Bayesian Learning in Oncology: A Case Study

Daniel Sabanés Bové

Basel Biometrics Society Seminar, Basel, 13th April 2015

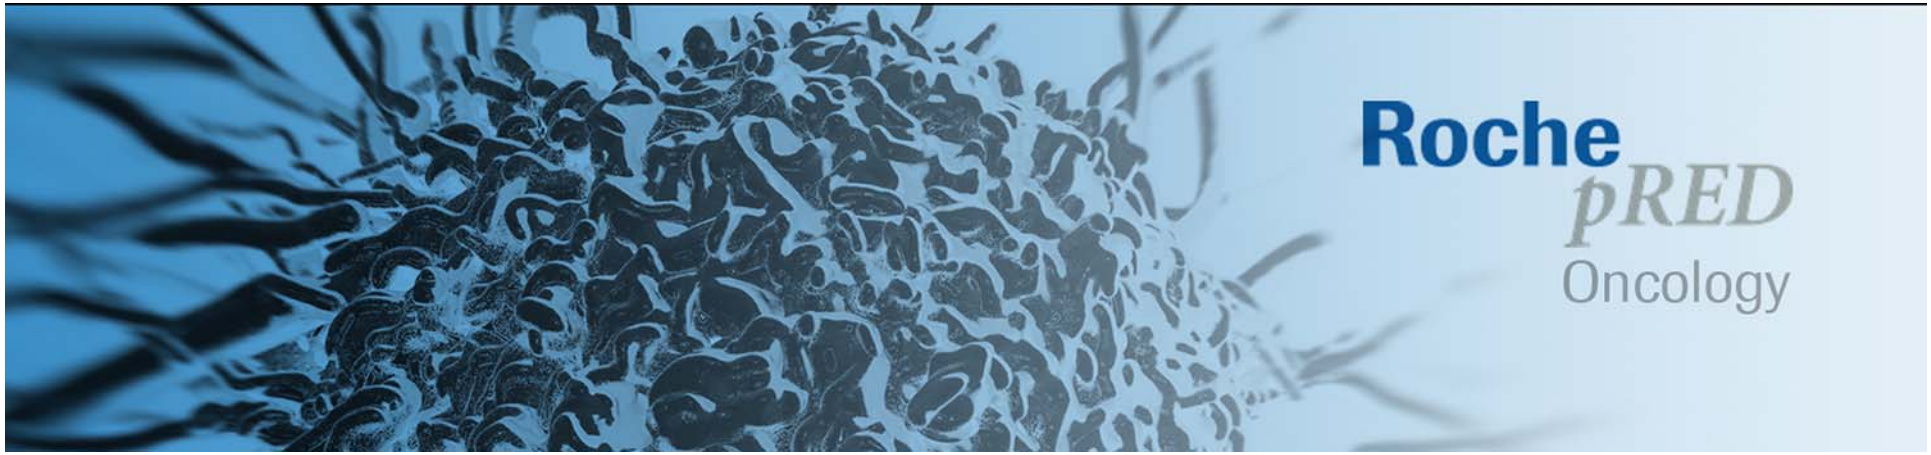

Case study on an NME (“A”) in Oncology
Introduction

- New biologic A for the treatment of solid tumors
- Mainly expected to work in combination with other drugs
- Frontrunner for a platform of similar biologics
- Entry-into-human phase 1 study started
- Combination A + B phase 1b study to start about one year later
- Follower NME C from same platform still to enter phase 1

This talk: How to use Bayesian statistics for answering the typical clinical development questions (in blue) in this project

What is the maximum tolerated dose (MTD)?

Approach used in phase 1 study

- Two-part dose escalation for our NME A
- First part: single patient cohorts in very low dose range
  - Safety rules for closing this part early
  - Rule-based escalation
- Second part: modified Continual Reassessment Method (mCRM) design
  - Dose-limiting toxicities (DLTs) defined in protocol (3 weeks period)
  - Logistic regression for modelling DLT rate in relation to dose
  - Start with biweekly (Q2W) schedule evaluation, afterwards Q3W
- Prior information which was used for mCRM setup:
  - In-vitro assays
  - In-vivo non-human primates study with a homologous molecule

