Guiding phase I dose escalation for modern cancer therapies: recent developments, opportunities and challenges

Lukas A. Widmer, Dr. sc. ETH Zürich
Statistical Planning of Translational Studies
Symposium, Göttingen
March 21st, 2023

Acknowledgements

Time-to-event modelling:

Sebastian Weber
Hans-Jochen Weber
Yunnan Xu

BLRM framework:

Andrew Bean
David Ohlssen
Beat Neuenschwander

The opinions expressed in this presentation and on the following slides are solely of the presenter, and not necessarily those of Novartis.

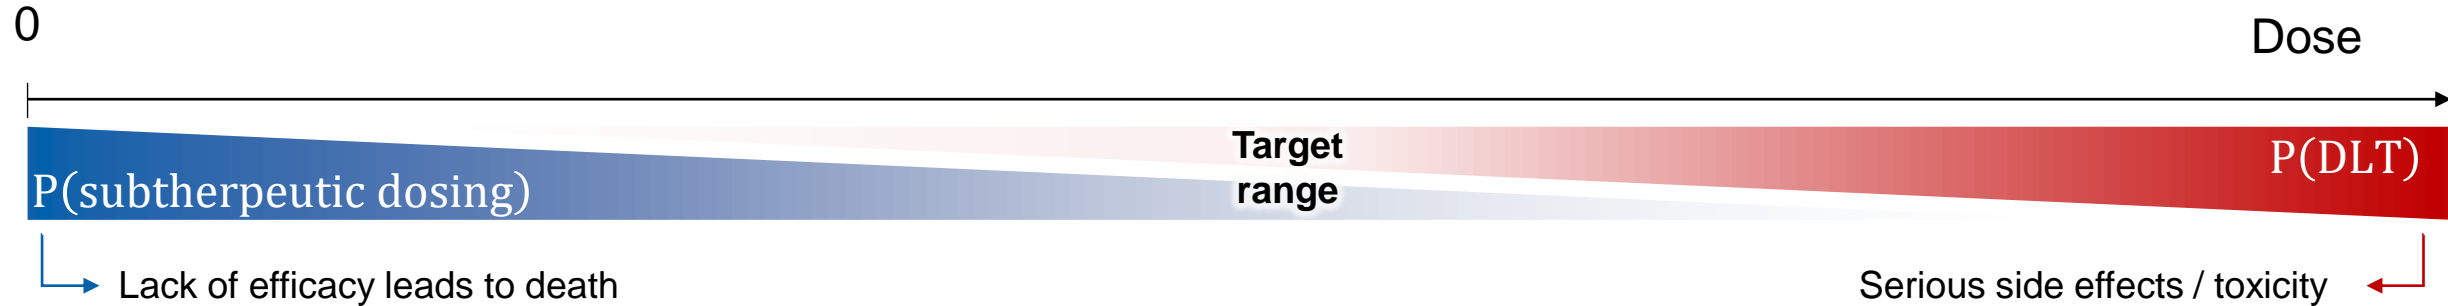# Phase I trials in Oncology: Small sample size & open-label trials

Volunteers are typically "patients ... whose cancers have progressed despite standard treatments"[1]

Not any data can be collected: ethical (a-c) & operational (d) challenges

a) Maximizing efficacy is mandated by the urgent need to have a treatment effect (due to the life-threatening nature of the disease).

b) One must not avoid giving knowingly working treatments to patients in order to study a new drug.

c) Toxicity risk to patients due to overdosing must be controlled given limited data.

d) Not so many patients qualify.

1. Le Tourneau, C. *et al*. Dose escalation methods in phase I cancer clinical trials. *J. Natl. Cancer Inst.* **101**, 708–720 (2009).

Oncology phase I dose escalation is a delicate balance between sub-therapeutic & toxic dosing

0

Dose

P(subtherpeutic dosing)

Target range

P(DLT)

→ Lack of efficacy leads to death

Serious side effects / toxicity ←

Initially, we have limited knowledge on toxicity

- Need to limit risk to current (and future) patients[1] ⇒ can initially only use small cohorts.

The goal of phase I dose escalation studies is to systematically increase the dosing

- as quickly as possible, to reach a biologically active and (hopefully) efficacious dose,

- as safely as possible, such that Dose-Limiting-Toxicity (DLT) events are controlled,

- to determine the safe dosing range for further development of the therapy (and ideally collect early efficacy data).

1. Babb, J *et al*. Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Stat. Med.* 17, 1103–1120 (1998).

For cytotoxic therapies, the maximum tolerated dose (MTD) in a treatment cycle (~ 4 weeks) is the optimal dose
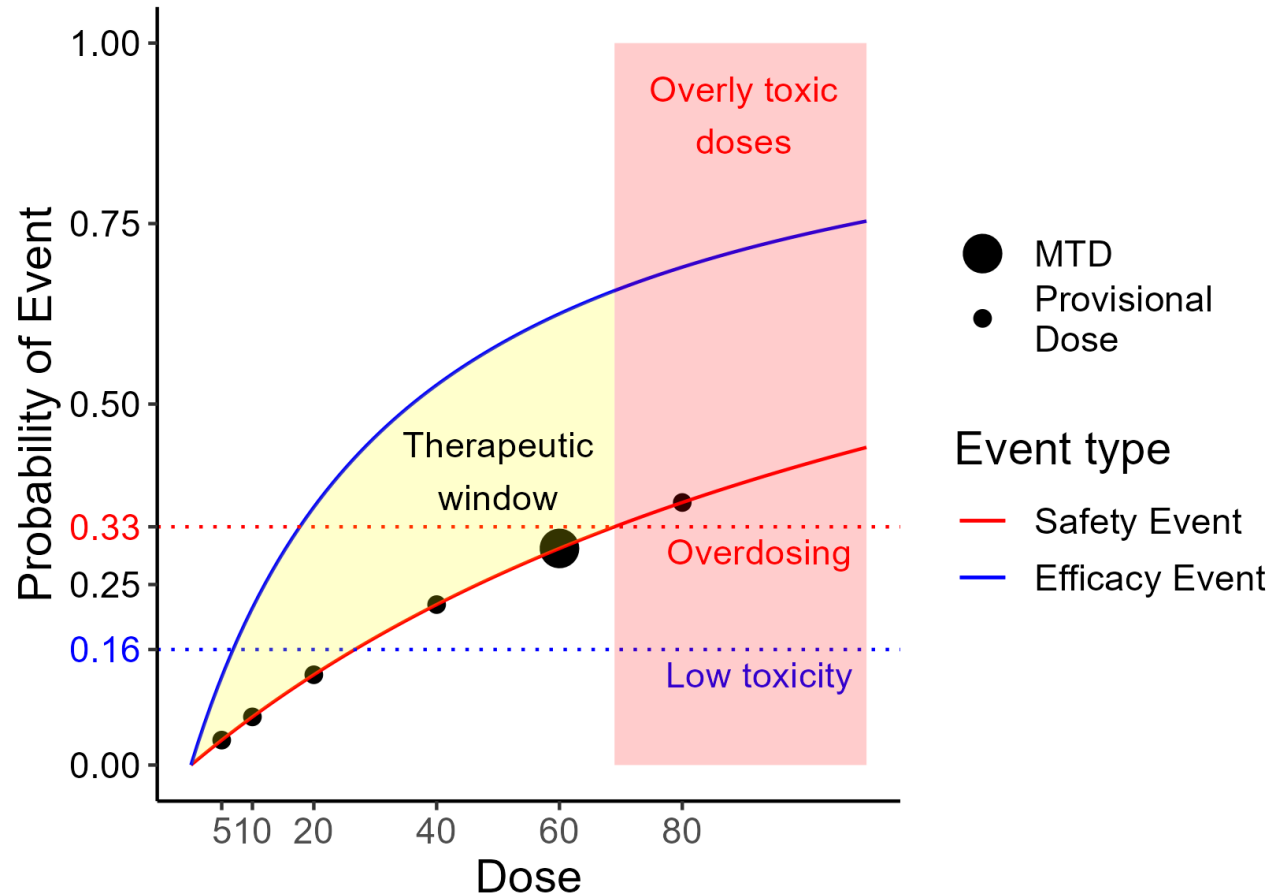
- For cytotoxic therapies, efficacy is due to the toxicity of the drug (especially to rapidly-dividing cells).
  → Efficacy ≈ Toxicity
  → Hit hard and fast & repeat in cycles

- Thus, given the severity of the disease, we want to go to as high a dose as possible to maximize the potential of killing all cancer cells, i.e., the maximum tolerated dose (MTD).

... however, many modern therapies are not cytotoxic.

