A Bayesian model-free approach to combination therapy phase I trials using censored time-to-toxicity data

Graham Wheeler[1,2] Michael Sweeting[3] Adrian Mander[2]

[1]Cancer Research UK & UCL Cancer Trials Centre, University College London, UK

[2]MRC Biostatistics Unit, University of Cambridge, UK

[3]Department of Health Sciences, University of Leicester, UK

Cambridge Statistics Discussion Group
19th March 2019

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

Phase I Clinical Trials

Phase I trials are the first investigation of a new treatment/therapy in humans

- In oncology, aim to find safe and (hopefully) beneficial dose/regimen

Typical phase I trial for cytotoxic anti-cancer drug

- Non-comparative, dose-escalation study, $15 - 50$ patients (exhausted standard treatments)
- Patients dosed sequentially (individuals or small groups)
- Based on whether dose is deemed safe or not, change dose level for next patient/group

Aim is to find the **Maximum Tolerated Dose (MTD)**.

Definition

MTD: *The dose expected to produce some degree of medically unacceptable, dose-limiting toxicity (DLT) in a specified proportion of patients (e.g.* 20%*).* (Babb and Rogatko, 2004)

Dose-Limiting Toxicity (DLT)

BLOOD AND LYMHPATIC SYSTEM DISORDERS					
Adverse Event	**Adverse Event Grade**				
	1	**2**	**3**	**4**	**5**
Anaemia	Haemoglobin (Hgb): 100 - LLN g/L	Hgb: 80 - 100 g/L	Hgb < 80 g/L; transfusion indicated	Life-threatening consequences; urgent intervention indicated	Death
Febrile Neutropenia	-	-	ANC<1000/mm^3 with a single temperature of > 38.3°C *or* sustained temperature of ≥ 38°C for more than one hour	Life-threatening consequences; urgent intervention indicated	Death

Dose-Limiting Toxicity (DLT)

BLOOD AND LYMHPATIC SYSTEM DISORDERS					
Adverse Event	**Adverse Event Grade**				
	1	**2**	**3**	**4**	**5**
Anaemia	Haemoglobin (Hgb): 100 - LLN g/L	Hgb: 80 - 100 g/L	Hgb < 80 g/L; transfusion indicated	Life-threatening consequences; urgent intervention indicated	Death
Febrile Neutropenia	-	-	ANC<1000/mm^3 with a single temperature of > 38.3°C *or* sustained temperature of ≥ 38°C for more than one hour	Life-threatening consequences; urgent intervention indicated	Death

Phase I Clinical Trials

DLT is usually recorded as binary response Y_i for patient i, where

$$Y_i = \begin{cases} 1 & \text{if patient } i \text{ has a DLT} \\ 0 & \text{otherwise} \end{cases}$$

Common assumptions for cytotoxic anti-cancer drugs:

- as dose increases, probability of experiencing DLT increases;
- toxicity is indicative of drug having an effect on body/disease

Aim to gradually increase dose of drug until we find a dose with an estimated risk of DLT close to our Target Toxicity Level (TTL), e.g. 20%.

Phase I Clinical Trials

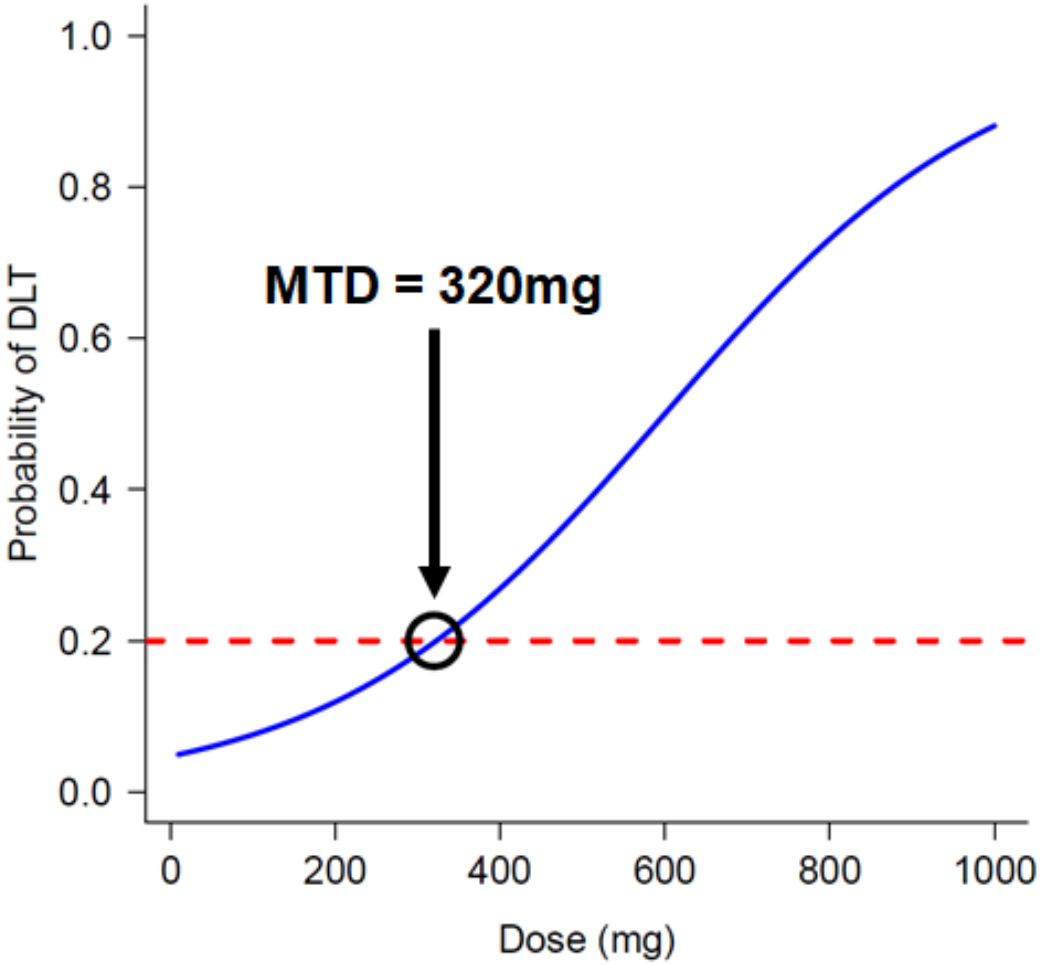

Target Toxicity Level (TTL)

TTL = 20%

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

ORCA-2 trial

Olaparib in high risk locally advanced squamous cell head and neck cancer

- Dose = $\{50, 100, 150, 200\}$ mg twice daily
- Weekly schedule = $\{3, 4, 5\}$ days
- Identify Maximum Tolerated Dose Combinations (MTDCs) of dose and schedule of olaparib
- Target Toxicity Level (TTL) = 33%
- DLT follow-up period is 14 weeks from beginning of treatment

