Bayesian adaptive methods for optimal dose finding studies : clinical development of molecularly targeted therapies in oncology

Pierre Colin

December 6th 2016

Supervisors:

Eric Parent & Maud Delattre, AgroParisTech, MIA
Sandrine Micallef & Pierre Mancini, Sanofi R&D, Biostatistics & Programming

Contents

1. Context
2. Local optimality
3. Global optimality
4. Modeling
5. Simulations
6. Conclusion

Phase I clinical trial

Possible path
Observed path

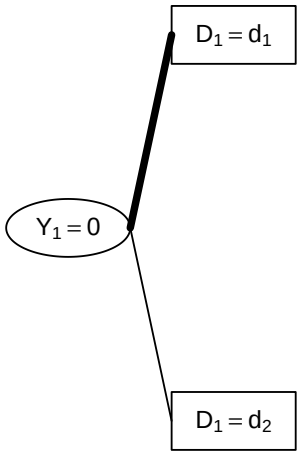

Phase I clinical trial

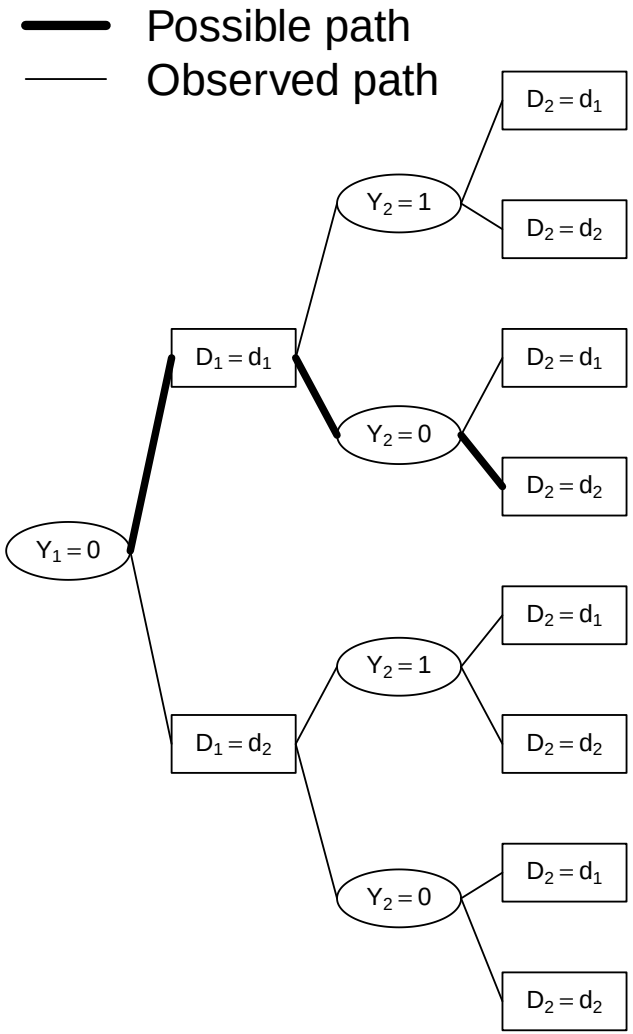

Phase I clinical trial

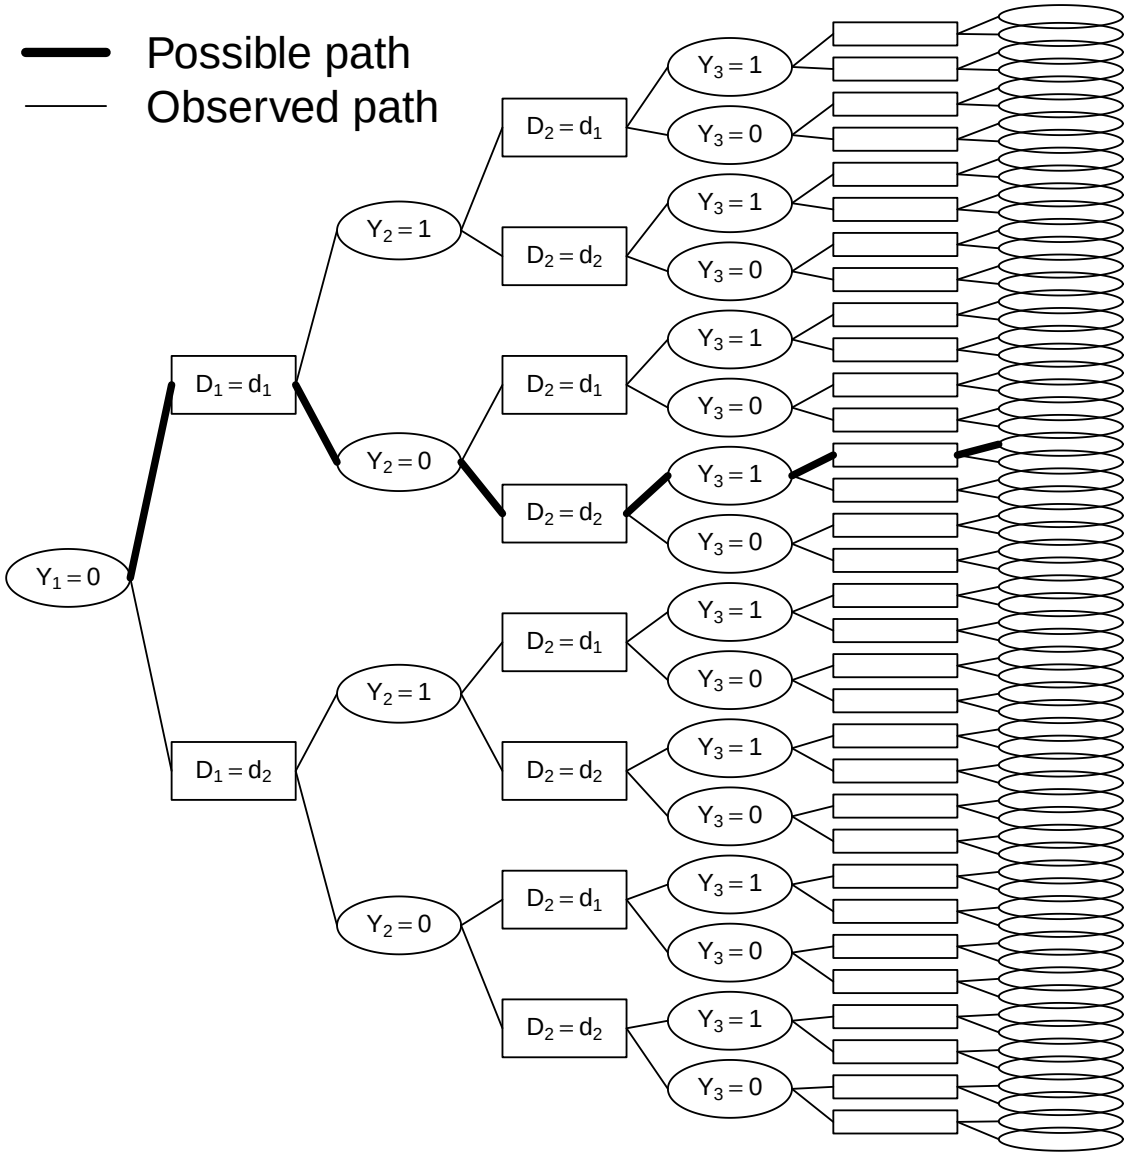

Phase I trial: a sequential dose recommendation

Design issues

- Sequential treatment of a small number of patients
  - Treat the patients according to up-to-date information
  - Challenge the dose level after each cohort of patients
- Several potential decisions (number of predefined dose levels)
- Inclusion of patients who have reached a therapeutic "dead end"
  - Looking for innovative drugs (potentially active)

Drug development issues

- Increasing information on drug toxicity and activity
- Paradigm shift due to the development of Molecularly Targeted Therapies (MTA)
  - Drug activity need to be assessed to recommend the dose level

Paradigm shift

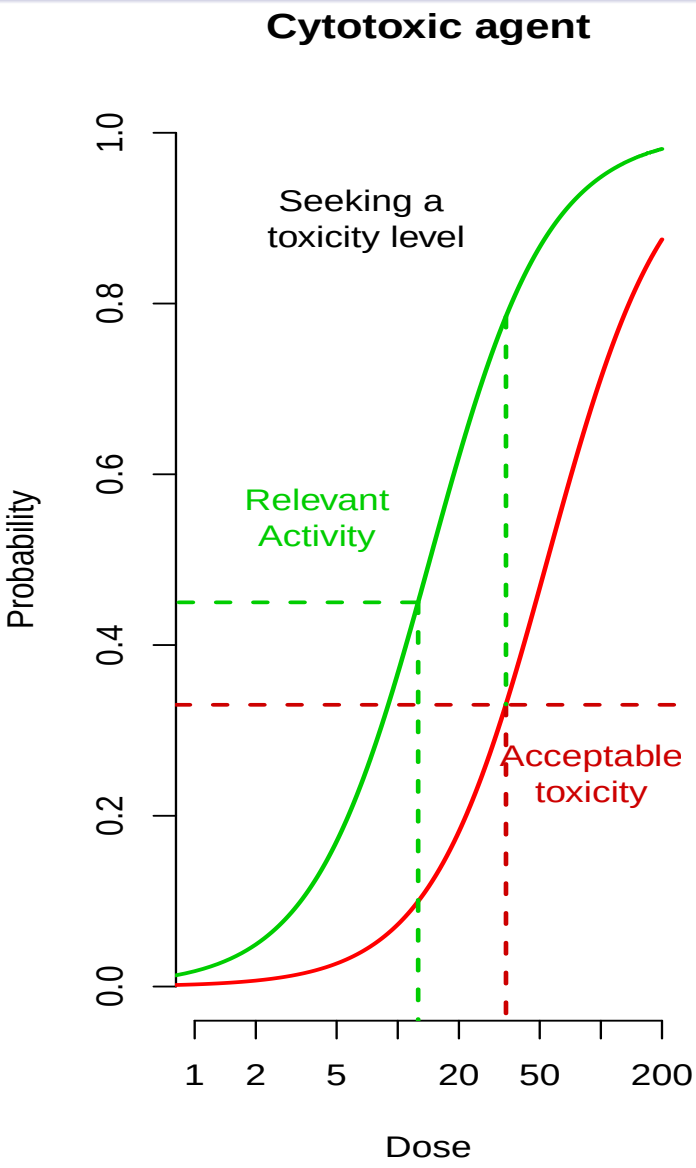

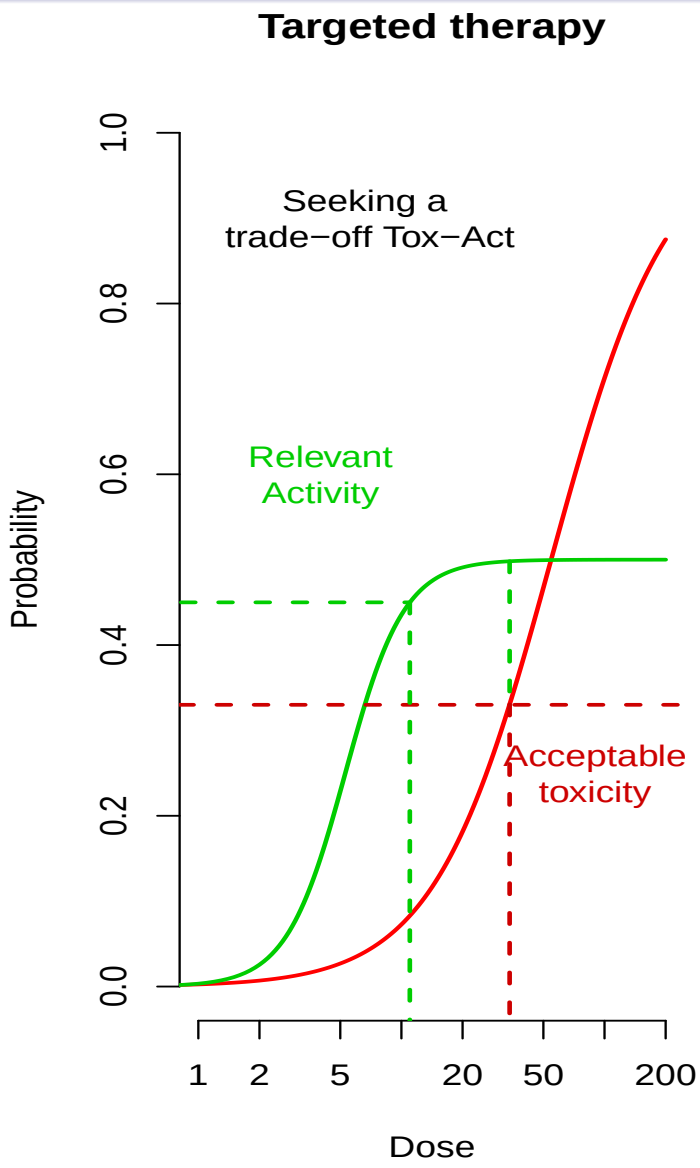

Objectives of a Phase I clinical trial

Dose finding

- Recommend the dose level with the highest chance of success for the following clinical trials
- Treat correctly future patients (long-term objective)
- Depends on the final estimation of the drug toxicity and activity (and the final design of the study)

Objectives of a Phase I clinical trial

Dose finding

- Recommend the dose level with the highest chance of success for the following clinical trials
- Treat correctly future patients (long-term objective)
- Depends on the final estimation of the drug toxicity and activity (and the final design of the study)

