# Beyond the 3+3 design: How to find the optimal dose

Dr Thomas Jaki[1] Dr Adrian Mander[2]

Medical and Pharmaceutical Statistics Research Unit

MRC Hubs for Trials Methodology Research
MRC Biostatistics Unit Hub

[1]Lancaster University

[2]MRC Biostatistics Unit, Cambridge

3. December 2014

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Introduction to Phase I trials

- Phase I studies involve the first experimentation of a new drug / clinical procedure in human subjects
- The emphasis is on finding a safe, yet potentially effective, dose for future Phase II experimentation
- Trials are small, typically 20-50 patients
- Due to ethical considerations patients are added sequentially after side-effects from previous patients have been assessed
- Subjects
  - Healthy volunteers for relatively non-toxic agents
  - Patients when drugs are highly toxic (e.g. cytotoxic agents in cancer)
- Aim: To seek the highest possible dose subject to toxicity constraints
  - This is known as the *maximum tolerated dose* (MTD)
  - Based on a monotonicity assumption that the benefit (efficacy) of treatment increases with dose
  - Ethically, we would like to treat every patient at a dose just

# Key elements

1. A starting dose that will be given to the first patient
   - Often chosen as $\frac{1}{10}LD_{10}$ in mice (one tenth of the lethal dose in 10% of mice)
2. A toxicity outcome
   - Often binary (e.g. occurrence of a *dose-limiting toxicity* (DLT) is used in cancer trials)
3. A *target toxicity level* (TTL)
   - The desired toxicity at the MTD (e.g. cancer trials often propose 30% prevalence of DLT at the MTD)
4. A dose-escalation design
   - Rule or model based
   - Cohort size:- No. individuals treated at each dose level
   - Possible dose levels for experimentation
   - Sample size / stopping rules

# Key elements

1. A starting dose that will be given to the first patient
   - Often chosen as $\frac{1}{10}LD_{10}$ in mice (one tenth of the lethal dose in 10% of mice)
2. A toxicity outcome
   - Often binary (e.g. occurrence of a *dose-limiting toxicity* (DLT) is used in cancer trials)
3. A *target toxicity level* (TTL)
   - The desired toxicity at the MTD (e.g. cancer trials often propose 30% prevalence of DLT at the MTD)
4. A dose-escalation design
   - Rule or model based
   - Cohort size:- No. individuals treated at each dose level
   - Possible dose levels for experimentation
   - Sample size / stopping rules

# Key elements

1. A starting dose that will be given to the first patient
   - Often chosen as $\frac{1}{10}LD_{10}$ in mice (one tenth of the lethal dose in 10% of mice)
2. A toxicity outcome
   - Often binary (e.g. occurrence of a *dose-limiting toxicity* (DLT) is used in cancer trials)
3. A *target toxicity level* (TTL)
   - The desired toxicity at the MTD (e.g. cancer trials often propose 30% prevalence of DLT at the MTD)
4. A dose-escalation design
   - Rule or model based
   - Cohort size:- No. individuals treated at each dose level
   - Possible dose levels for experimentation
   - Sample size / stopping rules

# Key elements

1. A starting dose that will be given to the first patient
   - Often chosen as $\frac{1}{10}LD_{10}$ in mice (one tenth of the lethal dose in 10% of mice)
2. A toxicity outcome
   - Often binary (e.g. occurrence of a *dose-limiting toxicity* (DLT) is used in cancer trials)
3. A *target toxicity level* (TTL)
   - The desired toxicity at the MTD (e.g. cancer trials often propose 30% prevalence of DLT at the MTD)
4. A dose-escalation design
   - Rule or model based
   - Cohort size:- No. individuals treated at each dose level
   - Possible dose levels for experimentation
   - Sample size / stopping rules

# Key elements

1. A starting dose that will be given to the first patient
   - Often chosen as $\frac{1}{10}LD_{10}$ in mice (one tenth of the lethal dose in 10% of mice)
2. A toxicity outcome
   - Often binary (e.g. occurrence of a *dose-limiting toxicity* (DLT) is used in cancer trials)
3. A *target toxicity level* (TTL)
   - The desired toxicity at the MTD (e.g. cancer trials often propose 30% prevalence of DLT at the MTD)
4. A dose-escalation design
   - Rule or model based
   - Cohort size:- No. individuals treated at each dose level
   - Possible dose levels for experimentation
   - Sample size / stopping rules

