Adaptive trial design for early phase trials

Model based dose-finding

Dr Thomas Jaki Dr Lisa Hampson

Lancaster University

Study specifications

Dose schedule: $d_1 < \cdots < d_k$

Response: $x = \begin{cases} 1 \text{ for toxic response} \\ 0 \text{ otherwise} \end{cases}$

Structure: treat successive cohorts of c subjects

Objective: find the “highest safe dose”

Phase I cancer trials

- Typically implemented as adaptive designs:

Doses	1	2.5	5	10	15	20	25
# patients	3	3	3				
# DLT	0	0	1				

- What next?
  - More patients on same dose (5)
  - More patients on lower dose (2.5)?
  - Patients on higher dose (10)?

$\Rightarrow$ Rule such that most patients on MTD, few overexposed

Standard Phase I cancer designs

- "3+3" designs
  - standard, simple up-and-down design
  - no statistical inference, simple data-based rules
  - very popular among clinicians
  - bad statistical properties ("operating characteristics")
- Continual Reassessment Method (CRM)
  - very popular among statisticians
  - good operating characteristics
  - bad on-study properties (non-intuitive dose recommendations)
- Bayesian adaptive dose-response escalation strategies
  - model based using flexible dose-response model
  - Usually very balanced dose-recommendations

Review of advantages Jaki *et al*. (2013)

Seeking a quantile

MTD – maximal dose acceptably tolerated by a particular patient population

$\rightarrow$ vague

$TD100\pi$ – dose at which the probability of toxicity is π (for $0 < \pi < 1$), e.g. TD20

$\rightarrow$ more specific

TD20

$$p(d) = P(\text{toxicity}|\text{dose } d)$$

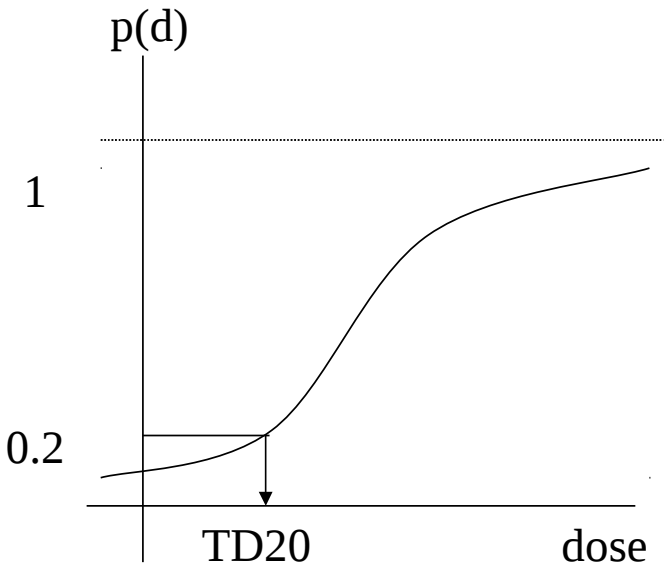

Assume that a 20% risk of toxicity is an acceptable risk to pay for a chance of benefit

General (Bayesian) approach

1. Make assumptions about the form of $p(d)$
2. Impose a prior distribution for the parameters that determine $p(d)$
3. Choose next dose to optimise treatment for this patient
4. Stop once target dose level can be estimated accurately enough

Prior distribution

Two possible purposes:

1. To include relevant additional information
   - Related trial data
   - Expert opinion
2. To control operating characteristics
   - Typically pessimistic - reflecting fears rather than beliefs
   - Often useful to use frequentist final analysis