ORCA-2 trial design

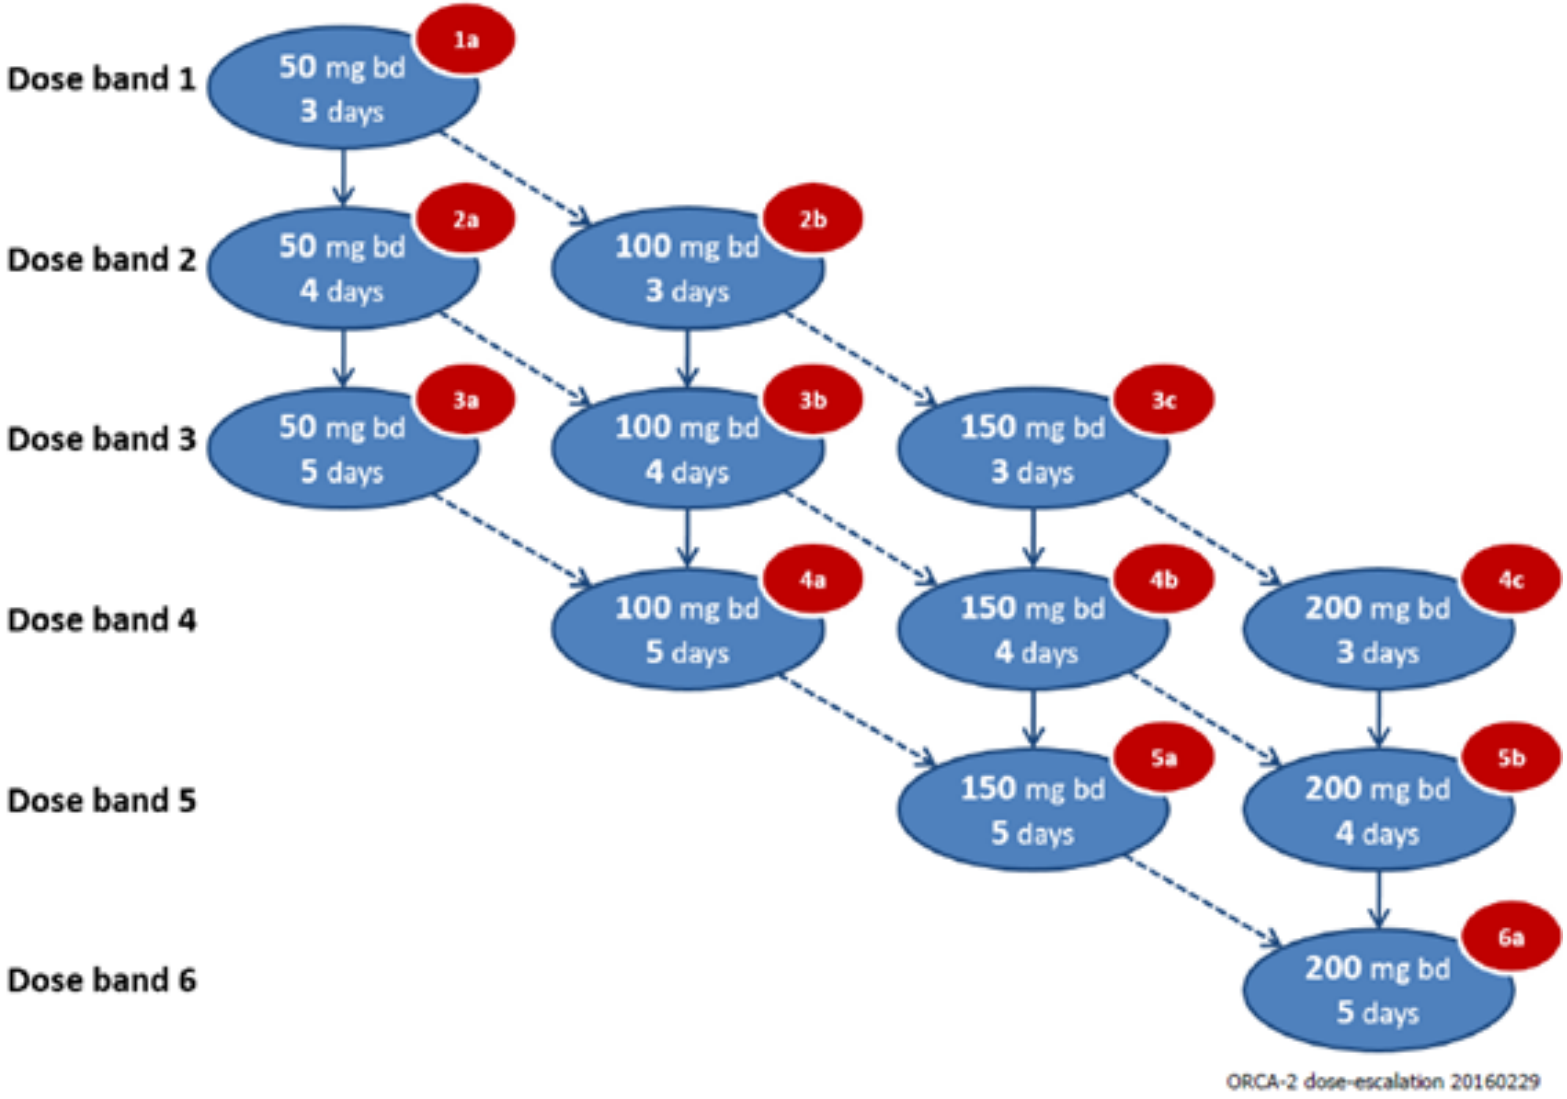

Dual agent dose escalation

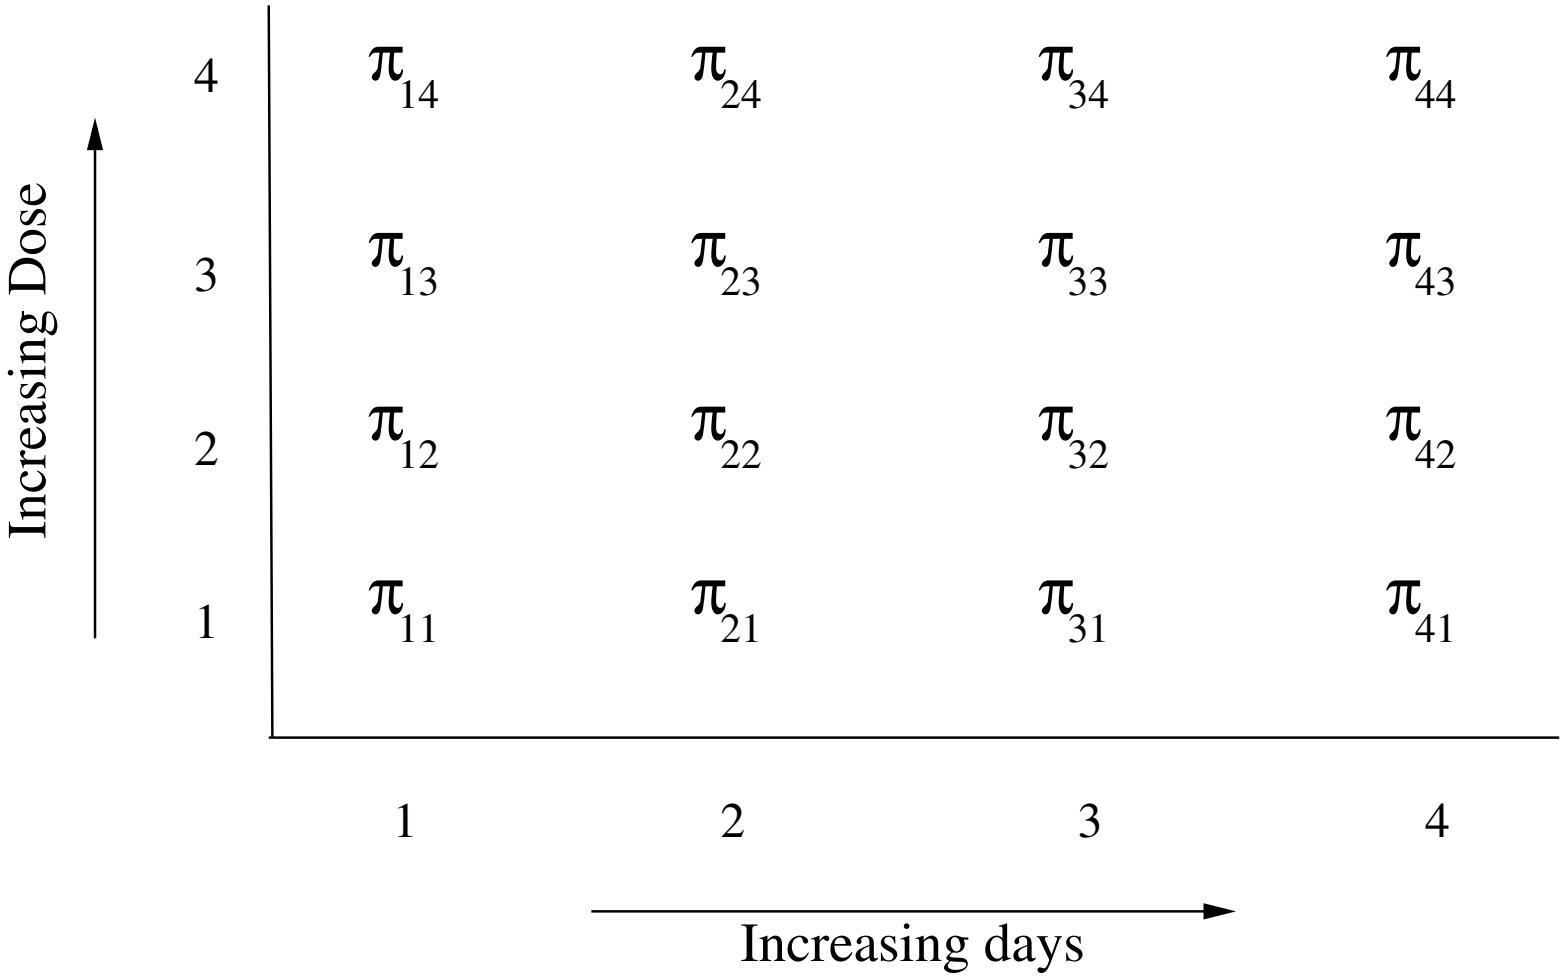

Monotonicity assumptions: 1) $\pi_{jk} \leqslant \pi_{(j+1)k}$ 2) $\pi_{jk} \leqslant \pi_{j(k+1)}$

