

Adaptive Bayesian Designs in Dose-finding Studies

VICC Cancer Biostatistics Workshop Series
2008

Leena Choi, Ph.D.

Department of Biostatistics

Vanderbilt University School of Medicine

May 16, 2008

Thomas Bayes – Bayes' Theorem



Diagnostic Testing

- A woman at age 40 had a positive mammography in a routine screening
- What is the probability that she actually has breast cancer?

Diagnostic Testing

- “The probability that a woman with a positive mammography has breast cancer”
- “The probability that a woman with breast cancer has a positive mammography”

Diagnostic Testing

- “The probability that a woman with a positive mammography has breast cancer”
- “The probability that a woman with breast cancer has a positive mammography”
- 3 pieces of information
 - the prior probability that a woman has breast cancer (prevalence)
 - the probability that a woman with breast cancer gets a positive mammography (sensitivity)
 - the probability that a woman without breast cancer gets a positive mammography (1-specificity)

Diagnostic Testing

- **Prevalence:** 1% of women at age forty who participate in routine screening have breast cancer
- **Sensitivity:** 80% of women with breast cancer will get positive mammographies
- **1- specificity:** 10% of women without breast cancer will also get positive mammographies
- **What is the probability that she actually has breast cancer?**

Diagnostic Testing

- 100 out of 10,000 women at age forty who participate in routine screening have breast cancer
- **Before** the mammography screening:
 - Breast Cancer
 - 100 women +
 - 9,900 women -
- **After** the mammography:
 - 100 women +
 - Group A: 80 women *with* breast cancer, and a *positive* mammography.
 - Group B: 20 women *with* breast cancer, and a *negative* mammography.
 - 9,900 women -
 - Group C: 990 women *without* breast cancer, and a *positive* mammography.
 - Group D: 8,910 women *without* breast cancer, and a *negative* mammography.
- $A/(A + C) = 80 / (80 + 990) = 80 / 1070 = 7.5\%$

Bayes' Rule

$P[\text{cancer} | T (+)]$

$P[\text{cancer}] \times P[T (+) | \text{cancer}]$

$$= \frac{P[\text{cancer}] \times P[T (+) | \text{cancer}] + P[\sim \text{cancer}] \times P[T (+) | \sim \text{cancer}]}$$

0.01×0.8

$$= \frac{0.01 \times 0.8 + (1-0.01) \times (1-0.9)}{0.01 \times 0.8 + (1-0.01) \times (1-0.9)} = 7.5\%$$

$0.01 \times 0.8 + (1-0.01) \times (1-0.9)$

Bayes' Rule

$P[\text{cancer} | T (+)]$

$P[\text{cancer}] \times P[T (+) | \text{cancer}]$

=

$P[\text{cancer}] \times P[T (+) | \text{cancer}] + P[\sim \text{cancer}] \times P[T (+) | \sim \text{cancer}]$

- **prior probability** - the original proportion of patients with breast cancer.
- **posterior probability** - the estimated probability that a patient has breast cancer, given that we know she has a positive result on her mammography - is known as the *revised probability*
- **Prior** \times **Data** \rightarrow **Posterior**

Bayes' Rule

$P[\text{cancer} | T (+)]$

$P[\text{cancer}] \times P[T (+) | \text{cancer}]$

=

$P[\text{cancer}] \times P[T (+) | \text{cancer}] + P[\sim \text{cancer}] \times P[T (+) | \sim \text{cancer}]$

- Posterior belief in H (hypothesis, cancer) given the data
 \propto Prior belief in H \times support for H by the data
 - The probability of H before the data (prior)
 - The likelihood of the data given H (likelihood)
 - The probability of H given the data (posterior)
- Posterior is a synthesis of Prior and Data using Bayes' rule

Prior, Likelihood and Posterior

- **Prior distribution**: represents the prior information associated with parameter
 - Information available in literature, clinical data base, expert opinion, or any other appropriate source
 - Non-informative representing very little or no relevant information
- **Data distribution (Likelihood)**
- **Posterior distribution**: updating the prior with the likelihood, formally done via Bayes' theorem
 - Prior is merged with the likelihood to give a final posterior

Example: Placenta Previa


- The proportion of female births in the general population = 0.485 (female births < male births)
- 437 out of total 980 births in *placenta previa* ($437/980=0.446$) were females in German

The sex of *placenta previa* births:

Female	437 (0.446)
--------	-------------

Male	543 (0.554)
------	-------------

Total	980 (1.0)
-------	-----------



Normal Placenta



Marginal Placenta Previa




Complete Placenta Previa


Hypothesis:

the proportion of *placenta previa* female births is less than 0.485


Prior

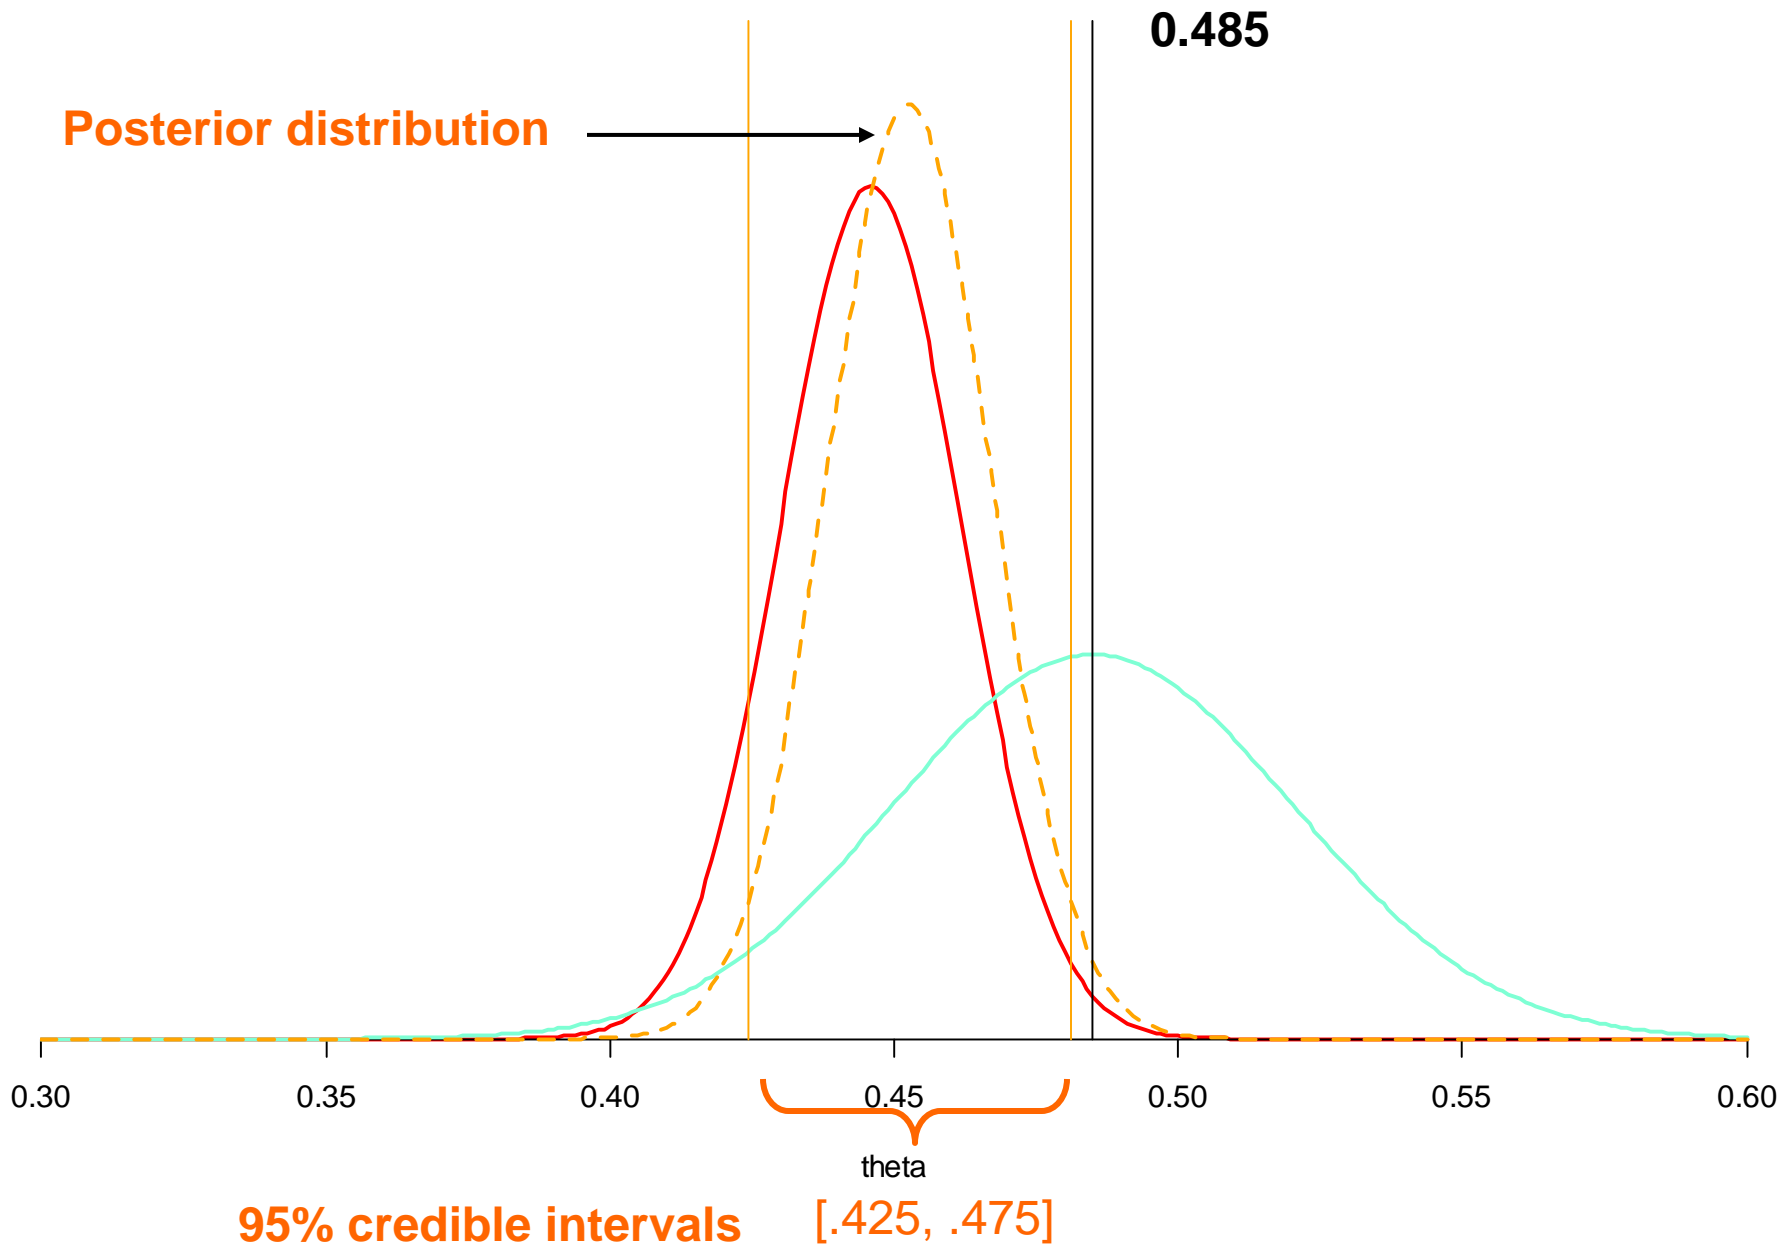


Likelihood



Posterior





Why Bayesian Adaptive Designs?

- Update beliefs in light of new evidence
- Dose finding (dose dropping)
- Stopping early, or late
 - Efficacy
 - Futility
- Seamless phases
- Add arms or drop arms
- Adaptive randomization
- Advantages
 - smaller trials (usually)
 - More precise conclusions, earlier decision
 - Increase the probabilities of success (faster, better drug development, reduce cost)
 - Better treatment of patients (right drug to right patient at right time)

Challenges in Bayesian Approach

- Need appropriate priors – define what we know based on external evidence (conclusions may depend priors)
- Define complexity of our models – what likelihood to use
- Require good knowledge of probability theory
- Computational challenge

Dose Finding Studies

Garrett-Mayer E. *Clin Trials* 2006

- Oncology compounds are cytotoxic: the rationale in cancer dose-finding trials: to find the highest dose that is also safe for use in a Phase II trial
- Maximum Tolerated Dose (MTD): “optimal” dose, “target”, relatively high dose with manageable side effects
- Dose-limiting toxicities (DLTs): 1 if toxicity occurs, 0 otherwise
- Standard “3+3” dose escalation design starts at dose k with fixed number of ordered dose levels