For modern therapies, the optimal dose is not necessarily the maximum tolerated dose (MTD)
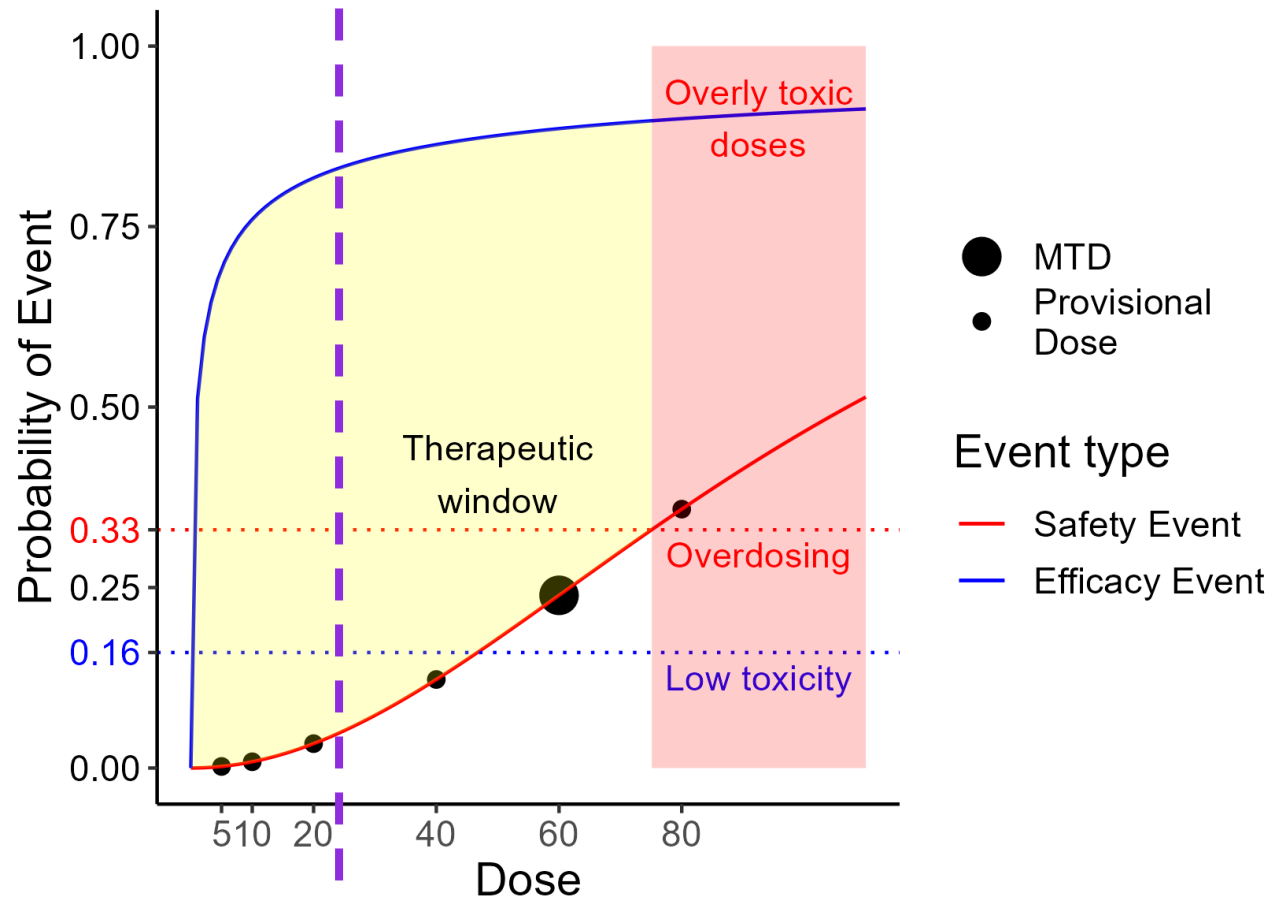
For modern therapies:

Efficacy and toxicity dose-response might be different, and the dose might not be as directly linked to efficacy & toxicity simultaneously.

E.g.: cell therapies are living "drugs" where efficacy of a given "dose" also depends on the patient & disease.

Also, these therapies must be tolerated over longer periods of time.

→ **The optimum benefit / risk (Optimal Biological Dose OBD) may no longer be at the MTD for one cycle of treatment!**

Consequences of not optimizing dosage

If treatment is poorly-tolerated, patients may stop taking an otherwise efficacious treatment or choose to try a different one

The treatment might put patients at excessive risk in downstream phases:

- Risk for development programs in phases II / III, e.g., if the investigational treatment is doing worse than standard-of-care due to tolerability issues (examples of this are shown in a recent FDA project Optimus seminar[1])

- If the drug is already on the market, it might have to be withdrawn

Dose optimization is more challenging to conduct post-approval

- Patients may not want to enroll in a trial for a commercially available treatment

- Focus may shift to novel (and potentially better) treatments

1. Project Optimus – FDA's New Dose Optimization & Selection Paradigm in Oncology Drug Development (youtube.com)

Phase I dose-escalation designs:
Standard safety models in the past and nowadays

Traditional designs aimed to quickly find the maximum tolerated dose (MTD)

- Ideal for cytotoxic treatments where efficacy ≈ toxicity (MTD close to optimal), and
- Monitoring cycle 1 toxicity to guide dose escalation is sufficient

Novel treatments differ

- **Short-term safety not sufficient** to assess efficacy
- Longer-term **tolerability of therapy** must be warranted
- **Timing of dosing** can be critical (e.g., ramp-up regimes avoid cytokine release syndrome)
- → **Novel designs should allow for separate safety modeling (which doses are too toxic?) and efficacy modeling (of the safe doses, which are efficacious?).**

FDA project Optimus challenges current standard paradigm fundamentally

Can we evolve current practice to develop novel therapies better with a focus on safety?

NOVARTIS | Reimagining Medicine

3+3 design:
stereotypically traditional rule-based dose escalation

3+3 design does not need a statistician

Originally introduced in the 1940s, described 1989 in the context of Phase I Oncology[1].

Recruits 3 patients at a given dose level, then, at

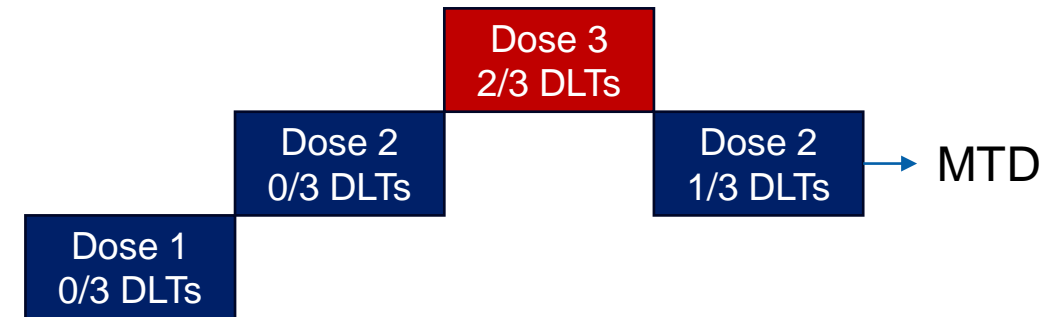0/3 DLTs: escalate

1/3 DLT: recruit another 3 patients, then:

      1/6 DLT: escalate

      2+   DLTs: de-escalate and
               never re-escalate.

2+ DLTs: de-escalate and never re-escalate.
        If 6 patients at lower dose level
        already tested, declare MTD.

Why is this a bad design?

- Does not use data within/across dose levels.
- Cohort size always 3
- Declares MTD based on 6 patients
- etc... too many reasons to list here, really.

1.   Storer, B. E. Design and Analysis of Phase I Clinical Trials. *Biometrics*, *45*(3), 925 (1989).

Here, we consider the Bayesian Logistic Regression Model (BLRM)[1] framework

Why? Because it can deal with

- arbitrary dose levels: it is a continuous model where distance from existing data is considered

- an arbitrary number of drugs being escalated simultaneously (N-dimensional)[2]

- historical data for single drug and/or combinations of arbitrary order[2]

- an arbitrary number of drug-drug interactions of arbitrary order[3]

Such a flexible and powerful framework comes with the cost of dealing with priors, proper model specification and computational complexity (model update / MCMC after each cohort)

1. Neuenschwander, B *et al*. Critical aspects of the Bayesian approach to phase I cancer trials. *Stat. Med.* 27(13), 2420–2439 (2008).
2. Weber S, Widmer L, Bean A. OncoBayes2: Bayesian Logistic Regression for Oncology Dose-Escalation Trials. R package version 0.8-9 (2023).
3. Widmer, Lukas A., *et al*. Principled Drug-Drug Interaction Terms for Bayesian Logistic Regression Models of Drug Safety in Oncology Phase I Combination Trials. arXiv preprint arXiv:2302.11437 (2023).

