDLM

Locally around z = Z_j a straight line with level θ_j and slope δ_j

Parameters of the straight lines change between doses by adding a (small) evolution noise.

$$\theta_{j+1} = \theta_j + \delta_j + \omega_j$$
$$\theta_{j+1} = \theta_j + \delta_j$$

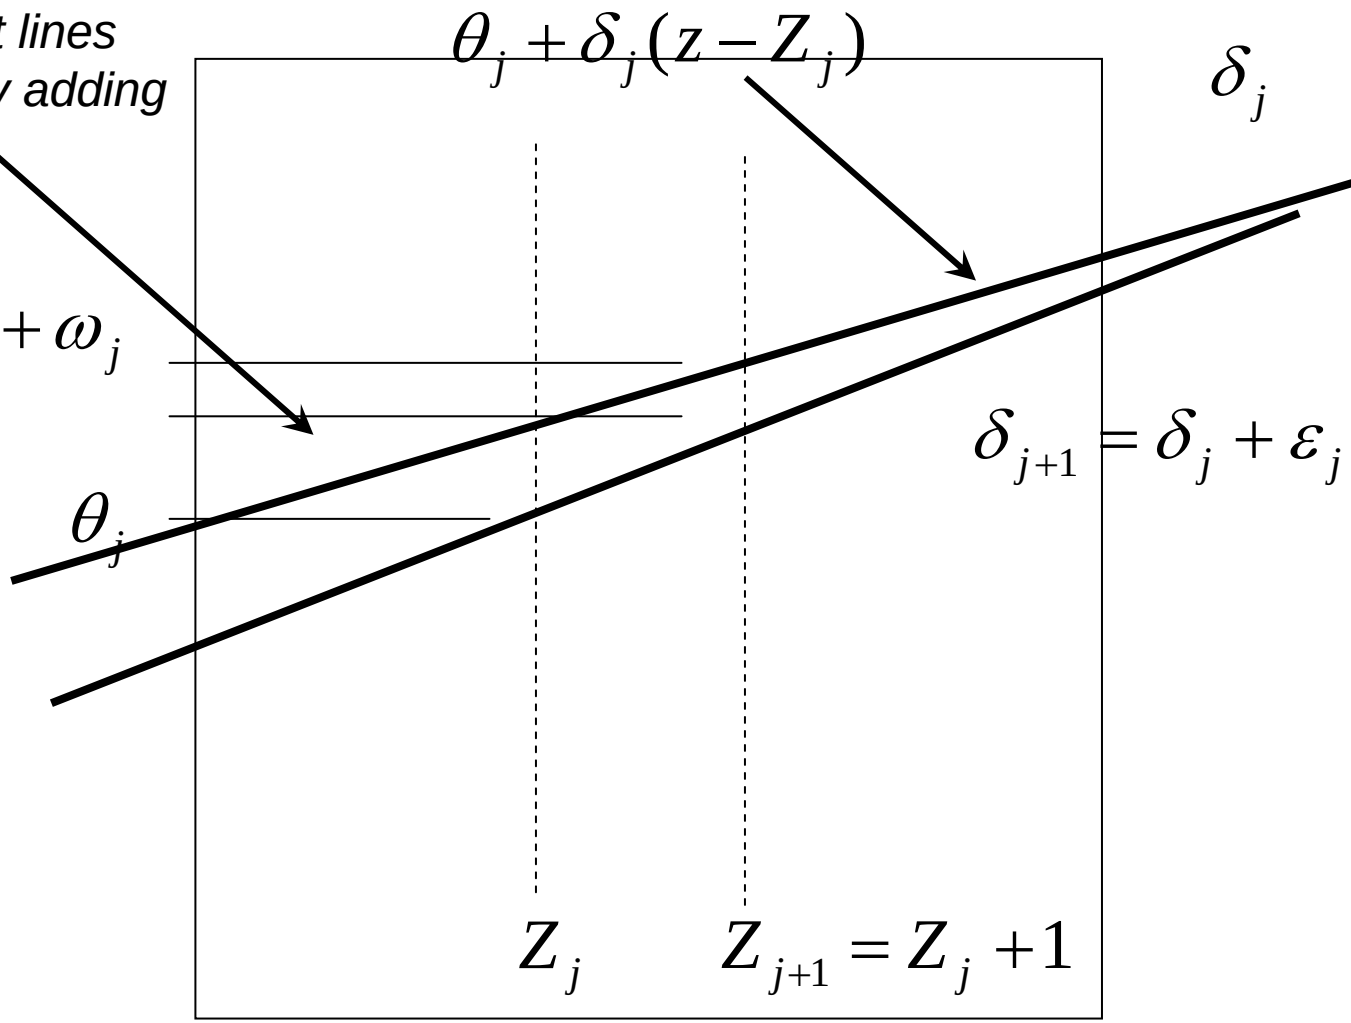

Evolution Variance = Smoother

Normal Dynamic Linear Model

- NDLM – 2nd Order polynomial
- Observation Equation : $Y_{jk} = \mu_j + \nu_{ik}$, $\nu_{ik} \sim N(0, V\sigma^2)$
- System Equation : $\mu_j = \mu_{j-1} + \delta_{j-1} + \omega_j$, $\omega_j \sim N(0, W_j\sigma^2)$