Treatment of patients

- Provide to most of included patients a dose level with:
  - Low toxicity
  - High activity
- Depends on the sequential dose recommendation
- Leads to the final design

The treatment versus experimentation dilemma

Dilemma

- Give priority to final recommendation
  - $\Rightarrow$ Included patients are poorly treated
- Give priority to patients' treatment
  - $\Rightarrow$ Poor final dose recommendation

Global cost of the clinical trial

- Trade-off between immediate and final cost functions
- Final objective of the clinical trial is clearly defined
- Patient's treatment is not only an objective, but also a source of information

Treatment and dose recommendation considered as a cost

Hazard-Decision system

- Decision step: $x_t \longmapsto d_t = \delta_t(x_t)$
- Hazard step: $X_{t+1} = f_t(x_t, \delta_t(x_t), Y_t), \quad Y_t \sim F_\theta(d_t)$

Treatment and dose recommendation considered as a cost

Hazard-Decision system

- Decision step: $x_t \longmapsto d_t = \delta_t(x_t)$
- Hazard step: $X_{t+1} = f_t(x_t, \delta_t(x_t), Y_t), \quad Y_t \sim F_\theta(d_t)$

Decision evaluation

Decisions $d_1, \ldots, d_{T+1}$: $\sum_{t=1}^{T} L(\theta, d_t, y_t, t) + L(\theta, d_{T+1}, T+1)$

Treatment and dose recommendation considered as a cost

Hazard-Decision system

- Decision step: $x_t \longmapsto d_t = \delta_t(x_t)$
- Hazard step: $X_{t+1} = f_t(x_t, \delta_t(x_t), Y_t), \quad Y_t \sim F_\theta(d_t)$

Decision evaluation

Decisions $d_1, \ldots, d_{T+1}$: $\sum_{t=1}^{T} L(\theta, d_t, y_t, t) + L(\theta, d_{T+1}, T+1)$

Strategy evaluation

$$\mathbb{E}_\theta\left(\mathbb{E}_{Y_{1:T}}\left(\sum_{t=1}^{T} L(\theta, \delta_t(X_t), Y_t, t) + L(\theta, \delta_{T+1}(X_{T+1}), T+1)\right)\right)$$

Contents

2 Local optimality
- Treatment versus experimentation dilemma
- Literature review

Treatment before anything else

Assumption of the optimization problem

- Treatment of the actual patient is the only one objective
- Final dose recommendation is ignored

Treatment before anything else

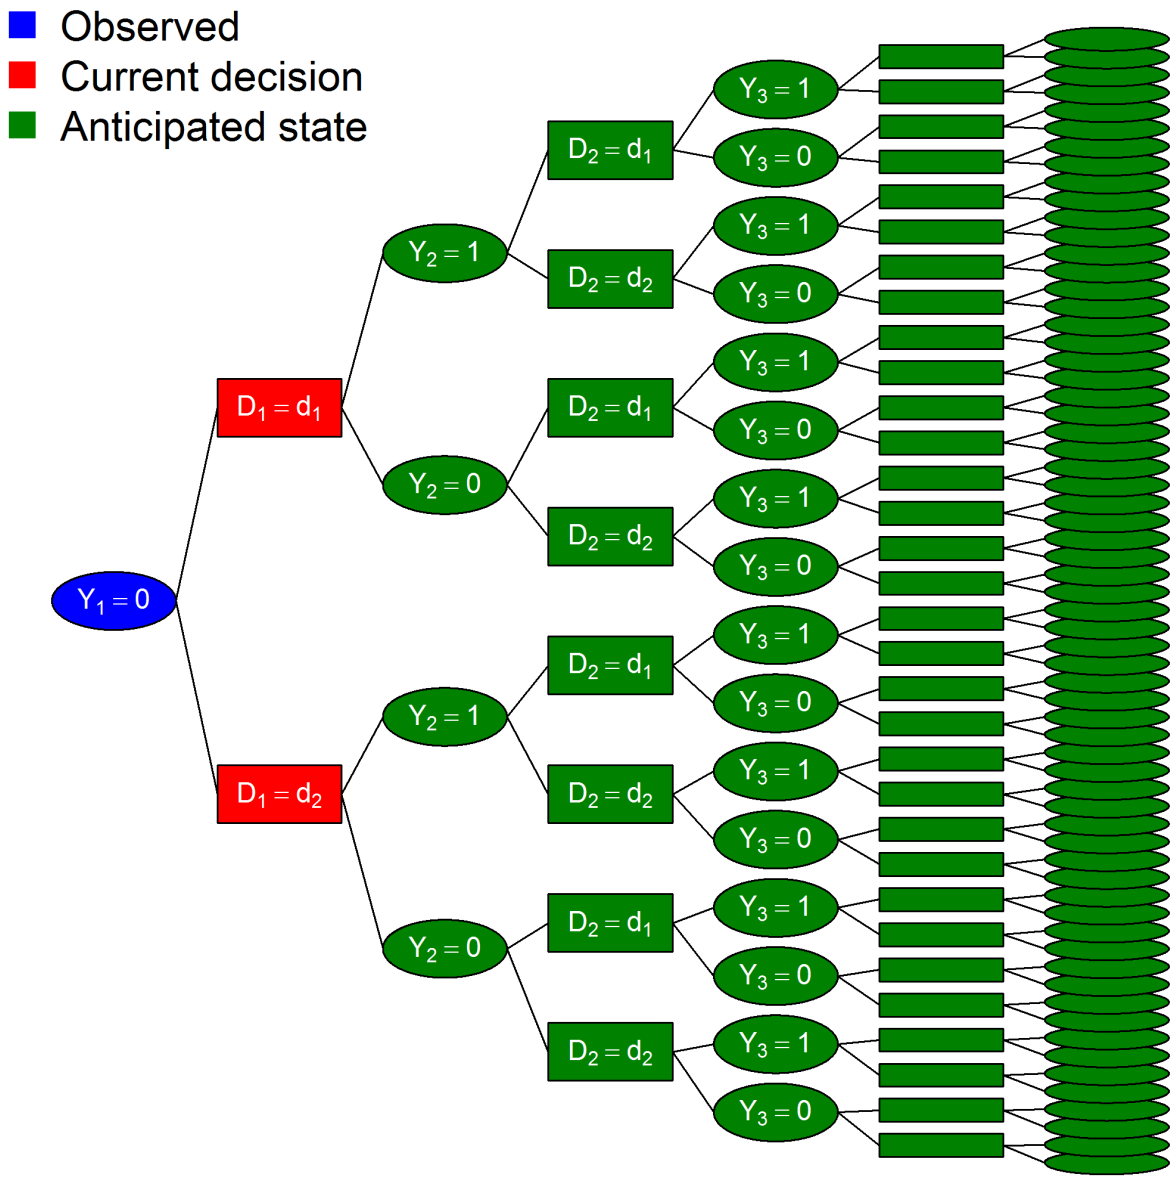

Treatment before anything else

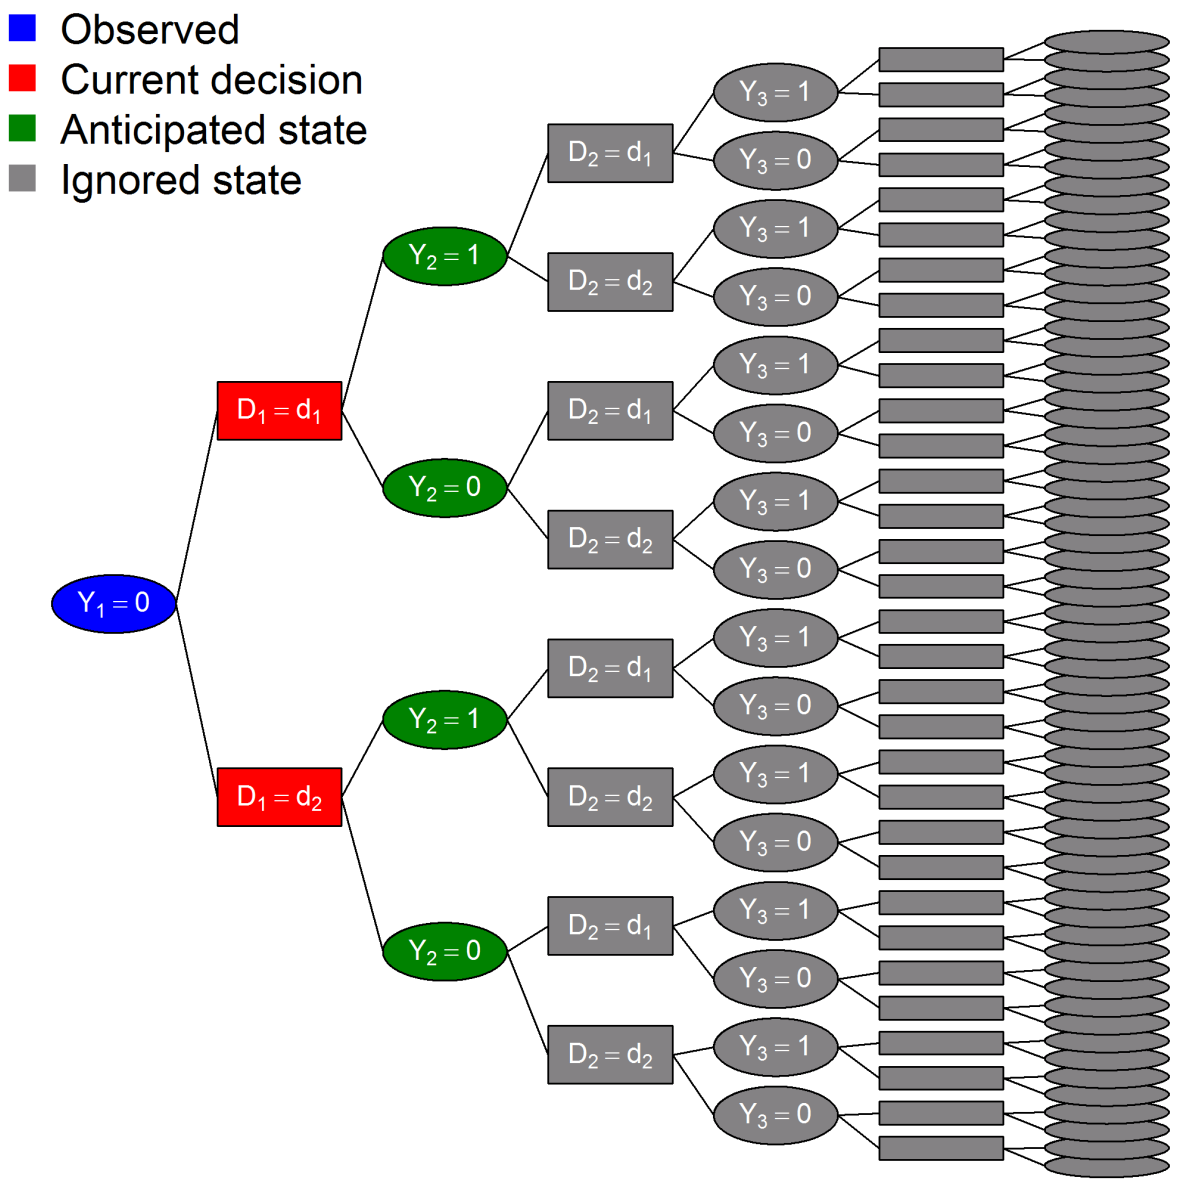

Treatment before anything else

Assumption of the optimization problem

- Treatment of the actual patient is the only one objective
- Final dose recommendation is ignored

Local optimization criterion

$$d_{t+1} = \arg\min_{d} \; \mathbb{E}_{\theta}\left[L\Big(\theta, d, t+1\Big) \middle| y_{1:t} \right]$$

Advantage/Drawback

- Guaranteed patient's treatment (at least for first patients)
- May converge to the wrong dose recommendation

Literature review

Clinical cost

- Targeting a toxicity level (MTD)
  - CRM (O'Quigley 1990)
  - EWOC (Babb 1998)
  - Tite-CRM (Cheung 2000)
- Targeting a toxicity and/or activity level(s)
  - EffTox (Thall 2004)
  - bCRM (Braun 2002)

Information cost

- Optimal design (Bretz 2010)
- SLOD (Roth 2009)