Given sparse phase I data, dose-safety curves are uncertain
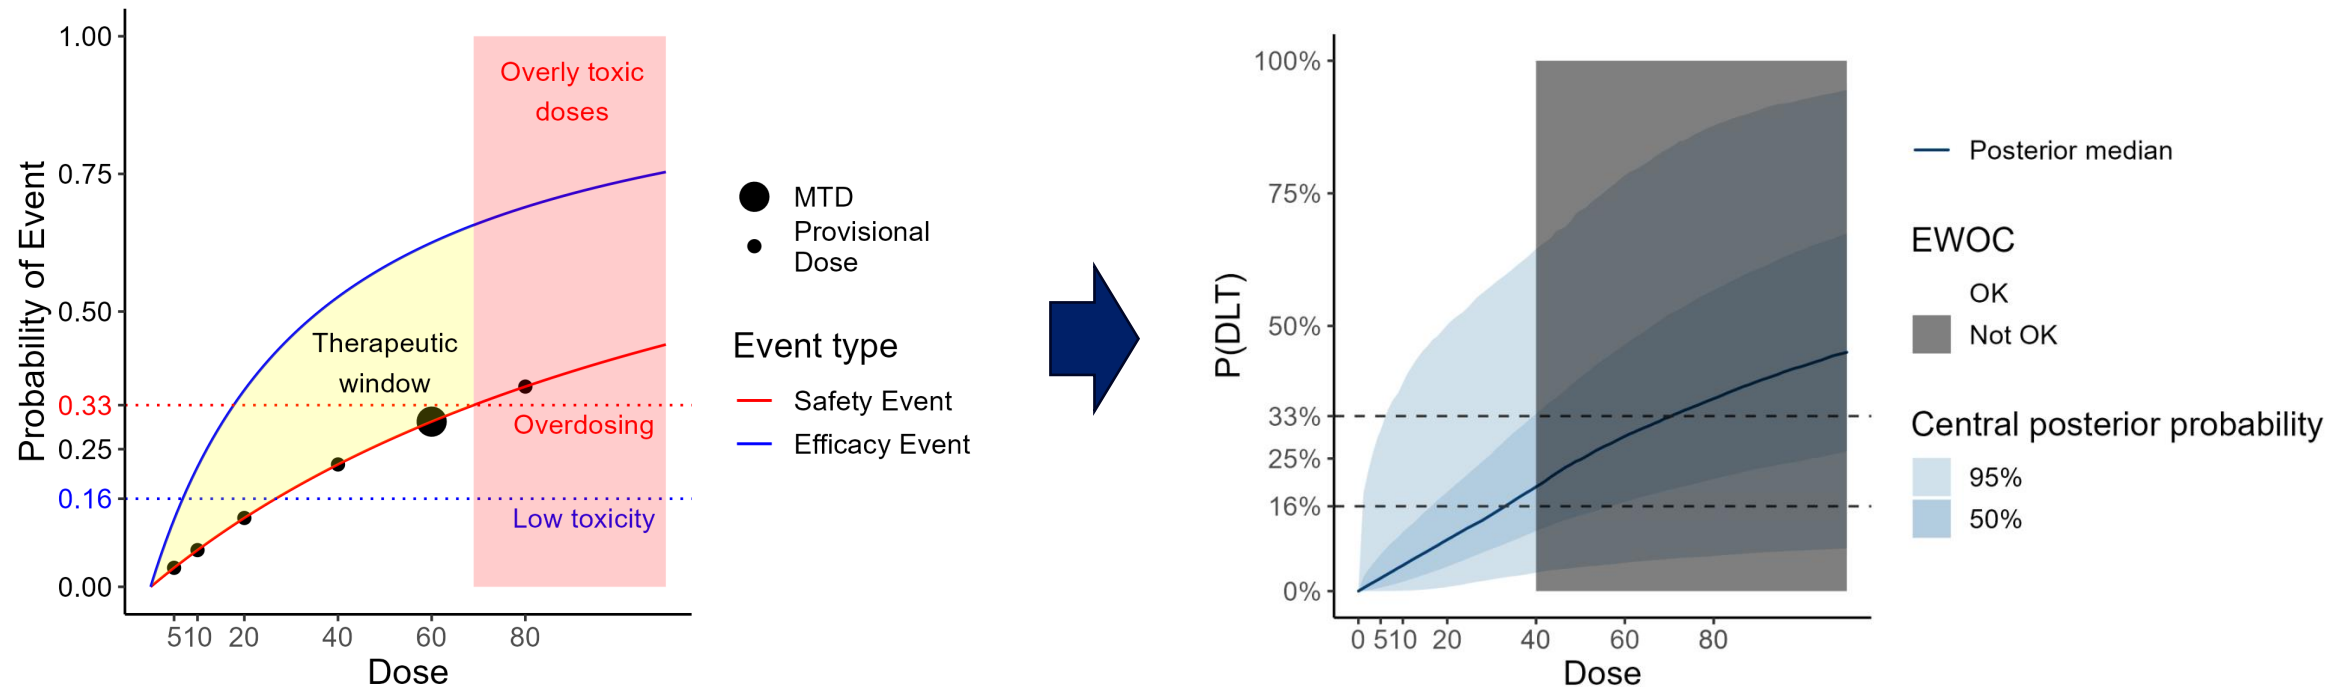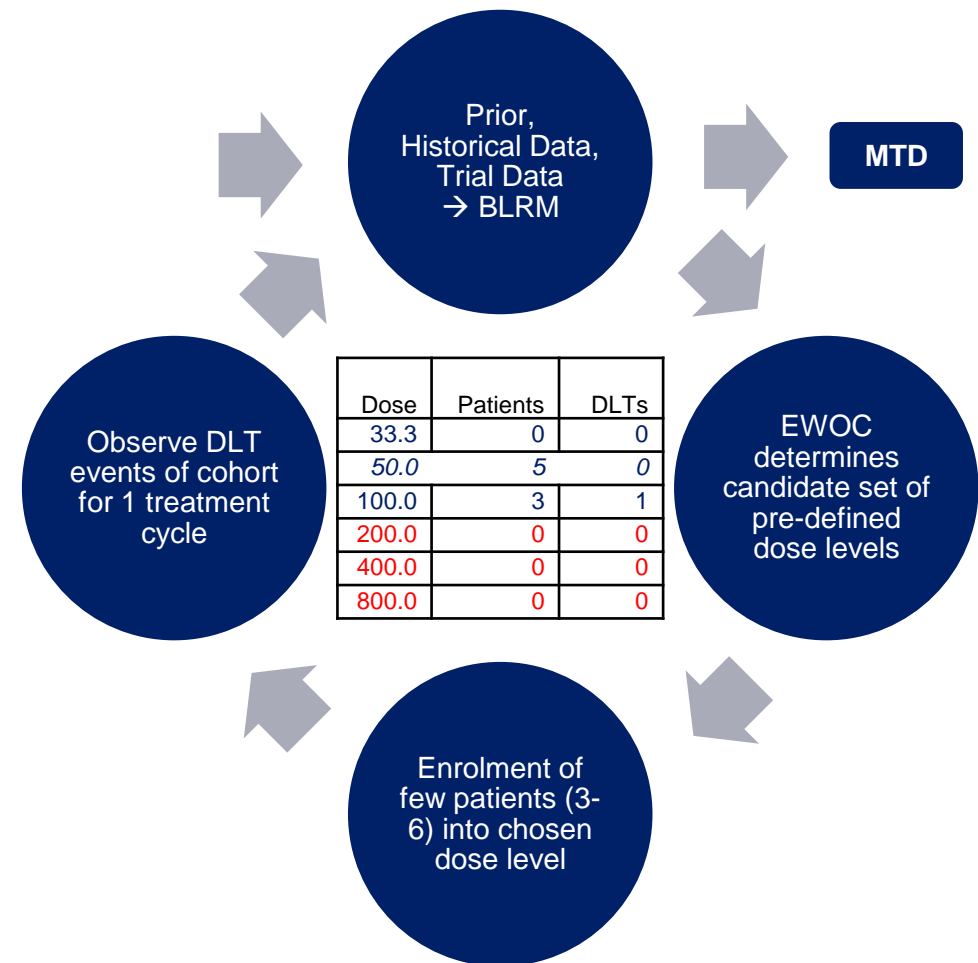
Running a trial using the BLRM with EWOC

We start with a small sample and perform adaptive **Escalation With Overdose Control (EWOC)**[1] step-by-step to warrant patient safety.

Probability π of dose limiting toxicity (DLT) event at dose d during one cycle.

EWOC for dose fulfilled
$$\Leftrightarrow P(\pi(d) \in \text{overdose}) < 0.25$$

- **Overdose: $\pi(d) > 1/3$**

- **Avoids too toxic (or too uncertain) doses!**

Dose	Patients	DLTs
33.3	0	0
50.0	5	0
100.0	3	1
200.0	0	0
400.0	0	0
800.0	0	0

Prior, Historical Data, Trial Data → BLRM

MTD

EWOC determines candidate set of pre-defined dose levels

Enrolment of few patients (3-6) into chosen dose level

Observe DLT events of cohort for 1 treatment cycle

1. Babb, J *et al.* Cancer phase I clinical trials: Efficient dose escalation with overdose control. *Stat. Med.* 17, 1103–1120 (1998).