Doses

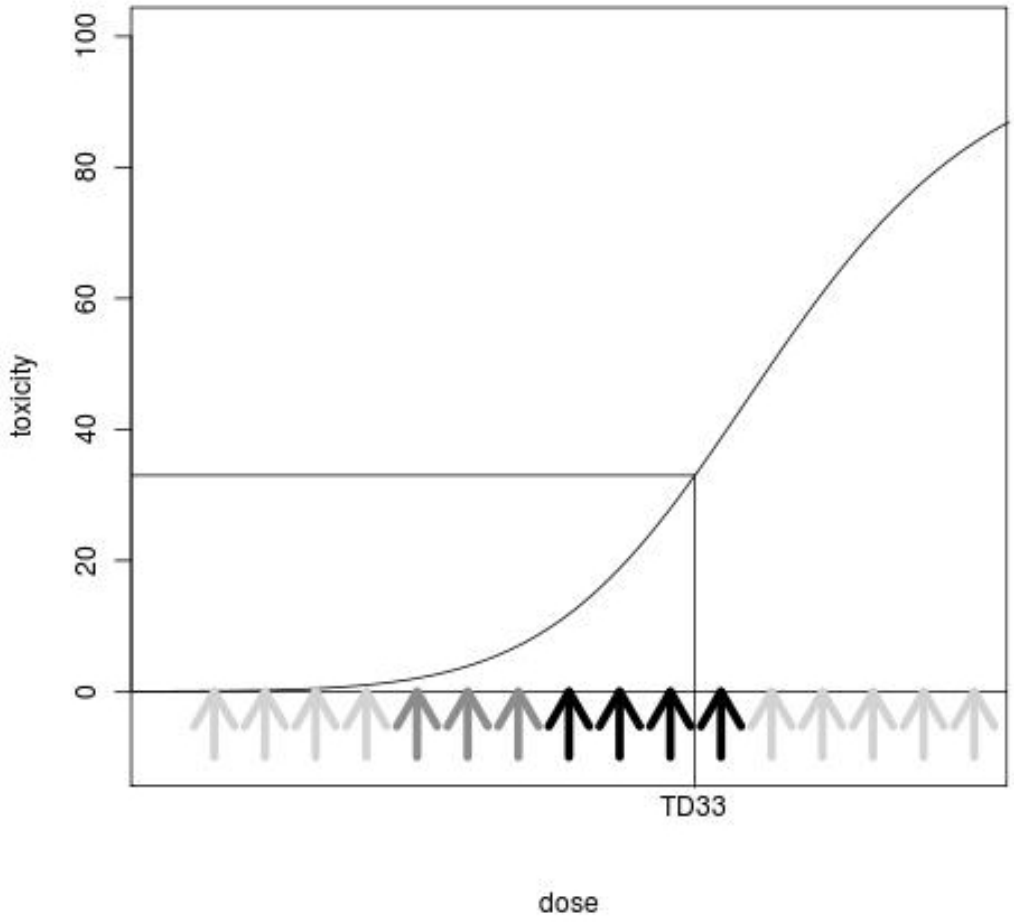

- Want most patients treated at doses with dark arrows
- Useful to have many doses available
- Skipping doses should be considered

Continual Reassessment method (CRM)

O'Quigley et al (1990)

Dose schedule: $d_1 < \cdots < d_k$

Response: $x = \begin{cases} 1 \text{ for toxic response} \\ 0 \text{ otherwise} \end{cases}$

Objective: find TDϑ

Cohort size: 1

Model

One parameter log-log model

$$p(d_i) = \pi_i^{\theta},\ i = 1, \ldots, k$$

where π_i is the fixed prior guess at the probability of toxicity at d_i, such that

$$\pi_1 < \pi_2 < \cdots < \pi_k$$

A Bayesian prior for θ is imposed: $\theta \sim Exp(1)$ so that $E_0(\theta) = 1$