PIPE (Mander and Sweeting, 2015)

ORCA-2 uses product of independent beta priors escalation (PIPE) approach

Prior distribution of probability of DLT at combination (a_j, b_k) is

$$\pi_{jk}|r_{jk}, s_{jk} \sim Beta(r_{jk}, s_{jk}). \tag{1}$$

Data after m cohorts = $\mathcal{D}^{(m)} = \left\{ R_{jk}^{(m)}, n_{jk}^{(m)} : j = 1, \ldots, J;\ k = 1, \ldots, K \right\}$

Posterior distribution of π_{jk} is also beta distributed, i.e.

$$\pi_{jk}^{(m)} = \pi_{jk}|\mathcal{D}^{(m)}, r_{jk}, s_{jk} \sim Beta(r_{jk} + R_{jk}^{(m)}, s_{jk} + n_{jk}^{(m)} - R_{jk}^{(m)}). \tag{2}$$

Posterior probability of DLT at (a_j, b_k)

Beta(prior + DLTs on (a_j, b_k), prior + non-DLTs on (a_j, b_k))

Maximum Tolerated Contour

Target toxicity limit (θ) is 0.33

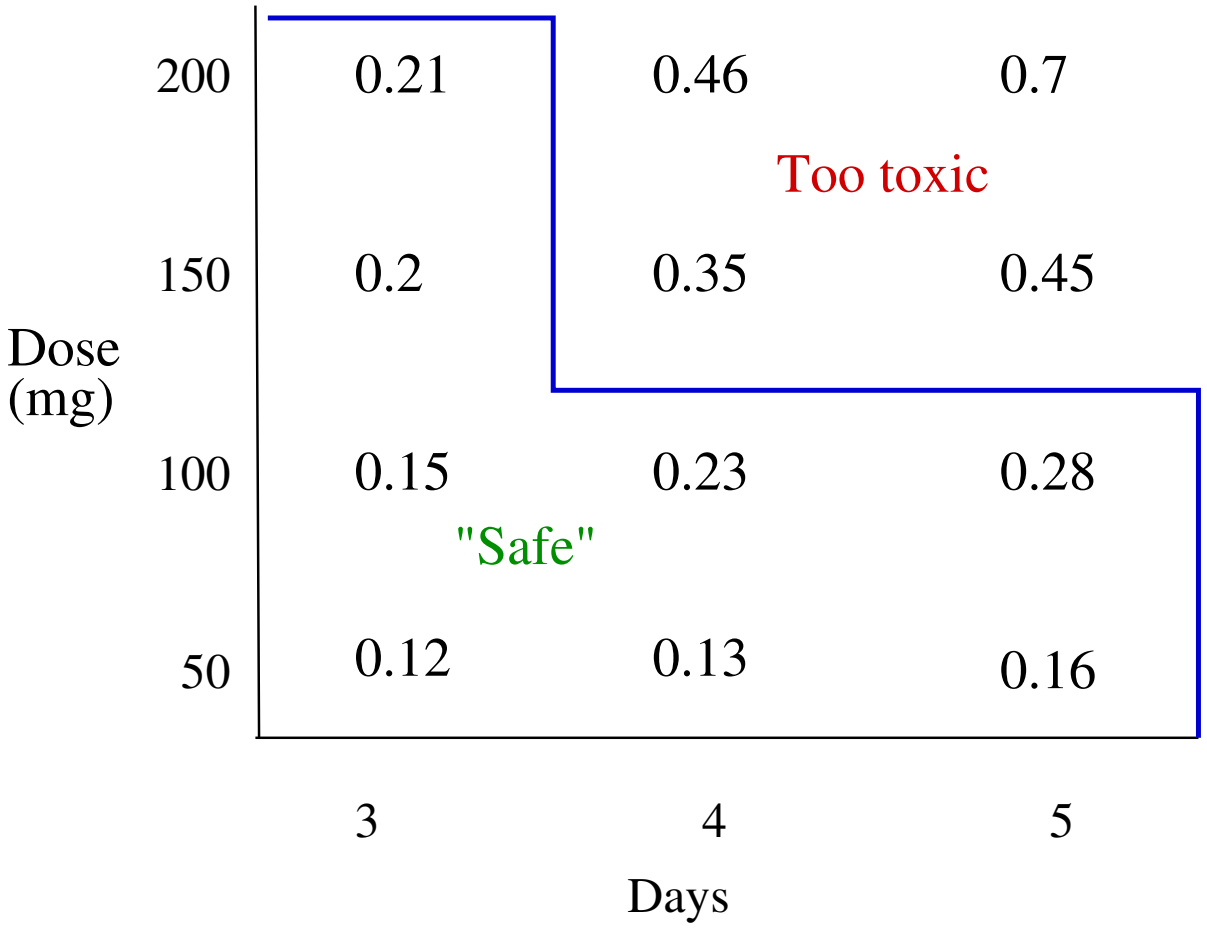

Binary matrix to define the contour

$\mathcal{C}$ = set of all contours satisfying monotonicity assumptions

$C_l \in \mathcal{C}$ is a contour defined by a $J \times K$ binary matrix

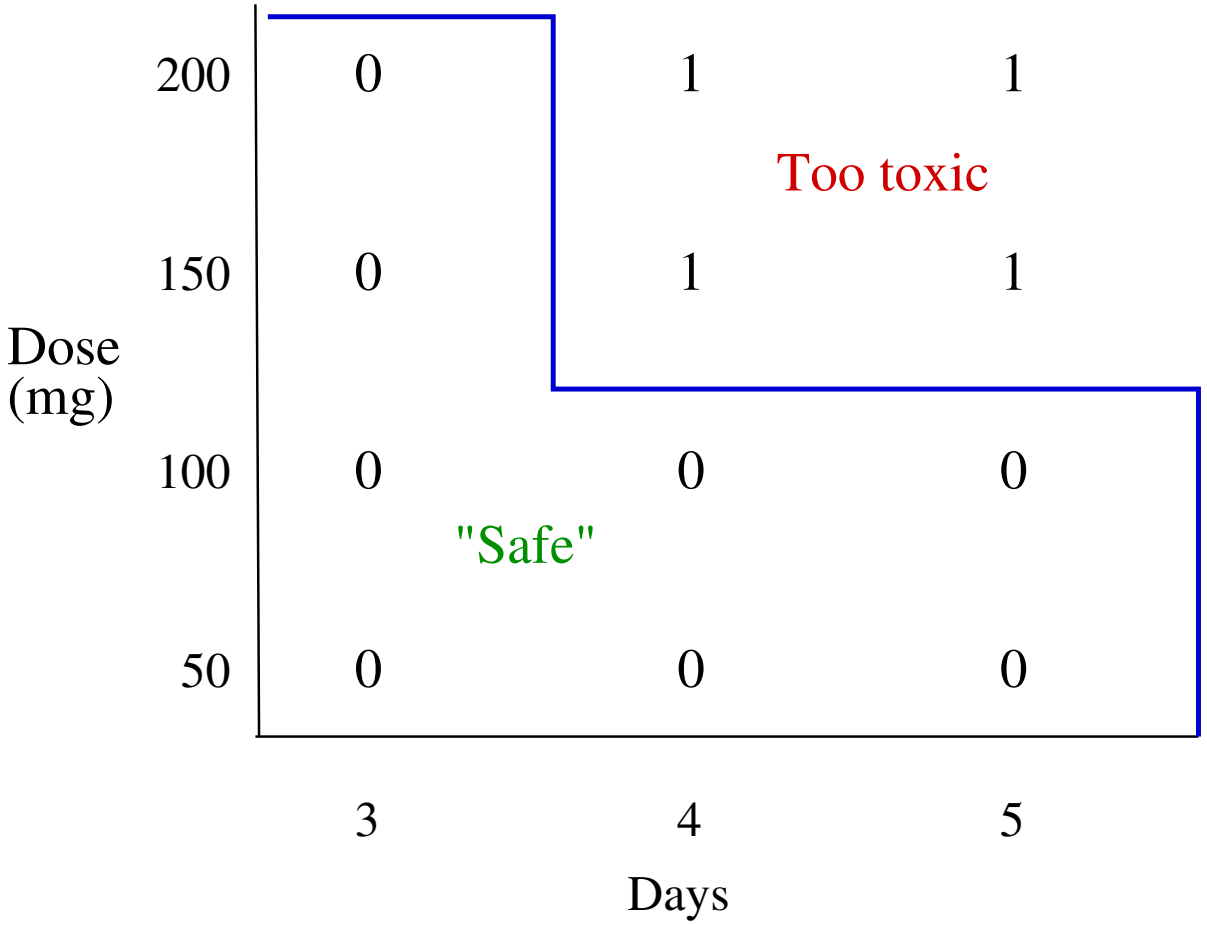

Using $\pi_{jk}^{(m)}$, define the tail probability

$$p_{jk}(\theta|\mathcal{D}^{(m)}) = \mathbb{P}(\pi_{jk}^{(m)} \leqslant \theta|R_{jk}^{(m)}, n_{jk}^{(m)}, r_{jk}, s_{jk})$$

and

$$\mathbb{P}(MTC_\theta = C_l|\mathcal{D}^{(m)}) = \prod_{j,k}\{1 - p_{jk}(\theta|\mathcal{D}^{(m)})\}^{C_l[j,k]} p_{jk}(\theta|\mathcal{D}^{(m)})^{1-C_l[j,k]}$$