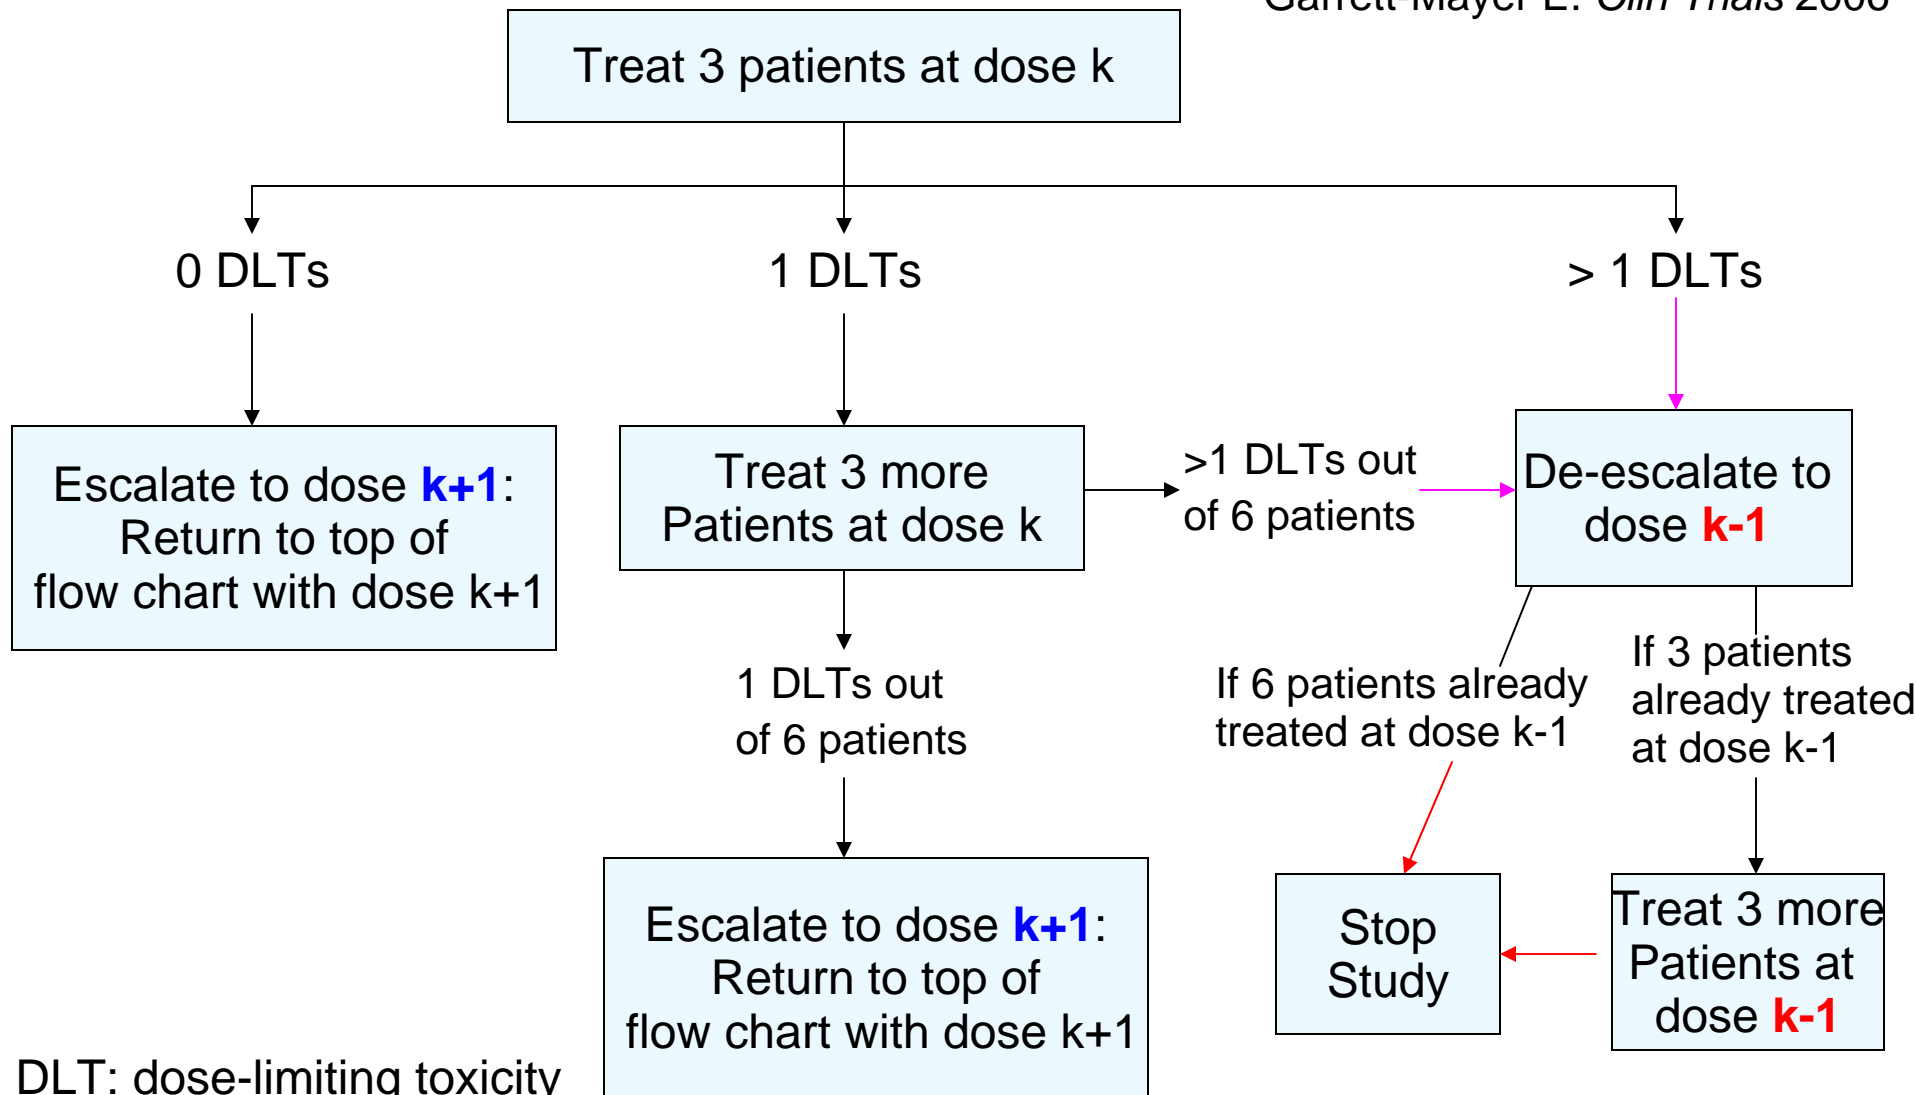
Level	$k-1$	k	$k+1$	$k+2$	$k+3$	$k+4$
Dose	100 mg	200 mg	400 mg	600 mg	800 mg	1000 mg