NOVARTIS | Reimagining Medicine

BLRM with Escalation With Overdose Control (EWOC)

Underdose: $\pi(d) \leq 1/6$
Target dose: $1/6 < \pi(d) \leq 1/3$
Overdose: $\pi(d) > 1/3$
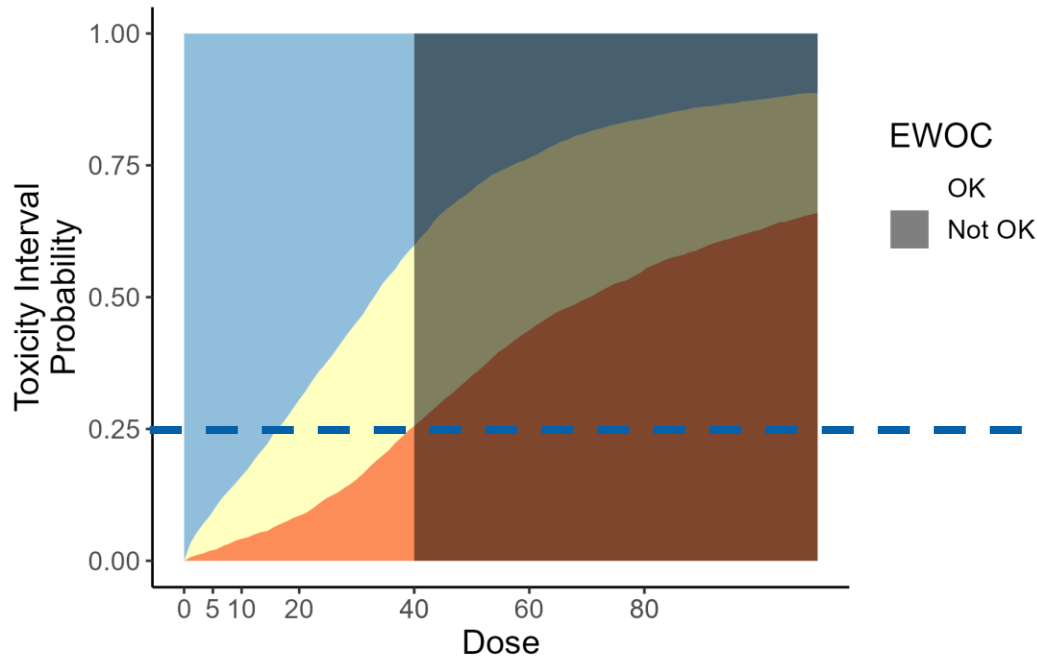
Bayesian logistic regression model (BLRM)[1]:
$$Y_i \,|\, \pi(d_i), n_i \sim \text{Binomial}(\pi(d_i), n_i)$$
$$\text{logit}\big(\pi(d_i)\big) = \alpha + \beta \log(d_i/d_{\text{ref}})$$

Probability π of dose limiting toxicity (DLT) event at dose d_i during one cycle:

- α log-odds for DLT when $d_i = d_{\text{ref}}$

- $\beta > 0$ slope (monotonically increasing)

EWOC for dose fulfilled
$\Leftrightarrow P(\pi(d) \in \text{overdose}) < 0.25$

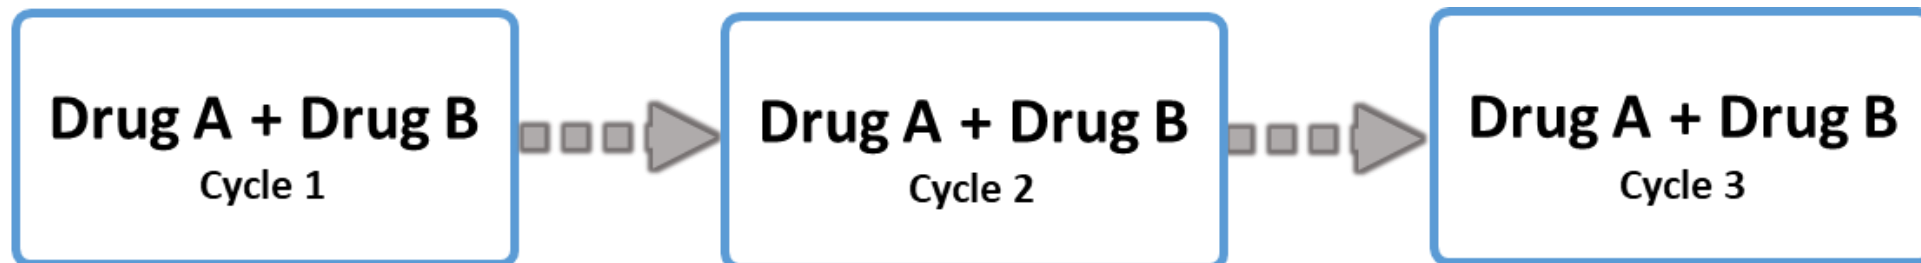Avoids too toxic (or too uncertain) doses!

EWOC
OK
Not OK

1. Neuenschwander, B *et al*. Critical aspects of the Bayesian approach to phase I cancer trials. *Stat. Med.* 27(13), 2420–2439 (2008).

NOVARTIS | Reimagining Medicine

Motivating indication: Post-transplant Acute Myeloid Leukemia (AML)

AML is eventually treated with a stem cell transplant

- Following transplantation, drug A is given for multiple cycles to enhance benefits of transplant.

- Can an additional drug B (at different doses) after the transplant reduce the risk of a relapse?

- How can we monitor safety of the entire therapy and maintain typical trial conduct with dose escalation decisions already after cycle 1 with partial exposure data of patients?

Turning to time-to-event (TTE) modeling…

Challenge	Approach
Need for absolute risk probabilities	Modeling of baseline hazard
Changes in baseline hazard	Time/cycle-varying hazard (cycle specific intercept/slope)
Data sparsity	- Coarse time units aligned with model focus on cycles (or weekly dosing regimens if needed) - Priors accounting for continuity in time - Structural modeling techniques, i.e., monotone increase/decrease of intercepts/slopes
Drug combinations	Additive hazard modeling (not discussed here)
Drug-drug interactions	Multiplicative hazard modeling (not discussed here)
Multi-cycle model prior	Extend reference concept for dose with time

NOVARTIS | Reimagining Medicine

BLRM & TTE model equivalence

BLRM – binomial endpoint

- Probability for DLT within 1 cycle: $\pi(d)$

- Linear model on log-odds scale:
$\text{logit}\big(\pi(d)\big) = \alpha + \beta \log(d)$

TTE – Poisson process (time-varying)

- Hazard for an event in 1 cycle: $h(d)$

- Linear model on log-hazard scale:
$\log\big(h(d)\big) = \alpha + \beta \log(d)$

Given $P(T \leq t|d) = 1 - \exp(-h(d)t)$, we get
$\text{cloglog}\big(P(T \leq t|d)\big) = \alpha + \beta \log(d) + \log(t)$

Same as the BLRM in cycle 1, except for the link function
(if time elapses in units of cycles, since then t=1)!

NOVARTIS | Reimagining Medicine

Model priors: *biologically plausible* ones for the intercept α and the slope β are important!
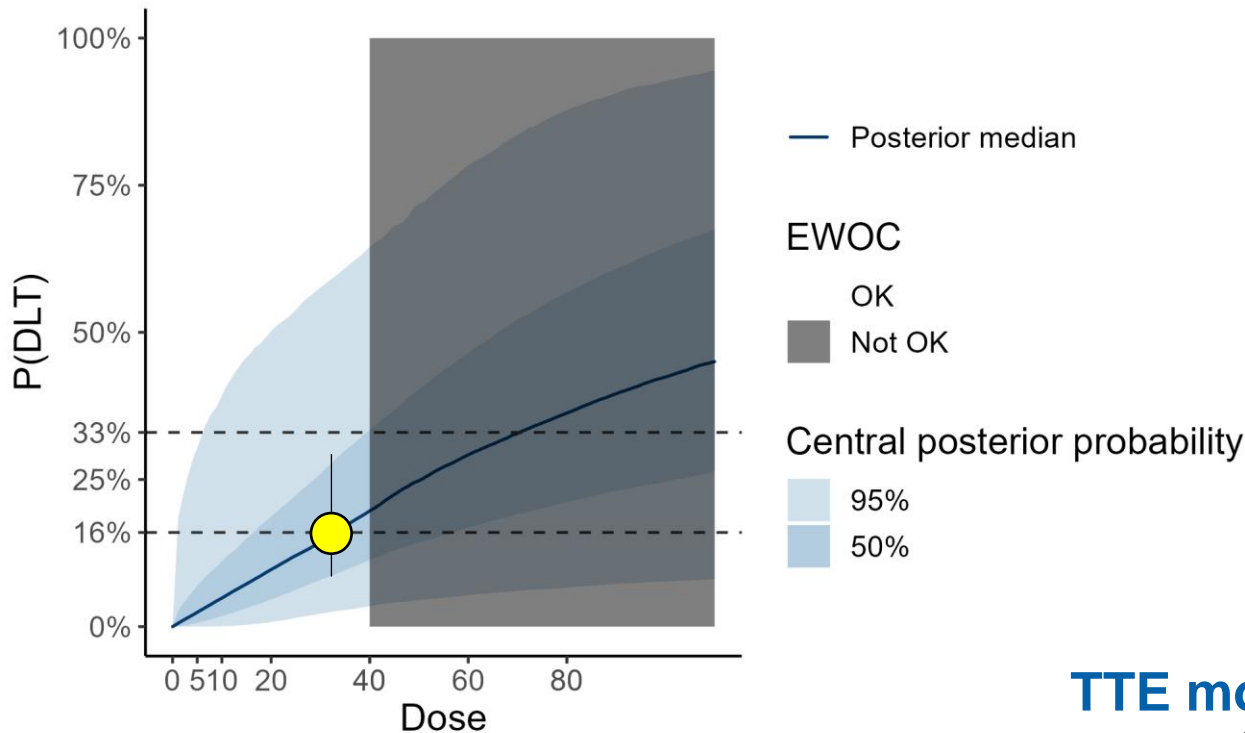
Logistic model:

$$\text{logit } \pi(d) = \alpha + \beta \log(d/d_{\text{ref}})$$

🟡 Reference dose d_{ref} determines meaning of prior

Intercept: $\alpha \sim \mathcal{N}(\text{logit}(\pi_{\text{ref}} = 0.20), 1^2)$

➡ Knowledge from preclinical studies, competitors with similar MoA, ...