Representation of the model

Starting values for the π_i

π_1	π_2	π_3	π_4	π_5	π_6
0.05	0.10	0.20	0.30	0.50	0.70

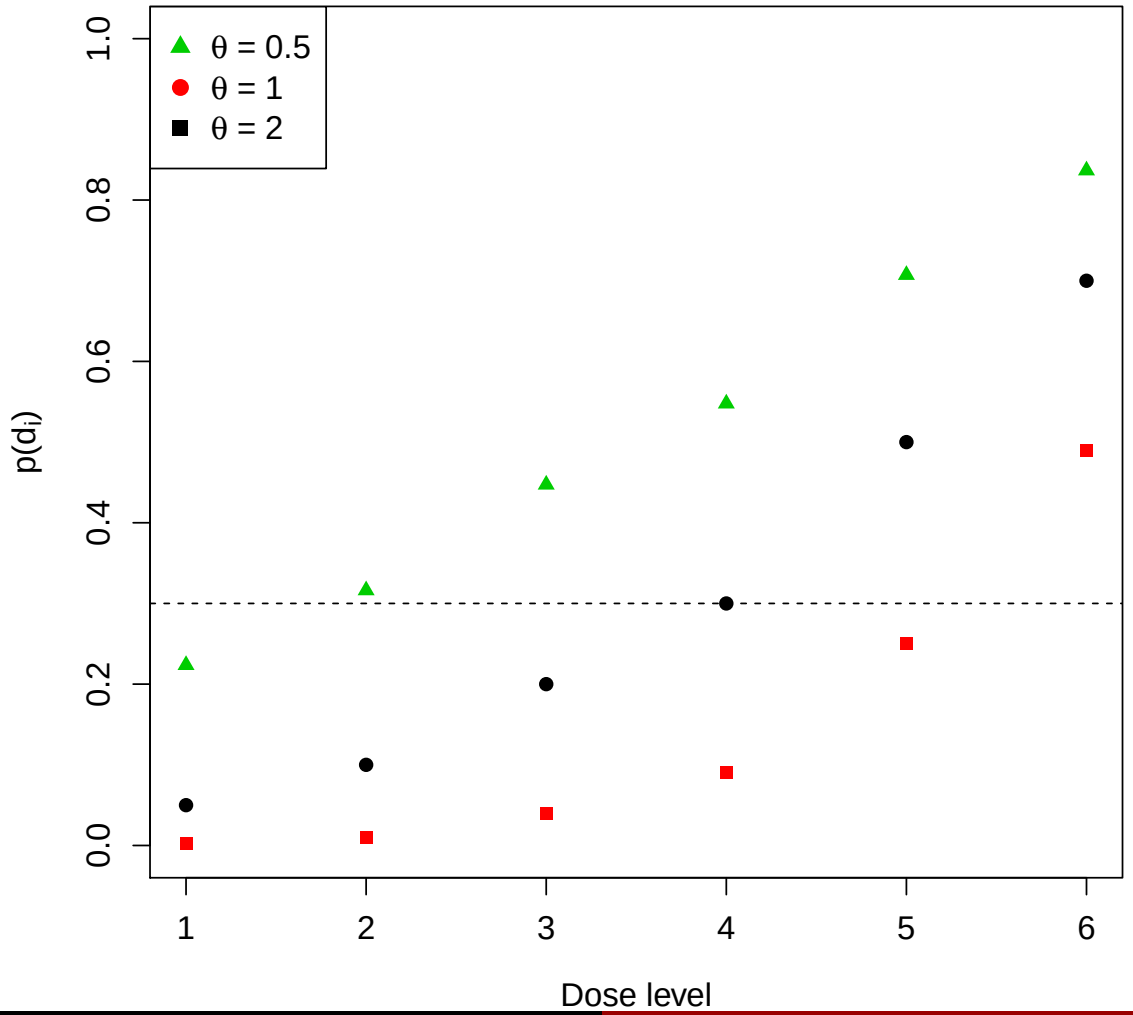

Example
O'Quigley et al (1990)