- MTD defined as the highest dose at which 0 or 1 DLTs are observed in 6 patients
- If de-escalation occurs at the first dose level, the study is discontinued

Dose Finding Studies:

Traditional “3 + 3” dose escalation designs

Garrett-Mayer E. *Clin Trials* 2006




Bayesian Dose-finding Method

Continual Reassessment Method

- Continual Reassessment method (CRM) developed by O'Quigley et al. (1990): a Bayesian Phase I design to estimate MTD
- Assumption: probabilities of both efficacy and toxicity increase with increasing dose
- Use dose-toxicity relationship (dose-efficacy relationship for non-cytotoxic): have rough idea at least
 - Starts with ***a priori*** dose response curve: parameters chosen based on investigators' prior belief
 - As data accumulate, the curve completely determined by the data, little like the *a priori* curve

a priori dose-toxicity



Bayesian Dose-finding Method

Continual Reassessment Method

- Advantages:
 - Superior to traditional dose-escalation designs because it “learns” from information gained at early time points
 - Less likely to treat patients at toxic doses, more likely to treat patients at efficacious doses – more ethical
 - Shown that CRM-based designs to be more efficient and safer (MTD is more precise, fewer cases of DLT – will show comparison)
- Criticism:
 - *a priori* curve could be dangerous due to large uncertainty
 - Large dose escalation could occur based on little information
 - Practical issue: duration of time for completion of study (original CRM evaluate every patient)
- Modified CRMs

a priori dose-toxicity




Modified CRM

Fraries, Goodman *et al.*, Möller

- Pre-defined dose levels for escalation as if for a “3+3” design
- Always start at the lowest dose level under consideration
- Any given dose escalation cannot increase by more than one level, although dose de-escalation can be large
- Enrol two or three patients at each prescribed cohort (not one)
- Proceed as a standard dose escalation design in the absence of dose-limiting toxicities

a priori dose-toxicity




Bayesian Dose-finding Method

Continual Reassessment Method

- Choice of dose level: practical consideration of preparation and packaging
- Target rate of toxicity (or response)
 - Chemotherapy given short time period: serious side effects, often set 0.2 – 0.3 (efficacy 0.8)
 - Choice not from statistician or single investigator, consultations with colleagues – dramatic difference, beneficial
- Choose *a priori* dose-toxicity curve
 - Visual display: clear understanding of the implications of various choices
 - One-parameter or two parameter models: hyperbolic tangent, one- or two-parameter logistic models
 - CRM robust in choosing MTD even if the model is incorrect
 - Good choice of model increase efficiency (smaller size of trial)
 - Informative: Pass through doses for high DLT (eg, 90%) and low DLT (eg, 5%)