For $C_l \in \mathcal{C}$, the normalised probability that $MTC_\theta = C_l$ is

$$\mathbb{P}(MTC_\theta = C_l|C_l \in \mathcal{C}, \mathcal{D}^{(m)}) = \frac{\mathbb{P}(MTC_\theta = C_l|\mathcal{D}^{(m)}) \cdot \mathbb{I}[C_l \in \mathcal{C}]}{\sum_{C_v \in \mathcal{C}} \mathbb{P}(MTC_\theta = C_v|\mathcal{D}^{(m)})}.$$

- Use the most likely contour for Decision making...
- ... subject to any safety constraints

Closest doses to MTC_θ

Define a set of dose combinations that are allowed to be given.

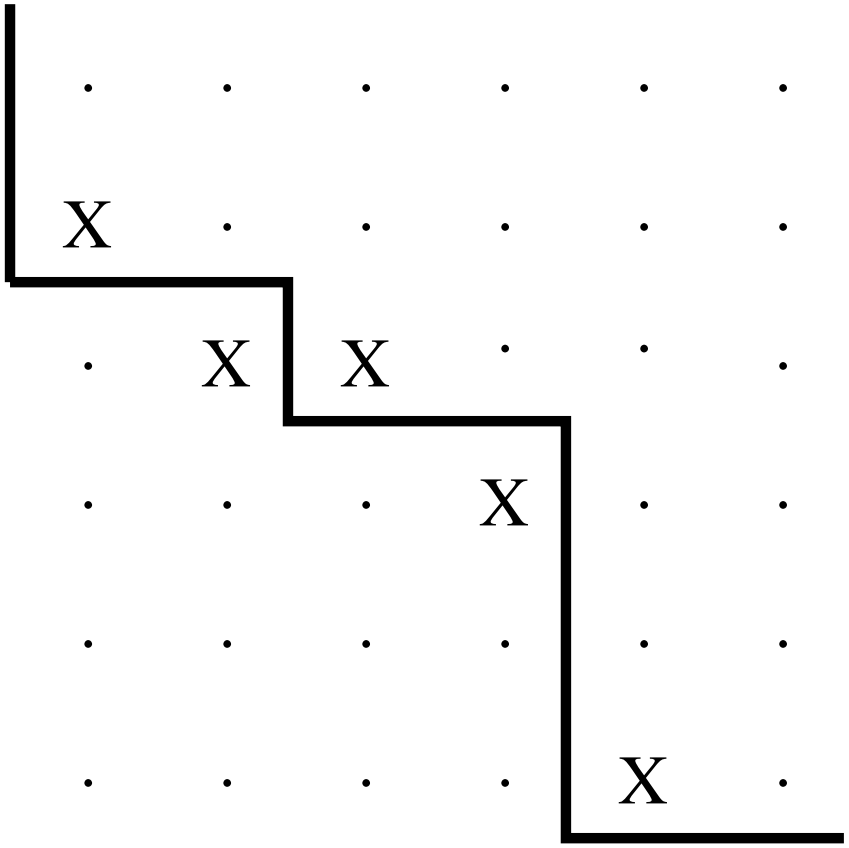

Dose skipping

To avoid skipping dose a boundary box defined by the current highest administered dose-combination (Hi) is used to restrict experimentation