Slope: $\log(\beta_i) \sim \mathcal{N}(0, (\log(4)/1.96)^2)$

➡ Knowledge of what is typically feasible in biological systems

TTE model:

$$\text{cloglog}\big(P(T \leq t_{\text{ref}}|d)\big) = \alpha + \beta \log(d/d_{\text{ref}}) + \log(t_{\text{ref}})$$

TTE model requires an additional **reference time-point** (often set to unity for end of cycle 1) for the intercept:

$$\alpha \sim \mathcal{N}(\text{cloglog}(\pi_{\text{ref}} = 0.20) - \log(t_{\text{ref}}), 1^2)$$

Multi-cycle overdose control metric

- Escalation with overdose control (EWOC) principle restricts candidate dose set

  $$\pi_c = 33\% \text{ and } p_c = 25\%$$

  $$P(\pi(d) > \pi_c) < p_c$$

- Per-cycle toxicity control
  - Toxicity within each cycle must be controlled

    $$\forall_{j=1}^{J} P\big(P\big(t \in I_j \big| t > I_{j-1}, d\big) > \pi_c\big) < p_c$$
  - Conditions on no event up to a given cycle

- Therapy-centric toxicity control
  - *Cumulative* toxicity over all J cycles must be controlled

    $$P\big(P\big(t \in \{I_1, \ldots, I_J\} \big| d\big) > \pi_c\big) < p_c$$
  - Due to longer exposure time more stringent (if the critical thresholds are kept constant)

Which method should we use for our 3-cycle case study? What do we gain from the individual-patient TTE data?

→ **Simulation study:**

How do the methods control toxicity of the entire therapy under different timing of DLT events for
- patients in the trial, and
- further development (future patients – MTD should not be overly toxic)?

Longer treatments often mean higher dropout - how do the methods perform then?

Operating characteristics scenarios

3-cycle therapy with toxicity profile: constant, decreasing or increasing

Dropout rates over 3 cycles:
0%, 33%, 55%

Accrual rate: 1 patient every 10 days

MTD declaration rules:
N = 6 on dose reached, and

- N ≥ 12 per trial, or
- Target probability ≥ 50%

1000 trials per scenario

Label	Model	Time	Overdose metric
B1	BLRM	1 cycle	standard EWOC
TCO	TTE	3 cycles	conditional by cycle
B3	BLRM	3 cycles	standard EWOC with 3 cycles
TCU	TTE	3 cycles	cumulative for 3 cycles