# 3+3 design with escalation only —*Storer(1989) Biometrics*

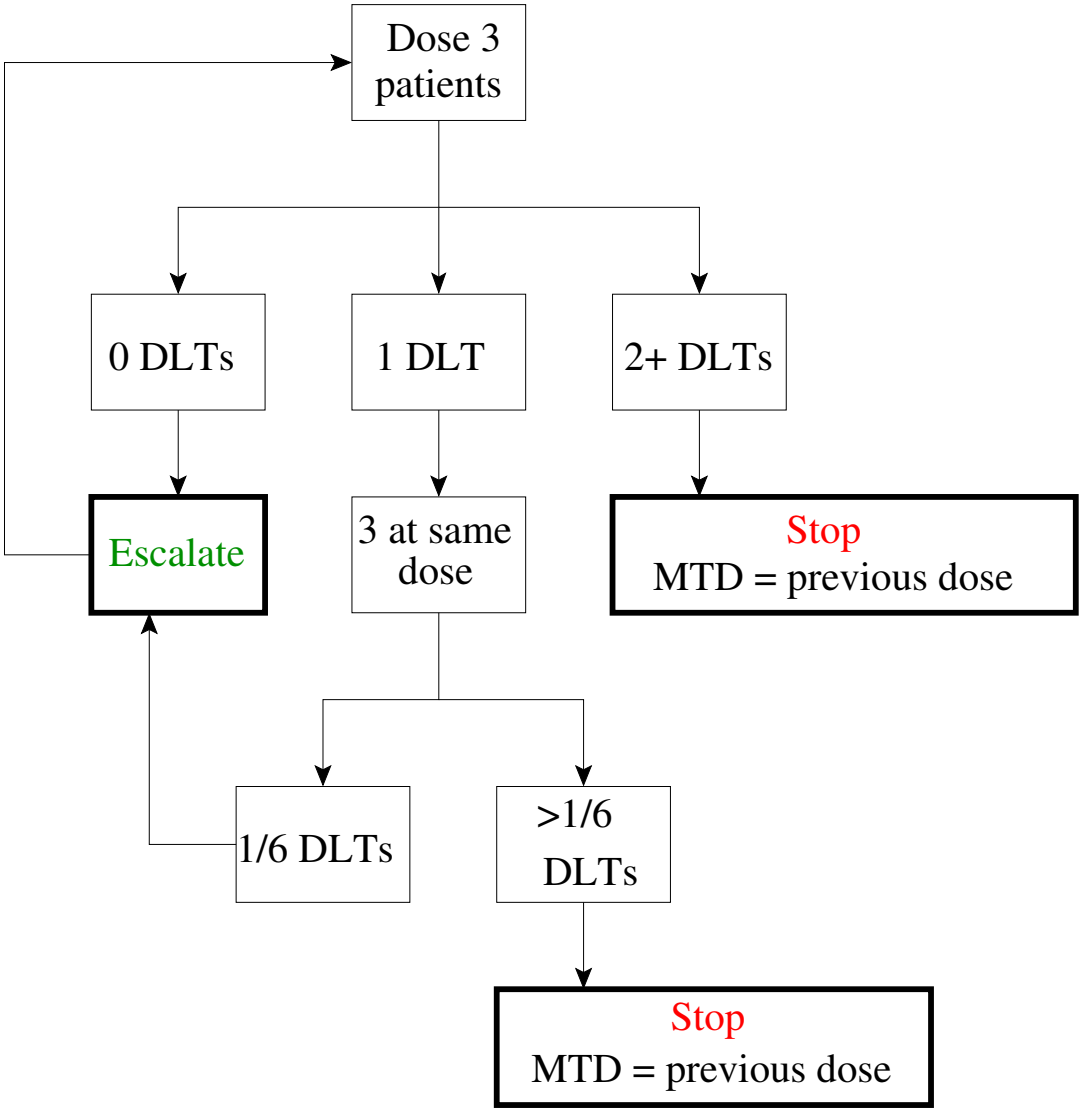


- Dose Limiting Toxicity (DLT)
- Simple rule based approach
- No need for a statistician
- Actual dose not used
- The data to declare an MTD are either $0/3$ or $1/6$