CRM example - John O'Quigley 1990

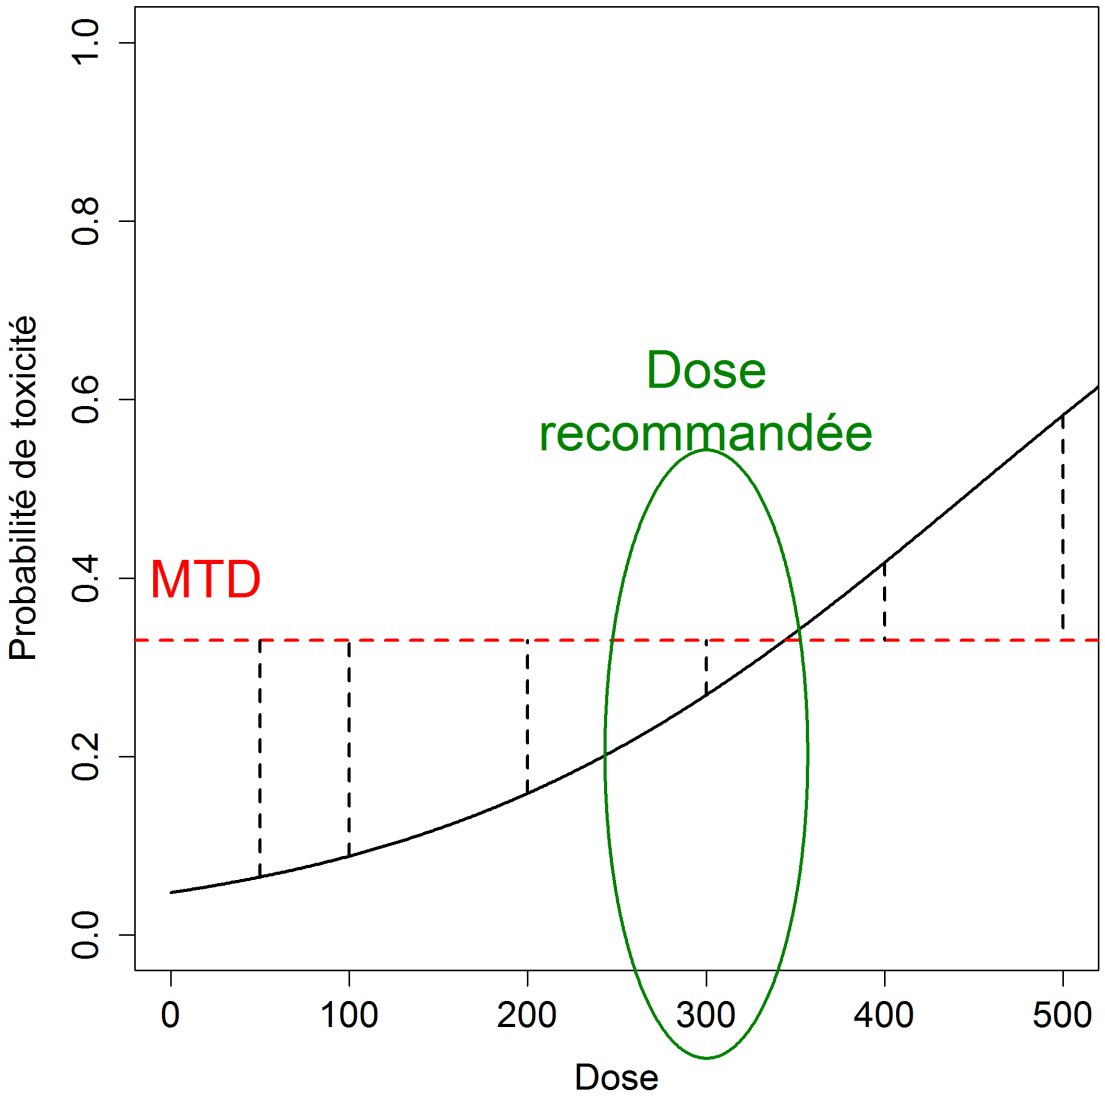

Contents

3 Global optimality
- Treatment versus experimentation dilemma
- Strategy optimization
- Approximated methods

Final recommendation is taken into account

Assumption of the optimization problem

- Treatment of the actual patient is the only a part of the clinical trial objectives
- Final dose recommendation must be assessed to evaluate the intermediate recommendation

Final recommendation is taken into account

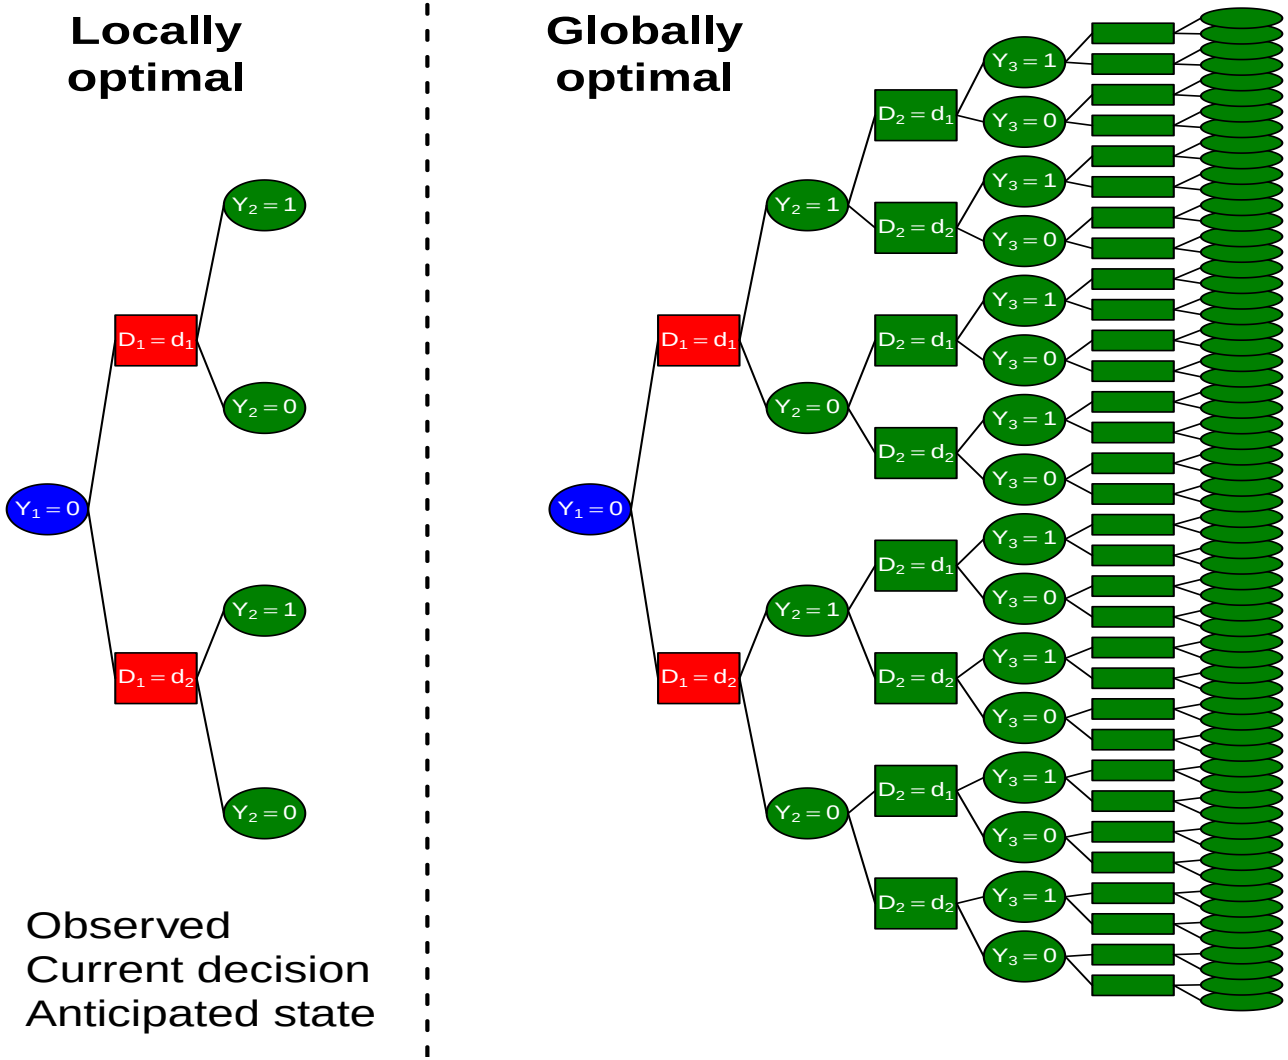

Final recommendation is taken into account

Global optimization criterion

$$d_{t+1} = \arg\min_{d} \; \mathbb{E}_{\theta}\left(\mathbb{E}_{Y_{1:T}}\left(\sum_{t=1}^{T} L(\theta, d, Y_t, t) + L(\theta, \delta_{T+1}(X_{T+1}), T+1)\right)\right)$$

Advantage/Drawback

- Look for a good treatment for all included patients
- Take into account the quality of the final dose recommendation

Inflation

- Future costs can be down weighted by: $\left(\frac{1}{1+\alpha}\right)^t$

Exact optimization (Dynamic programming)

Final recommendation

- No more expected data
- Final decision is based on all available information ($x_{t+1} = \{y_{1:t}, d_{1:t}\}$)

$$d_{T+1} = \arg\min_{d} \left\{ E_{\theta|x_{T+1}} \left\{ L(\theta, d, T+1) \right\} \right\}$$

Exact optimization (Dynamic programming)

Final recommendation

- No more expected data
- Final decision is based on all available information ($x_{t+1} = \{y_{1:t}, d_{1:t}\}$)

$$d_{T+1} = \arg\min_d \left\{ E_{\theta|x_{T+1}} \left\{ L(\theta, d, T+1) \right\} \right\}$$

Bellman equation

- Use the future optimal decision to evaluate the current decision
- Anticipate the future optimal decision (and cost) according to available data

$$d_{t+1} = \arg\min_d \left\{ E_{Y_t} \left\{ E_{\theta|y_t} L(\theta, d, t) + V_{t+1}(f_t(x, d, Y_t)) \right\} \right\}$$

Exact optimization (Dynamic programming)

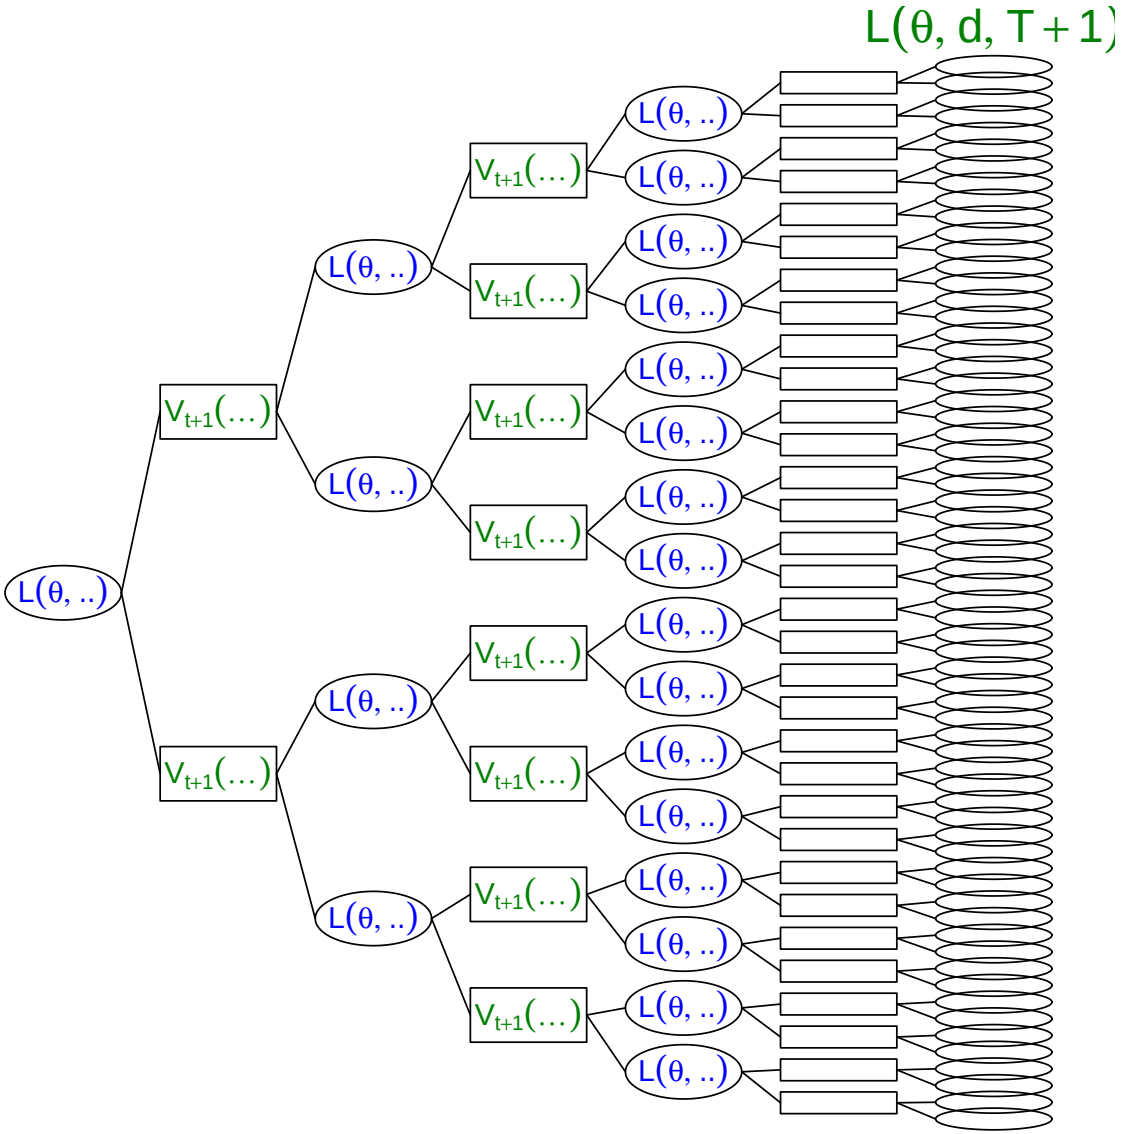

Curse of the dimension

Notations

- N: Total number of patients
- K: Number of prespecified dose levels
- M: Number of potential patient's response (categorical variable)