Prior dose-DLT rate model

Logistic regression with bivariate normal prior

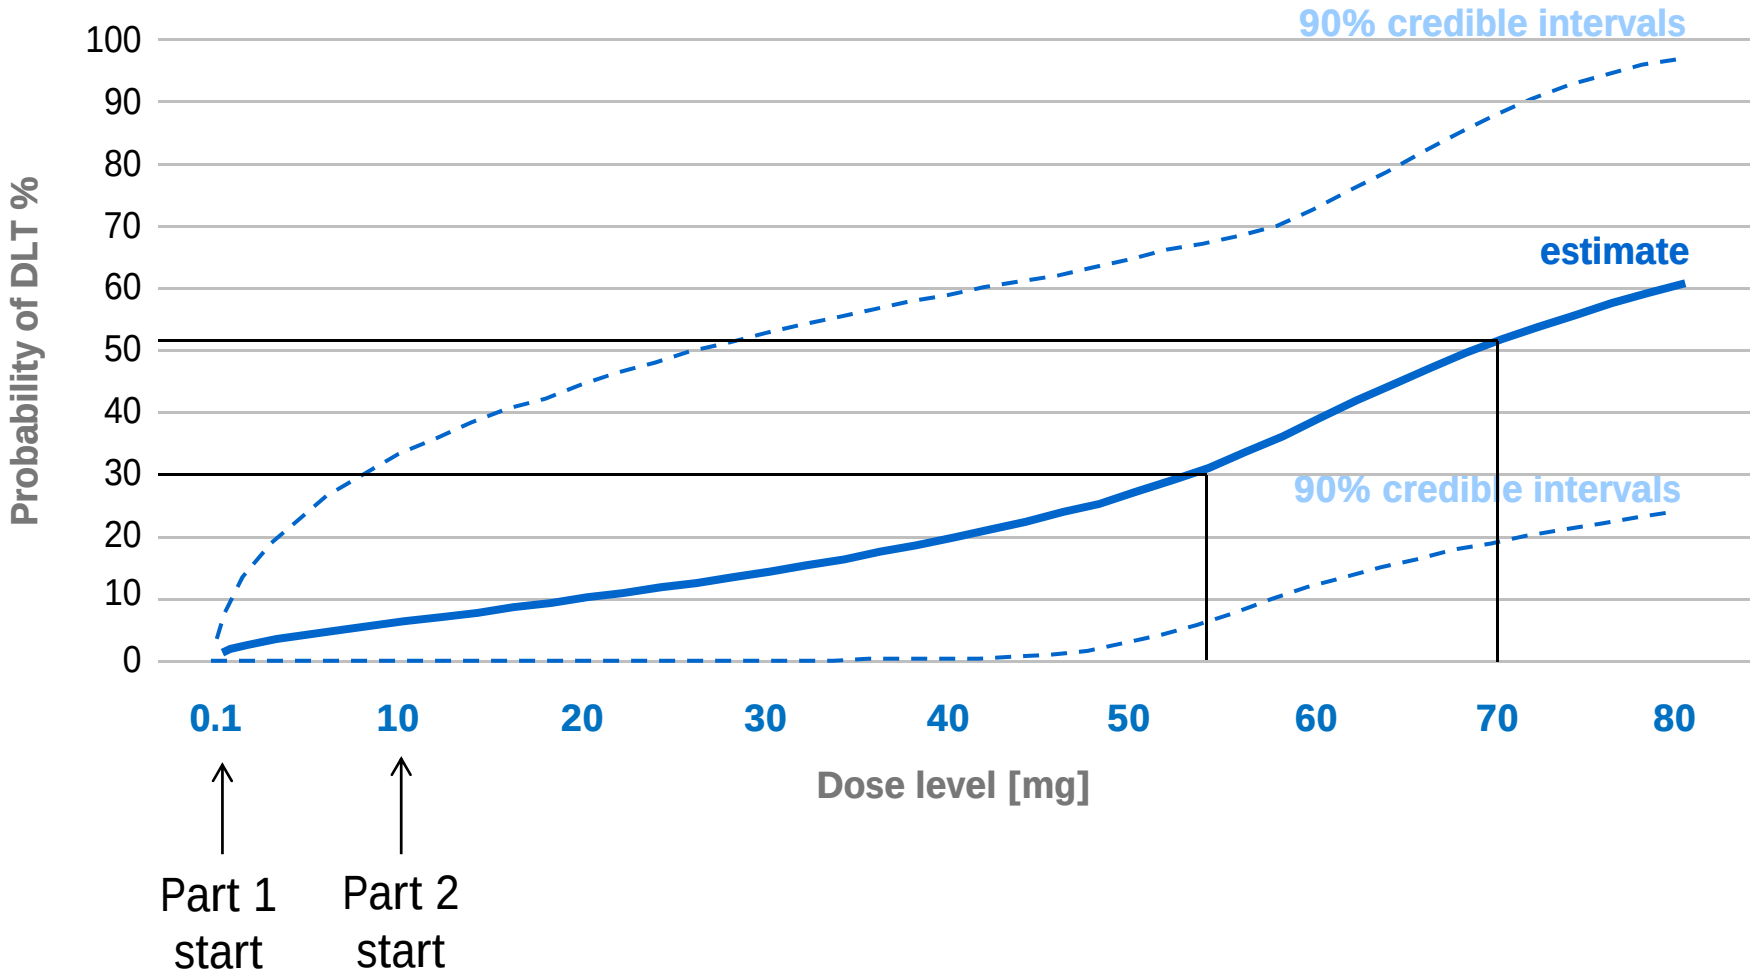

Dose escalation with mCRM
Standard design to find MTD

1. Gather DLT data (binary) for each patient
2. Update the dose-DLT rate model with Bayesian inference (get posterior distribution of parameters with MCMC)
3. Use updated model to estimate the next best dose:
   - **Priority No. 1: overdose control**
     Dose must have < 25% risk that DLT rate is > 35%
   - **Priority No. 2: "target toxicity window"**
     Dose should maximize chance that DLT rate $\in (20\%, 35\%)$
4. Are prespecified stopping criteria met? (relating to precision reached)
   - Then define MTD = next best dose.
   - Otherwise continue to step 5.
5. Treat next cohort of 3 patients at next best or a lower dose (clinical judgement has always final call!)