# 3+3 design with escalation only —*Storer(1989) Biometrics*

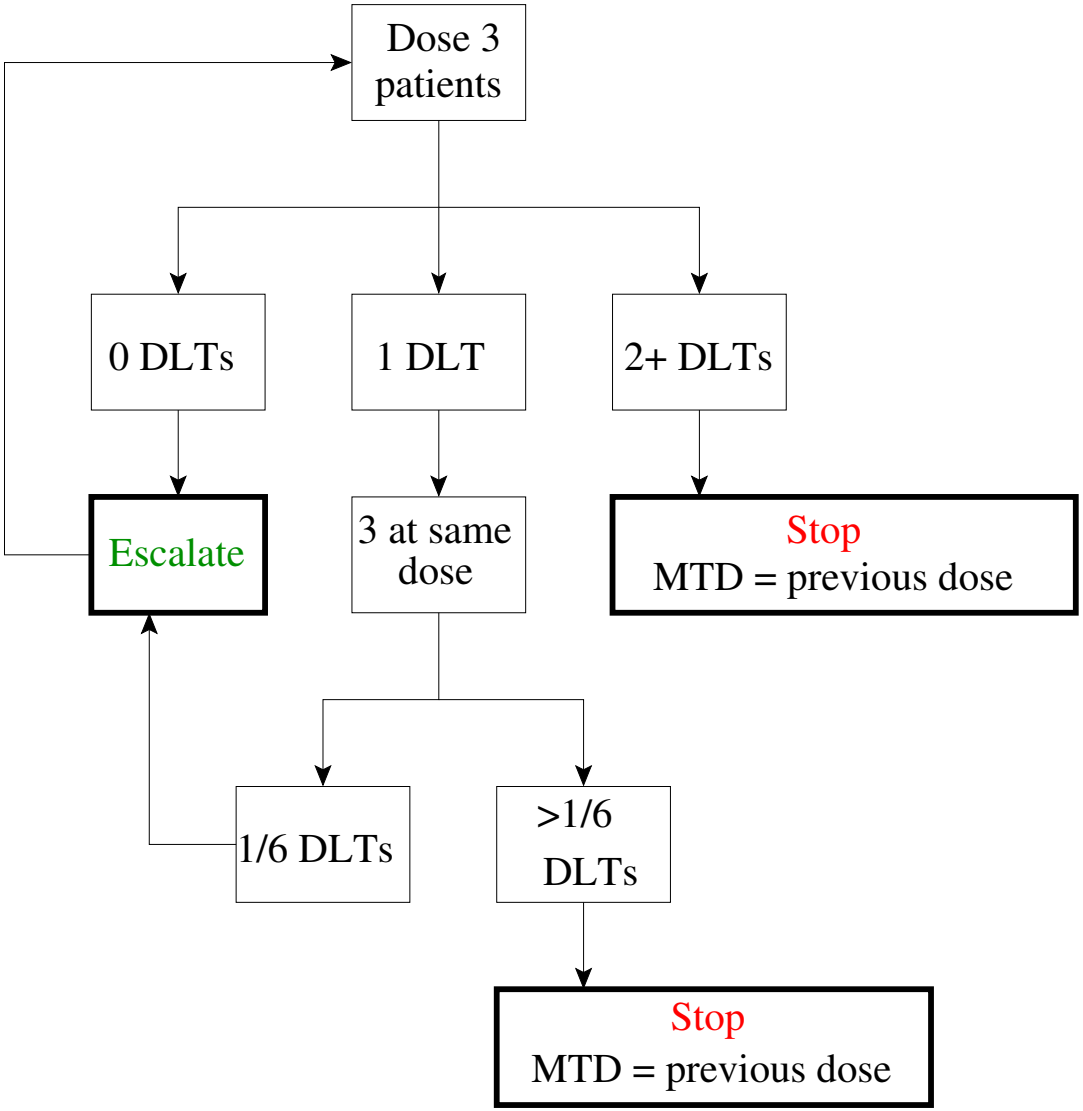


- Dose Limiting Toxicity (DLT)
- Simple rule based approach
- No need for a statistician
- Actual dose not used
- The data to declare an MTD are either $0/3$ or $1/6$

# 3+3 design with escalation only —*Storer(1989) Biometrics*

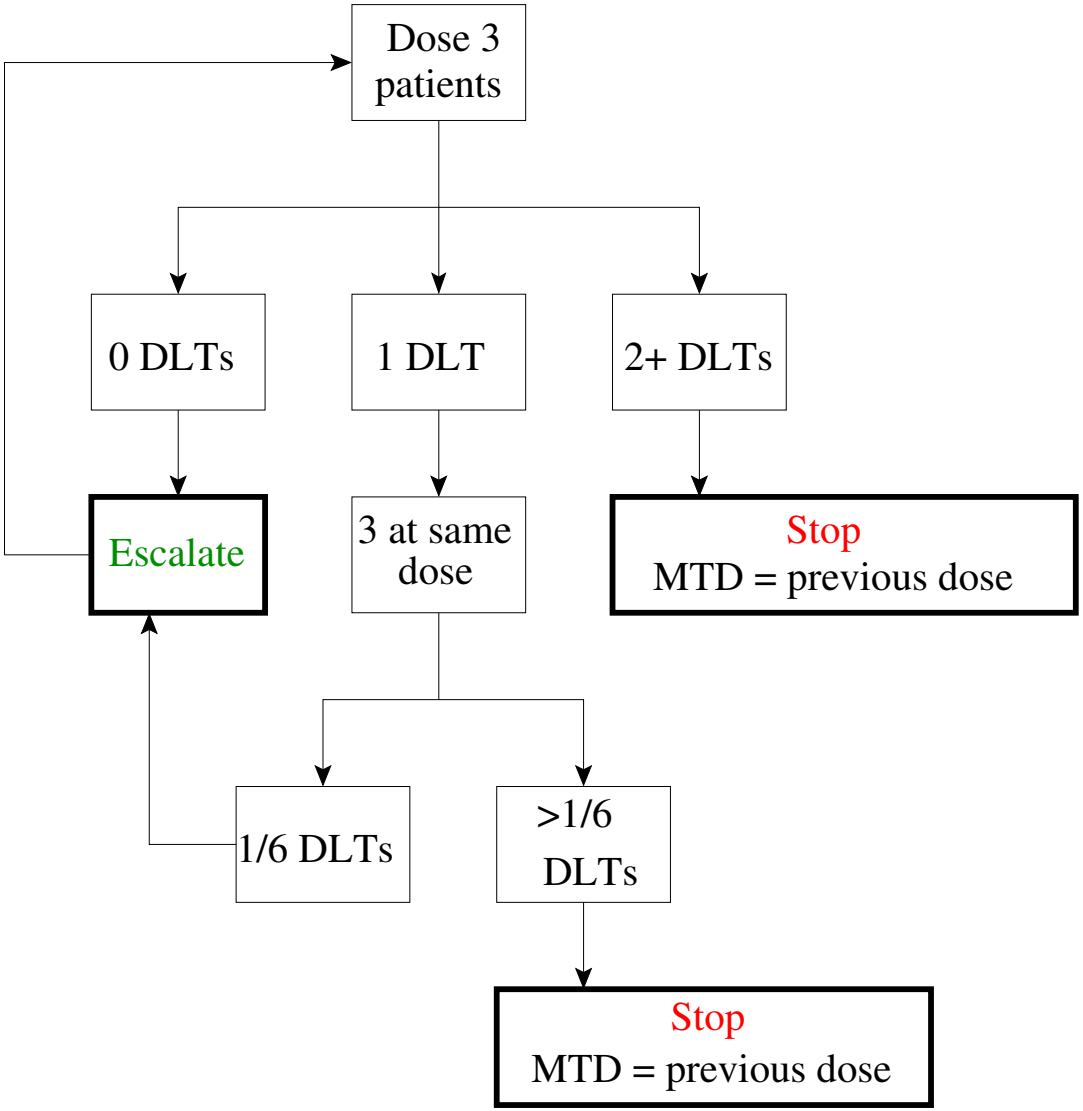


- Dose Limiting Toxicity (DLT)
- Simple rule based approach
- No need for a statistician
- Actual dose not used
- The data to declare an MTD are either $0/3$ or $1/6$

# Current opinion about the 3+3 design

Phase I trial design: Is 3+3 the best? — *Hansen et al.(2014) Cancer Control*

> *The evidence from this review suggests that the 3+3 design identifies the recommended phase 2 dose and toxicities with an acceptable level of precision in some circumstances*

> *Novel trial designs demonstrating superiority over the 3+3 method in statistical simulations without corroborating clinical evidence are of theoretical value alone*

What comes first the simulations (chicken) or the practice (egg)?