Assign starting values for the π_i

π_1	π_2	π_3	π_4	π_5	π_6
0.05	0.10	0.20	0.30	0.50	0.70

Thus first subject receives d_3

Simulated data

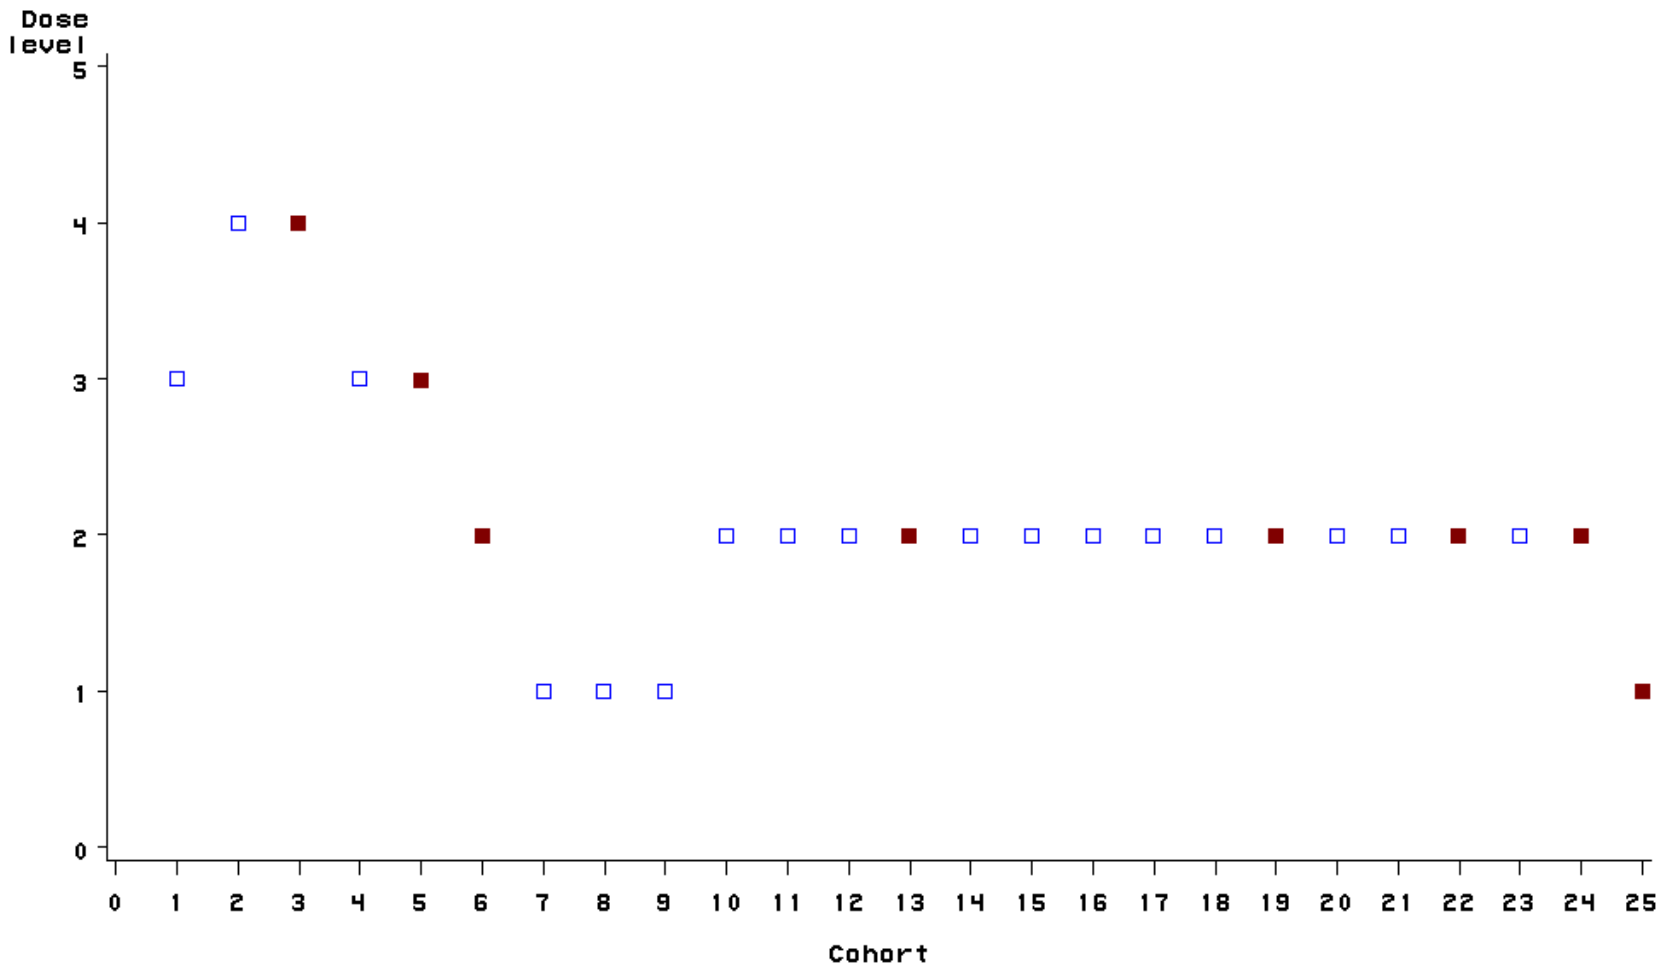

Criticisms and modifications

Criticisms:

- Starting dose is usually high
- Treats too many subjects on high doses
- Doses can be skipped
- No appropriate stopping rule

Modifications:

- Start from the lowest dose
- Safety constraints
- No dose skipping
- Formal stopping rules

Bayesian Logistic Regression Model (BLRM)

Neuenschwander *et al.* (2008)