Practical experience in the case study
Benefits, challenges and opportunities

- High flexibility of the mCRM design:
  - Number of patients per cohort could be increased as needed (e.g. to get additional pharmacodynamic data)
  - Parallel enrollment at different doses possible (e.g. imaging substudy)
  - Later detected DLTs could be incorporated
- Some practical challenges encountered:
  - Clinicians were more conservative due to non-DLT adverse events (AEs)
  - Change from originally planned Q3W to QW schedule
  - Newly proposed escalation in cycle 2 or cycle 3 instead of only in cycle 1
- Opportunities for extensions of mCRM:
  - Ordinal AEs / multiple schedules / stepwise escalation / …

Posterior dose-DLT rate model

Current status after including 58 patients

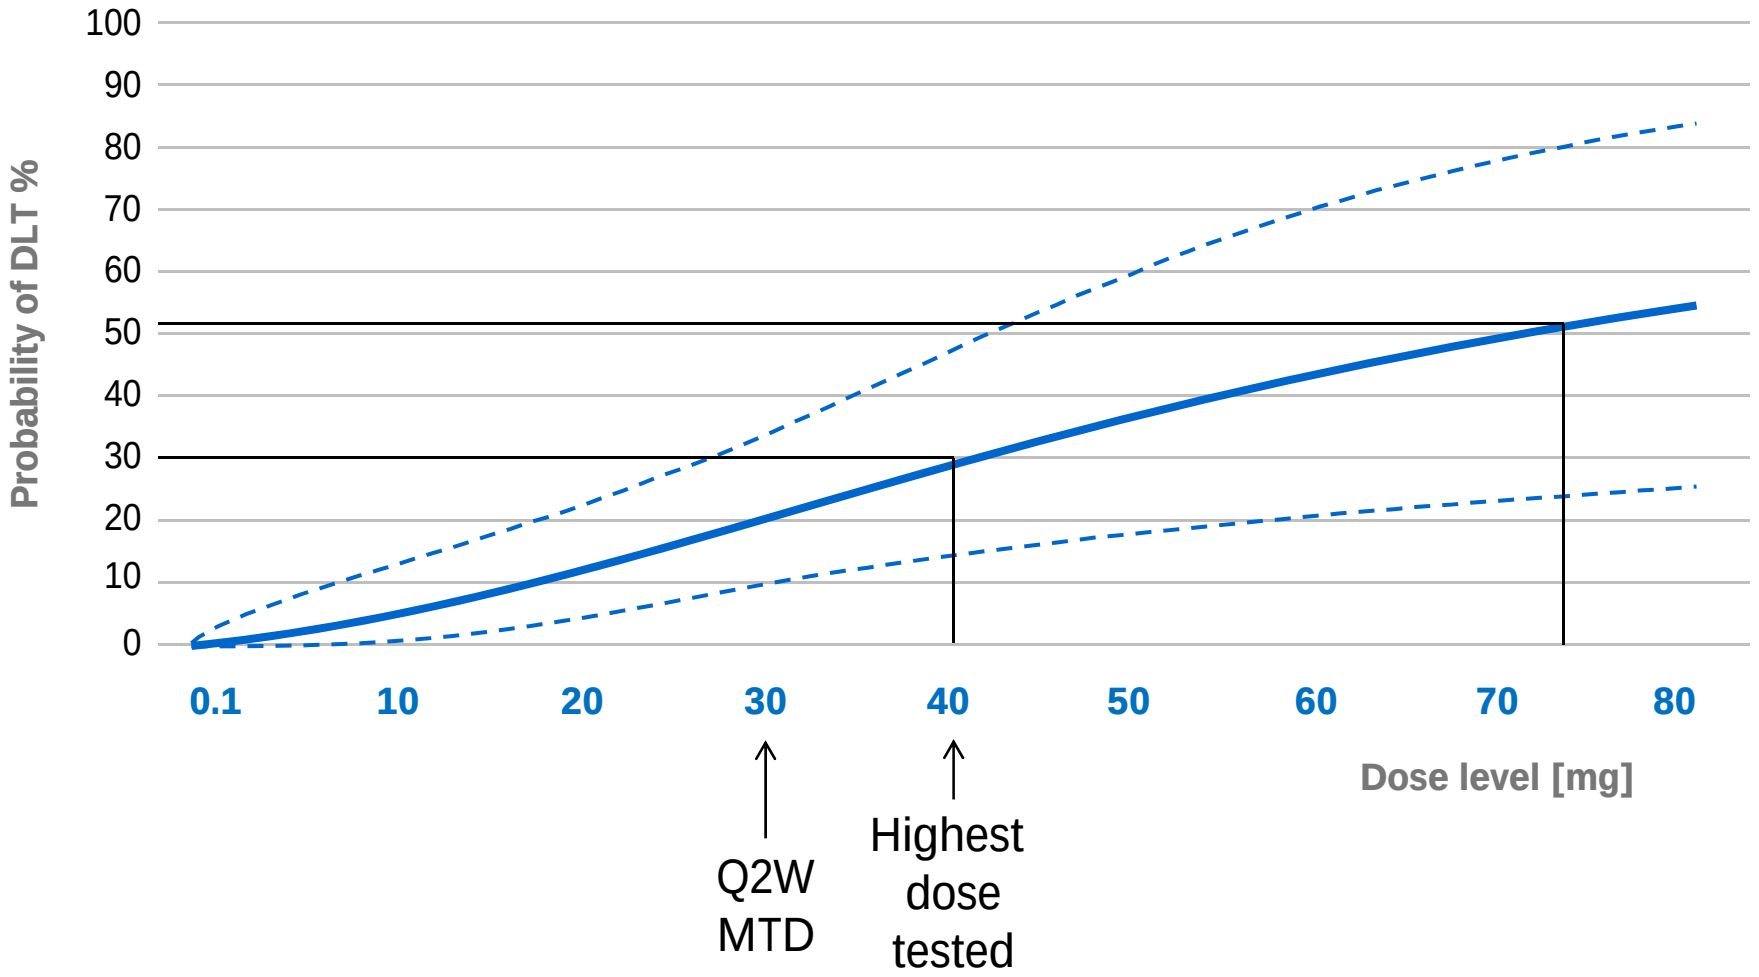

Current target and overdosing probabilities

mCRM design would give higher MTD

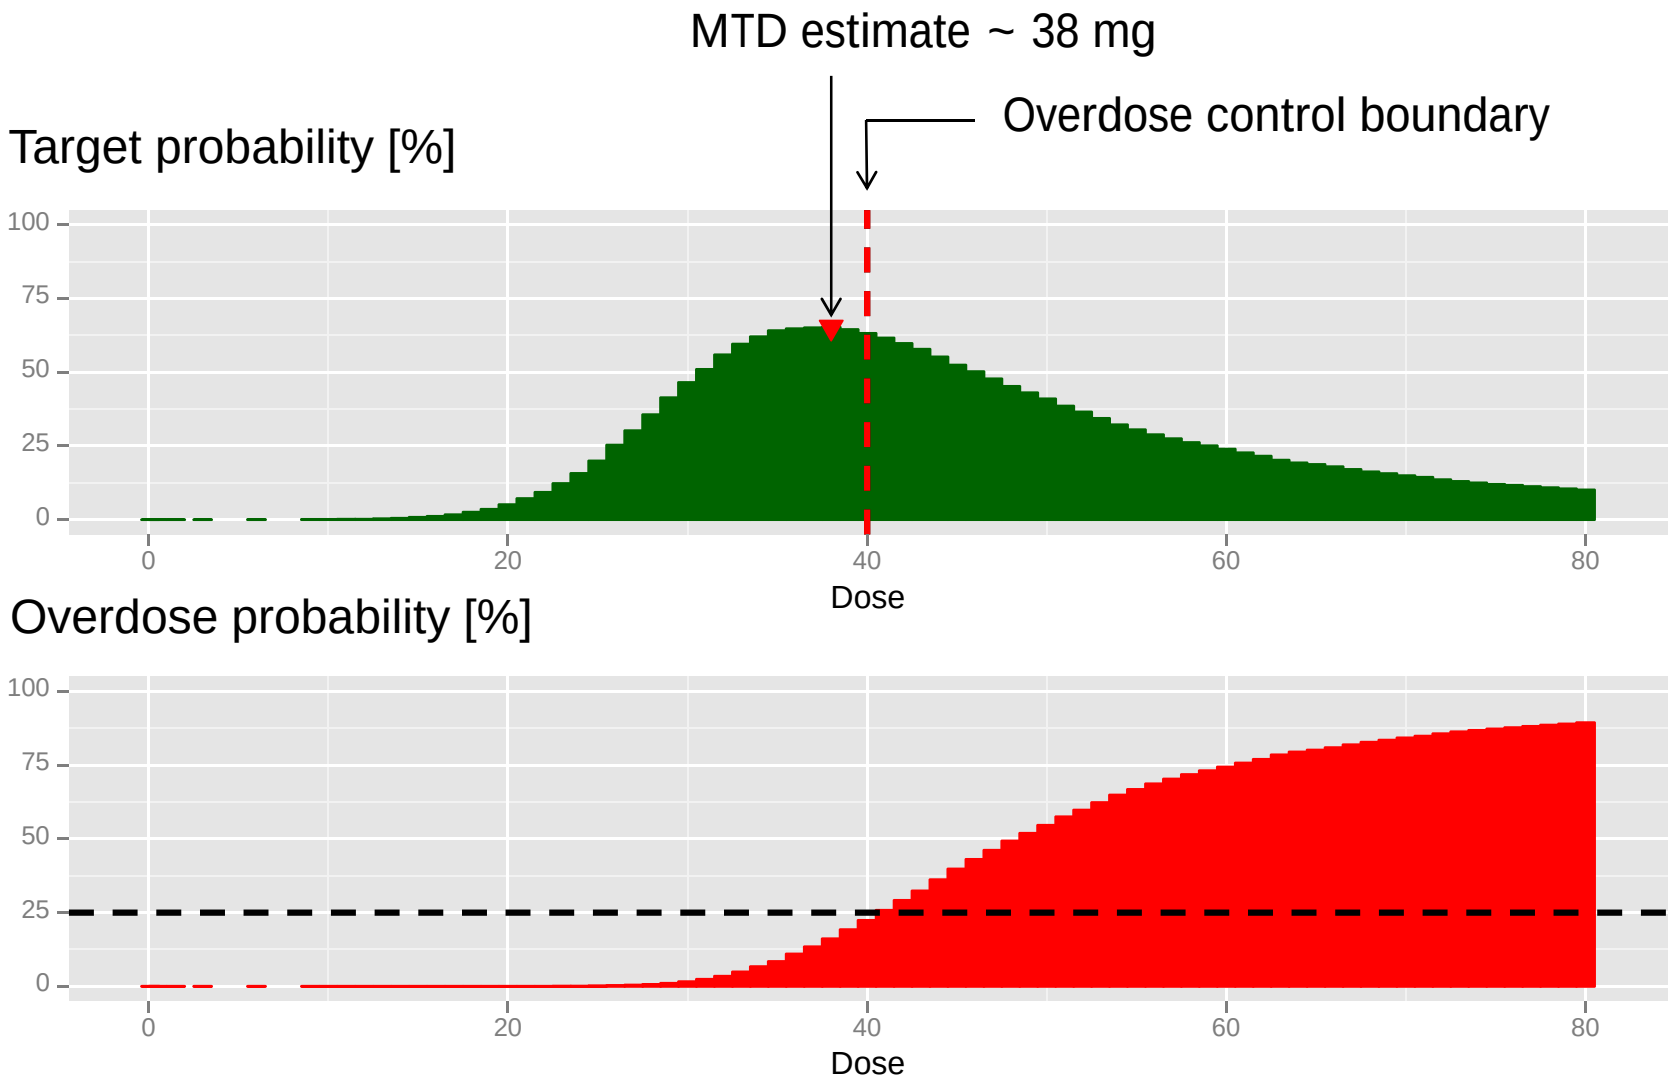

Combination study with NME B
Introduction

- No tumor shrinkage was observed so far in the phase 1 of our NME A.
- As per clinical development plan, proceed to 2NME combination A + B
- Synergistic toxicity cannot be excluded, concomitant administration
- Dose and schedule of NME B are given and should not be changed
- Therefore do single agent dose escalation only on our NME A
- How can we include the prior information from the phase 1 of A and the existing safety data from B?
- How can we decide whether the safety and efficacy of the 2NME combination A + B warrants further development?

How to include prior information from A and B?

Pragmatic approach

- Take the posterior model from phase 1 for NME A at the cutoff date, approximate parameter posterior with bivariate normal
- Add additional toxicity to be expected from the combination with NME B, by shifting the prior mean of the intercept
- Here: expect 5% additional DLT probability at dose of 10 mg
- Some uncertainty about it taken into account,
  by increasing the prior variance of the intercept
- This leads to the prior dose-DLT rate model for the combo A + B
- With target toxicity window of 20-30%, the prior MTD estimate is 28 mg (22 mg with overdose control)