a priori dose-toxicity




A priori dose-toxicity curve

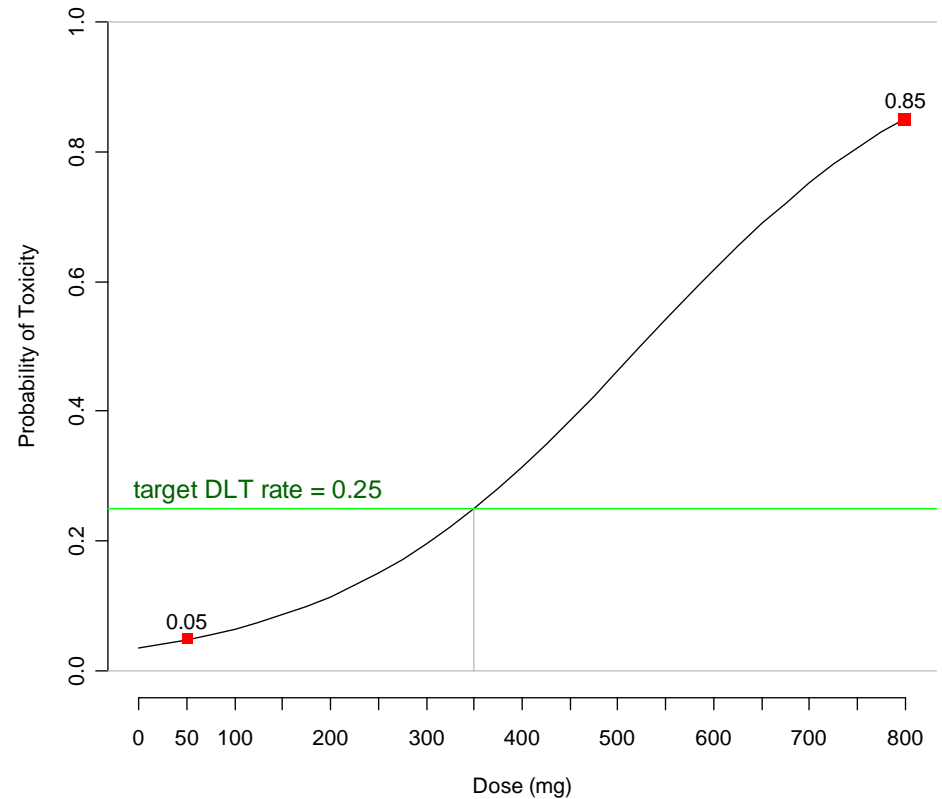


A priori dose-toxicity curve

non-informative



informative



Bayesian Dose-finding Method

Continual Reassessment Method

- Number of patients per dose level (cohort size)
 - 2-3 /cohort
 - Logistical issues: how long, accrual rate, total number of patients available, maximum number of patients, number of dose levels, seriousness of DLT, etc.
- Stopping and sample size: fixed number of patients (continue until total sample size reached for continuous dose levels, infusion), fixed number of patients/dose (6-8 treated at the MTD for discrete dose)


Example: *A priori* curve dose-toxicity with until first DLT observed

Garrett-Mayer E. *Clin Trials* 2006

Cohort	Dose	Outcomes (0 if no DLT, 1 if DLT)	
--------	------	----------------------------------	--

1	50 mg	0	0
2	100 mg	0	0
3	200 mg	0	0
4	350 mg	0	1
	400 mg		
	450 mg		
5	500 mg		
	550 mg		
	800 mg		

- Cohort size = 2
- Dose levels: 50 – 800 by 50 mg
- DLT rate = 0.25
- *A priori* curve: two parameter logistic model
- Stopping rule: 10 patients treated at the same dose


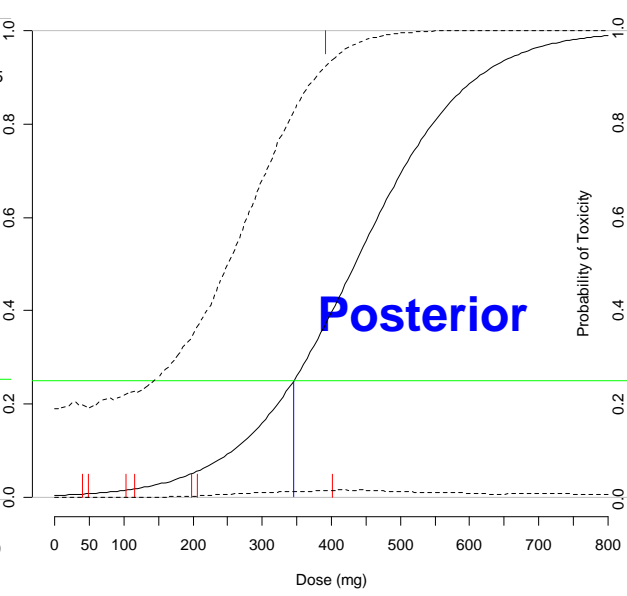
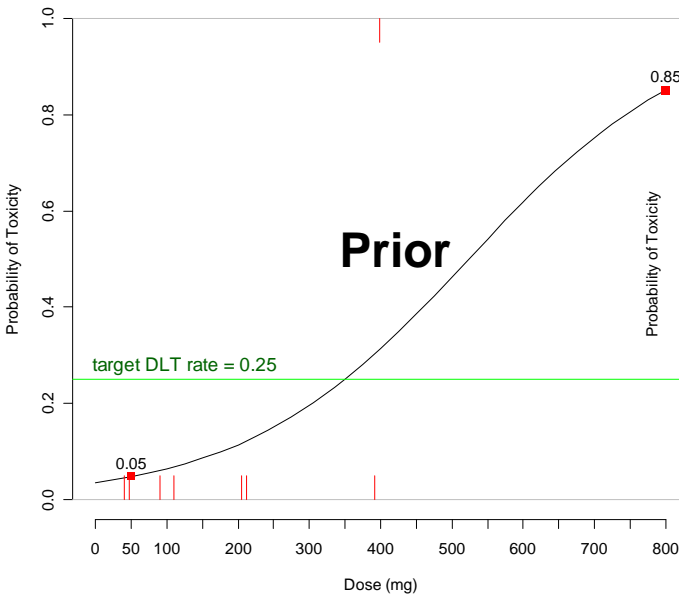