2-parameter logistic regression model

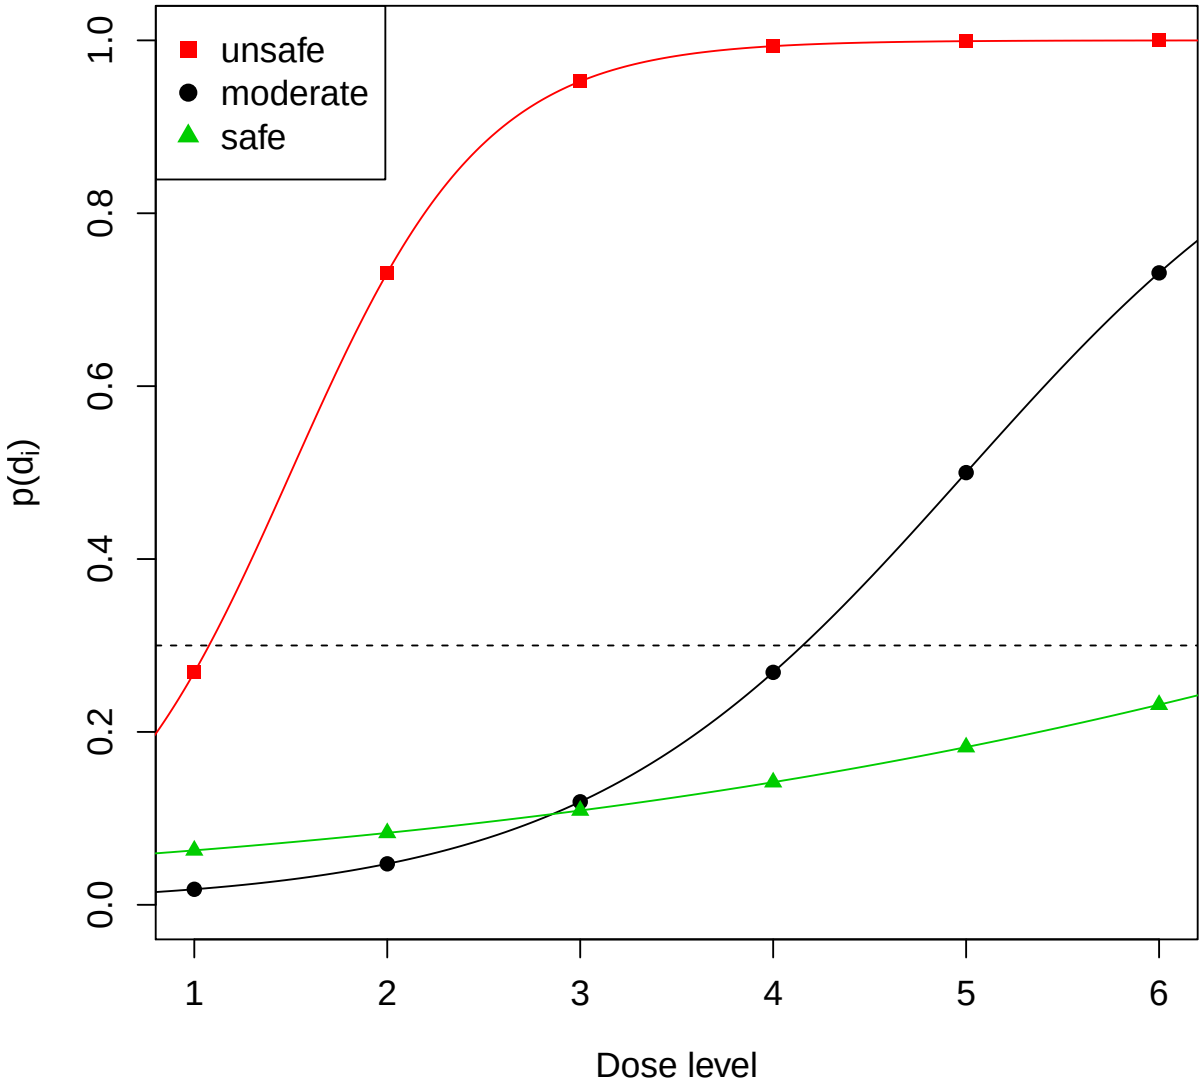

Specifying priors

- Specify two quantiles for probability of toxicity at each dose level
- Define prior distribution for the model parameters such that they are in close agreement with the above
- Requires Markov Chain Monte Carlo (MCMC)

Dose escalation and stopping

Choose recommended dose, d, such that

- probability of overdosing
  $P(\text{DLT rate} > 0.33 \mid d) < 0.2$
- probability of target toxicity
  $P(\text{DLT rate} \in (0.16, 0.33) \mid d) \geq 0.5$
- probability of underdosing
  $P(\text{DLT rate} < 0.16 \mid d) < 0.3$

is controlled.