  $\delta_j = \delta_{j-1} + \varepsilon_j$, $\varepsilon_j \sim N(0, W_j\sigma^2)$
- Issues
  - Choice of W_i - in our study fixed
    - can learn about it
  - Covariates can be included by making the expected responses depend linearly on the covariates
  - $E(y_{jk} \mid z = Z_j , x_k) = \theta_j + \beta \times x_k$
  - The NDLM is then applied to these θ_j's

Decision Problem 1 : Dose Allocation

- df (z,θ) advantage over placebo at dose z, using a curve parameterized by θ.
- z^* dose which achieves 95% of possible improvement over placebo (ED95)
- Utility $u(z, \tilde{y}, x, \tilde{D}, D) = -Var[df(z^*, \theta) \mid D, \tilde{D}, \tilde{y}, x, z]$
- Where :
  - x : covariates of a new patient
  - z : the assigned dose
  - $\tilde{y}$: predicted response of a new patient
  - D : data
  - $\tilde{D}$: missing data (missing final response)

- Expected Utility :

$$U(z, x, D) = \int_{\tilde{y}, \tilde{D}} u[z, \tilde{y}, x, \tilde{D}, D] \times p(\tilde{D} \mid D)\, p[\tilde{y} \mid D, z]\, d\tilde{D}\, d\tilde{y}$$

- Substitute average value : $x \equiv \bar{x}$
- Approximate Expected Utilities :

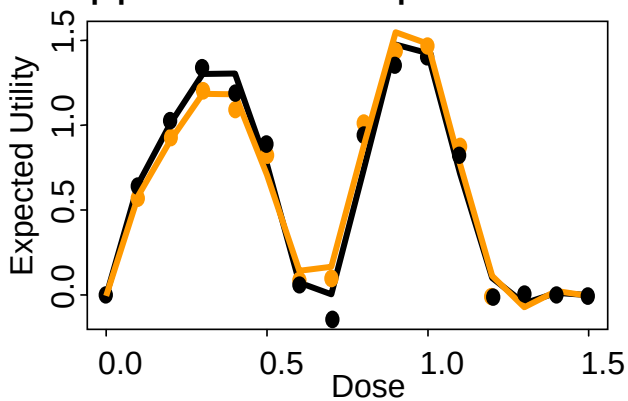

- Maximise expected utilities as a function of dose

Randomisation

- Based on current information the “optimal” dose d* is chosen to :
  - minimise expected variance of ED95
  - minimise expected variance of response at ED95
- If d* not placebo then placebo is assigned with some minimum probability :
  - 10%, 15%, 20%
- The assigned dose is selected randomly from within all doses for which the expected response is within
  - 5%, 10% that of d*

Decision Problem 2 : Early Stopping

- Formal Bayesian decision theoretic approach

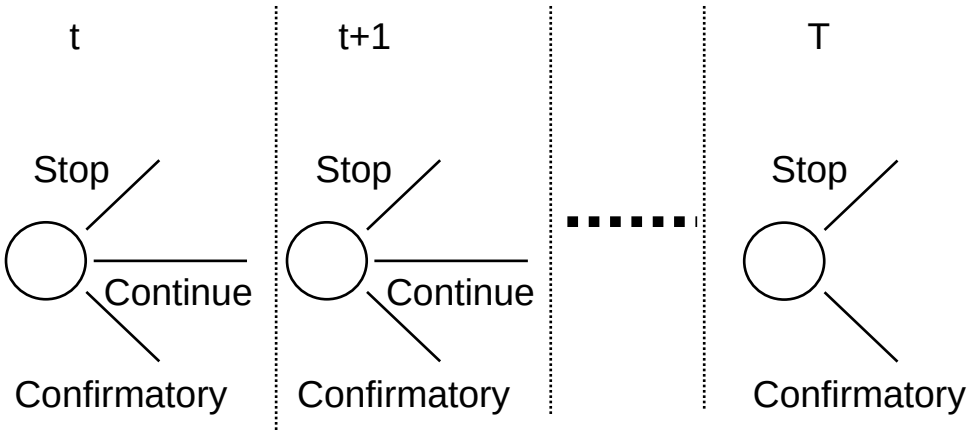

"... if one decision leads to another, then to analyse the former, one first needs to analyse the latter, since the outcome of the former depends on the choice of the latter."