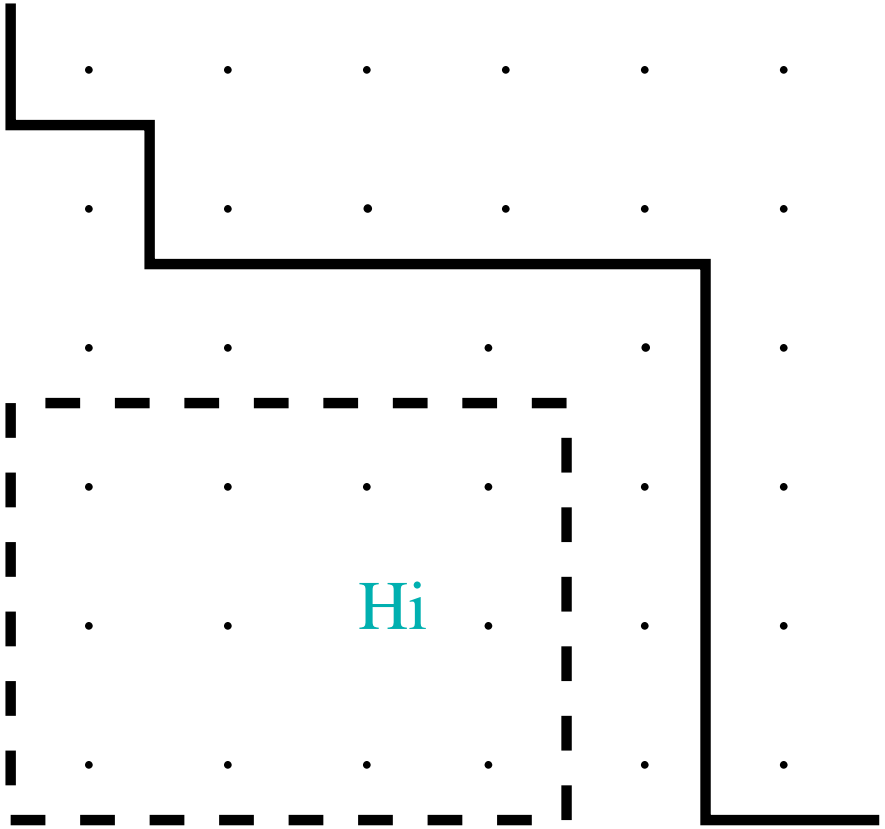

Set of monotonic contours $\mathcal{C} = \{C_1, C_2, \ldots\}$

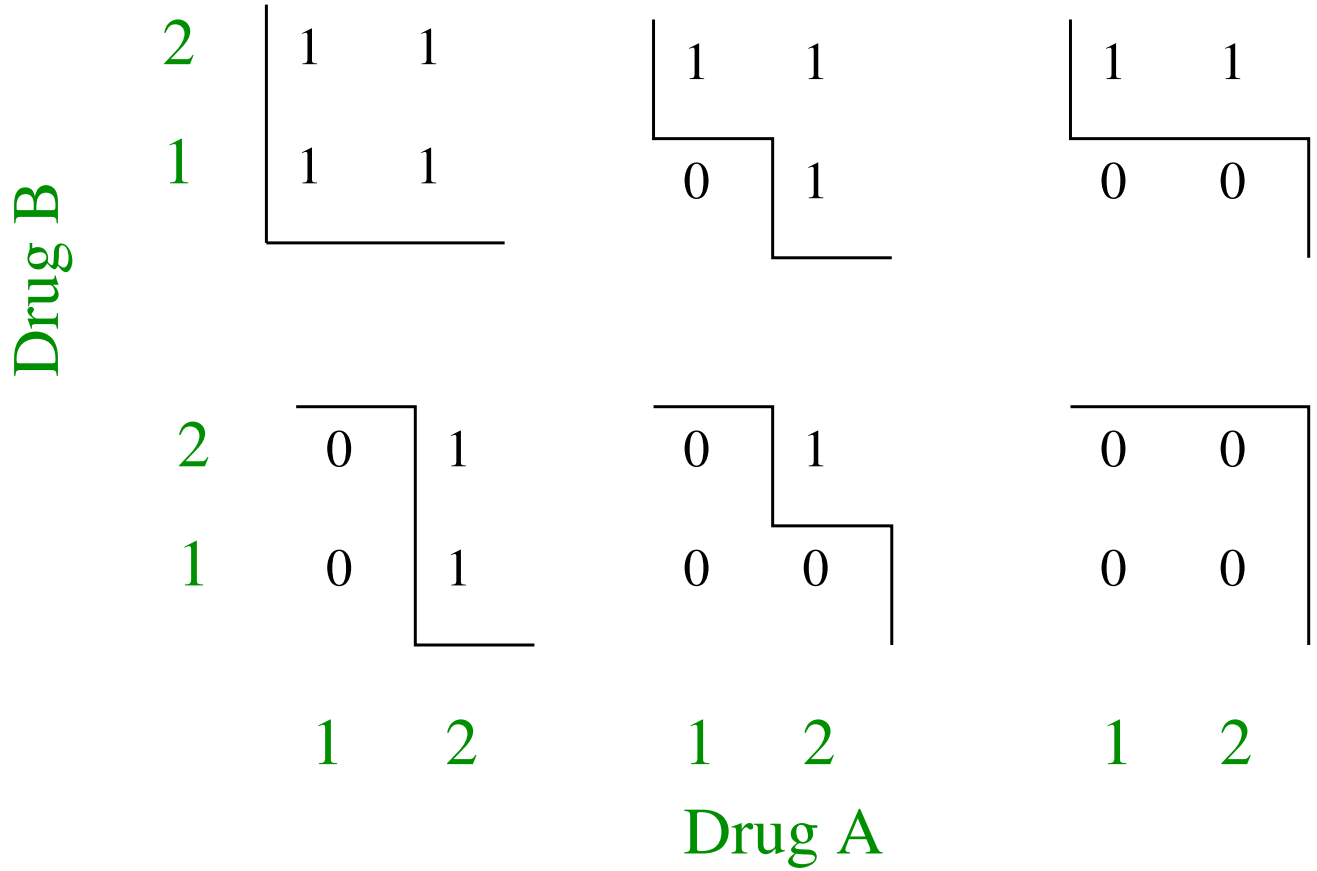

For an $J \times K$ dose-combination grid there are:

- 2^{J+K} contours — $\binom{J+K}{K}$ monotonic contours

Safety constraint

Weak Prior

0.2	0.6
0.05	0.3

TTL=0.33

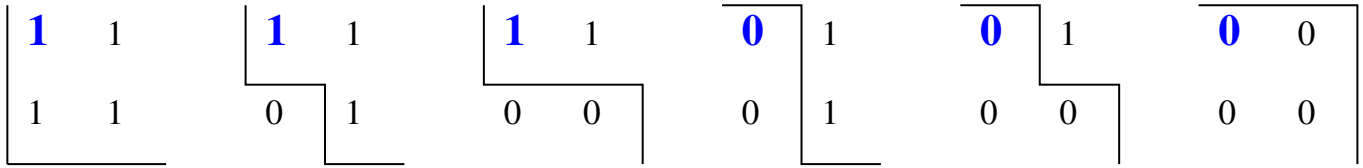

Contour Probs	0.11	0.17	0.17	0.19	0.20	0.16

Average prob above TTL

0.45	**0.84**
0.11	0.47

Dose selection for next cohort

- Identify the most likely MTC
- Given dose skipping restrictions, list set of dose/day combinations closest to MTC
- Select closest dose/day combination with **smallest** sample size
  - In event of tie, randomly select
- If no dose combinations are available due to violating safety constraint, trial is terminated early (no MTDC recommended).

At the end of the trial, the modal MTC is estimated.

- All combinations closest to MTC from below that have been experimented on are chosen as MTDCs (Mander and Sweeting, 2015).

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

Arrival of New Patients

Clinicians on ORCA-2 wanted option to enrol new patients when current patients still in DLT follow-up period

- How to decide where new patient should be allocated given partial data?
- Consider Time-to-Event (TITE) approach

Terminology

- J doses of drug A ($\{a_j : j = 1 \ldots, J\}$)
- K doses of drug B ($\{b_k : k = 1 \ldots, K\}$)
- Dose combination (a_j, b_k)
- $n_{jk,t}$ = number of people on dose combination (a_j, b_k) at time t
- $y_{i,t}$ = DLT outcome for patient i at time t
  - if $y_{i,t} = 1$, then $y_{i,t'} = 1 \; \forall t' \geqslant t$
- $\pi_{jk,t}$ = probability of DLT on dose combination (a_j, b_k) at time t

Weight functions (Cheung and Chappell, 2000)

$w_{i,t}$ = "partial" outcome for patient i given (a_j, b_k) at time t_{i0} and observed at time $t \in [t_{i0}, T + t_{i0}]$.

Linear

$$w_{i,t} = \begin{cases} 1 & \text{if } y_{i,t} = 1 \text{ and } t - t_{i0} \leqslant T \\ 1 - \frac{t - t_{i0}}{T} & \text{if } y_{i,t} = 0 \text{ and } t - t_{i0} \leqslant T \end{cases} \tag{3}$$

Adaptive (Cheung and Chappell, 2000)

$$w_{i,t} = 1 - \frac{1}{z+1}\left(\kappa + \frac{t - t_{(\kappa)}}{t_{(\kappa+1)} - t_{(\kappa)}}\right) \tag{4}$$

- DLT times $t_{(1)}, t_{(2)}, \ldots, t_{(z)}$ $(0 \equiv t_{(0)} < t_{(1)} \leqslant \cdots \leqslant t_{(z)} < t_{(z+1)} \equiv T)$
- $\kappa = \max_{0 \leqslant i \leqslant z}\{i : t \geqslant t_{(i)}\}$