Comparison: posterior for A vs. prior for A + B

At cutoff date included in the protocol

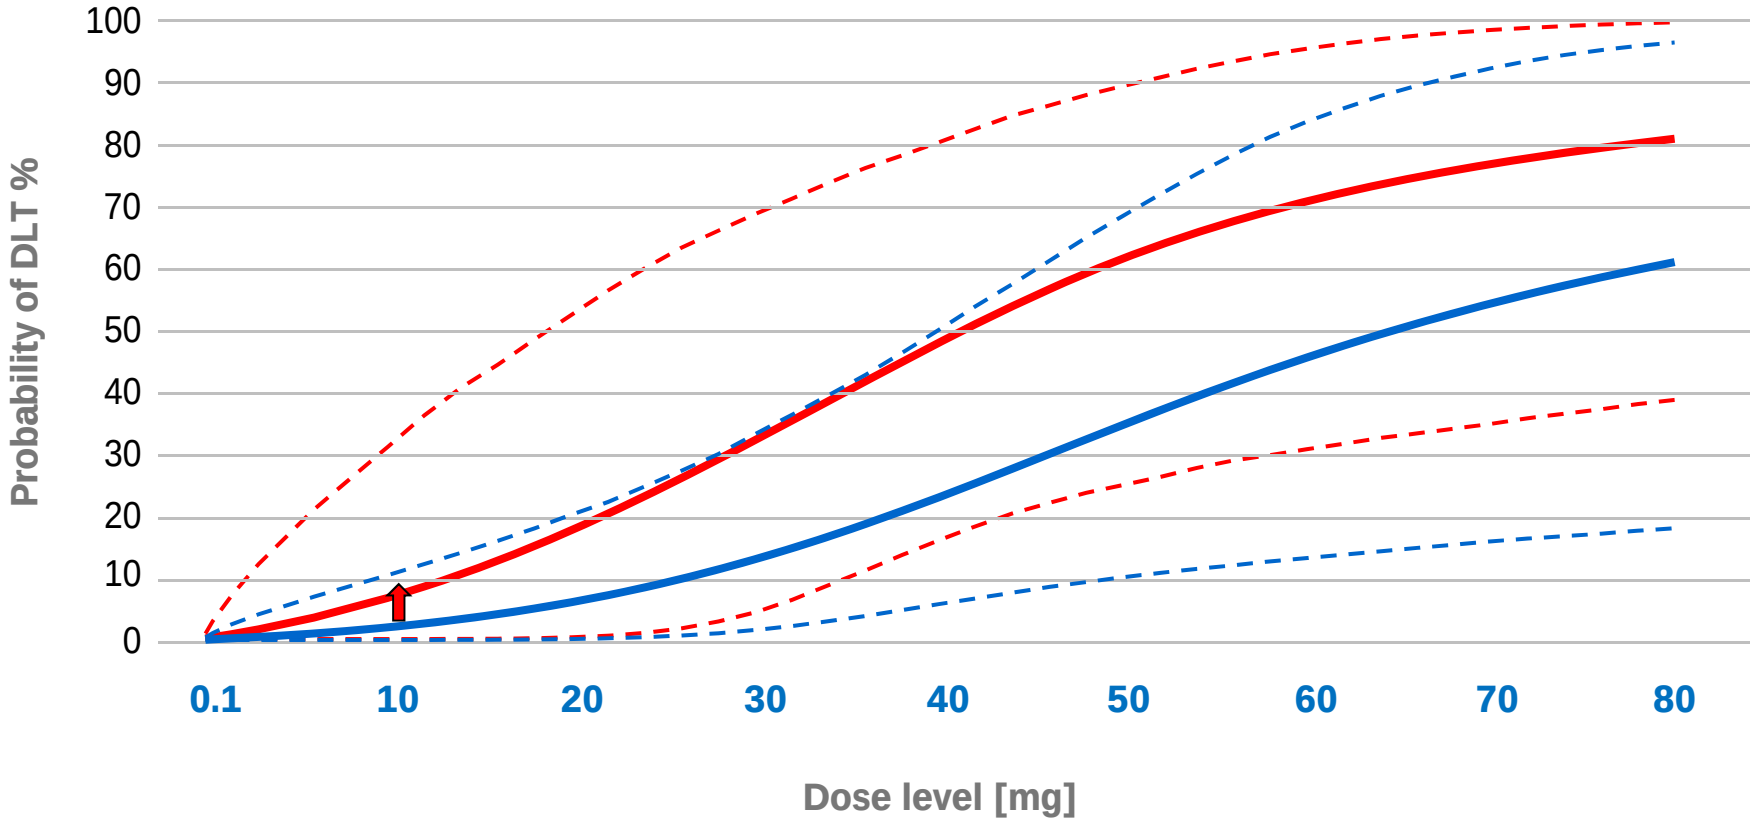

Should we continue developing combo A + B ?