Simon French - Readings in Decision Analysis
Chapman & Hall, 1989

- Numbers of scenarios exponentially increasing
- Expectations analytically intractable
- Computationally intensive
- Approximations have been developed
- Regulatory attitude
  - FDA : "Our regulations state that we are only to consider safety and effectiveness(efficacy) in determining whether a medical device (drug) can be marketed."

    "Therefore, the cost per observation, because it involves the cost per device, cannot be considered in our evaluation."
- Posterior Probabilities of Clinical Importance

Dose Effect Curve

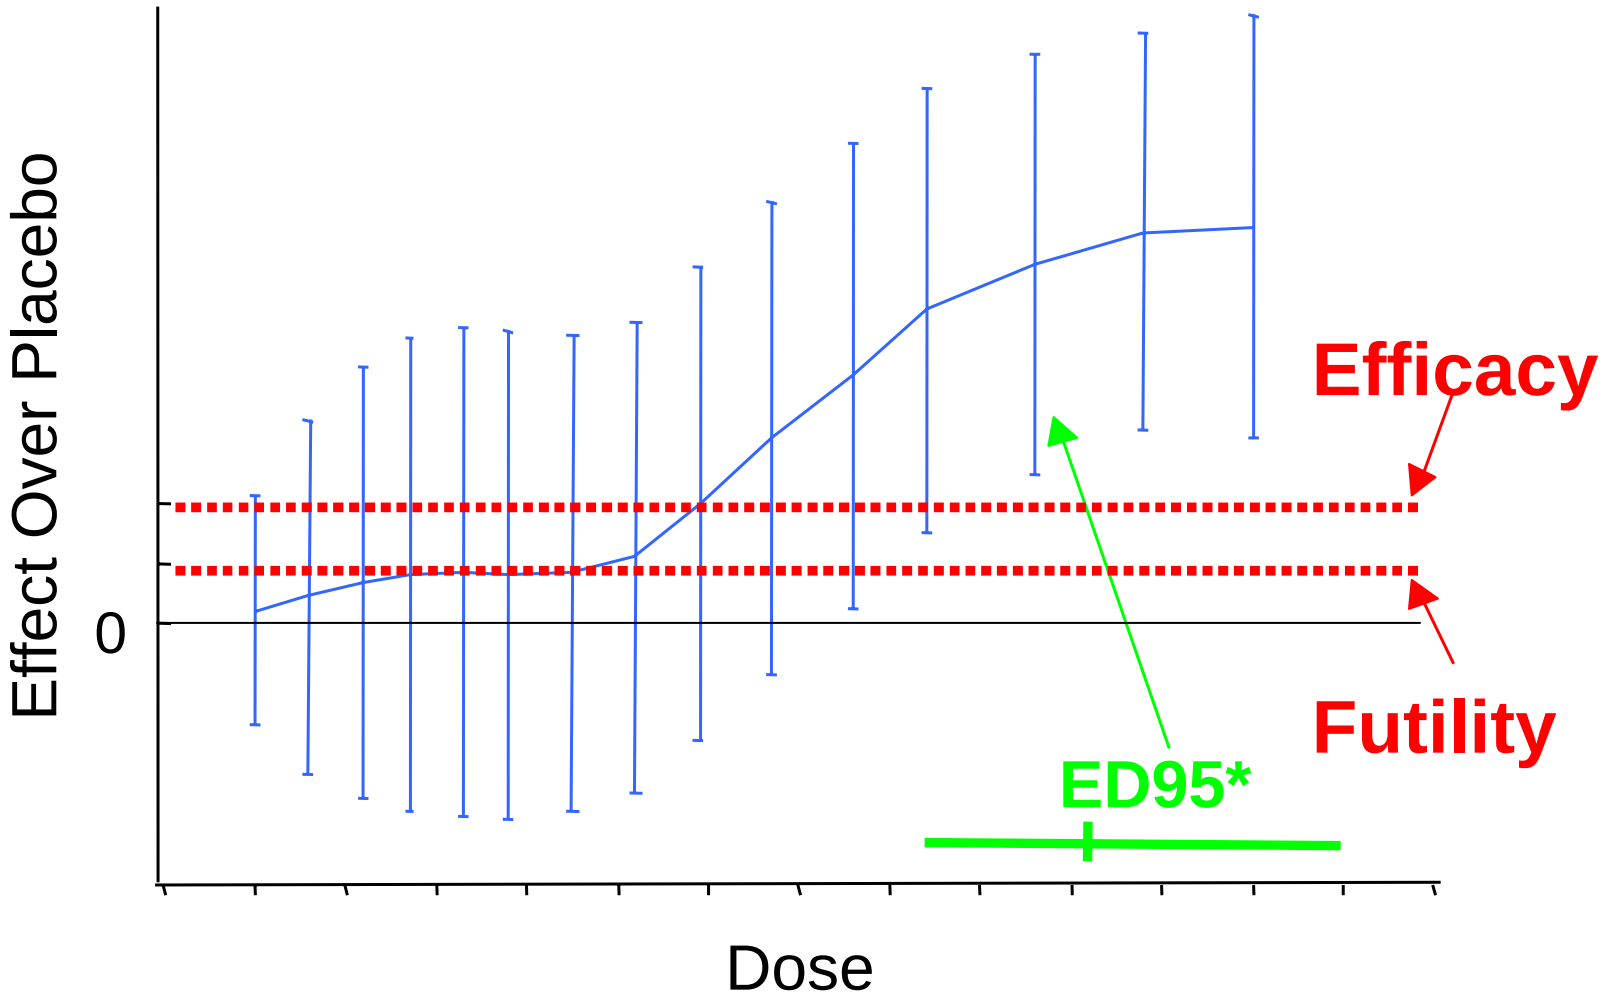

Next Steps

- Design Developed in theory
- Can we run it?
- Not Yet !
  - Sell it
    - Management
    - Regulators
  - Validation
    - Computer System
    - Algorithm

Statistical Research & Consulting Center

Astin Results

Dose Selection Designs Issues / Generalisations

- Time to response
  - Long-time to response - surrogate measures ?
- Group of patients vs individuals
  - Possible to allocate cohorts of patients
- Response Type : Binomial, Poisson, Ordered categorical
- Covariates : age/gender/severity
  - not used for allocation
- Dose interval restriction
  - removed
- "Accrual" Bias
  - Chronic diseases
  - Knowledge of design can cause bias
  - Later patients > prob of optimal treatment
    - Delayed entry into study
- Time Bias
  - Population Drift
- Practical set-up costs
- Need data quickly
- More regulatory experience
- More need to be tried before more general acceptance

Statistical Research & Consulting Center

SKB Bayesian Approach to FIM Studies

First in Man

- Placebo-controlled
- 4-Period Crossover
- 2-6 cohorts of 4 healthy volunteers
- 1st cohort :