Linear and adaptive weights

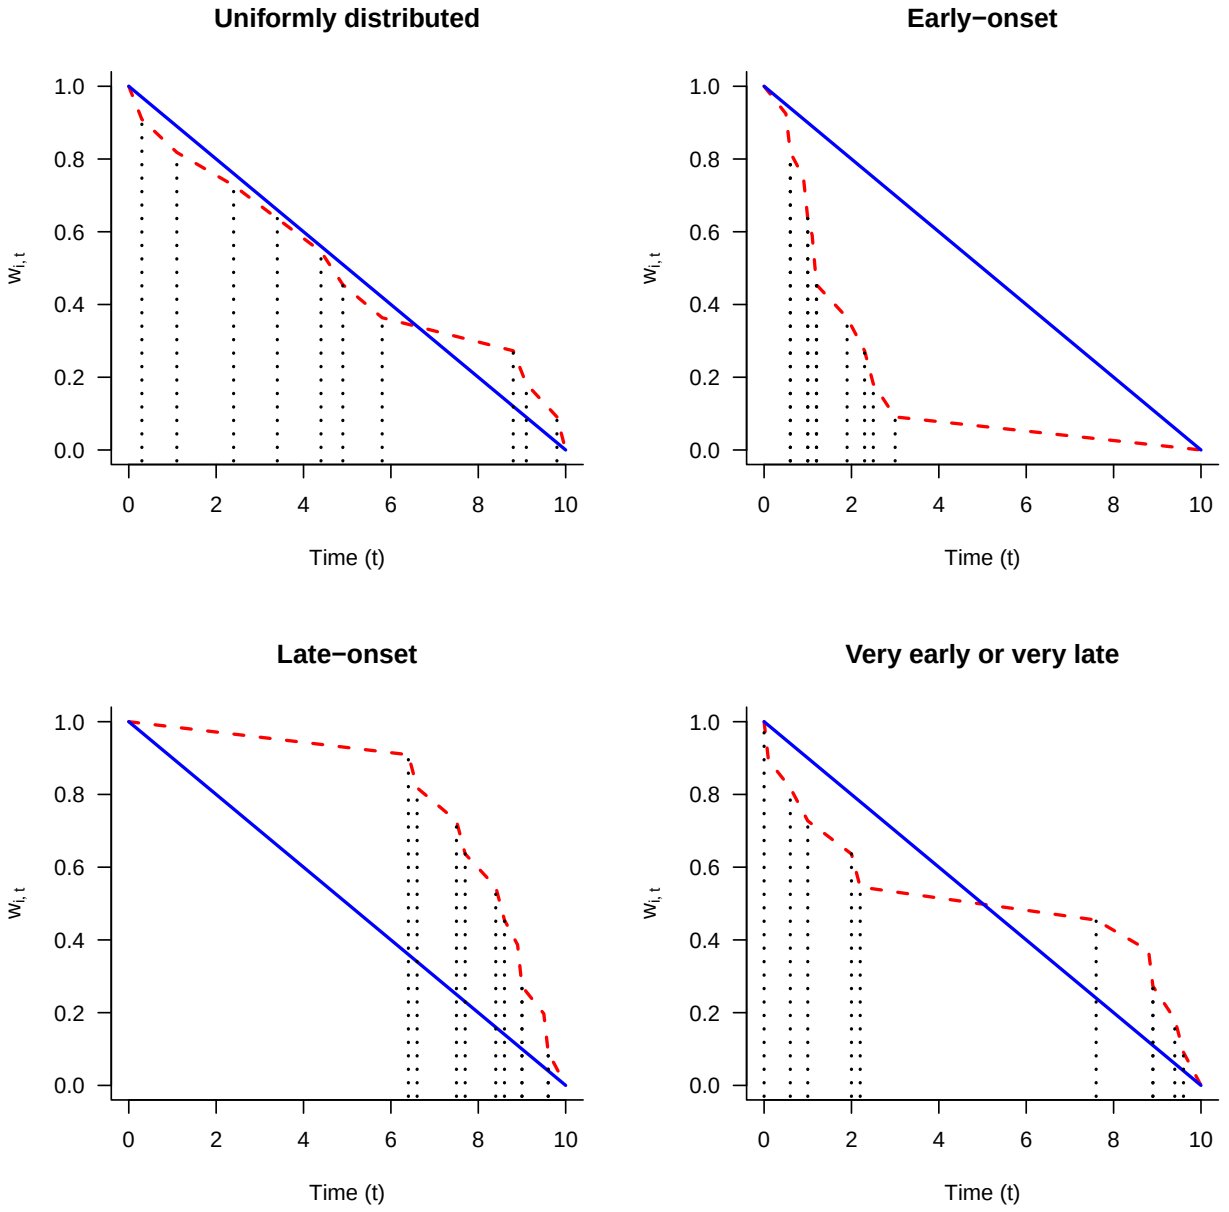

TITE-PIPE

A priori, $\pi_{jk,0} \sim \text{Beta}(r_{jk,0}, s_{jk,0})$.

- $r_{jk,0}$ and $s_{jk,0}$ chosen s.t. centered on prior medians and $\sum_{j=1}^{J} \sum_{k=1}^{K} (r_{jk,0} + s_{jk,0}) = 1$

For dose combination (a_j, b_k), at time t:

- $R_{jk,t} = \sum_{i=1}^{n_{jk,t}} w_{i,t}$ = number of DLTs
- $S_{jk,t} = \sum_{i=1}^{n_{jk,t}} (1 - w_{i,t}) = n_{jk,t} - R_{jk,t}$ = number of non-DLTs
- $\pi_{jk,t} \sim \text{Beta}(r_{jk,0} + R_{jk,t}, s_{jk,0} + S_{jk,t})$.

Posterior probability for dose $(a_j, b_k)_t$

Beta(prior + DLTs on (a_j, b_k) at time t, prior + non-DLTs on (a_j, b_k) at time t)

Arrival Times

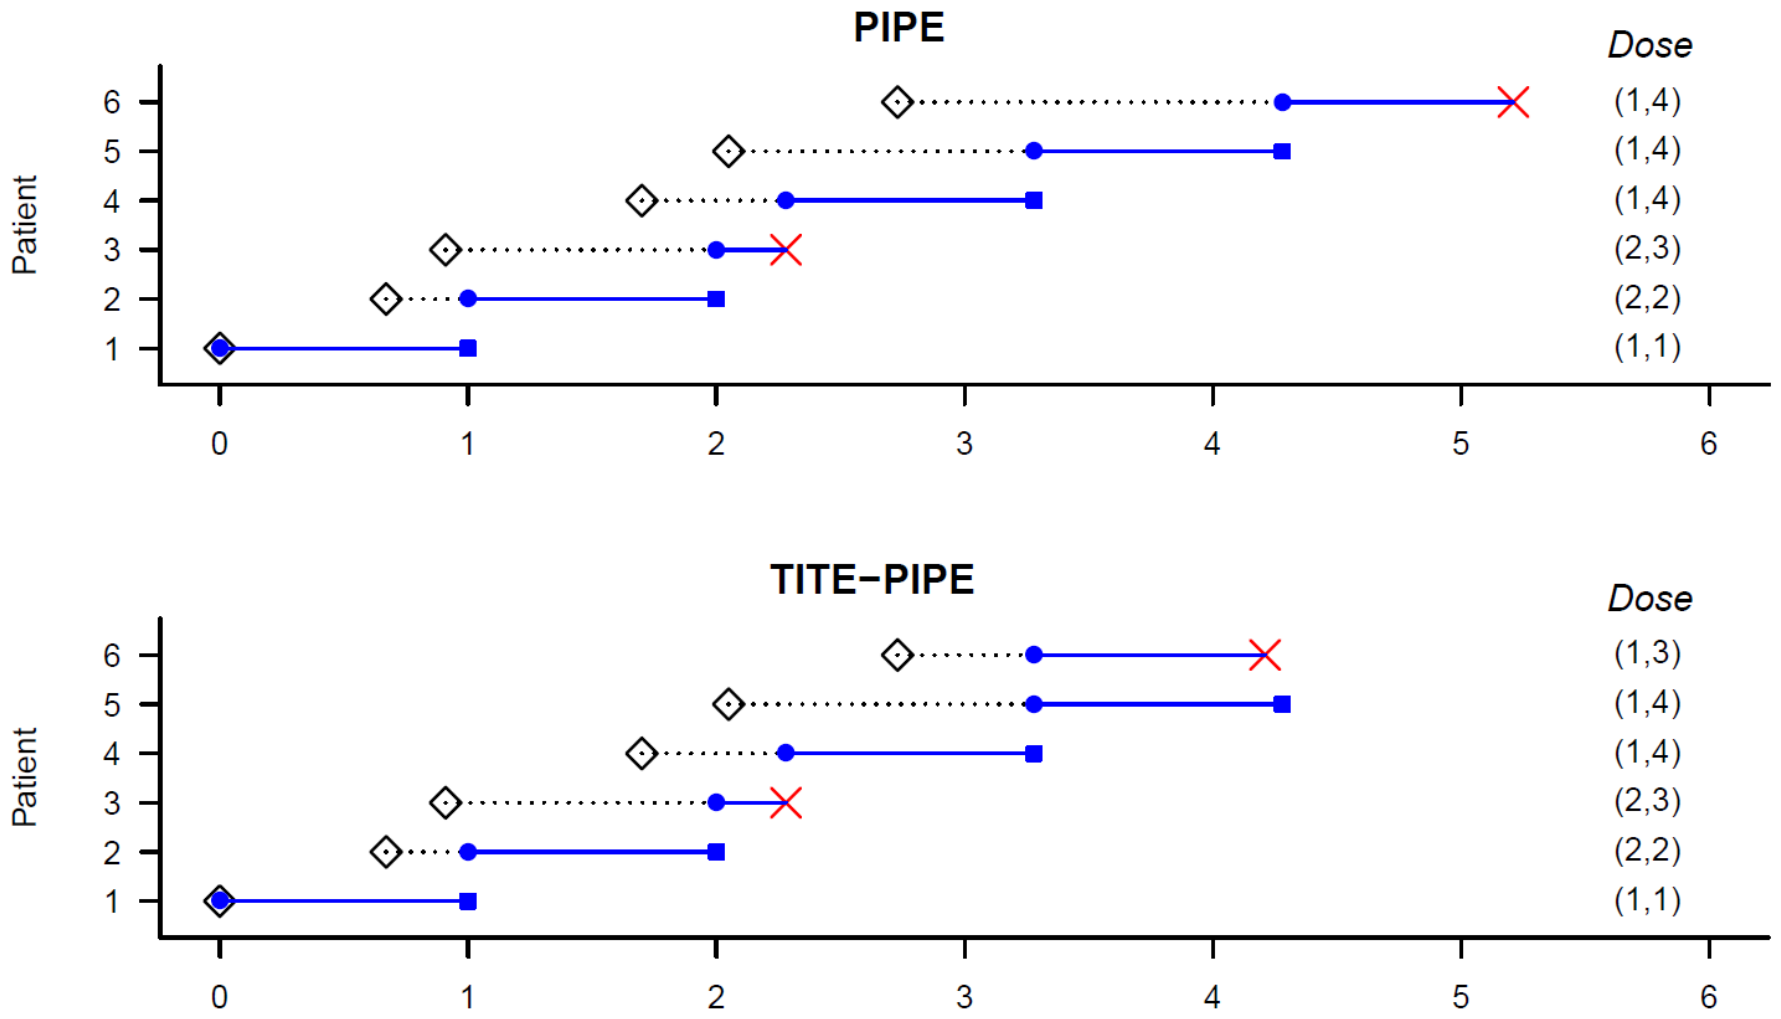

Safety constraints and early stopping

Stop trial when maximum sample size is reached, or

- $\mathbb{P}((a_1, b_1)_t > MTC_\theta) \geqslant \epsilon$ (computed using completed follow-up data only)
- if current data (complete and partial) give no dose as admissible for next cohort, wait until all patients have completed follow-up before potentially enrolling future patients (Ivanova et al. (2016)).