Maximal number of paths

$$\sum_{n=2}^{N} \binom{K \times M + n - 1}{n}$$

- $O(N^{KM-1})$ when $N \to +\infty$
- $O(M^N)$ when $M \to +\infty$

Curse of the dimension

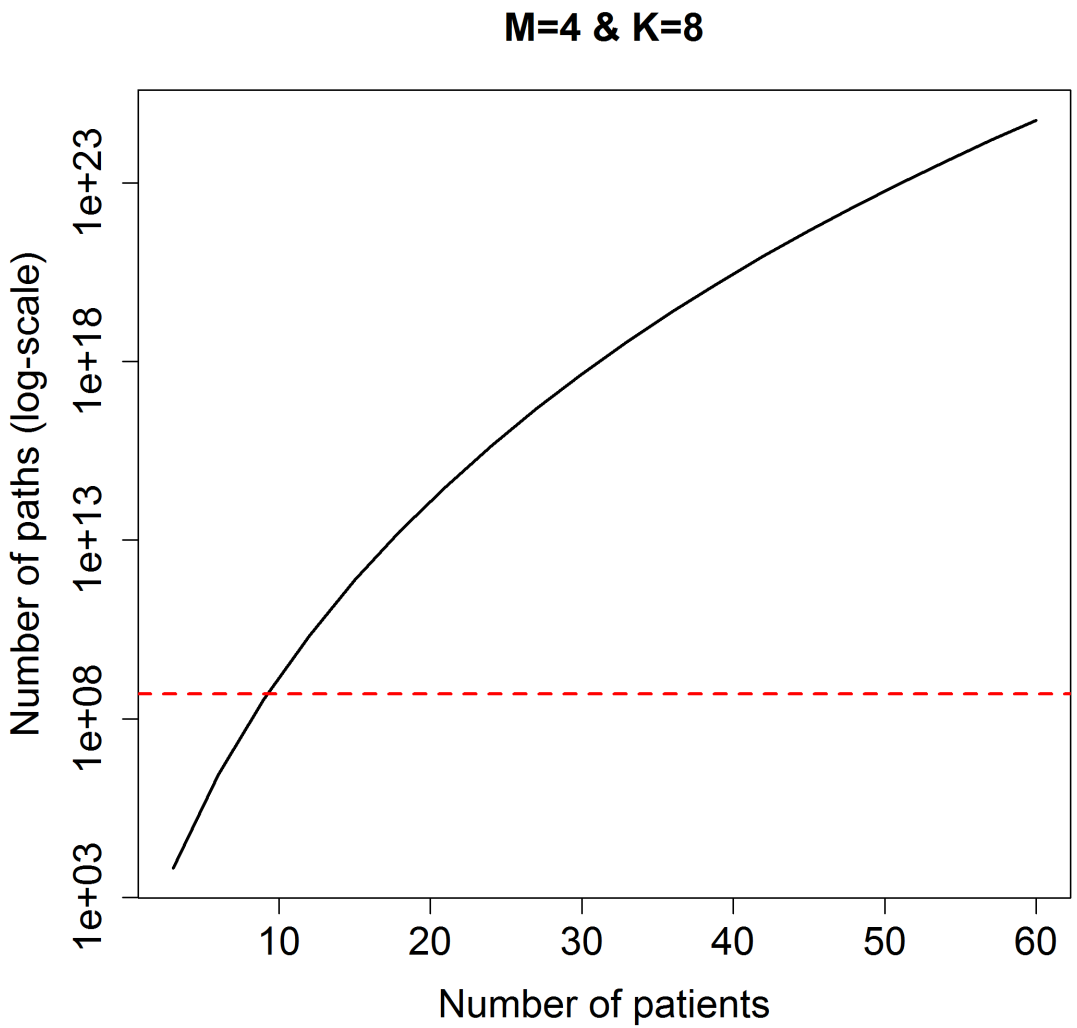

Approximate optimization (1/3)

Approximation by truncation

- Current decision does not impact later costs
- Trial is assumed to be stopped within $\tilde{T}$ steps ($\ll T$)

Approximate optimization (1/3)

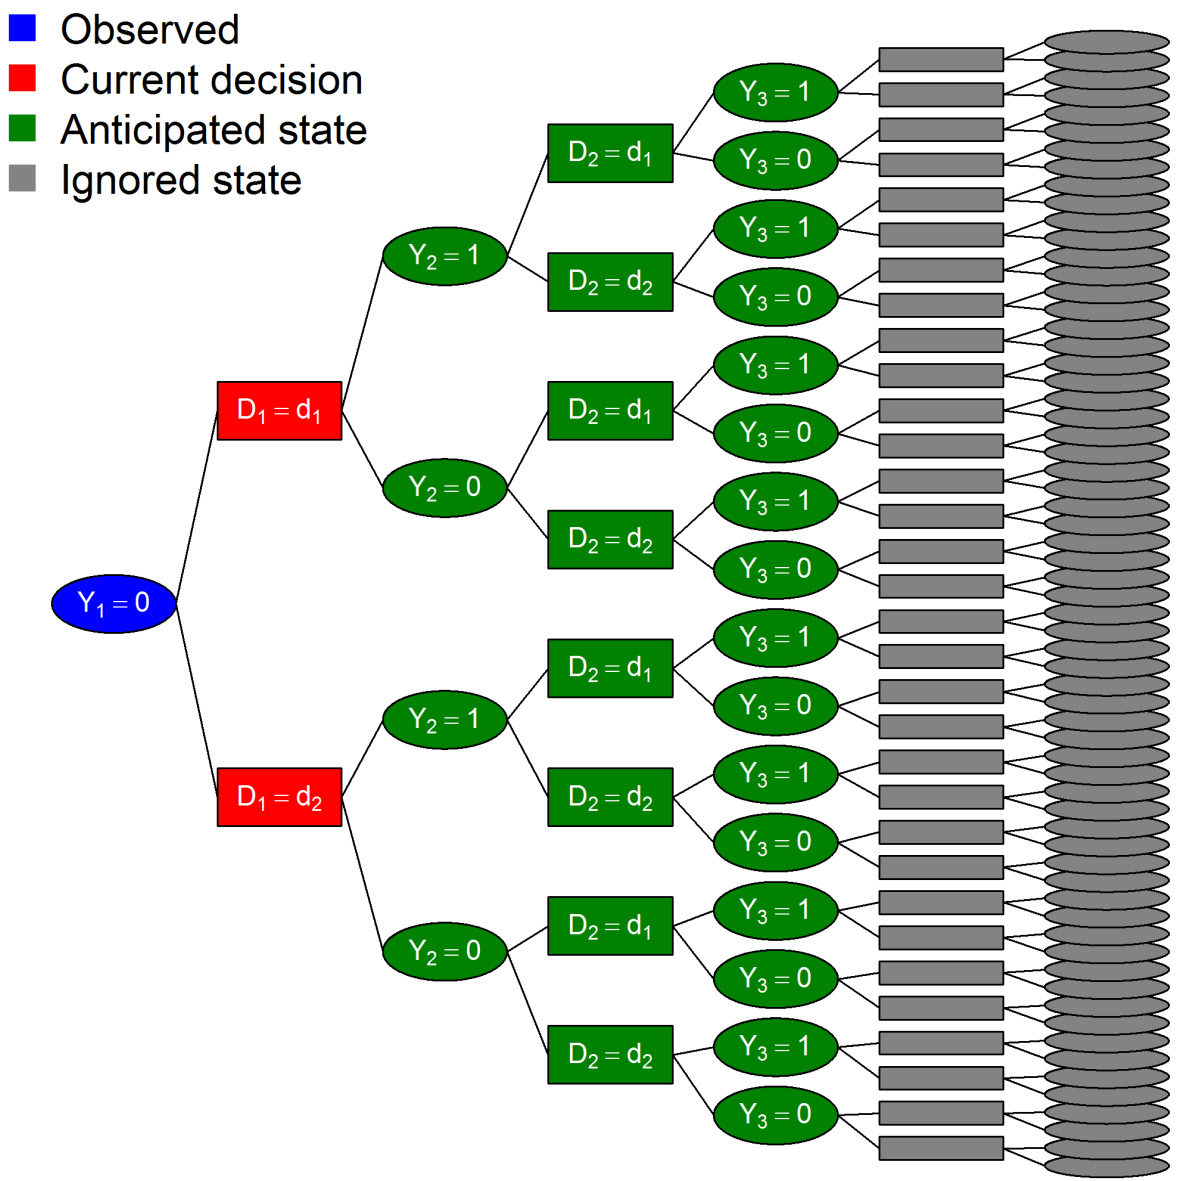

Approximate optimization (1/3)

Approximation by truncation

- Current decision does not impact later costs
- Trial is assumed to be stopped within $\tilde{T}$ steps ($\ll T$)

Optimization criterion

$$\mathbb{E}_{Y_{t:t+\widetilde{T}}}\left[\sum_{\tau=t}^{t+\widetilde{T}} \mathbb{E}_{\theta}\left[L(\theta, \delta_{\tau}(X_{\tau-1}), Y_{\tau}, \tau)\middle| y_{1:\tau-1}\right]\right].$$

Advantage/Drawback

- Less complex dynamic programming: $O(N^{KM-1}) \rightarrow O(\tilde{T}^{K\tilde{M}-1})$
- Biased expected cost

Approximate optimization (2/3)

Approximation by Monte Carlo simulations

- Expectation is approximated by averaged simulations
- Simulations are performed through any strategy (algorithmic or locally optimal strategy)

Approximate optimization (2/3)

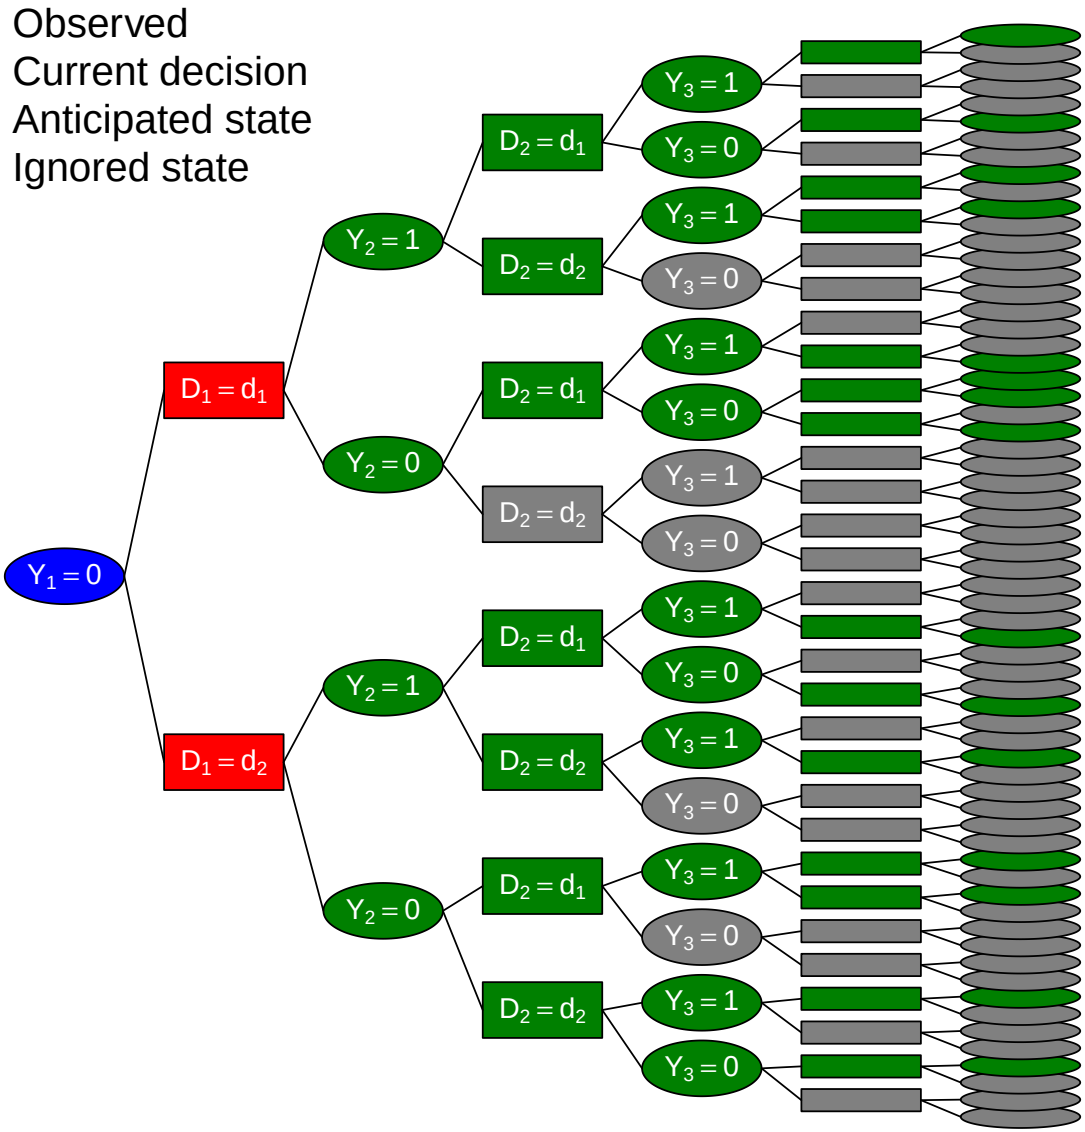

Approximate optimization (2/3)

Approximation by Monte Carlo simulations

- Expectation is approximated by averaged simulations
- Simulations are performed through any strategy (algorithmic or locally optimal strategy)