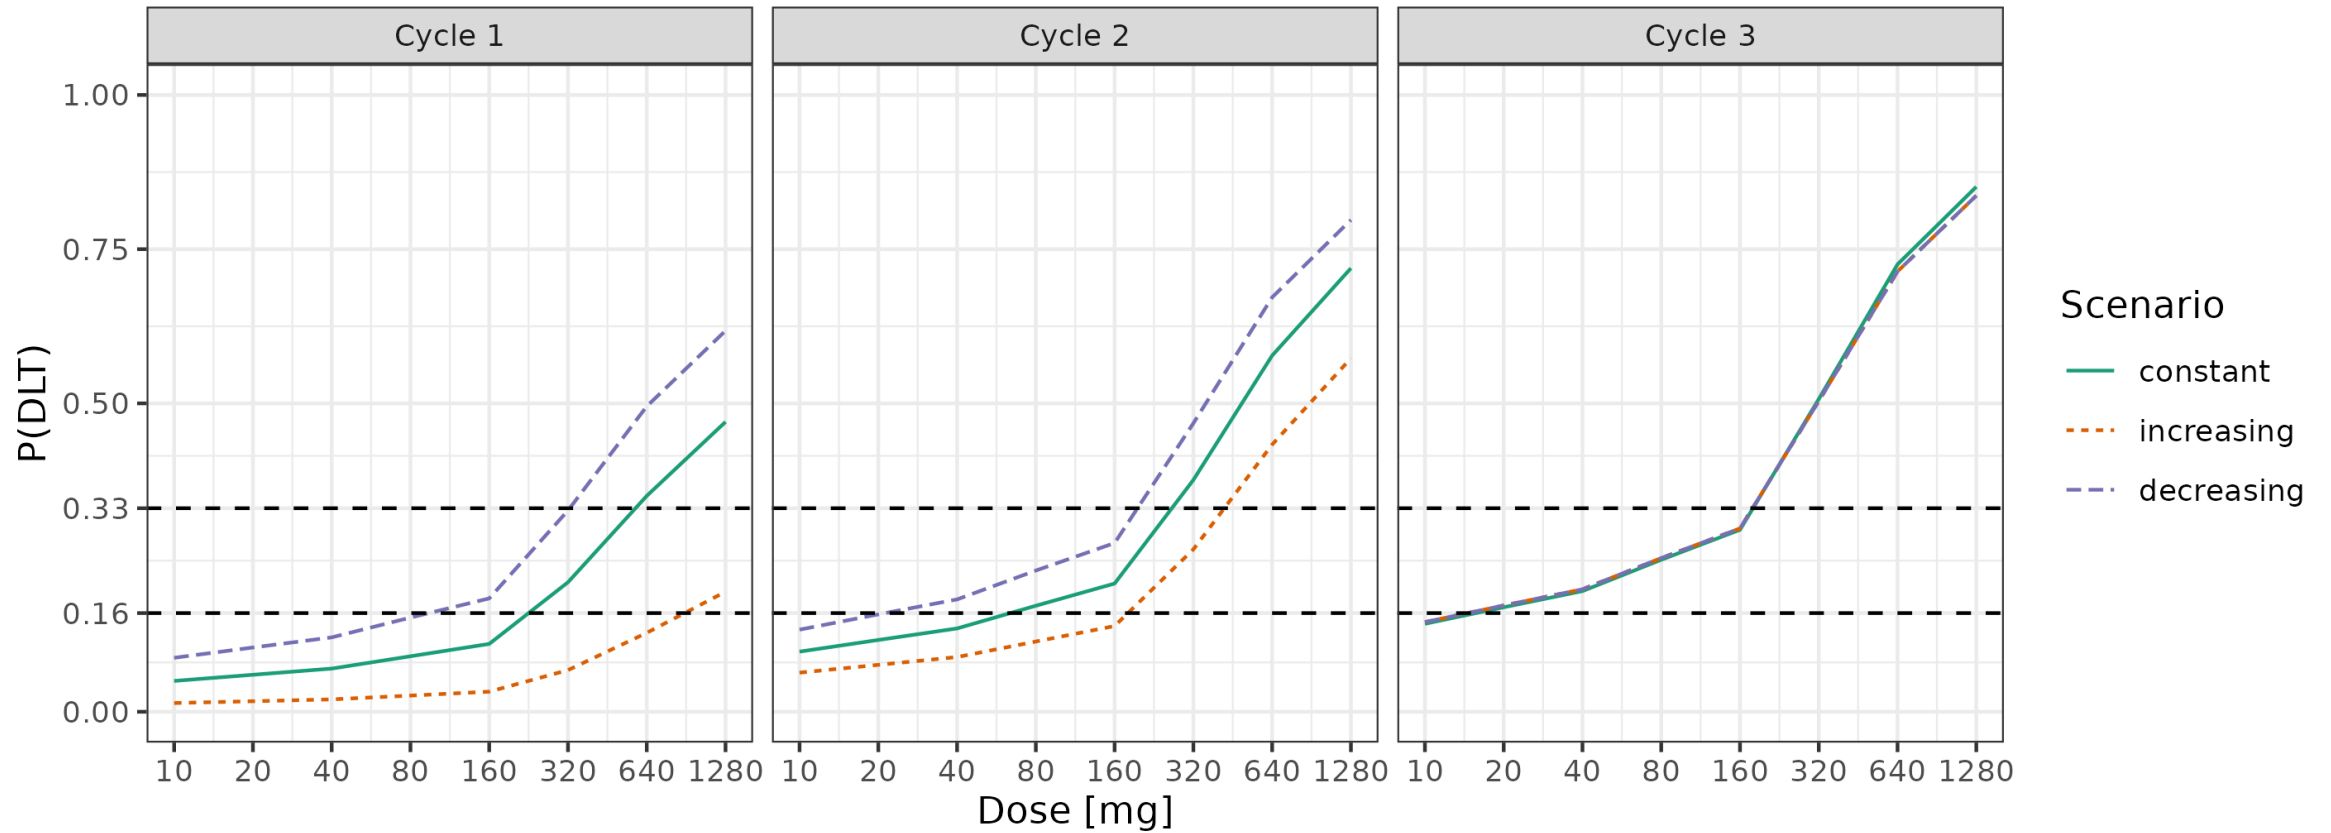NOVARTIS | Reimagining Medicine

Scenarios: constant & in-/decreasing per-cycle toxicity

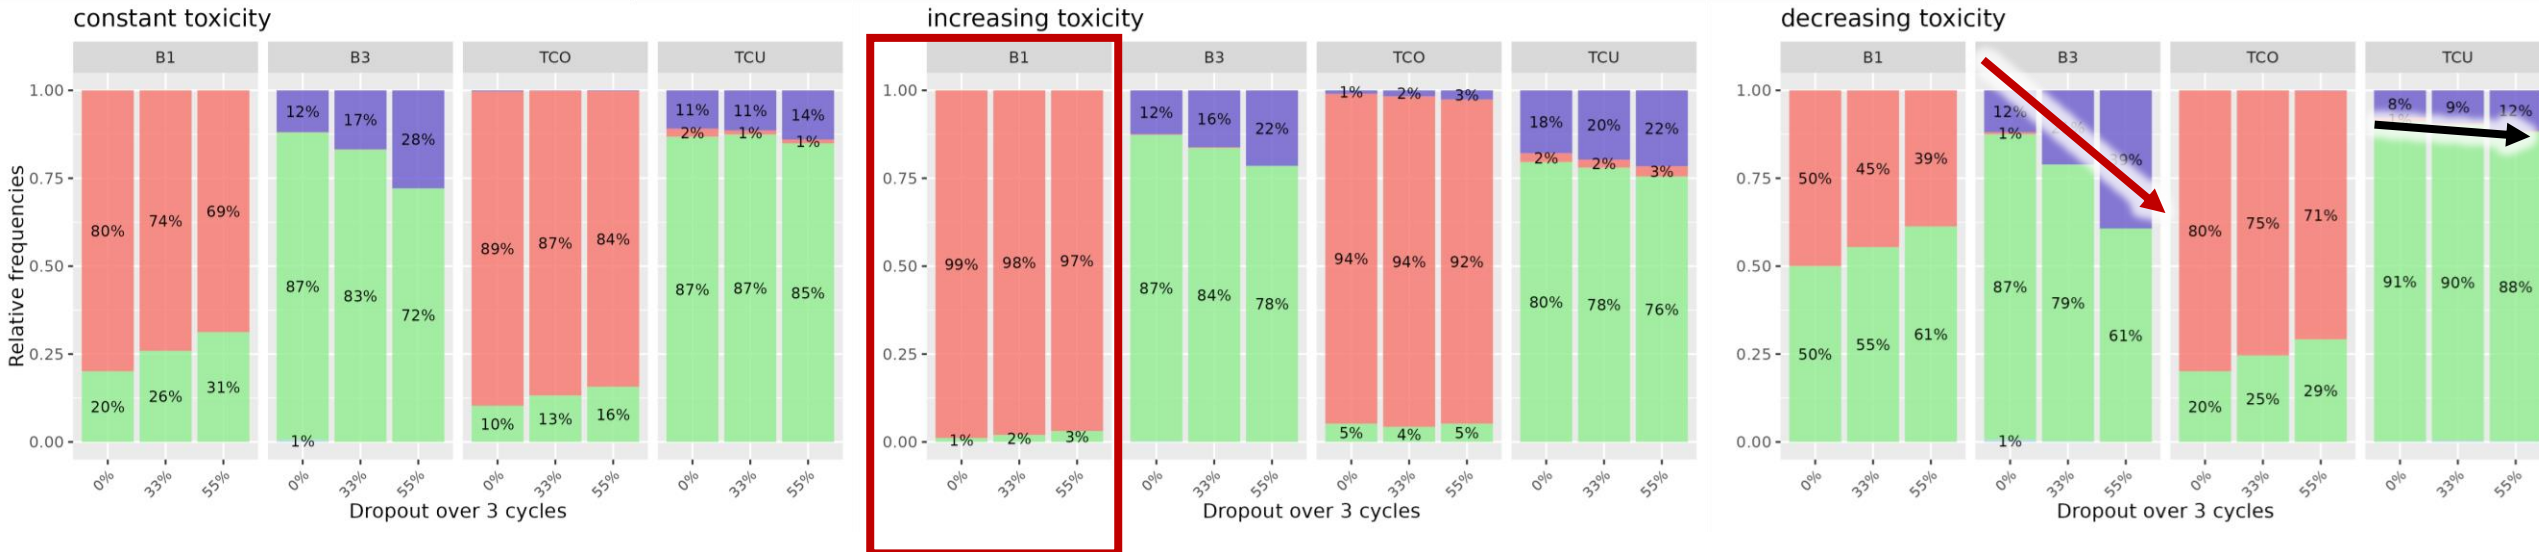Scenario risk for a DLT, per-cycle / conditional
True probability
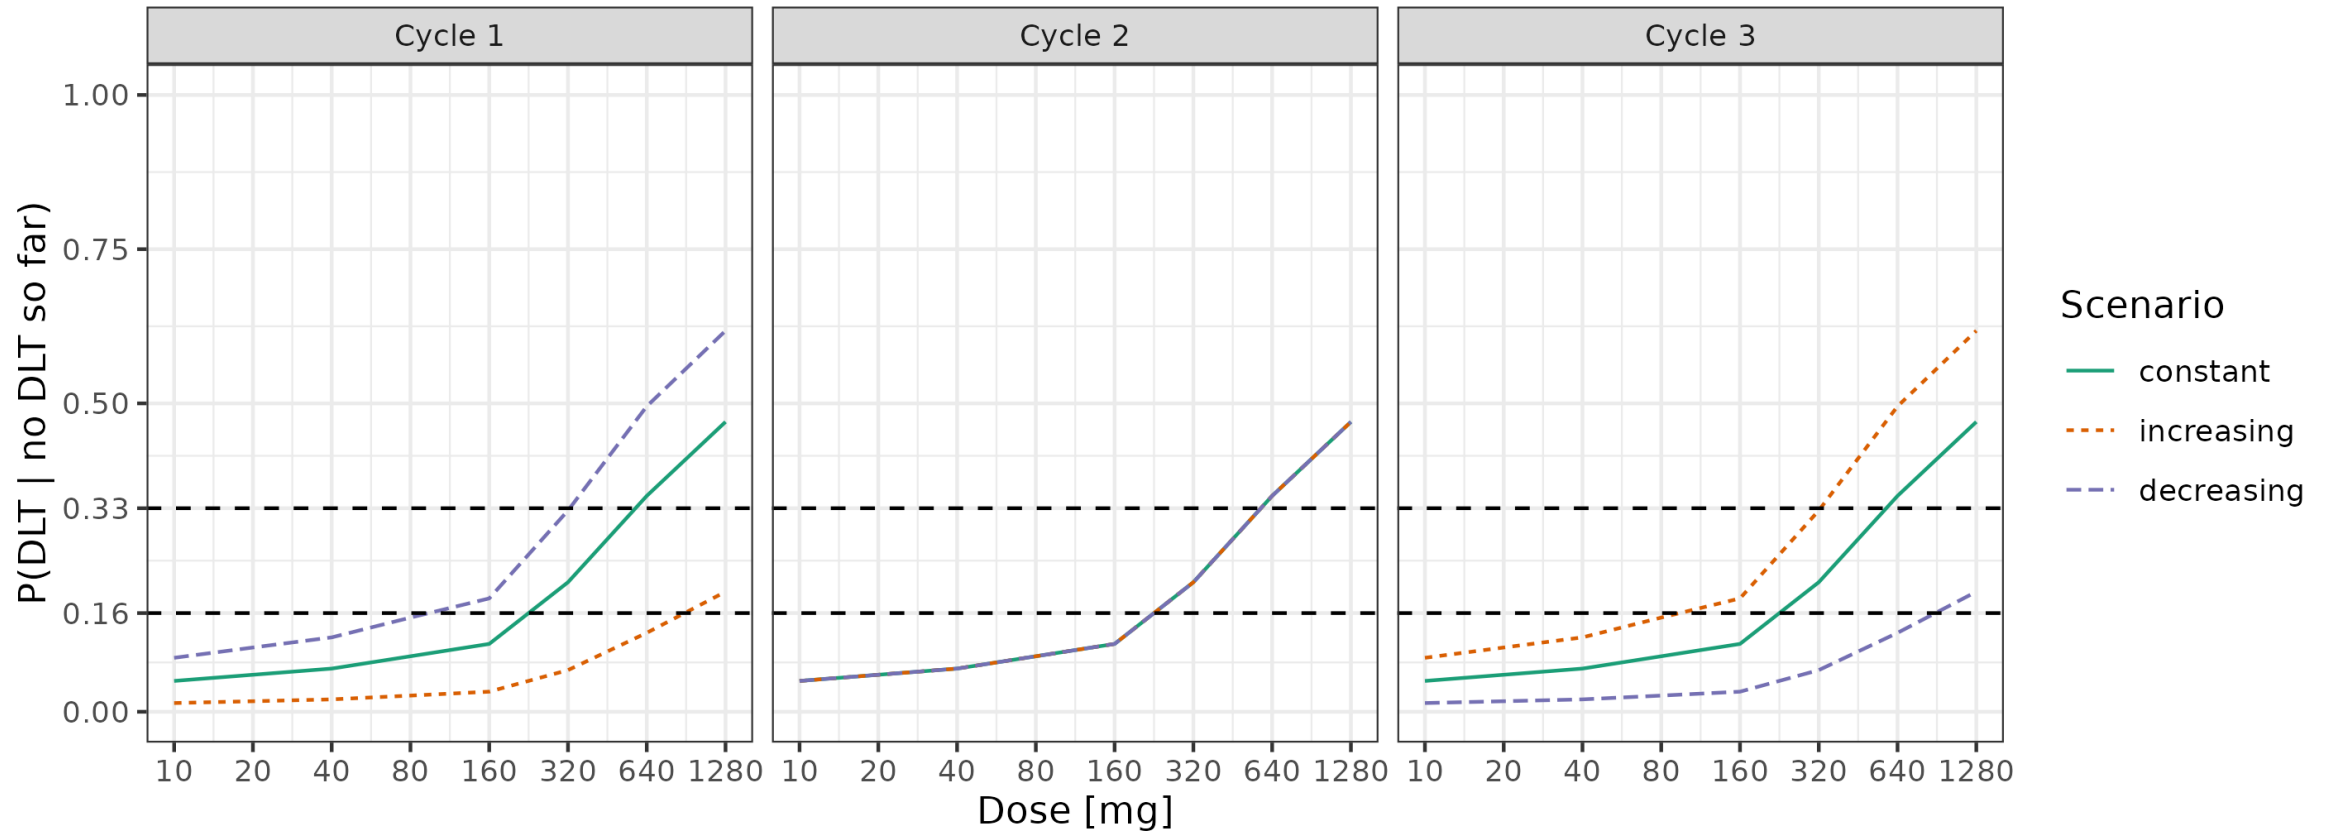
Scenarios: same cumulative toxicity at end of cycle 3