Multiple gates approach

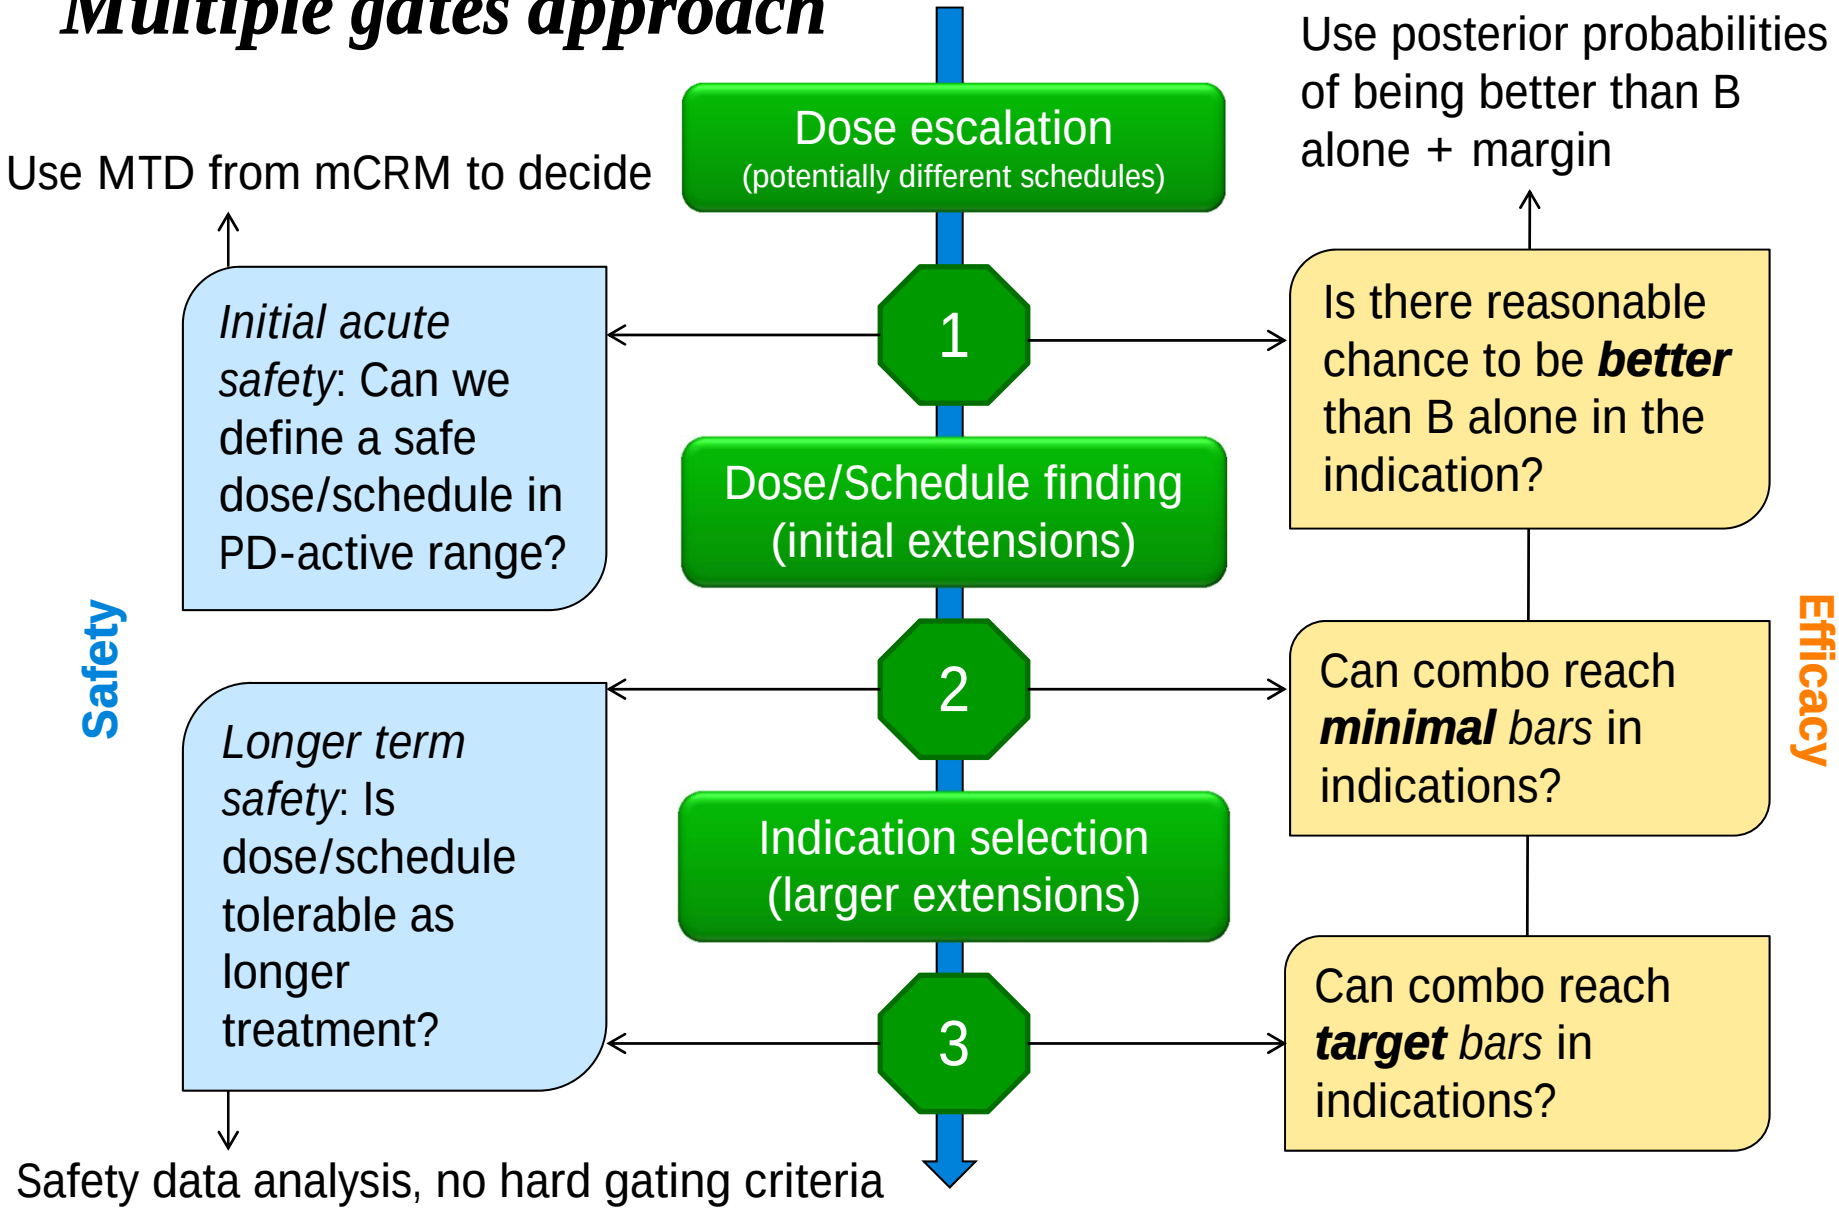

Use posterior probabilities for gating decisions

Example at the first efficacy gate

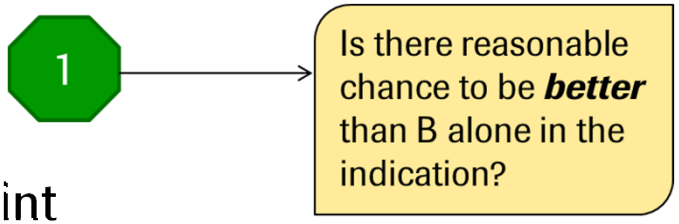

- Use objective response rate (ORR) as endpoint
- Define “reasonable chance” = “at least 30% posterior probability”.
  Use uniform prior, and conjugate beta-binomial model for computations.
- Assume we have n=9 patients in an indication during dose escalation, (across all doses and schedule to avoid «cherry picking»)
- Assume with NME B alone there were 7/43 responses (obs. ORR: 16%) in this indication
- → Possible ORR / probability / decision outcomes for this indication:

Responses:	0	1	2	3	4	5	6	7	8	9
Obs. ORR [%]:	0	11	22	33	44	55	67	78	89	100
Probability to be better than B alone [%]:	17	47	73	89	97	99	100	100	100	100

STOP 30% threshold **GO**

Immunogenicity of A and pretreatment with X
Introduction

- Many of the phase 1 patients developed anti-drug antibodies (ADAs)
- However, no safety impact and most of them no PD impact
- Does pretreatment with drug X reduce the proportion of ADAs ?
- Proof-of-concept study with n=20 patients and binary ADA endpoint
- Patients will be randomized 3:1 to pretreatment, and control arm will be enlarged / have informative prior from the relevant phase 1 patients.
- Compute posterior probability $\mathbb{P}(\pi_1 < (1-\delta)\pi_0 \mid \mathcal{D})$

ADA prop. with pretreatment

Target reduction of ADA prop.

ADA prop. in control

- If > 70% → success; if < 30% → failure; otherwise → inconclusive.
- Look here for large effect $\delta = 50\%$; reasonable power for detecting this.

How to decide if pretreatment is effective?

Illustration of the Bayesian answer with example

Data from n=20 patients

+

ADA results in suitable control patients from EIH study: 12/21 patients ADA+

>70% probability that ADA proportion is at least 50% lower	→	**GO**: Conclude sufficient pretreatment effect
> 30% & < 70% probability	→	**Evaluate**: Consider ADA specificity and biomarkers
<30% probability that ADA proportion is at least 50% lower	→	**STOP**: Conclude insufficient pretreatment effect

May consider further evaluation of pretreatment concept

Overview of potential outcomes with probabilities $\mathbb{P}(\pi_1 < (1-0.5)\pi_0 \mid \mathcal{D})$ and implied decisions

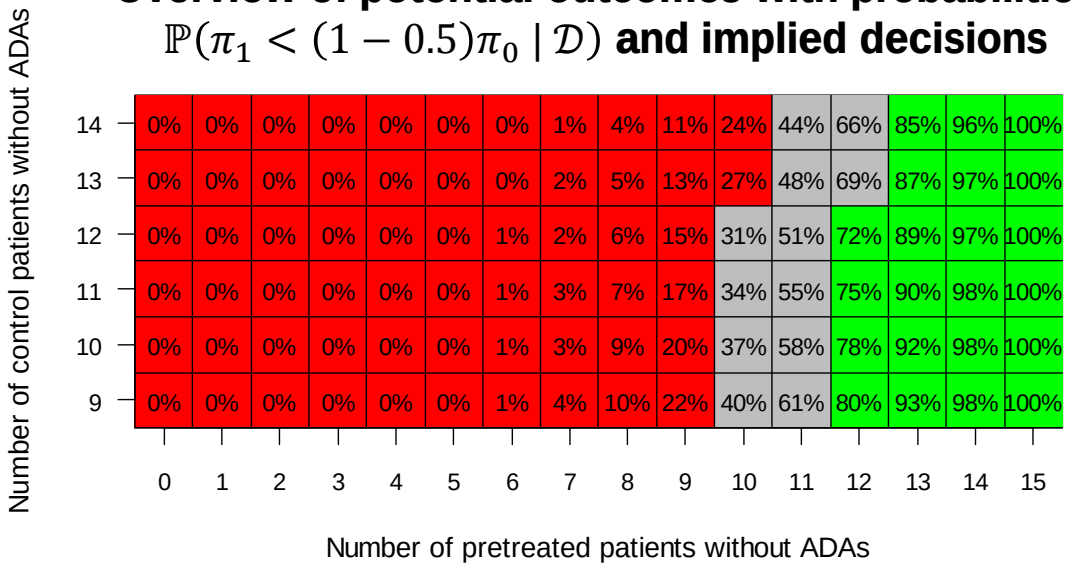

How to use the data for the follower NME C ?

General and statistical aspects

- Follower NME C from the same platform still to enter-into-human.
- Safety and PK profile expected to be very similar to that of our NME A.
- Therefore, we should learn as much as possible from A to design the dose escalation trial for C as efficient as possible, for example general aspects:
  - Starting dose?
  - Schedule?
  - DLT period: could it be shortened? / account for later DLTs? Etc…
- Specifically, for the mCRM design, we would like to use the DLT data from A to specify an informative prior for C.
- Still, the design should be robust against surprises.

Robust mixture prior approach
Overview of the concept

- Take the posterior distribution of the model parameters and approximate it with a bivariate normal → informative component
- Construct a minimally informative prior → «neutral» component
- Mix the two components in a ratio such that the prior sample size is adequate (note: pragmatic approximations here)
- For example: If the informative prior is constructed from 60 patients, and prior sample size should be 10 patients, then choose 1:5 ratio (1/6 weight for informative component)
- The neutral component makes the design robust against surprising deviations from the informative component.
- Still, the dose escalation can be more efficient using the previous data.

Thank you! Questions?

Doing now what patients need next