Period	Sub 1	Sub 2	Sub 3	Sub 4
1	d_1	d_1	d_1	Pla
2	d_2	d_2	Pla	d_1
3	d_3	Pla	d_2	d_2
4	Pla	d_3	d_3	d_3

Statistical Research & Consulting Center

First in Man

- Study may be terminated, or dosing regimen altered
  - if volunteers exceed ***pre-specified*** exposure level (AUC or CMAX)
  - an unacceptable adverse event profile seen
- If safe to continue next cohort with doses d_3,d_4 and d_5 + Pla.
- Continuation until completion of planned cohorts or unacceptable safety
- d_1 and d_{max} based on tox and pre-clinical data

First in Man Model for Data

- yij [log(AUC or CMAX)] = q1+q2log(dij) +si+eij
  - i - subject
  - j - dose (to ith subject)
  - si - random subject effect
  - ei - random error term
  - q1 - intercept
  - q2 - slope

First in Man
Model for Data

- y_{ij} [log(AUC or CMAX)] = $\theta_1+\theta_2 \log(d_{ij}) + s_i+\varepsilon_{ij}$
  - i - subject
  - j - dose (to ith subject)
  - s_i - random subject effect
  - ε_i - random error term
  - θ_1 - intercept
  - θ_2 - slope

First in Man Bayesian Analysis

- Prior - based on imaginary data - choice determines speed of escalation
- Uses PROC MIXED in SAS - point estimates are posterior modes
- Dose escalation
  - given doses $d_1, d_2, \ldots, d_k$
  - prior information, data from previous cohorts, data from previous periods in current cohort
  - choose 3 real doses to administer in next period

Use of Safety Constraint

- Suppose a limiting level L [log(AUC or CMAX)] pre-specified - larger values to be avoided
- A candidate dose d_f should satisfy
  - $P(y_{if}>L) <= c_0$
  - where y_{if} is the future response corresponding to d_f
  - this gives a set of acceptable doses
  - The dose d_f^* which gives equality is the MTD