Scenario risk for a DLT, cumulative
True probability

Simulation results: MTD for three-cycle therapy (cumulative toxicity)

Approximately equivalent models:
B1 ⇔ TCO, B3 ⇔ TCU
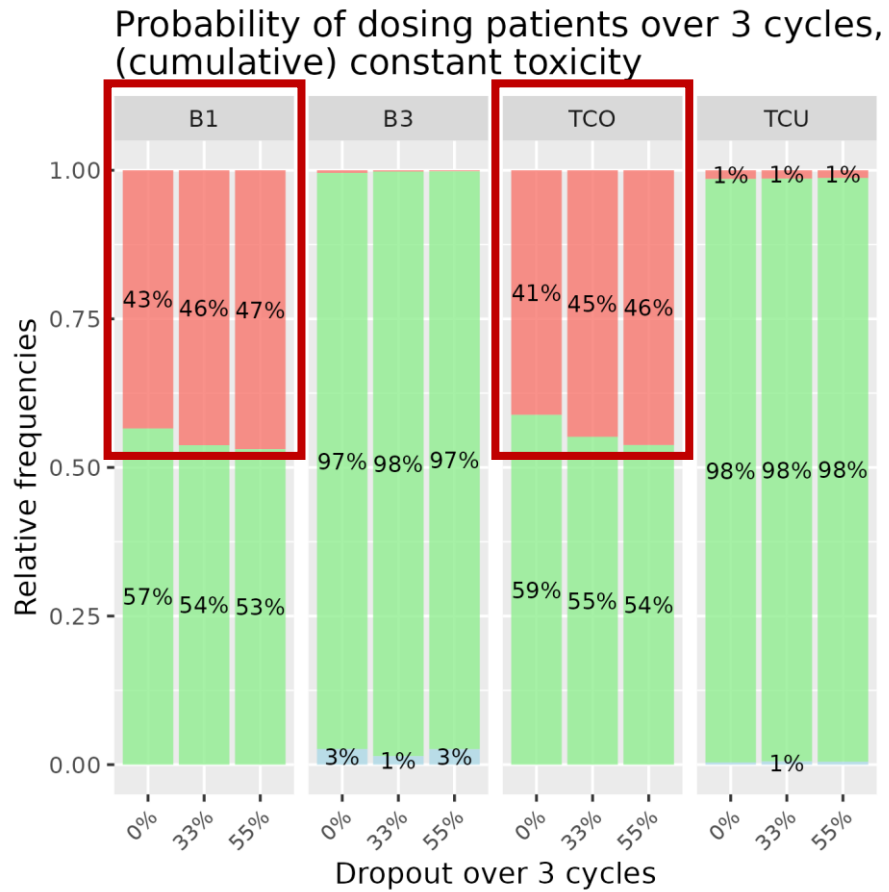

Trial stopping reason:

- 🟪 All doses predicted to be toxic
- 🟥 MTD ∈ overdose
- 🟩 MTD ∈ target dose
- 🟦 MTD ∈ underdose

If the treatment has more late / long-term toxicity, a method that considers only cycle 1 does *very* poorly!

In the face of early toxicity and increasing dropout, binary methods get very conservative. TTE method does not!

NOVARTIS | Reimagining Medicine

Simulation results:
patient allocation and trial length

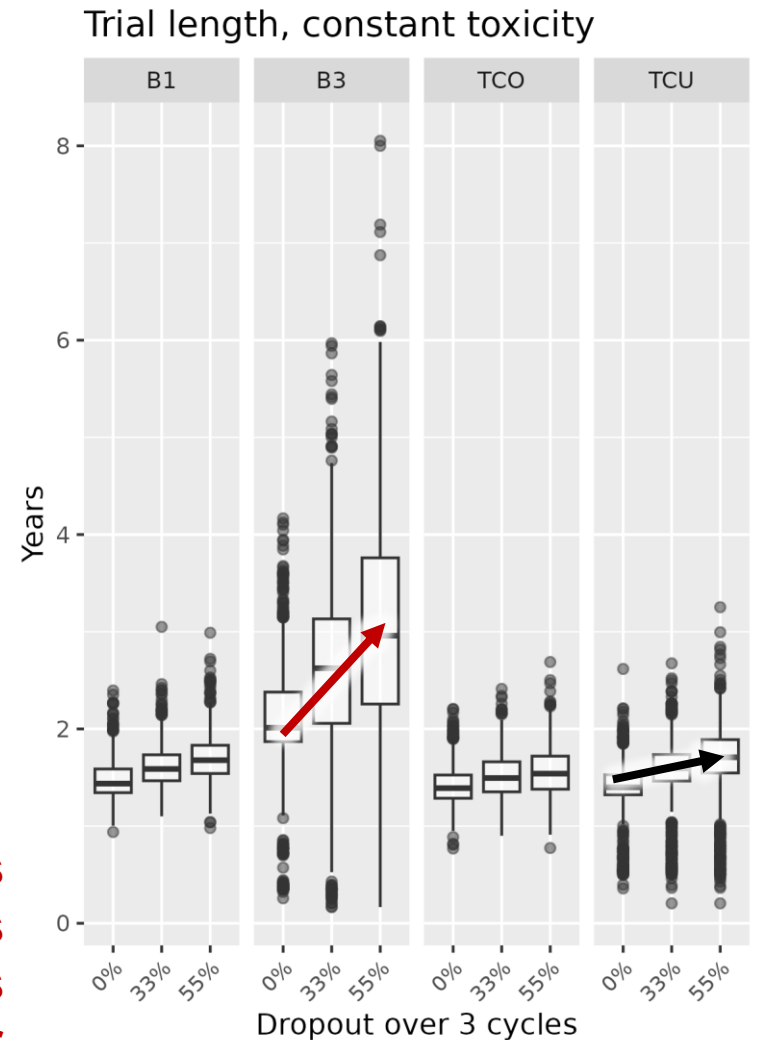Probability of dosing patients over 3 cycles, (cumulative) constant toxicity

Patient allocation:
- Overdose
- Target dose
- Underdose

Cumulative toxicity to patients not well-controlled by per-cycle methods

Waiting for 3 cycles before escalating is slow, dropout makes MTD declaration slower

Trial length, constant toxicity

NOVARTIS | Reimagining Medicine

Time-to-event modeling: summary

Novel treatments in Oncology often require safety monitoring of entire therapies extending beyond cycle 1

Simulation study demonstrates substantial benefits of time-to-event (TTE) over Bayesian logistic regression models (BLRM) 1/3 cycle design at the cost of a more complex model, protocol & data collection:

- **TTE design comparable in execution duration & safer**
- Cumulative TTE & BLRM 3-cycle very similar if no dropout (but escalation with BLRM is slow!)
- **Presence of dropout grows advantages of TTE over BLRM**

Simplified TTE model can seamlessly replace BLRM with minor change of the link function

- Could also consider the timing of dosing, not just the dose alone (time-varying exposure)

Outlook – challenges in Phase I Oncology

For cytotoxic agents, toxicity ≈ efficacy and going for the MTD was "easy" – modern treatments require modelling of safety *and* efficacy for finding the optimal biological dose.