Optimization criterion

$$\frac{1}{M}\sum_{m=1}^{M}\left[\sum_{\tau=t+1}^{T}\mathbb{E}_{\theta}\left[L\left(\theta,\tilde{d}_{\tau}^{(m)},\tau\right)\middle|y_{1:t-1},y_{t:\tau-1}^{(m)}\right]\right]$$

Advantage/Drawback

- Clinical trial design is preserved
- Monte Carlo approximation
- Biased future distribution (due to the simulated strategy)

Approximate optimization (3/3)

Combined approximated methods

- Exact expectation is performed within $\tilde{T}$ steps
- Further expected costs are approximated by Monte Carlo simulations

Approximate optimization (3/3)

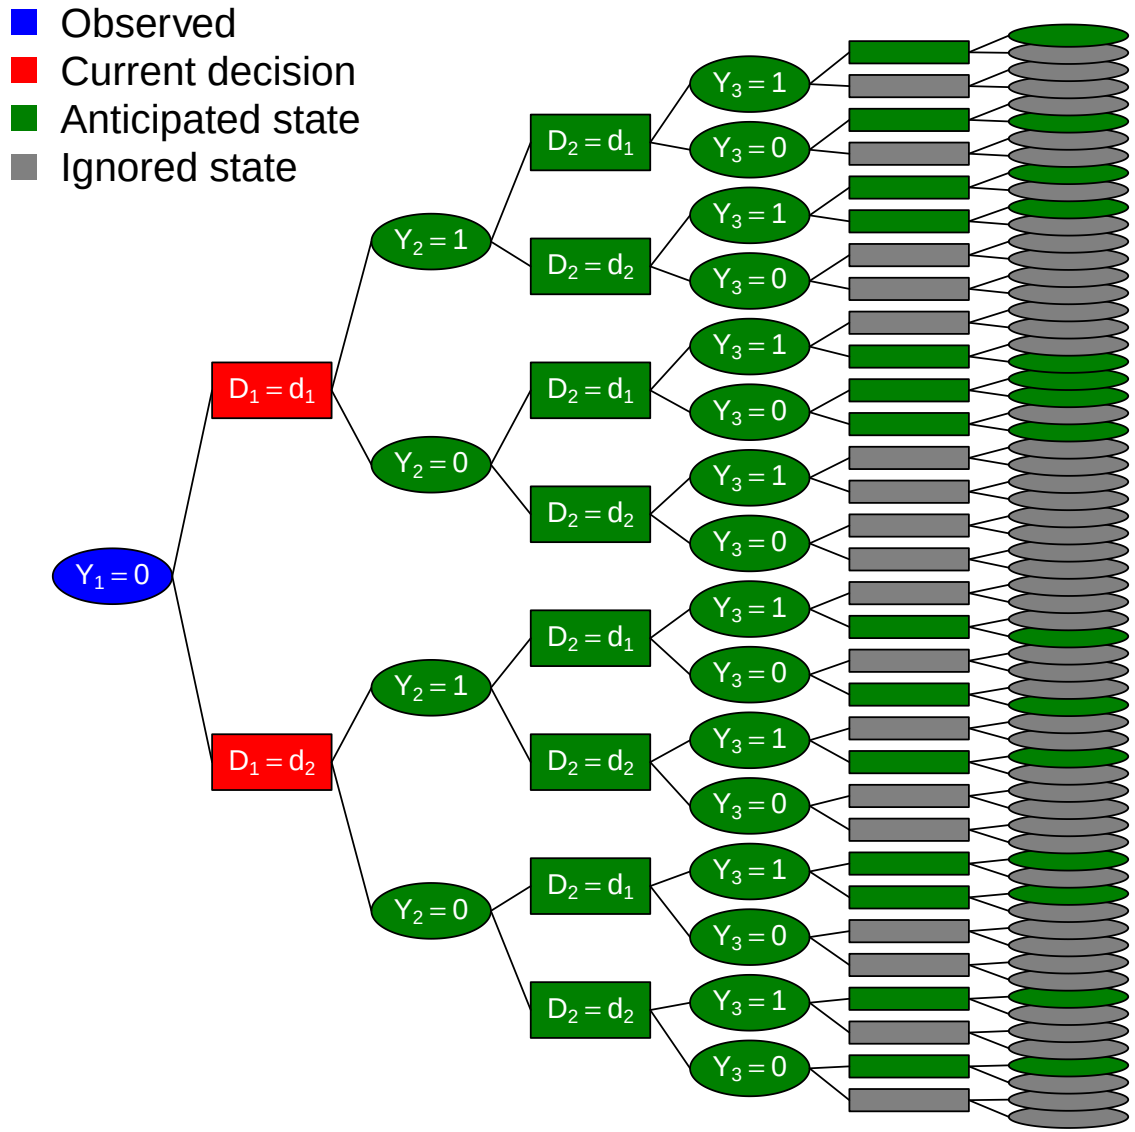

Approximate optimization (3/3)

Combined approximated methods

- Exact expectation is performed within $\tilde{T}$ steps
- Further expected costs are approximated by Monte Carlo simulations

Advantage/Drawback

- Clinical trial design is preserved
- Less biased future distribution

Conclusion

- Approximated method used for subsequent designs
- Fully globally optimal designs A_α
- Alternative design A^*

Contents

4 Modeling
- Paradigm shift
- Model
- Bayesian inference
- Cost functions

Paradigm shift

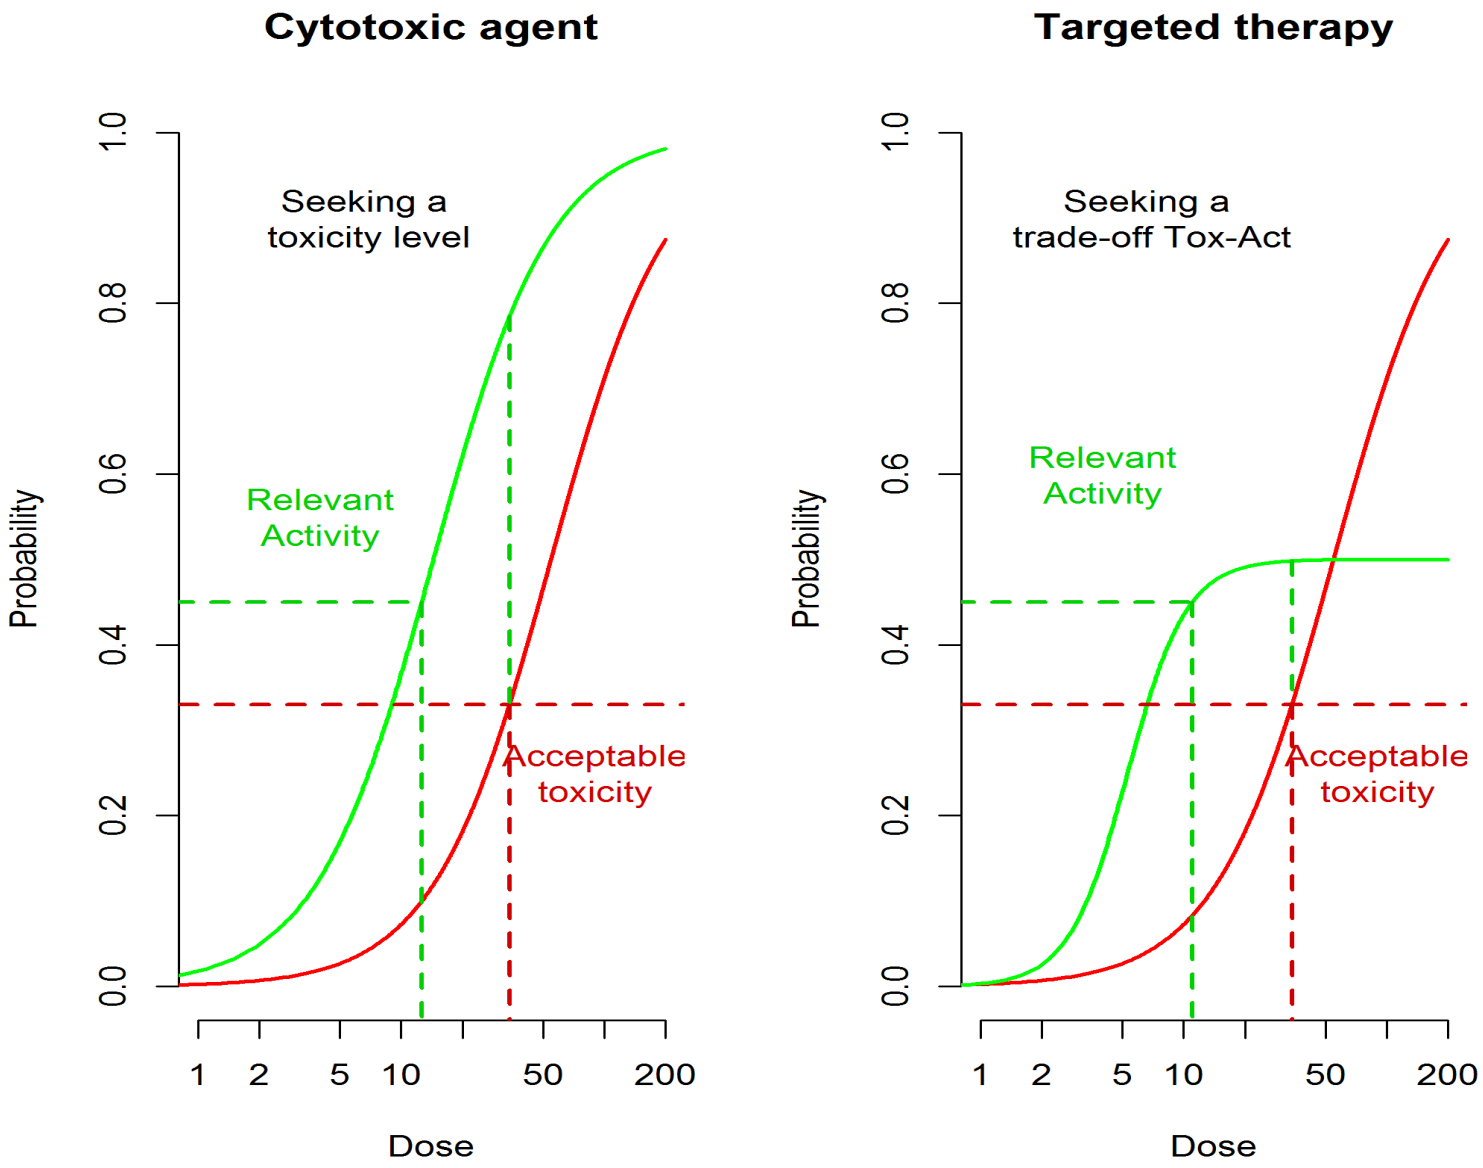

Multivariate probit model

Binary endpoints

$$Y^{Tox} = \mathbb{1}(Z^{Tox} \geq 0) \quad \pi_{Tox}(d, \theta) = \mathbb{P}(Y^{Tox} = 1)$$
$$Y^{Act} = \mathbb{1}(Z^{Act} \geq 0) \quad \pi_{Act}(d, \theta) = \mathbb{P}(Y^{Act} = 1)$$
$$Z = \begin{pmatrix} Z^{Tox} \\ Z^{Act} \end{pmatrix}$$

Multivariate probit model

$$\eta_{Tox}(d) = \alpha_1 + \beta_1 \times \log(d)$$
$$\eta_{Act}(d) = \alpha_2 + \beta_2 \times \log(d) + \gamma_2 \times \log(d)^2$$
$$\begin{pmatrix} Z^{Tox} \\ Z^{Act} \end{pmatrix} \sim \mathcal{N}_{mv}\left(\begin{bmatrix} \eta_{Tox}(d) \\ \eta_{Act}(d) \end{bmatrix}, \Sigma \right)$$
$$\Sigma = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$$

Inference

Prior distributions

- Gaussian distributions for regression parameters
- Wishart distribution for the latent variance matrix
- Prior information is controlled by the Effective Sample Size
  - The final results must be *data-driven*

Particle filter

- Suitable to sequential analyses
- Re-use of the previous *a posteriori* distributions
- Faster than MCMC algorithms
- Prior distributions are used as instrumental distributions

Cost functions

Target dose levels

- Safe dose level: $\pi_{Tox}(d, \theta) < R_2$
- Active dose level: $\pi_{Act}(d, \theta) \geq R_1$
- Target doses = Safe doses $\cap$ Active doses

Immediate {0,-1} cost function

$$L(\theta, d_t, t) = -\mathbb{1}(\pi_{Act}(d_t, \theta) \geq R_1) \times \mathbb{1}(\pi_{Tox}(d_t, \theta) < R_2)$$