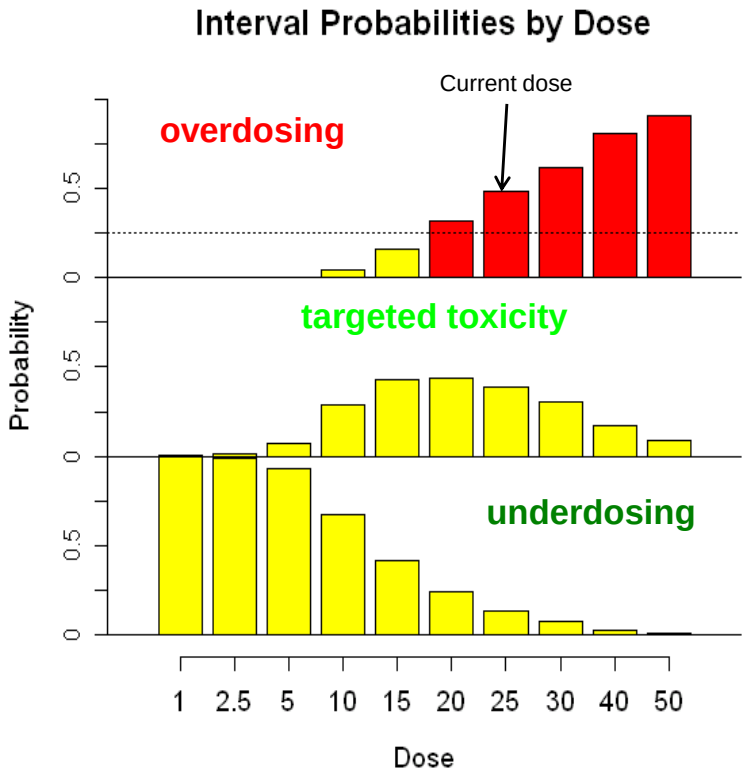

Comments

- Widely used in industry now
- Specifying priors can be time consuming
- Requires MCMC
- Very intuitive dose-selection

Example
Neuenschwander et al (2008)

- Open-label, multicenter, dose-escalation cancer trial
- Find dose that has 30% risk of toxicity, the TD30.
- Use CRM but do not allow for skipping dose levels
- After 4 cohorts (4 dose levels) no DLTs
- Team decides to skip 2 dose levels
- Two DLTs in two patients

Example

	Dose in mg									
	1	2.5	5	10	15	20	25	30	40	50
No. of patients	3	4	5	4	—	—	2	—	—	—
No. of DLTs	0	0	0	0	—	—	2	—	—	—
Posterior summaries:										
Mean	0.069	0.085	0.099	0.111	0.123	0.144	0.163	0.242	**0.330**	0.465
Std. dev.	0.055	0.062	0.068	0.072	0.076	0.082	0.087	0.101	0.109	0.108

Dose recommendation for next cohort:

- 3+3: Unclear
- CRM: 40mg
- BLRM: 15mg (from previous figure)

A simple BLRM
Whitehead & Williamson (1998)