# Current opinion about the 3+3 design

Phase I trial design: Is 3+3 the best? — *Hansen et al.(2014) Cancer Control*

> *The evidence from this review suggests that the 3+3 design identifies the recommended phase 2 dose and toxicities with an acceptable level of precision in some circumstances*
>
> *Novel trial designs demonstrating superiority over the 3+3 method in statistical simulations without corroborating clinical evidence are of theoretical value alone*

What comes first the simulations (chicken) or the practice (egg)?

# Current opinion about the 3+3 design

Phase I trial design: Is 3+3 the best? — *Hansen et al.(2014) Cancer Control*

> *The evidence from this review suggests that the 3+3 design identifies the recommended phase 2 dose and toxicities with an acceptable level of precision in some circumstances*
>
> *Novel trial designs demonstrating superiority over the 3+3 method in statistical simulations without corroborating clinical evidence are of theoretical value alone*

What comes first the simulations (chicken) or the practice (egg)?

# The truth about the 3+3 design

> Given such a simple system of rules there is no need for simulations
>
> *Lin and Shih (2001) Biostatistics*

- Take one example with $4$ doses
  - Let the true toxicity probabilities be $(0.04, 0.29, 0.36, 0.74)$
  - The percentage of patients experimented on each dose are $(35\%, 43\%, 17\%, 5\%)$ —averaged over all possible trials
  - The recommended MTD probabilities are $(48\%, 31\%, 19\%, 0\%)$, $2\%$ no recommended doses
- The 3+3 design
  - is conservative if the TTL is $33\%$
  - can recommend MTDs with minimal toxicity
  - is memoryless — *O'Quigley and Zohar (2006) BJC*

# The truth about the 3+3 design

> Given such a simple system of rules there is no need for simulations
>
> *Lin and Shih (2001) Biostatistics*

- Take one example with $4$ doses
  - Let the true toxicity probabilities be $(0.04, 0.29, 0.36, 0.74)$
  - The percentage of patients experimented on each dose are $(35\%, 43\%, 17\%, 5\%)$ —averaged over all possible trials
  - The recommended MTD probabilities are $(48\%, 31\%, 19\%, 0\%)$, $2\%$ no recommended doses
- The 3+3 design
  - is conservative if the TTL is $33\%$
  - can recommend MTDs with minimal toxicity
  - is memoryless — *O'Quigley and Zohar (2006) BJC*

# The truth about the 3+3 design

> Given such a simple system of rules there is no need for simulations
>
> *Lin and Shih (2001) Biostatistics*

- Take one example with $4$ doses
  - Let the true toxicity probabilities be $(0.04, 0.29, 0.36, 0.74)$
  - The percentage of patients experimented on each dose are $(35\%, 43\%, 17\%, 5\%)$ —averaged over all possible trials
  - The recommended MTD probabilities are $(48\%, 31\%, 19\%, 0\%)$, $2\%$ no recommended doses
- The 3+3 design
  - is conservative if the TTL is $33\%$
  - can recommend MTDs with minimal toxicity
  - is memoryless — *O'Quigley and Zohar (2006) BJC*

# The truth about the 3+3 design

> Given such a simple system of rules there is no need for simulations
>
> *Lin and Shih (2001) Biostatistics*

- Take one example with $4$ doses
  - Let the true toxicity probabilities be $(0.04, 0.29, 0.36, 0.74)$
  - The percentage of patients experimented on each dose are $(35\%, 43\%, 17\%, 5\%)$ —averaged over all possible trials
  - The recommended MTD probabilities are $(48\%, 31\%, 19\%, 0\%)$, $2\%$ no recommended doses
- The 3+3 design
  - is conservative if the TTL is $33\%$
  - can recommend MTDs with minimal toxicity
  - is memoryless — *O'Quigley and Zohar (2006) BJC*

# The tipping point - 0.297 (Maximum TTL)

For any true toxicity probability for a single dose — the exact chance of escalating or stopping the $3 + 3$ design