Final cost function

$$L\Big(\theta, d, T+1\Big) = N_2 \times L\Big(\theta, d, T\Big)$$

This function assures a common scale for both cost functions

Cost functions

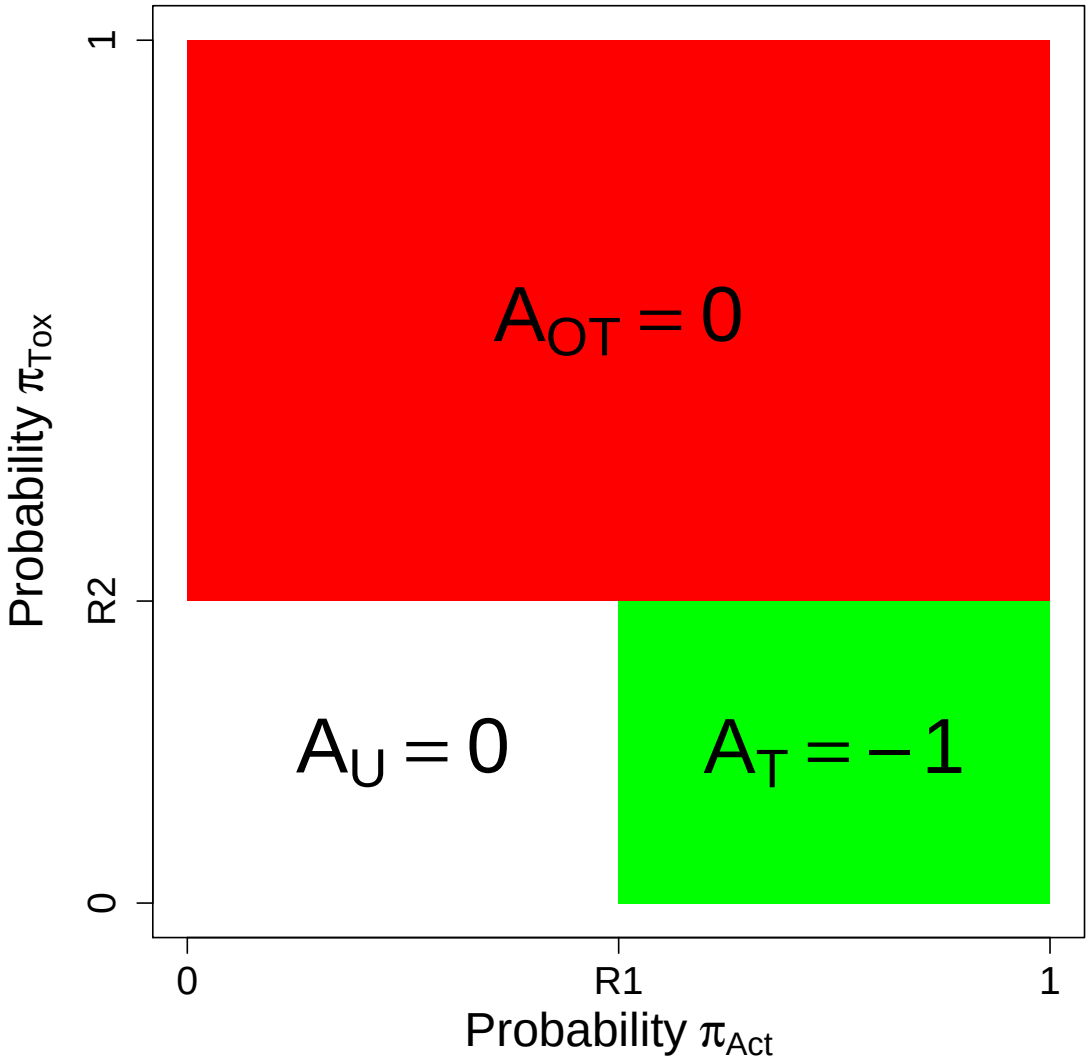

Contents

5 Simulations
- Objectives
- Designs
- Simulations
- Results

Objectives and criteria

Objectives

- Compare designs' operating characteristics
- Simulations with fixed number of patients (30 & 45)

Criteria

- Average strategy cost
- Dose recommendation
- Number of events (activity & toxicity)
- Patient allocation (patients treated with active or toxic dose levels)

Design comparison

Algorithmic design

- $3 + 3$

Locally optimal designs

- EffTox (continuous cost function)
- EBE-CRO (discrete cost function)

Globally optimal designs

- Adaptive design with inflation
  - A_3
  - A_{10}
  - A_{20}
- A^* is a $3 + 3$ followed by an adaptive design

Simulation plan

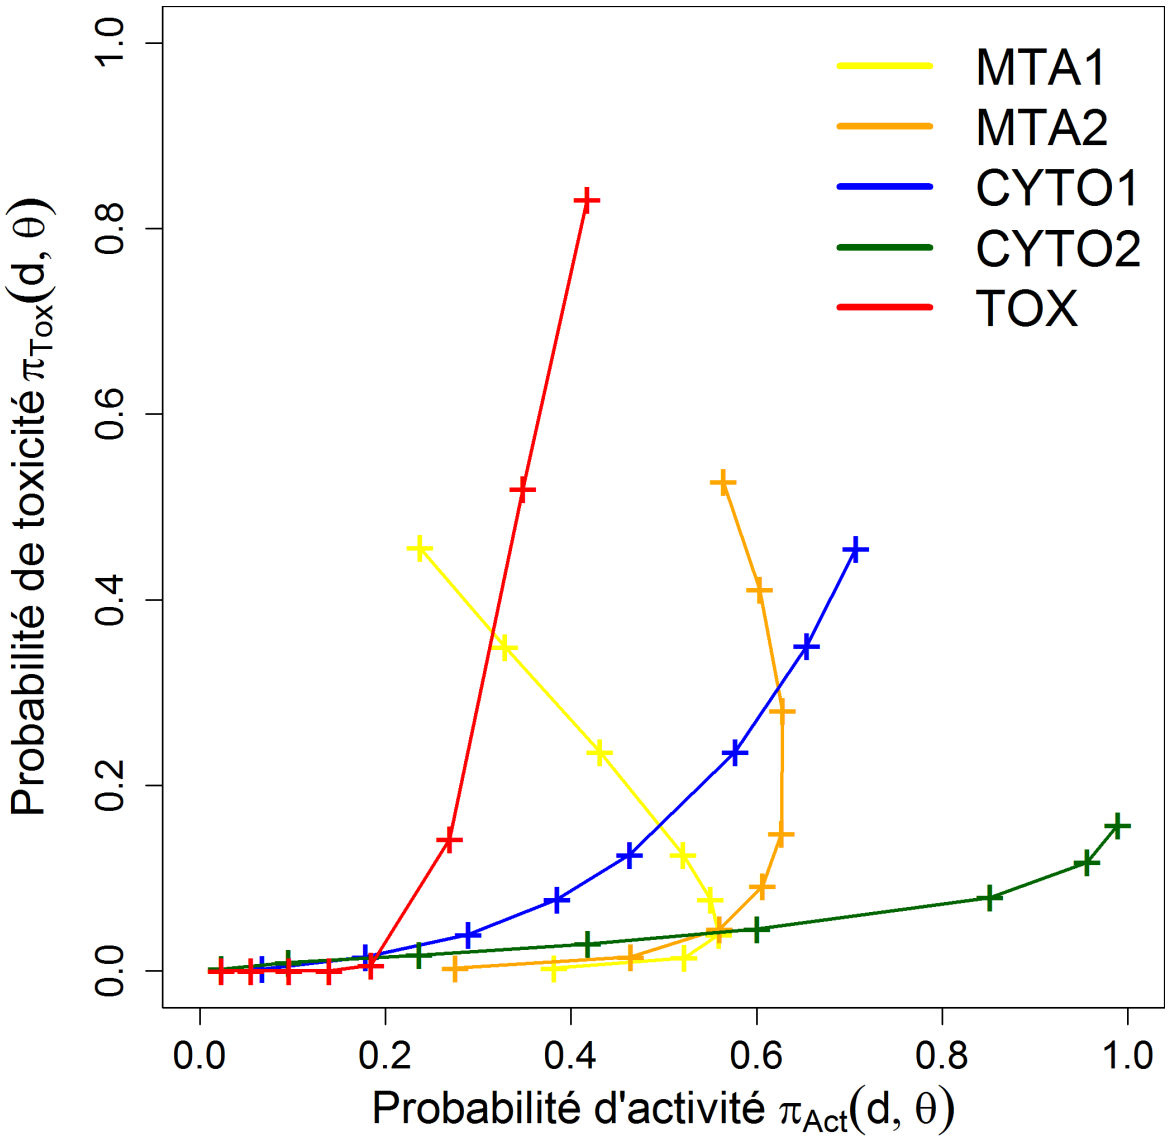

Results - Strategy cost (1/3)

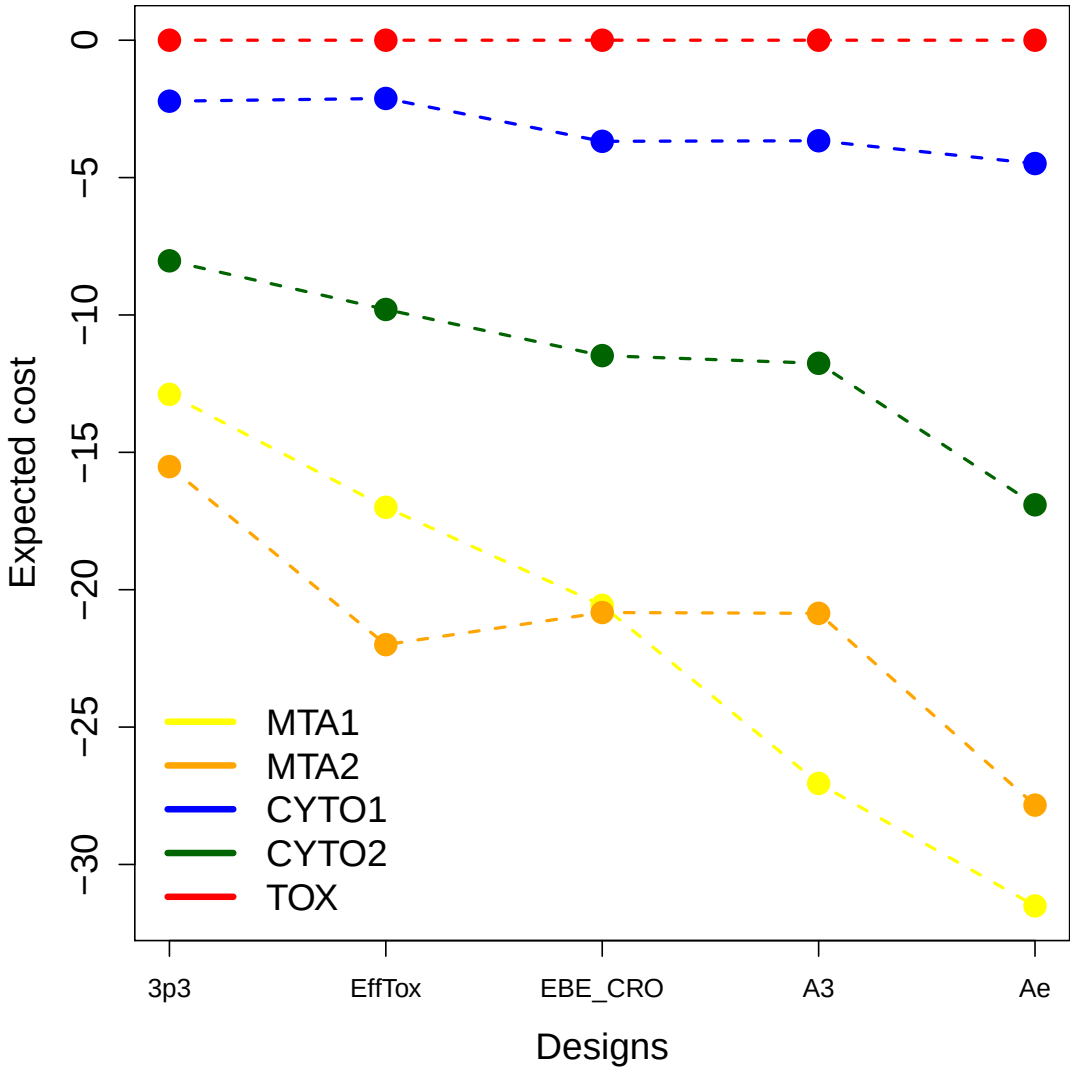

Results - Final dose recommendation (2/3)