2-parameter logistic regression model

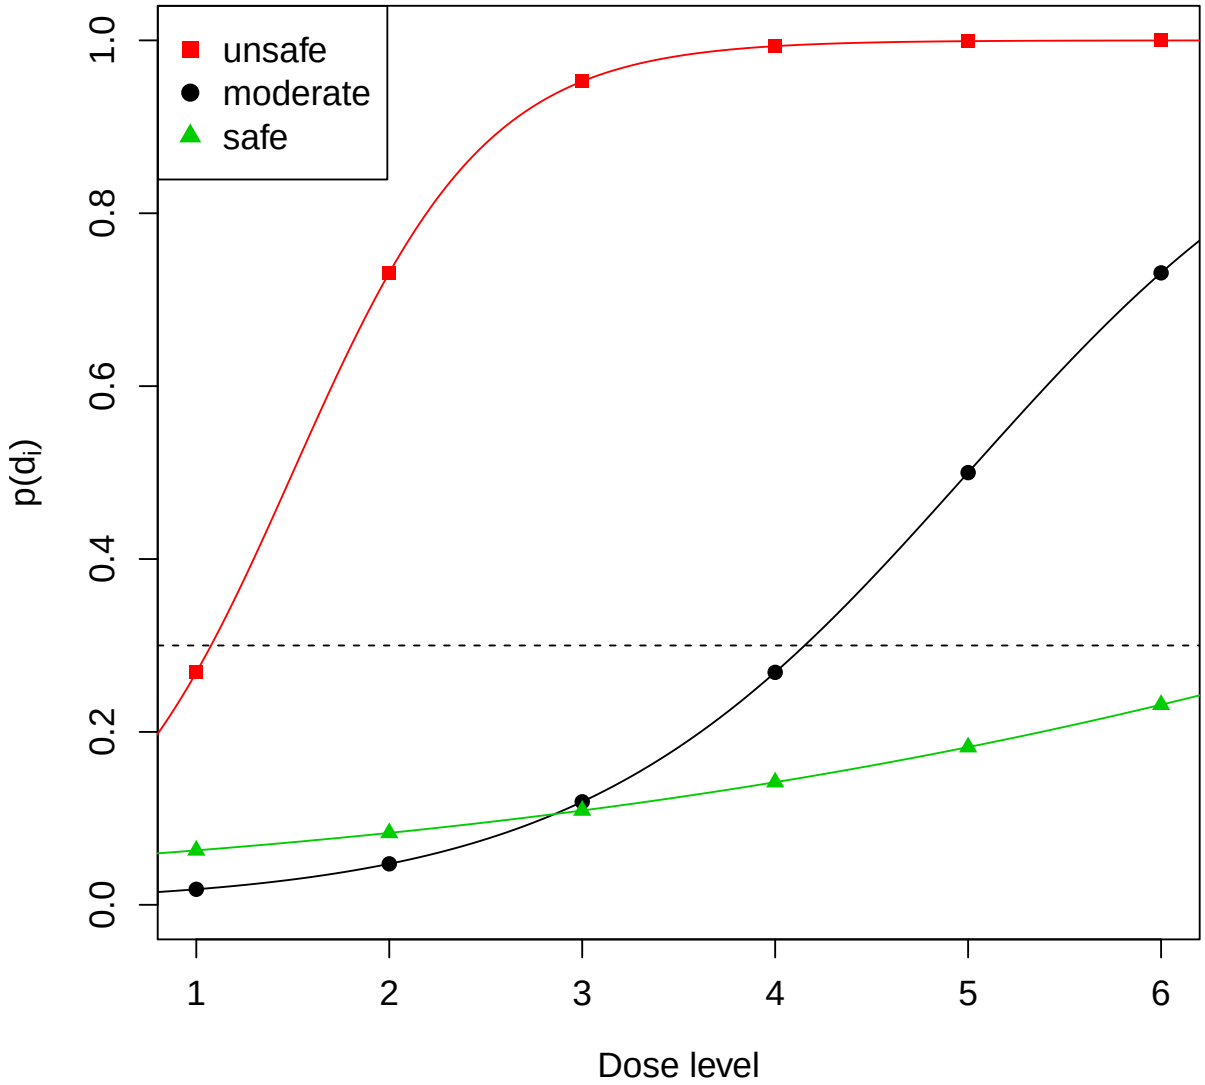

Specifying the prior

- Specify two dose levels (low and high)
- Elicit probability of toxicity at these levels from experts
- Determine how many patients this information is worth
- Adjust to start escalation at lowest dose
- Include “pseudo-patients” in analysis based on above

Note: This corresponds to using a beta-prior on $p(d)$.

Inference

- Treat pseudo-patients as real patients
- Find parameter estimates for logistic model
- Any software that can fit logistic models can be used

Dose escalation and stopping

- Next dose:
  - current estimated target toxicity level
  - Usually subject to some additional safety rule
- Stop
  - When maximum number of patients has been recruited
  - When the ratio of credibility interval limits is small enough

Comments

- Easy to elicit priors from experts
- Any software that can fit a logistic model can be used
- Useful to allow higher TDϑ^* during trial when seeking TDϑ
- Not possible to use more complex rules for dose selection without MCMC

A comparison

- 3+3 design as discussed in Lecture 1
- Simple BLRM
  - Cohort size 1
  - Operational prior worth 6 patients
  - Accuracy stop if ratio of CI < 4
- True probability of toxicity at each dose level

π_1	π_2	π_3	π_4	π_5	π_6
0.05	0.10	0.20	0.30	0.50	0.70

A comparison