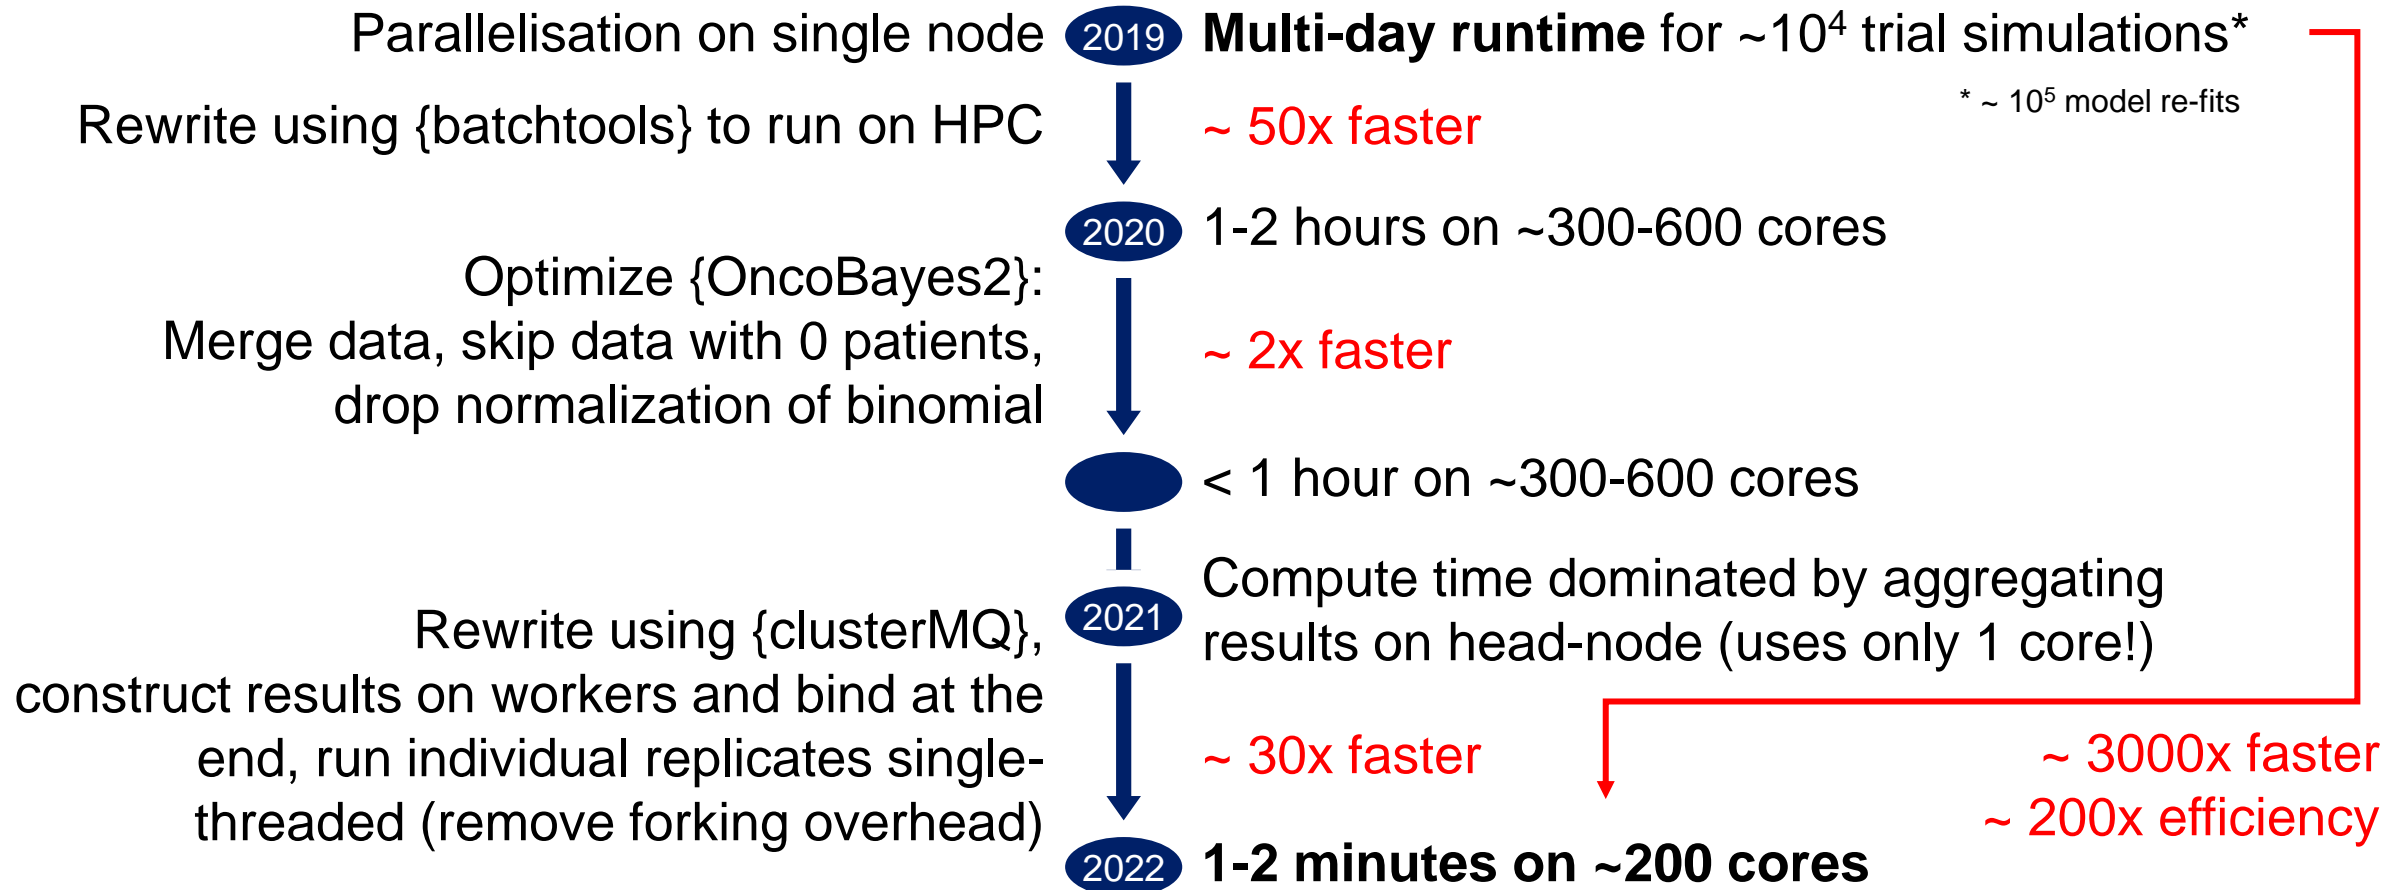**Optimizing the benefit / risk profile early requires an early idea of efficacy**

- We're typically interested in longer-term efficacy (survival) – while there are earlier biomarkers, they may only give quite an uncertain dose-efficacy curve.

→ Longer-term safety of therapy is often important for novel therapies.

Challenges in safety models: dosing interruptions/reductions & different regimens

- What *was* administered (rather than what was *planned* to be administered)?

- Zero dose in cycle 2 ≠ zero DLT probability (if dose in cycle 1 was > 0 – long-term effects).
  ⇒ Exposure metric needed (K-PD / PK-PD modelling with incomplete data).

Addressing the compute requirements in simulation studies – making model-based methods go fast

Parallelisation on single node **2019** **Multi-day runtime** for ~10^4 trial simulations*

* ~ 10^5 model re-fits

Rewrite using {batchtools} to run on HPC — ~ 50x faster

2020 1-2 hours on ~300-600 cores

Optimize {OncoBayes2}:
Merge data, skip data with 0 patients,
drop normalization of binomial — ~ 2x faster

< 1 hour on ~300-600 cores

2021 Compute time dominated by aggregating results on head-node (uses only 1 core!)

Rewrite using {clusterMQ},
construct results on workers and bind at the
end, run individual replicates single-
threaded (remove forking overhead) — ~ 30x faster

~ 3000x faster
~ 200x efficiency

2022 **1-2 minutes on ~200 cores**

U **NOVARTIS** | Reimagining Medicine

The principles and techniques are now a course: go fastR!

The course covers the following learning goals:

- Be able to debug R code and identify & optimize bottlenecks

- Basics of R parallelization on high performance compute environments

- Know how to apply this knowledge on relevant case studies

Find our free open-source course at https://luwidmer.github.io/fastR-website/

- Developed together with Michael Mayer (Posit)

- Same techniques were applied for the presented simulation study

Lukas A. Widmer
lukas_andreas.widmer@novartis.com

Thank you