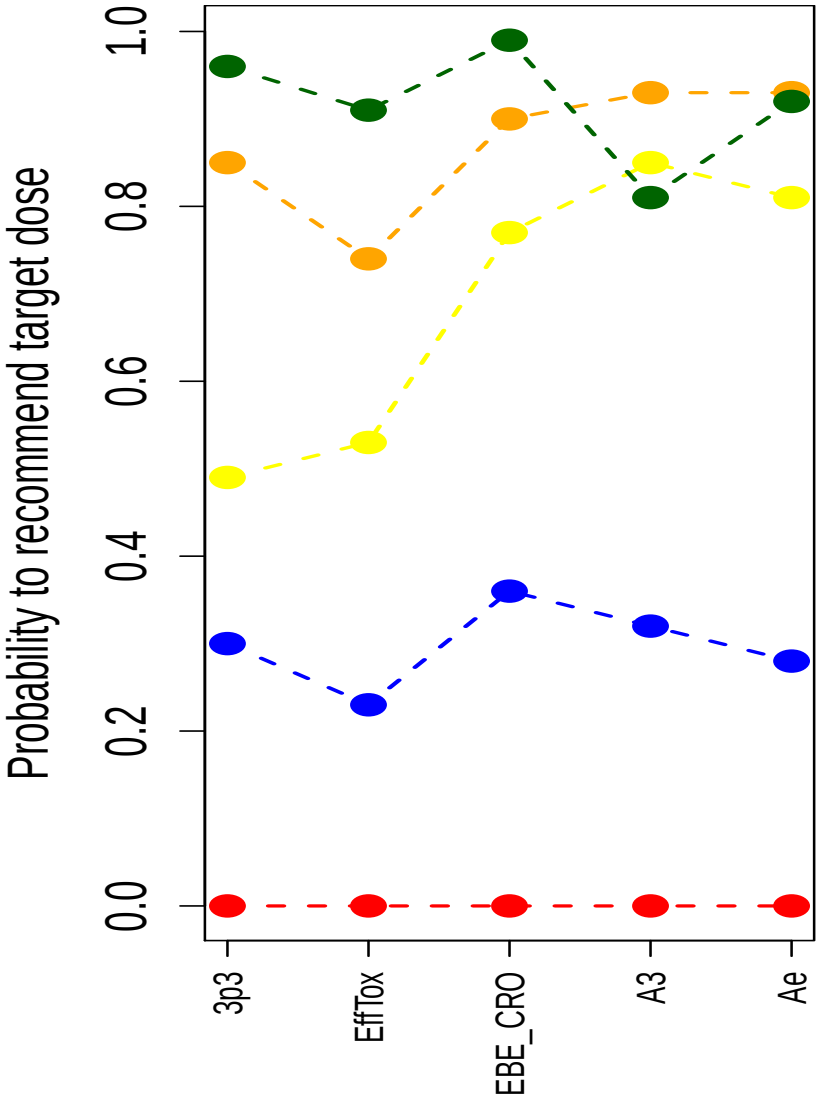

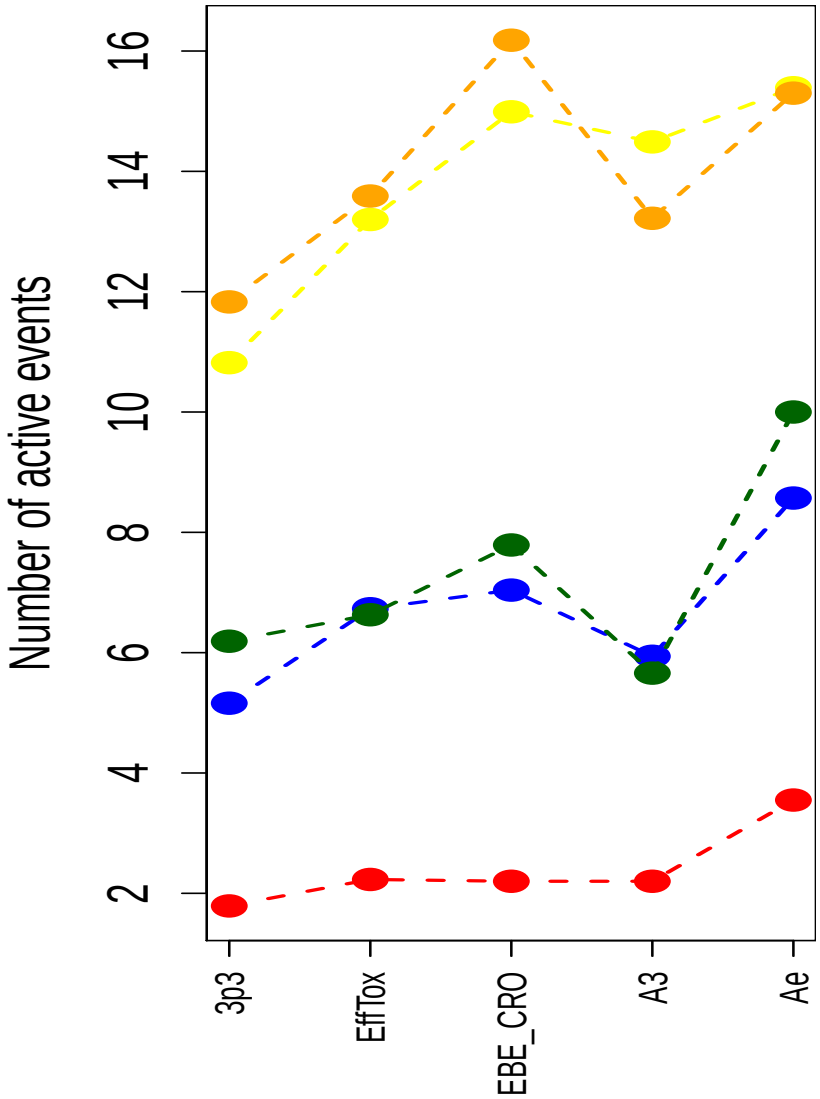

Results - Patient allocation (3/3)

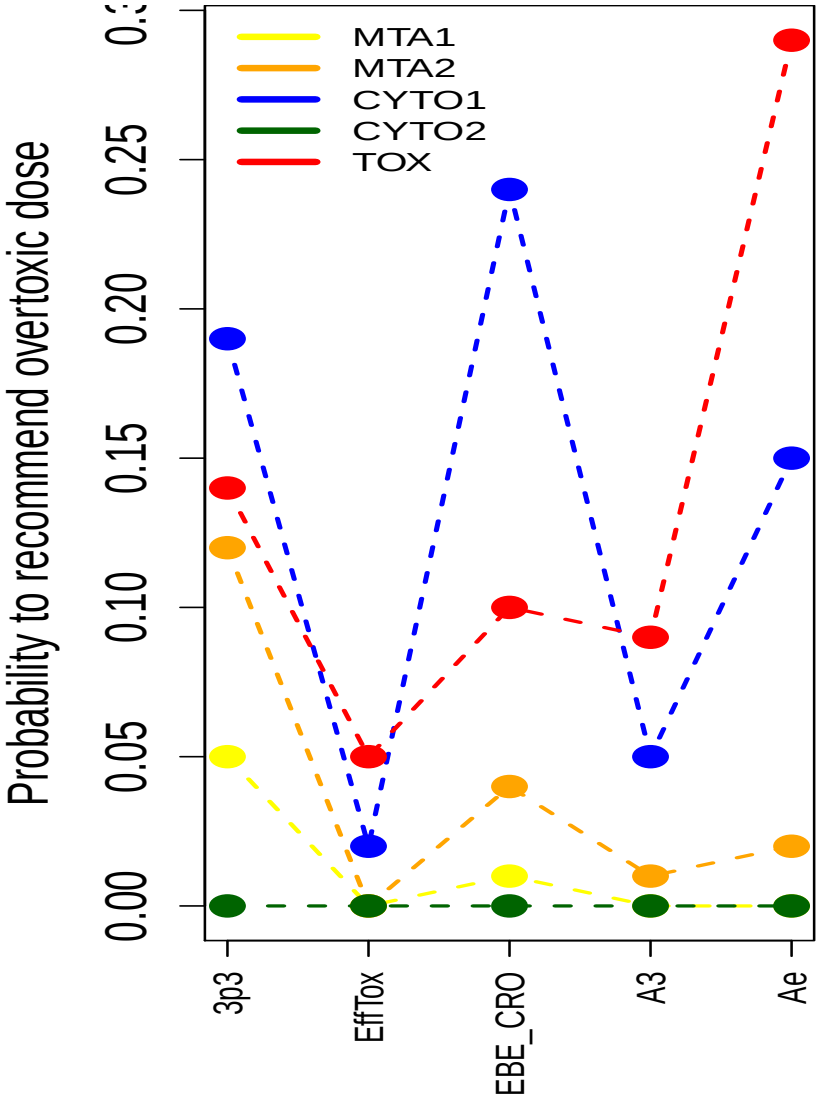

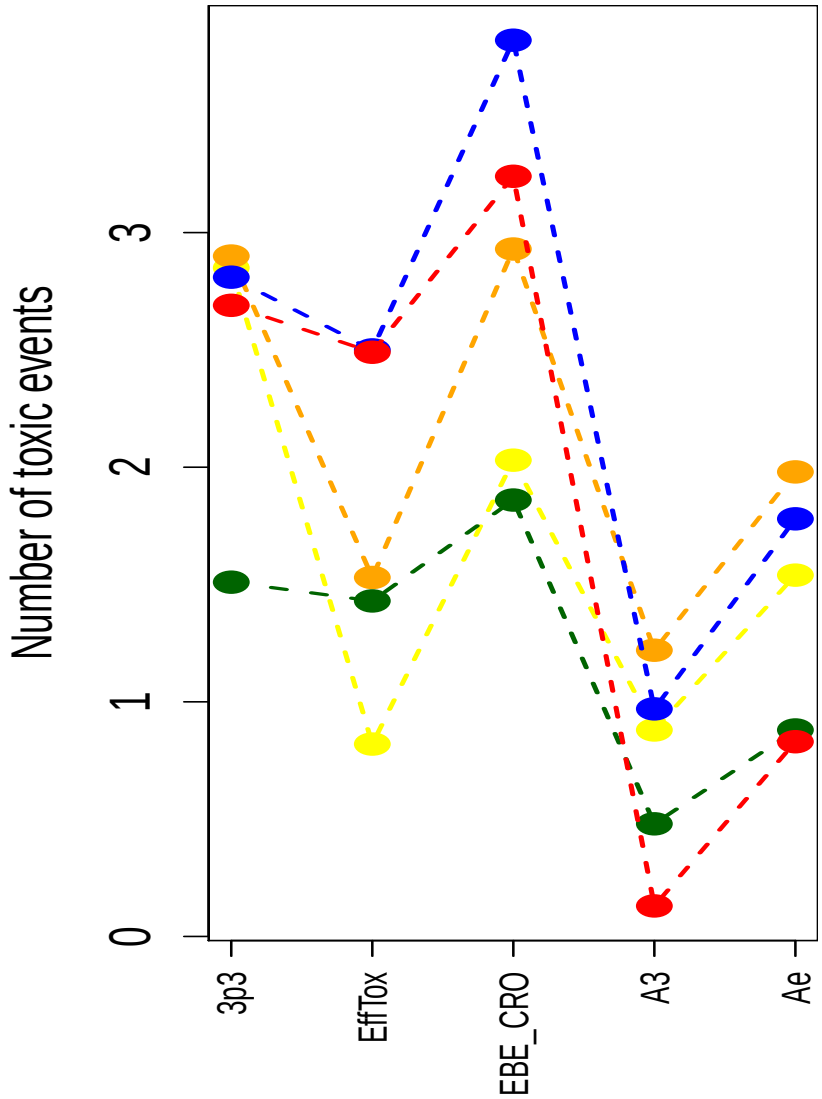

Conclusion

Advantages of globally optimal designs

- Provide a better overall cost of the clinical trial
- A better final dose recommendation & patient allocation

Global conclusions

- Operating characteristics depend mainly on the cost functions (both locally and globally optimal designs)
- Compromise between drug activity and toxicity risk

Drawbacks

- Globally optimal designs need a minimum prior information to perform well
- Complex to implement
- Utility to be chosed carefully

Perspectives on Phase II design (1/2)

Context

- Parallel groups of treatment (several dose levels)
- Interim analysis and choice of allocation ratio (based on efficacy)

Objective

- Optimize the allocation ratio at the interim analysis according to Phase 2 and Phase 3 potential results
- Cost functions
  - Immediate cost on the stage 2 of Phase 2 (can be null)
  - Final cost on the result of Phase 3 clinical trial

Perspectives on Phase II design (2/2)

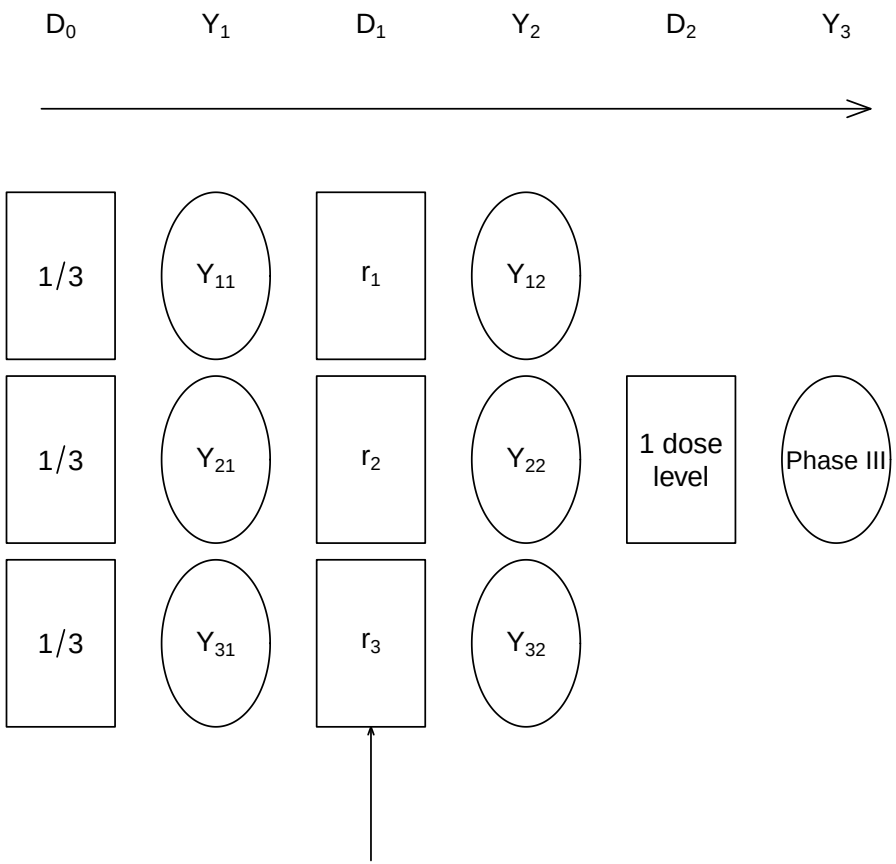

Main references

- Locally optimal designs
  - O'Quigley J. Continual reassessment method: a practical design for phase I clinical trials in cancer
  - Neuenschwander B. Critical aspects of the Bayesian approach to phase I cancer trials
- Globally optimal designs
  - Azriel D. The treatment versus experimentation dilemma in dose finding studies
  - Azriel D. Optimal sequential designs in phase I studies
- Proposed articles
  - (Published) Towards using a full spectrum of early clinical trial data: a retrospective analysis to compare potential longitudinal categorical models for molecular targeted, Statistics in Medicine therapies in oncology
  - (Submitted) An Escalation for Bivariate Binary Endpoints Controlling the Risk of Overtoxicity (EBE-CRO)