Data

CRM

Cohort 4




Cohort 5




Cohort 6

Cohort 7


Cohort 8




Cohort 9




Cohort 10




Cohort 11



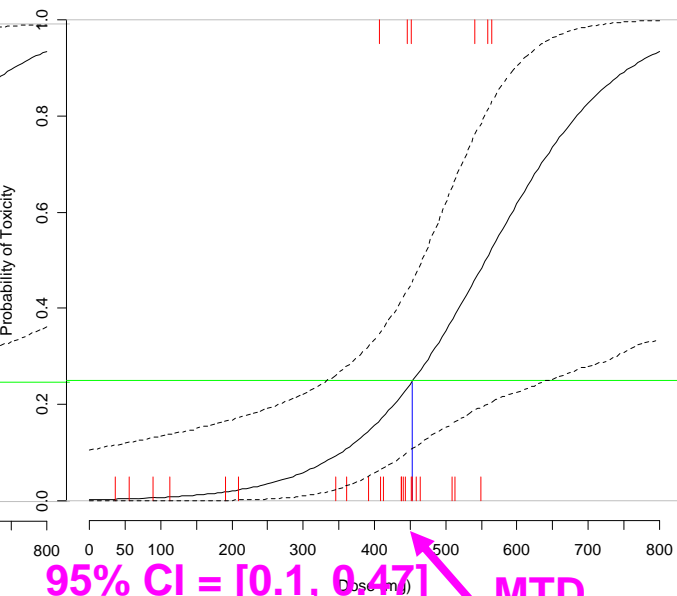
Cohort 12



Cohort 13



Cohort 14



95% CI = [0.1, 0.47] MTD

Accrual of Data

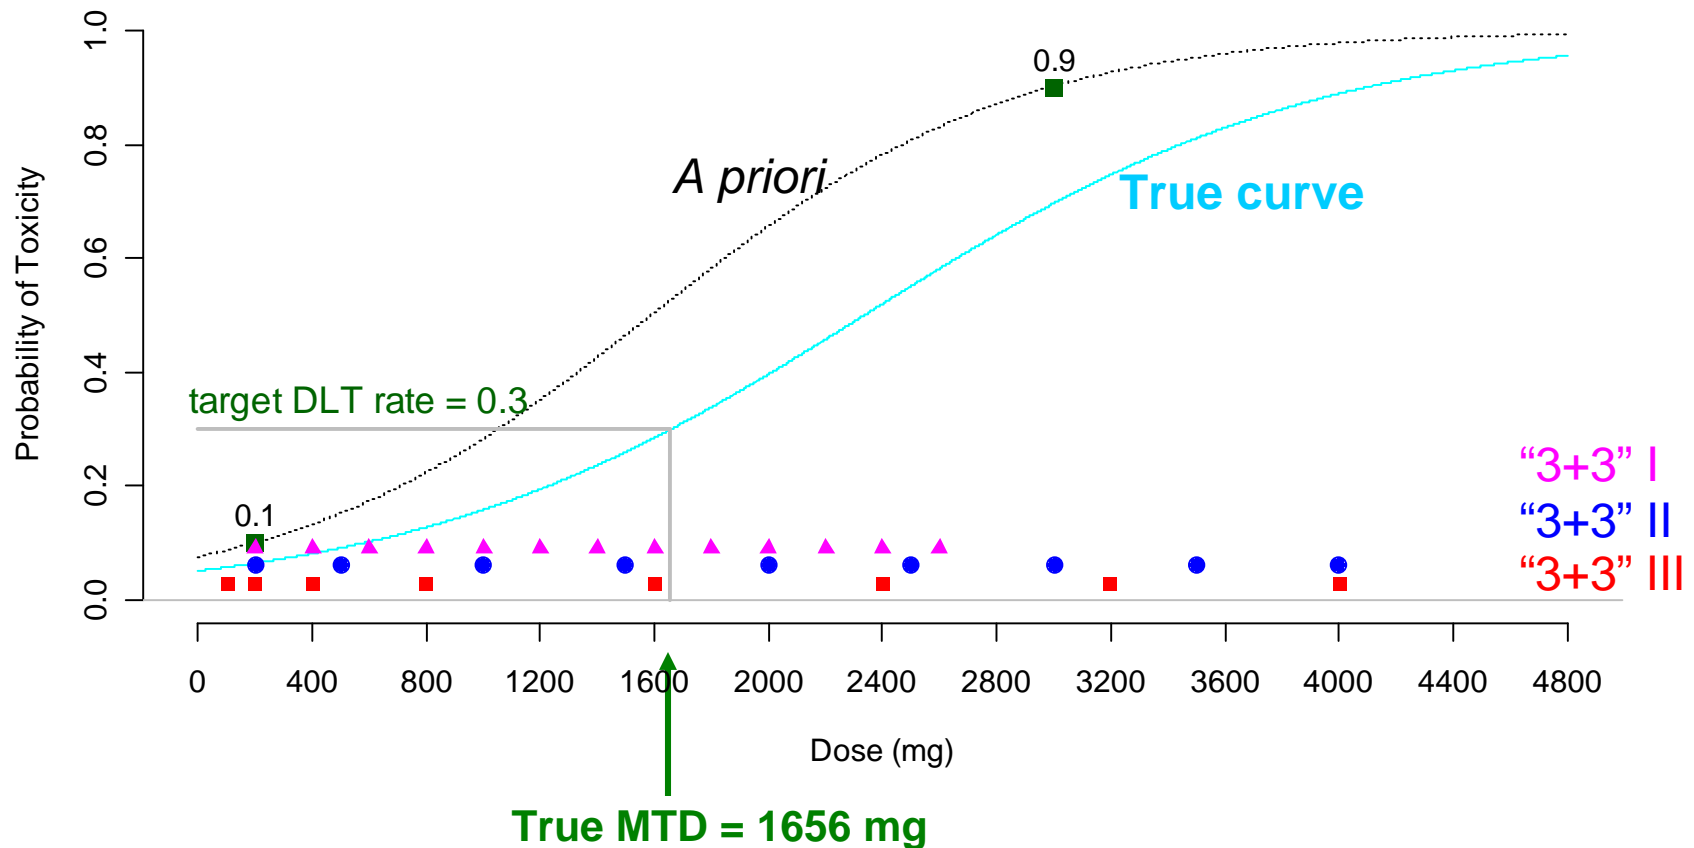
	Cohort	Dose	Outcomes (0 if no DLT, 1 if DLT)	
	1	50	0	0
	2	100	0	0
	3	200	0	0
CRM →	4	400	0	1
	5	350	0	0
	6	400	0	0
	7	550	1	0
	8	450	0	0
	9	500	0	0
	10	550	1	1
	11	450	0	1
	12	450	0	0
	13	450	1	0
	14	450	0	0
MTD →	15	450	STOP	

Stopping rule: 10 patients treated at the same dose

Comparison of CRM and “3+3” Designs using Simulations

Garrett-Mayer E. *Clin Trials* 2006

		CRM I	CRM II	CRM III	“3+3” I	“3+3” II	“3+3” III
1	Total sample size	30	50	60	27	27	39
2	Patients per cohort	3	5	3	3	3	3
3	Number of cohorts	10	10	20	9	9	13



Comparison of CRM and “3+3” Designs using Simulations

Garrett-Mayer E. *Clin Trials* 2006

		CRM I	CRM II	CRM III	“3+3” I	“3+3” II	“3+3” III
4	% of trials with recommend. dose within 250 mg of true (1656 mg)	57%	72%	71%	41%	35%	32%
5	% of trials with recommend. dose within 400 mg of true (1656 mg)	80%	91%	89%	41%	54%	58%
6	% of trials with recommend. dose DLT rate > 40%	9.5%	5.8%	5.9%	7.1%	21%	12%
7	% of trials with recommend. dose DLT rate > 50%	0.9%	0.2%	0.6%	7.1%	2.0%	0.6%
8	% of trials with recommend. dose DLT rate < 20%	13%	5.7%	6.2%	52%	44%	38%
9	% of trials with recommend. dose DLT rate < 10%	0.0%	0.1%	0.0%	11%	16%	6.9%
10	Average % of patients treated at doses with 40% or greater DLT rate	7.6%	7.8%	5.7%	17%	23%	7.5%
11	Average % of patients treated at doses with 20% or less DLT rate	32%	19%	24%	62%	53%	64%
12	Average % of patients treated at doses with DLT rate	26%	28%	26%	21%	22%	19%

Discussions

- Real collaboration between clinician and statistician
 - Clinician: dose levels, DLT rate, range of sample sizes, accrual rate
 - Statistician: cohort size, stopping rule
 - Together: dose-response curve