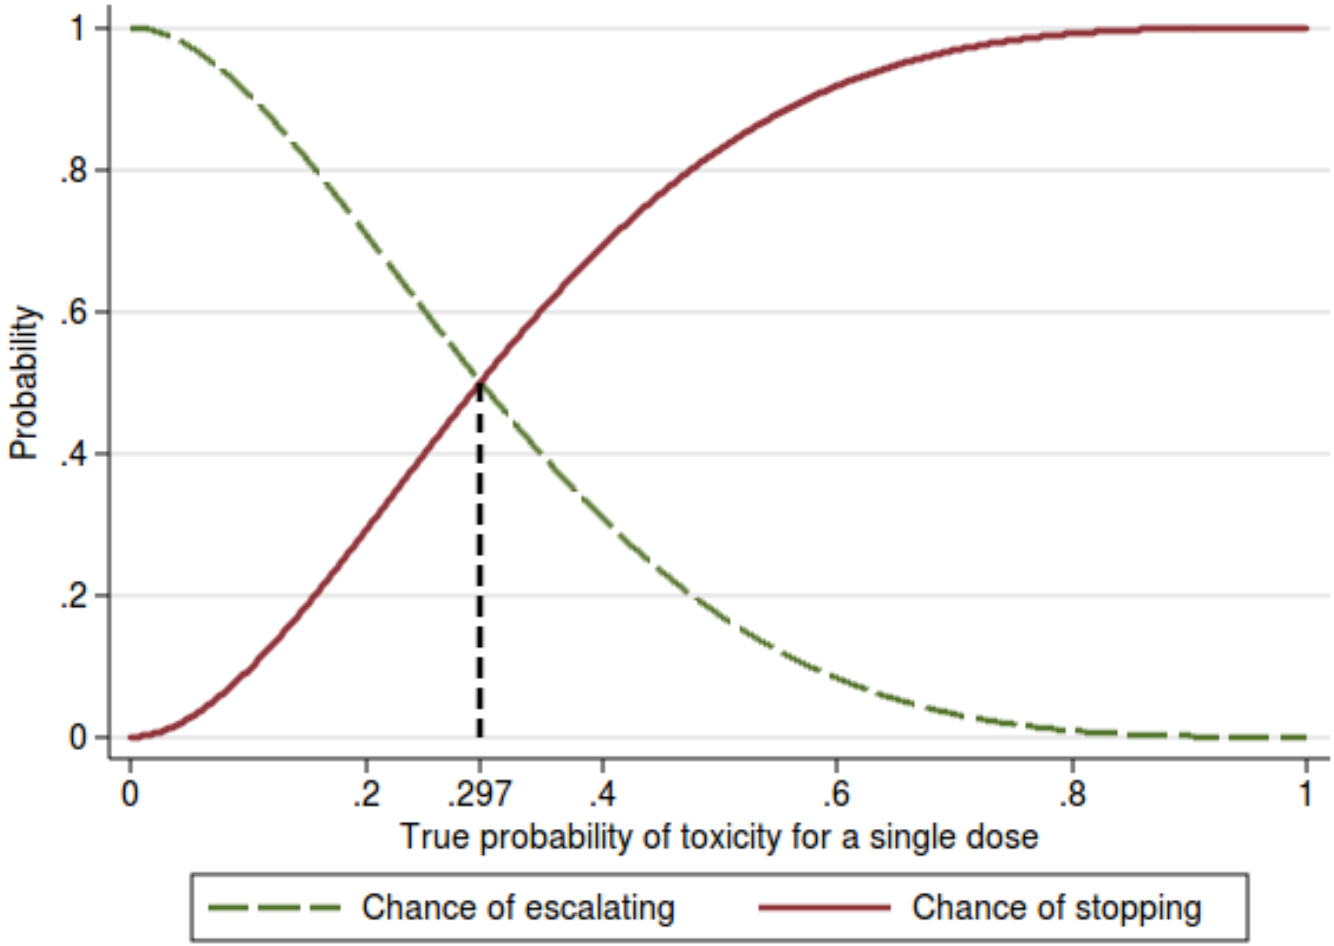

# Final thought about the 3+3

The 3+3 design is about finding the unknown toxicity probabilities with an unknown target toxicity limit.

> *"... there are also unknown unknowns – the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is* **the latter category that tend to be the difficult ones**.*" Donald Rumsfeld*

# Final thought about the 3+3

The 3+3 design is about finding the unknown toxicity probabilities with an unknown target toxicity limit.

> *"... there are also unknown unknowns – the ones we don't know we don't know. And if one looks throughout the history of our country and other free countries, it is* **the latter category that tend to be the difficult ones**.*" Donald Rumsfeld*

# 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 $\pi$ (for $0 < \pi < 1$), e.g. TD20

$\rightarrow$ more specific

# TD20

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

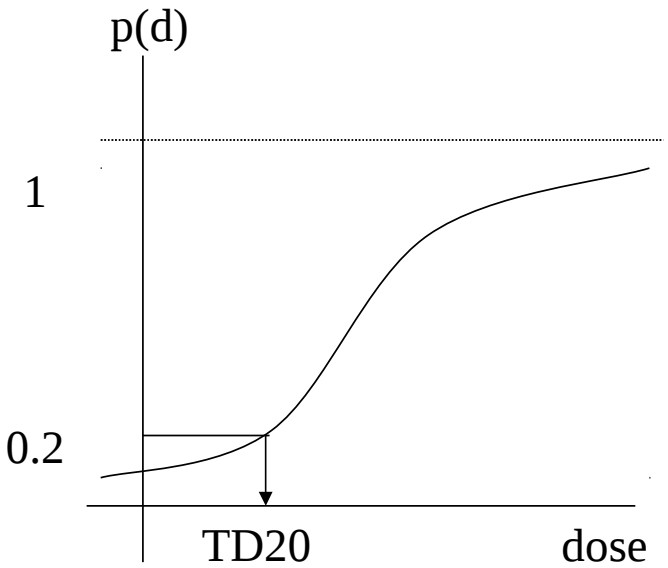


Assume that a 20% risk of toxicity is an acceptable risk to pay for a chance of benefit

# 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 some form of expected gain
4. Stop once target dose level can be estimated accurately enough

# 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 TD20

Cohort size: 1

Prior guess of corresponding probabilities of toxicity at $d_i$

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

were $\pi_i = p(d_i)$ and the $\pi_i$ are treated as fixed

# Model

One parameter log-log model

$$p(d_i) = \pi_i^{\theta}, \; i = 1, \ldots, k \quad \text{only } \theta \text{ is unknown}$$

That is

$$\log[-\log\{p(d_i)\}] = \log(\theta) + \log(-\log(\pi_i))$$

A Bayesian prior for $\theta$ is imposed: $\theta \sim Exp(1)$

so that $E_0(\theta) = 1$

# Representation of the model

Starting values for the $\pi_i$

| $\pi_1$ | $\pi_2$ | $\pi_3$ | $\pi_4$ | $\pi_5$ | $\pi_6$ |
|---|---|---|---|---|---|
| 0.05 | 0.10 | 0.20 | 0.30 | 0.50 | 0.70 |