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

Scenarios (Braun and Jia, 2013)

Scenario A

Drug B \ Drug A	1	2	3	4
1	4	8	12	16
2	10	14	**18**	**22**
3	16	**20**	24	28
4	**22**	26	30	34

Scenario B

Drug B \ Drug A	1	2	3	4
1	2	4	6	8
2	5	7	9	11
3	8	10	12	14
4	11	13	15	17

Scenario C

Drug B \ Drug A	1	2	3	4
1	10	**20**	30	40
2	25	35	45	55
3	40	50	60	70
4	55	65	75	85

Scenario D

Drug B \ Drug A	1	2	3	4
1	44	48	52	56
2	50	54	58	62
3	56	60	64	68
4	62	66	70	74

Scenario E

Drug B \ Drug A	1	2	3	4
1	8	**18**	28	29
2	9	**19**	29	30
3	10	**20**	30	31
4	11	**21**	31	41

Scenario F

Drug B \ Drug A	1	2	3	4
1	12	13	14	15
2	16	**18**	**20**	**22**
3	44	45	46	47
4	50	52	54	55

Scenario G

Drug B \ Drug A	1	2	3	4
1	1	2	3	4
2	4	10	15	**20**
3	6	15	30	45
4	10	30	50	80

Simulation setup

Study starts when first patient given combination (a_1, b_1) ($t = t_{1,0} = 0$).

Patient followed up for $T = 1$ unit, or until onset of DLT, whichever occurs first.

Patients arrive as a Poisson process of rate λ:

- $\lambda = \{0.5, 1, 2\}$

For both PIPE and TITE-PIPE, require minimum of 2 patients to have completed treatment on each open dose combination before new decisions are made

- **PIPE**: Require complete follow-up from $\geqslant 2$ patients before dosing next cohort (never use partial data)
- **TITE-PIPE**: Require complete follow-up on two patients per cohort before allowing partial data to be used

2000 simulations per scenario

- Maximum sample size of 40 patients
- Early stopping: $\epsilon = 0.80$
- Use `pipe.design` package in R (modified code)

Time to toxicity distributions

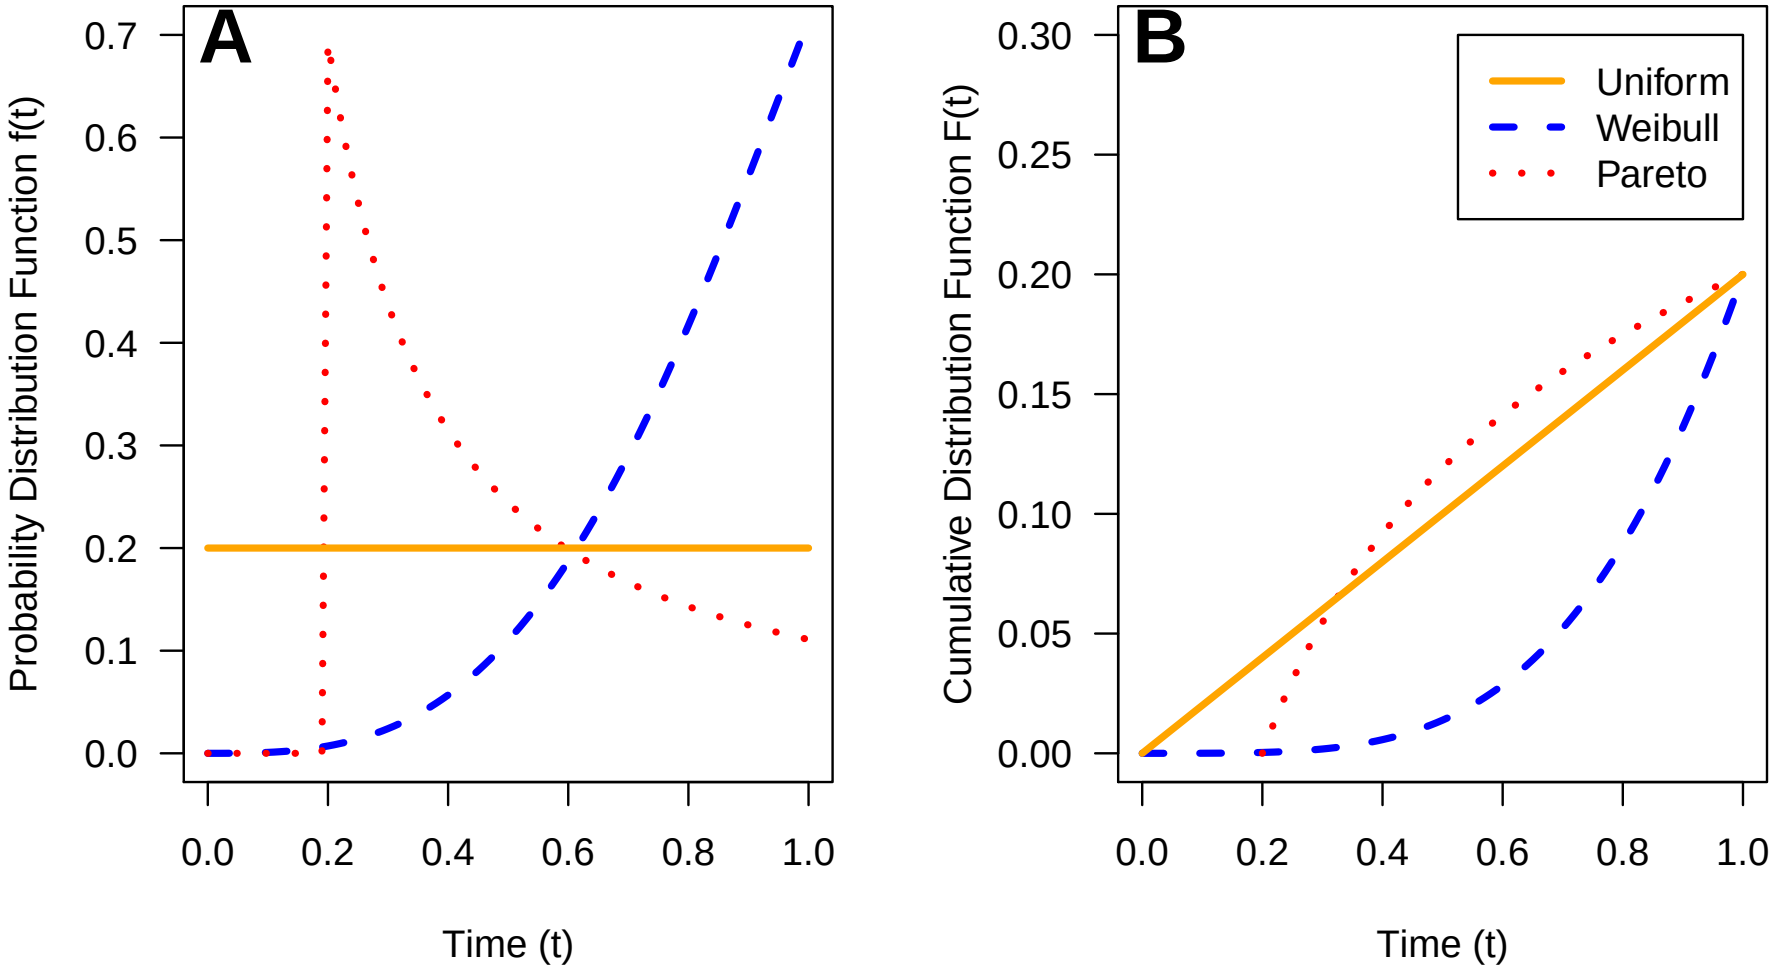

Figure: PDF (A) and CDF (B) of time to toxicity distributions (overall DLT risk = 0.20).