First in Man
Use of Safety Constraint

- Suppose a limiting level L [log(AUC or CMAX)] pre-specified - larger values to be avoided

- A candidate dose d_f should satisfy
  - $P(y_{if}>L) <= c_0$
  - where y_{if} is the future response corresponding to d_f
  - this gives a set of acceptable doses
  - The dose d_f^* which gives equality is the MTD

First in Man
Choice of Dose

- Amongst acceptable doses choose
  - maxsafe : give each subject maximum safe dose
  - optsafe : give that combination which optimses learning about θ_1 and θ_2

Adaptive Randomisation Giles et al, JCO(2003)

- Troxacitabine (T) in acute myeloid leukemia (AML) combined with cytarabine (A) or idarubicin (I)
- Adaptive randomization to:
  IA vs TA vs TI
- Max n = 75
- End point: Time to CR (< 50 days)

Adaptive Randomization

- Assign 1/3 to IA (standard) throughout (unless only 2 arms)
- Adaptive to TA and TI based on current results

  - Time to success : Exponential
  - Prior(Median : m_i)=Gamma(2.001,4.624) (i=0,1,2)
  - Initial randomisation : $\pi_0=\pi_1=\pi_2=1/3$
  - Define : $q_1=P(m_1<m_0|data)$, $q_2=P(m_2<m_0|data)$, $r=P(m_1<m_2|data)$

Adaptive Randomization

- Assign 1/3 to IA (standard) throughout (unless only 2 arms)
- Adaptive to TA and TI based on current results
- Results →

Statistical Research & Consulting Center

Adaptive Randomization

	Probability Assign to:				
Pat.	IA	TA	TI	Arm	CR<50
1	0.33	0.33	0.33	TI	NOT
2	0.33	0.34	0.32	IA	CR
3	0.33	0.35	0.32	TI	NOT
4	0.33	0.37	0.30	IA	NOT
5	0.33	0.38	0.28	IA	NOT
6	0.33	0.39	0.28	IA	CR
7	0.33	0.39	0.27	IA	NOT
8	0.33	0.44	0.23	TI	NOT
9	0.33	0.47	0.20	TI	NOT
10	0.33	0.43	0.24	TA	CR
11	0.33	0.50	0.17	TA	NOT
12	0.33	0.50	0.17	TA	NOT
13	0.33	0.47	0.20	TA	NOT
14	0.33	0.57	0.10	TI	NOT
15	0.33	0.57	0.10	TA	CR
16	0.33	0.56	0.11	IA	NOT
17	0.33	0.56	0.11	TA	CR

	Probability Assign to:				
Pat.	IA	TA	TI	Arm	CR<50
18	0.33	0.33	0.33	TA	NOT
19	0.33	0.34	0.32	TA	NOT
20	0.33	0.35	0.32	IA	CR
21	0.33	0.37	0.30	IA	CR
22	0.33	0.38	0.28	IA	CR
23	0.33	0.39	0.28	IA	CR
24	0.33	0.39	0.27	IA	CR
25	0.87	0.13	0	IA	NOT
26	0.87	0.13	0	TA	NOT
27	0.96	0.04	0	TA	NOT
28	0.96	0.04	0	IA	CR
29	0.96	0.04	0	IA	NOT
30	0.96	0.04	0	IA	CR
31	0.96	0.04	0	IA	NOT
32	0.96	0.04	0	TA	NOT
33	0.96	0.04	0	IA	NOT
34	0.96	0.04	0	IA	CR

Summary of results

CR < 50 days:

- IA: 10/18 = 56%
- TA: 3/11 = 27%
- TI: 0/5 = 0%

RPW Depression

- Tamura et al (1994) - JASA
- Primary Endpoint : Δ HAMD (8 weeks)
  - Is an adaptive design feasible ?
  - Surrogate end-point : > 3 weeks therapy, 50% ↓ HAMD in 2 consecutive visits
- Stratification Factor (2 levels)
- Ind. Urns within each strata
- Rand. block : 1st 6 pts in each stratum
- Data Collection by telephone
- Pat. Status determined by ind. CRA
- Urn updated by 2nd ind. CRA
- Randomisation schedule generated by CRA
- Allowed for multiple patients before next update

RPW Depression

- Results Important Stratum - Mean(se)

Endpoint	Placebo	Active
⊿HAMD	-5.5(1.6)	-11.4(1.2)

Bayes Posterior Prob (Active > pla) =0.003

- Interestingly : nearly equal allocation

Tamura et al - Conclusions

- Experience generally +ve
- Investigator enthusiasm - rapid accrual
- Need for automation to reduce burden of administrating design
- Encourage others to try