- $\theta = 1$
- $\theta = 2$
- $\theta = 0.5$

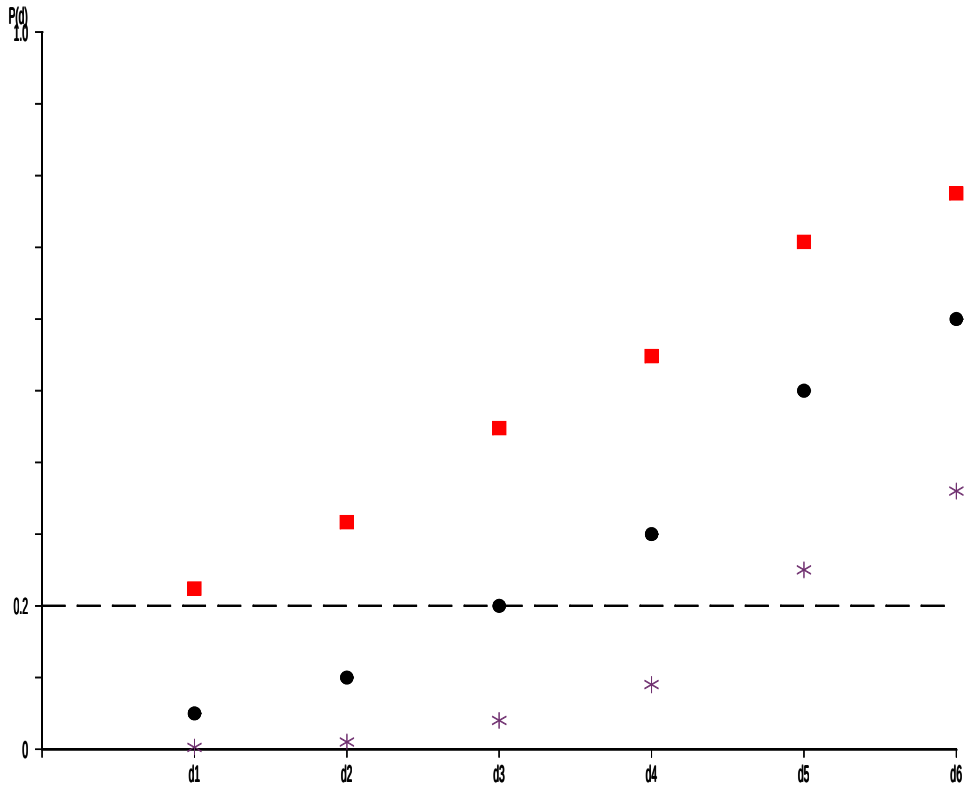

# Bayesian updating

Let $\boldsymbol{x} = (x_1, \ldots, x_s)$ be the data from the first $s$ patients

$$x_j = \begin{cases} 0 \text{ if no toxicity} \\ 1 \text{ if toxicity} \end{cases}$$

PRIOR $\rightarrow$ BAYES THEOREM $\rightarrow$ POSTERIOR

CRM then uses

$$\pi_i(x) = \pi_i^{E(\theta|\boldsymbol{x})}$$

as the updated estimate of $p(d_i)$, identifies which $\pi_i(x)$ is closest to 0.2, and uses corresponding dose next

# Comments

- one parameter model rather than two parameters: the dose-response relationship is well estimated close to the TD20, but not elsewhere
- non-standard log-log model: more usual one is a complementary log-log model CLOGLOG
- pre-defined $\pi_i, i = 1, \ldots, k$ are decided by investigators
- the prior is imposed not elicited
- A non-Bayesian version exists (O'Quigley J, Shen LZ, 1998)

# Neuenschwander et al (2008)

- Dose levels $d_{(1)} < \cdots < d_{(k)}$

$$p(d_{(j)}) = \frac{\exp\{\phi_1 + \phi_2 \log(d_{(j)})\}}{1 + \exp\{\phi_1 + \phi_2 \log(d_{(j)})\}}$$

  for all $j = 1, ..., k$.

- Logistic regression model

# Specifying priors

- Specify two quantiles for probability of toxicity at each dose level
- Define prior distribution for $(\phi_1;\ \phi_2)$ such that they are in close agreement with the above

# 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.

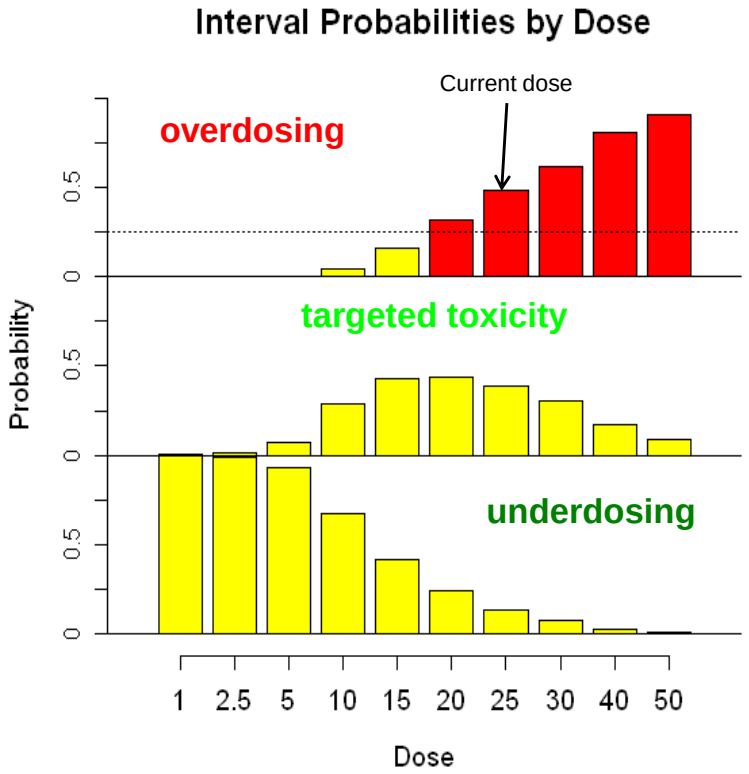

# Comments

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

# A simple approach (Whitehead & Williamson, 1998)

- Dose levels $d_{(1)} < \cdots < d_{(k)}$

$$p(d_{(j)}) = \frac{\exp\{\phi_1 + \phi_2 \log(d_{(j)})\}}{1 + \exp\{\phi_1 + \phi_2 \log(d_{(j)})\}}$$

  for all $j = 1, ..., k$.

- Logistic regression model

# Specifying the prior

- Specify two dose levels (low and high)
- Illicit 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 $p(d)$.