Experimentation (A-D)

Arrival rate (λ)	Design	Probability of DLT (%)					Mean Sample Size	Mean DLTs (%)
		0-14	15-24	25-34	35-45	46+		
	Scenario A							
	PIPE	20	63	17	0	0	40.0	19
0.5	TITE-PIPE	21	63	16	0	0	40.0	19
1	TITE-PIPE	22	62	15	0	0	39.9	19
2	TITE-PIPE	27	60	13	0	0	39.9	18
	Scenario B							
	PIPE	76	24	0	0	0	40.0	12
0.5	TITE-PIPE	77	23	0	0	0	40.0	12
1	TITE-PIPE	79	21	0	0	0	40.0	11
2	TITE-PIPE	83	17	0	0	0	40.0	11
	Scenario C							
	PIPE	14	15	27	33	11	39.5	31
0.5	TITE-PIPE	14	16	27	32	10	39.4	31
1	TITE-PIPE	15	16	27	32	10	39.4	31
2	TITE-PIPE	18	16	27	29	9	39.3	30
	Scenario D							
	PIPE	0	0	0	57	43	19.4	61
0.5	TITE-PIPE	0	0	0	60	40	19.4	61
1	TITE-PIPE	0	0	0	63	37	19.5	61
2	TITE-PIPE	0	0	0	66	34	19.5	61

Experimentation (E-G)

Arrival rate (λ)	Design	Probability of DLT (%)					Mean Sample Size	Mean DLTs (%)
		0-14	15-24	25-34	35-45	46+		
	Scenario E							
	PIPE	30	31	39	1	0	39.8	21
0.5	TITE-PIPE	30	31	39	0	0	39.8	21
1	TITE-PIPE	31	31	38	0	0	39.8	21
2	TITE-PIPE	33	30	36	0	0	39.8	20
	Scenario F							
	PIPE	20	55	0	14	11	39.5	25
0.5	TITE-PIPE	21	53	0	14	11	39.5	25
1	TITE-PIPE	23	52	0	14	11	39.4	25
2	TITE-PIPE	28	48	0	14	10	39.4	24
	Scenario G							
	PIPE	38	35	21	3	3	40.0	17
0.5	TITE-PIPE	40	35	20	3	2	40.0	17
1	TITE-PIPE	41	34	19	3	2	40.0	17
2	TITE-PIPE	46	33	17	2	2	40.0	16

Recommendation (A-D)

Arrival rate (λ)	Design	Probability of DLT (%)					Mean No. MTDCs	Trials with no MTDC (%)	Early Stop (%)
		0-14	15-24	25-34	35-45	46+			
	Scenario A								
	PIPE	12	73	15	0	0	2.3	0.2	0.2
0.5	TITE-PIPE	11	74	15	0	0	2.2	0.2	0.2
1	TITE-PIPE	12	73	15	0	0	2.2	0.4	0.4
2	TITE-PIPE	12	74	14	0	0	2.2	0.3	0.3
	Scenario B								
	PIPE	73	27	0	0	0	1.9	0	0
0.5	TITE-PIPE	73	27	0	0	0	1.9	0	0
1	TITE-PIPE	74	26	0	0	0	1.9	0	0
2	TITE-PIPE	77	23	0	0	0	2.0	0	0
	Scenario C								
	PIPE	16	24	35	19	1	1.3	4.5	2.1
0.5	TITE-PIPE	16	25	34	20	1	1.3	4.8	2.6
1	TITE-PIPE	14	23	35	20	2	1.3	5.6	2.8
2	TITE-PIPE	13	23	34	22	2	1.3	6.7	3.5
	Scenario D								
	PIPE	0	0	0	3	2	0	95.8	87.2
0.5	TITE-PIPE	0	0	0	2	2	0	96.2	87.0
1	TITE-PIPE	0	0	0	2	2	0	96.0	86.8
2	TITE-PIPE	0	0	0	2	2	0	96.3	86.4

Recommendation (E-G)

Arrival rate (λ)	Design	Probability of DLT (%)					Mean No. MTDCs	Trials with no MTDC (%)	Early Stop (%)
		0-14	15-24	25-34	35-45	46+			
	Scenario E								
	PIPE	30	32	37	0	0	2	0.9	0.7
0.5	TITE-PIPE	30	32	37	0	0	2	0.8	0.7
1	TITE-PIPE	30	31	38	0	0	2	0.8	0.6
2	TITE-PIPE	30	32	36	0	0	2	1.2	1.1
	Scenario F								
	PIPE	13	70	0	11	4	1.7	2.1	1.6
0.5	TITE-PIPE	12	71	0	11	4	1.7	1.9	1.5
1	TITE-PIPE	12	71	0	10	4	1.7	2.8	1.8
2	TITE-PIPE	12	70	0	11	4	1.7	3.5	2.0
	Scenario G								
	PIPE	44	38	17	1	0	2.7	0	0
0.5	TITE-PIPE	45	37	17	1	0	2.8	0	0
1	TITE-PIPE	44	37	18	1	1	2.7	0	0
2	TITE-PIPE	44	36	18	1	0	2.7	0	0

Trial duration (A-G)

Arrival rate (λ)	Design	Scenario						
		A	B	C	D	E	F	G
0.5	PIPE	79.3	79.5	78.4	38.7	79.0	78.4	79.5
	TITE-PIPE	79.0	79.1	78.0	38.7	78.6	78.0	79.1
1	PIPE	41.8	42.2	41.1	20.0	41.6	41.3	42.0
	TITE-PIPE	40.0	40.1	39.4	19.6	39.8	39.5	40.0
2	PIPE	29.8	30.8	29.1	13.6	29.7	29.5	30.1
	TITE-PIPE	20.5	20.6	20.2	10.3	20.4	20.2	20.5

- Experimentation more conservative under TITE-PIPE when recruitment faster than expected
- Recommendation similar between PIPE and TITE-PIPE approaches
- Savings are in trial duration (and thus cost)

Overview

1 Phase I clinical trials

2 ORCA-2 Trial and PIPE Design

3 Time-to-Event PIPE Design

4 Simulation Work

5 Summary

Summary

Partial outcomes can easily be incorporated into PIPE

- Potential for savings in time and cost
- comparable experimentation and recommendation performance, even with early and late-onset toxicity
- From our work, weight function choice does not matter
- TITE-PIPE code to be incorporated into `pipe.design` package (R)

For future research on TITE-PIPE, consider :

- having a “must-observe” observation window before new patients enrolled - e.g. WISTERIA trial (Birmingham CTC)
- larger and/or non-square dose-toxicity grids
- Ensure each contour has a uniform prior weight of being the MTC
- Comparison to model-based approaches (e.g. Wages et al. (2013))
- Use efficacy endpoint data to find a biologically optimum dose.

Publication

Appl. Statist. (2019)
68, *Part* 2, *pp.* 309–329

A Bayesian model-free approach to combination therapy phase I trials using censored time-to-toxicity data

Graham M. Wheeler,

University College London, UK

Michael J. Sweeting

University of Cambridge and University of Leicester, UK

and Adrian P. Mander

University of Cambridge, UK

Code: https://rss.onlinelibrary.wiley.com/hub/journal/14679876/series-c-datasets/68_2

Acknowledgements

- Michael Sweeting (University of Leicester)
- Adrian Mander (MRC Biostatistics Unit, Cambridge)

Funding from:

- MRC (AM, MS)
- NIHR (MS)
- British Heart Foundation (MS)
- Cancer Research UK (GW)

References

Babb, J. S., Rogatko, A., 2004. Bayesian methods for phase I cancer clinical trials. In: Geller, N. L. (Ed.), Advances in Clinical Trial Biostatistics. Marcel Dekker, New York, NY, pp. 1–40.

Braun, T. M., Jia, N., 2013. A Generalized Continual Reassessment Method for Two-Agent Phase I Trials. Statistics in Biopharmaceutical Research 5 (2), 105–115.

Cheung, Y. K., Chappell, R., 2000. Sequential designs for phase I clinical trials with late-onset toxicities. Biometrics 56, 1177–1182.

Ivanova, A., Wang, Y., Foster, M. C., 2016. The rapid enrollment design for Phase I clinical trials. Statistics in Medicine 35, 2516–2524.

Mander, A. P., Sweeting, M. J., 2015. A product of independent beta probabilities dose escalation design for dual-agent phase I trials. Statistics in Medicine 34 (8), 1261–1276.

Wages, N. A., Conaway, M. R., O'Quigley, J., 2013. Using the time-to-event continual reassessment method in the presence of partial orders. Statistics in Medicine 32 (1), 131–141.