# An example

- Dose-levels 25, 50, . . . , 500mg
- probability of toxicity at 25mg = 0.02
- probability of toxicity at 500mg = 0.50
- Each worth 3 patients

$\Rightarrow$ Include 0.02*3=0.06 toxicities on 25mg and 0.5*3 on 500mg

$\Rightarrow$ Include 0.2*3 toxicities on 25mg and 0.68*3=2.04 on 500mg

# Dose escalation

Allocate next dose that

- is current TD20 (Patient gain)
- allows learning most about dose-toxicity relationship (Variance gain)
- other

Note: Usually subject to some additional safety rule

# Stopping escalation

- When maximum number of patients has been recruited
- When we can estimate TD20 accurately enough
  - Determine current TD20
  - Find corresponding 95% credibility interval
  - Stop if the ratio of the upper and lower bound $< 5$

# Comments

- Easy to elicit priors from experts
- Any software that can fit logistic model can be used
- Prior (typically) pessimistic to ensure no additional rules necessary
- Not possible to use more complex rules for dose selection without MCMC

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- Binary outcomes (toxicity / no toxicity) commonly used in Phase I trials
- Adverse events defined over a time-horizon $T$ in which the drug is assumed to act
- For each cohort we must wait period $T$ before choosing dose for next cohort
- The TITE-CRM (Cheung, Chappell, *Biometrics* 2000) accelerates dose-escalation by using information on time-to-toxicity
  - Original CRM likelihood, $n$ patients

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_i} \left(1 - \pi\left(d(i); \alpha\right)\right)^{(1-y_i)}$$

  - TITE-CRM likelihood uses a weighted dose-response model. At time $t$ after start of recruitment

$$L(\alpha; \mathbf{y}_n) = \prod_{i=1}^{n} \pi\left(d(i); \alpha\right)^{y_{i,t}} \left(1 - w_{i,t}\pi\left(d(i); \alpha\right)\right)^{(1-y_{i,t})}$$

# Time-to-event (TITE) monitoring

- $y_{i,t} = \begin{cases} 1 & \text{if patient } i \text{ has experienced a toxicity by time } t \\ 0 & \text{otherwise} \end{cases}$
- $w_{i,t}$ is a weight assigned to patient $i$ at time $t$
  - Can be thought of as the probability that a toxicity occurs in patient $i$ by time $t$, conditional on it occurring within time horizon $T$
  - The function used for $w$ is as follows:

$$w_{i,t} = \begin{cases} (t - t_{i,0})/T & \text{if } y_{i,t} = 0 \text{ and } t - t_{i,0} < T \\ 1 & \text{otherwise} \end{cases}$$

    where $t_{i,0}$ is the time patient $i$ first receives treatment
- TITE-CRM has been shown to dramatically improve speed of recruitment, whilst still being comparatively safe
  - Safety can be compromised if many toxicities occur near end of time horizon $T$

# Time-to-event (TITE) monitoring

- $y_{i,t} = \begin{cases} 1 & \text{if patient } i \text{ has experienced a toxicity by time } t \\ 0 & \text{otherwise} \end{cases}$
- $w_{i,t}$ is a weight assigned to patient $i$ at time $t$
  - Can be thought of as the probability that a toxicity occurs in patient $i$ by time $t$, conditional on it occurring within time horizon $T$
  - The function used for $w$ is as follows:

$$w_{i,t} = \begin{cases} (t - t_{i,0})/T & \text{if } y_{i,t} = 0 \text{ and } t - t_{i,0} < T \\ 1 & \text{otherwise} \end{cases}$$

    where $t_{i,0}$ is the time patient $i$ first receives treatment
- TITE-CRM has been shown to dramatically improve speed of recruitment, whilst still being comparatively safe
  - Safety can be compromised if many toxicities occur near end of time horizon $T$

# Time-to-event (TITE) monitoring

- $y_{i,t} = \begin{cases} 1 & \text{if patient } i \text{ has experienced a toxicity by time } t \\ 0 & \text{otherwise} \end{cases}$
- $w_{i,t}$ is a weight assigned to patient $i$ at time $t$
  - Can be thought of as the probability that a toxicity occurs in patient $i$ by time $t$, conditional on it occurring within time horizon $T$
  - The function used for $w$ is as follows:

$$w_{i,t} = \begin{cases} (t - t_{i,0})/T & \text{if } y_{i,t} = 0 \text{ and } t - t_{i,0} < T \\ 1 & \text{otherwise} \end{cases}$$

    where $t_{i,0}$ is the time patient $i$ first receives treatment
- TITE-CRM has been shown to dramatically improve speed of recruitment, whilst still being comparatively safe
  - Safety can be compromised if many toxicities occur near end of time horizon $T$

# Time-to-event (TITE) monitoring

- $y_{i,t} = \begin{cases} 1 & \text{if patient } i \text{ has experienced a toxicity by time } t \\ 0 & \text{otherwise} \end{cases}$
- $w_{i,t}$ is a weight assigned to patient $i$ at time $t$
  - Can be thought of as the probability that a toxicity occurs in patient $i$ by time $t$, conditional on it occurring within time horizon $T$
  - The function used for $w$ is as follows:

$$w_{i,t} = \begin{cases} (t - t_{i,0})/T & \text{if } y_{i,t} = 0 \text{ and } t - t_{i,0} < T \\ 1 & \text{otherwise} \end{cases}$$

    where $t_{i,0}$ is the time patient $i$ first receives treatment
- TITE-CRM has been shown to dramatically improve speed of recruitment, whilst still being comparatively safe
  - Safety can be compromised if many toxicities occur near end of time horizon $T$

# Time-to-event (TITE) monitoring

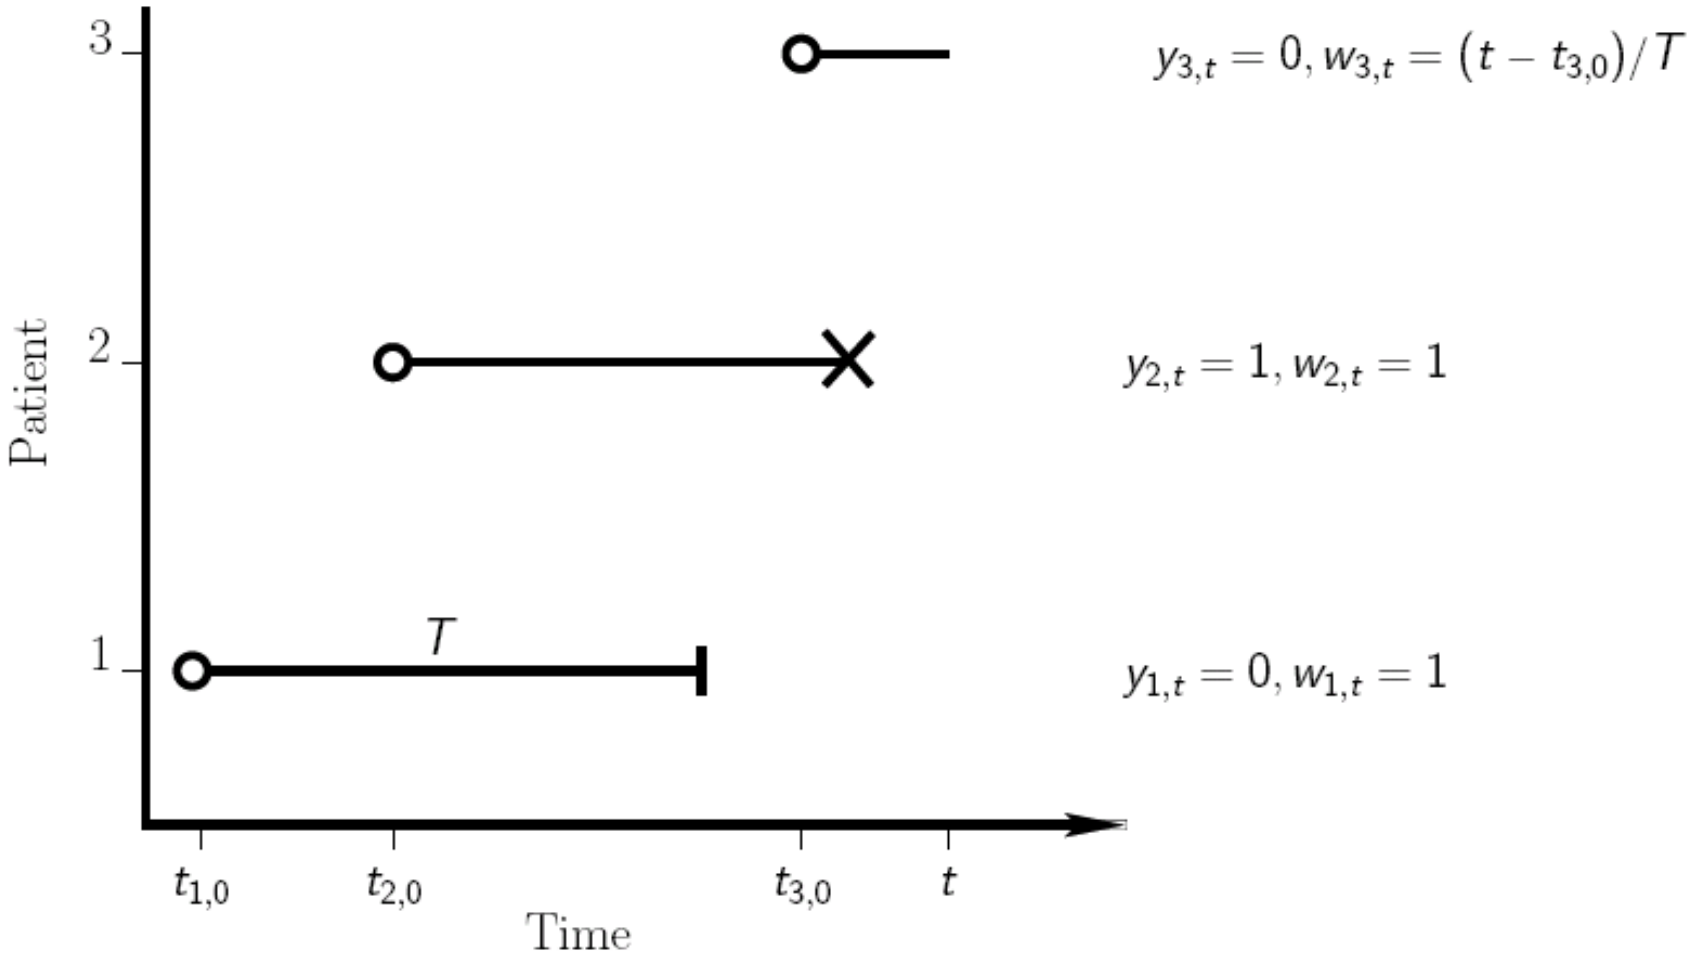

# References

Neuenschwander B, Branson M, Gsponer T. (2008) Critical aspects of the Bayesian approach to phase I cancer trials. *Statistics in Medicine*; 27:2420-2439

O'Quigley J, Pepe M, Fisher L. (1990) Continual Reassessment Method: A Practical Design for Phase 1 Clinical Trials in Cancer. *Biometrics*; 46(1):33-48.

O'Quigley J, Shen LZ (1998). Continual reassessment method: A likelihood approach. *Biometrics*; 52:673-684.

Whitehead J, Williamson D (1998). An evaluation of Bayesian decision procedures for dose-finding studies. *Journal of Biopharmaceutical Medicine*; 8, 445-467.