Main references

Thank you for your attention

Main references

BACKUP

Contents

7 References

8 Late & Repeated Toxicities

9 Globally optimal design

10 Decision theory

11 Utility function

12 Particule filter

13 Literature review

14 mTPI design

15 Standard designs

16 Cost functions

References on background (1/5)

Back

- Arrowsmith, John. Phase II failures: 2008–2010. Drug Discovery, 2011; 10:328-329. DOI: 10.1038/nrd3439
- Kola, Ismail and Landis, John. Can the pharmaceutical industry reduce attrition rates? Drug Discovery, 2004; 3:711-715. DOI: 10.1038/nrd1470
- Le Tourneau, Christophe *et al*. Current challenges for the early clinical development of anticancer drugs in the era of molecularly targeted agents. Targeted Oncology, 2010; 5:65-72. DOI: 10.1007/s11523-010-0137-6
- Postel-Vinay, Sophie *et al*. Phase I trials of molecularly targeted agents : should we pay more attention to late toxicities ? Journal of Clinical Oncology, 2011; 29:1728-1735. DOI: 10.1200/JCO.2010.31.9236

References on algorithmic designs (2/5)

Back

- Storer, B. Design and analysis of phase I clinical trials. Biometrics, 1989; 45:925-937. http://www.jstor.org/stable/2531693
- Skolnik, Jeffrey M. *et al*. Shortening the timeline of pediatric phase I trials : the rolling six design. Journal of Clinical Oncology, 2008; 26:190-195. DOI: 10.1200/JCO.2007.12.7712
- Ji, Yuan, Liu, Ping, Li, Yisheng and Bekele, Nbiyou. A modified toxicity probability interval method for dose-finding trials. Clinical Trials, 2010; 7:653-663. DOI: 10.1177/1740774510382799
- Simon, Richard *et al*. Accelerated titration designs for phase I clinical trials in oncology. Journal of the National Cancer Institute, 1997; 89:1138-1147. DOI: 10.1093/jnci/89.15.1138

References on locally optimal designs (3/5)

Back

- Babb, James *et al*. Cancer phase I clinical trials: efficient dose escalation with overdose control. Statistics in Medicine, 1998; 17:1103–1120.
- Bretz, Frank *et al*. Practical considerations for optimal designs in clinical dose finding studies. Statistics in Medicine, 2010; 29:731-742. DOI: 10.1002/sim.3802
- Chaloner, Kathryn and Larntz, Kinley. Optimal bayesian design applied to logistic regression experiments. Journal of Statistical Planning and Inference, 1989; 21:191-208. DOI: 10.1016/0378-3758(89)90004-9
- Neuenschwander, Beat *et al*. Critical aspects of the bayesian approach to phase I cancer trials. Statistics in Medicine, 2008; 27:2420-2439
- O'Quigley, John *et al*. Continual reassessment method: a practical design for phase I clinical trials in cancer. Biometrics, 1990; 46:33-48

References on gobally optimal designs (4/5)

Back

- Azriel, D., Mander, M. and Rinott, Y. The treatment versus experimentation dilemma in dose finding studies. Journal of Statistical Planning and Inference , 2011; 141:2759-2768. DOI: 10.1016/j.jspi.2011.03.001
- Azriel, David. Optimal sequential designs in phase I studies. Computational Statistics & Data Analysis , 2014; 71:288-297. DOI: 10.1016/j.csda.2013.05.010
- Bartroff, Jay and Lai, Tze Leung. Incorporating individual and collective ethics into phase I cancer trial designs. Biometrics, 2011; 67:596-603. DOI: 10.1111/j.1541-0420.2010.01471.x
- Bellman, Richard. Dynamic programming. 1957. Princeton University Press
- Bertsekas, Dimitri P. Dynamic programming and optimal control. 2000. Athena Scientific

References on modeling (5/5)

Back

- Albert, James H. and Chib, Siddhartha. Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association, 1993; 88:669-679. http://www.jstor.org/stable/2290350

Problèmes soulevés par les données

Back

- Toxicité
  - Données catégorielles ordonnées longitudinales (Gr0 à Gr4)
  - Événements tardifs, répétés ou persistants
  - Toxicités concentrées sur peu de patients (26/106)
  - Plusieurs types de toxicités (critère multivarié)
    - Neuropathies périphériques
    - Toxicités oculaires
  - Temps d'apparition variable (jusqu'à 4 mois après le début du traitement)
- Activité
  - Mesures répétées très espacées
  - 2 types d'observations (catégorielles et continues)

Questions statistiques posées

Back

- Comment évaluer le niveau de toxicité/d'activité (grade binaire ou catégoriel)?
- Combien de temps (de cycles de traitement) doit-on évaluer la réponse toxique et/ou active?
- Comment prendre en compte la durée des événements (non négligeable par rapport à la durée de l'étude)?
- Faut-il prendre en compte une corrélation entre les événements?

Modèle probit catégoriel multivarié (1/5)

Back

Données (unique par patient)

- X_i: Données disponibles pour le patient i
- Y_i: Données catégorielles observées chez le patient i

Modèle multivarié

$$Y_i = \sum_{k=1}^{K-1} k \times \mathbb{1}(c_k \leq Z_i < c_{k+1})$$
$$Z_i = X_i\beta + \varepsilon_t$$
$$\varepsilon_t \sim \mathcal{N}(0, \Sigma)$$

Cadres d'inférence

- Algorithme de Gibbs
- Algorithme Espérance-Maximisation

Modèle probit catégoriel multivarié (2/5)

Back

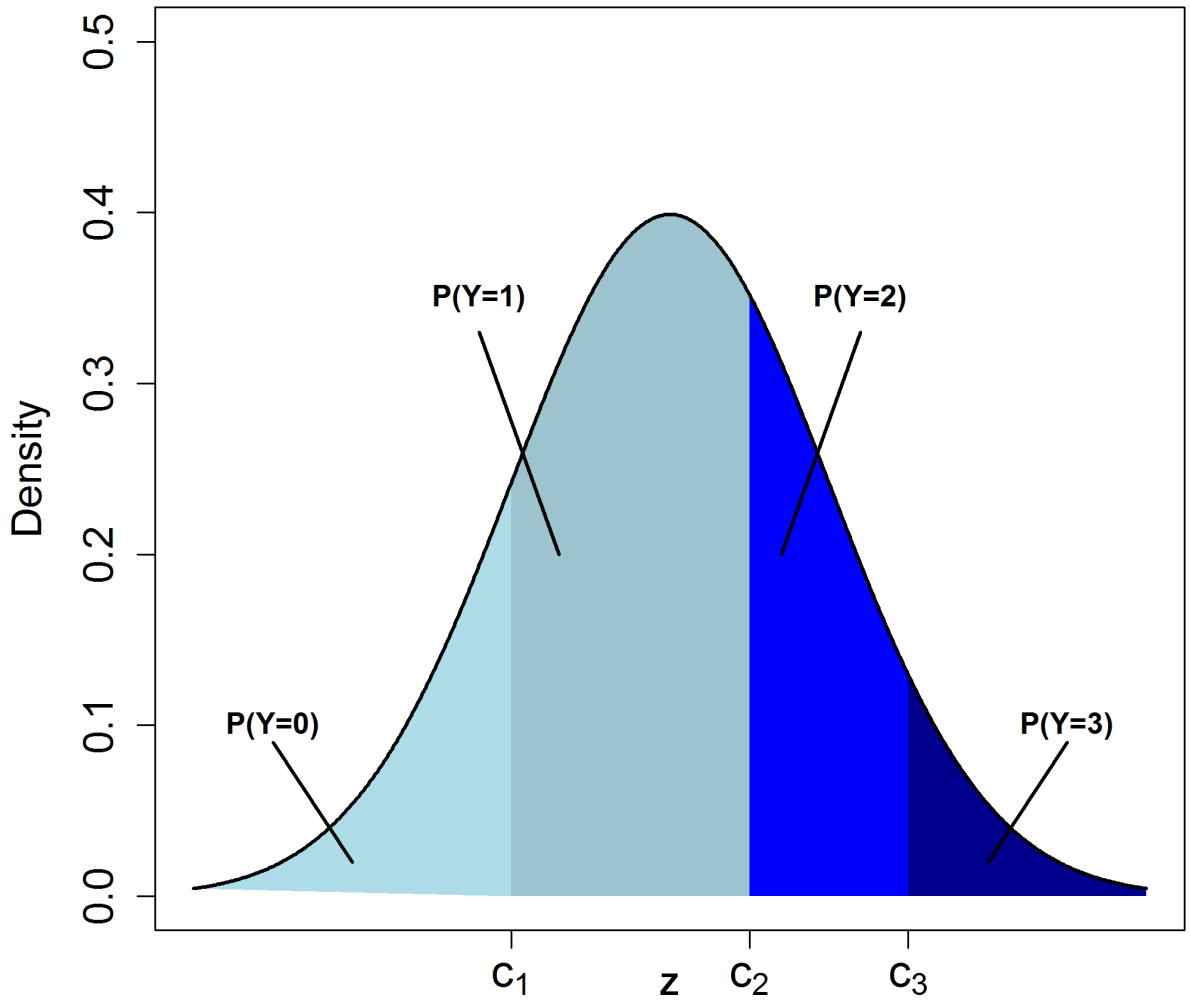

Modèle probit catégoriel multivarié - Gibbs (3/5)

Back

Principe

$$\begin{aligned}\text{Loi jointe} &= \text{Vraisemblance complète} \times \text{Lois a priori} \\ &= \text{Loi conditionnelle complète} \times \text{Loi marginale}\end{aligned}$$

Tirages successifs dans les lois conditionnelles complètes

- $[Z|Y, \beta, C, \Sigma] \sim$ Normale tronquée
- $[\beta|Y, Z, C, \Sigma] \sim$ Normale
- $[C|Y, Z, \beta, \Sigma] \sim$ Normale tronquée
- $[\Sigma^{-1}|Y, Z, \beta, C] \sim$ Wishart

Avantages

- Conjugaison des lois normales et Wishart
- Extension de paramètre pour déterminer Σ

Modèle pr

Back

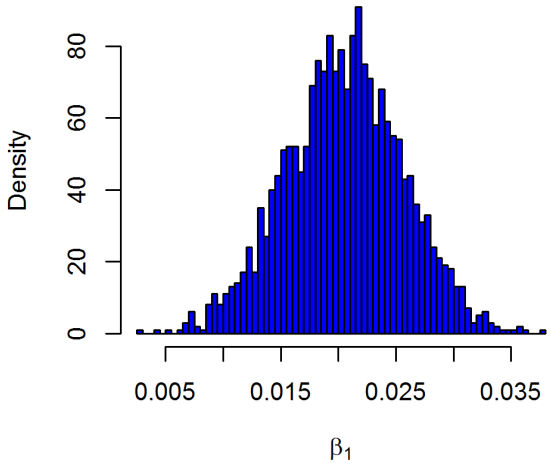

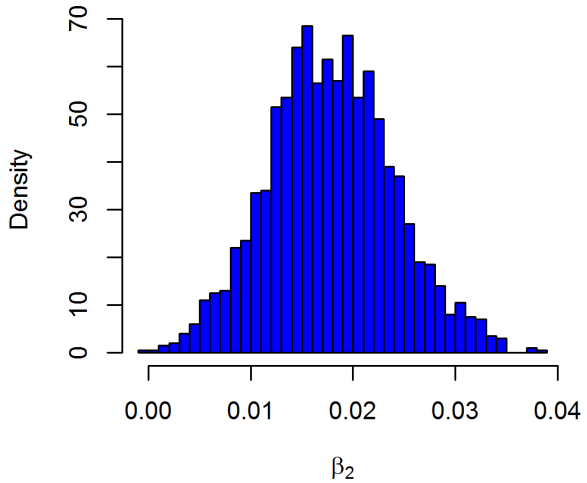

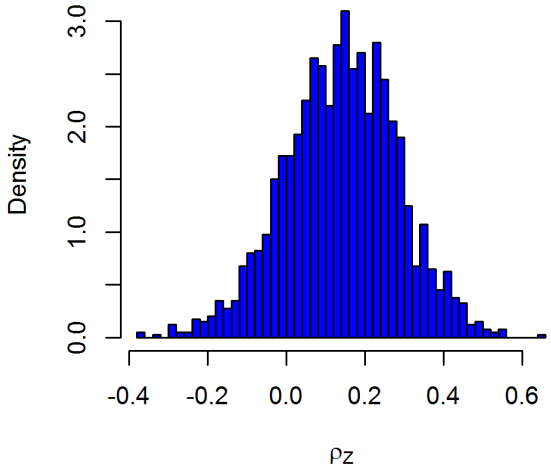

β_1 : paramètre de régression des toxicités oculaires

β_2 : paramètre de régression des neuropathies

ρ_Z : paramètre de corrélation

Modèle probit catégoriel multivarié - Gibbs (5/5)

Back

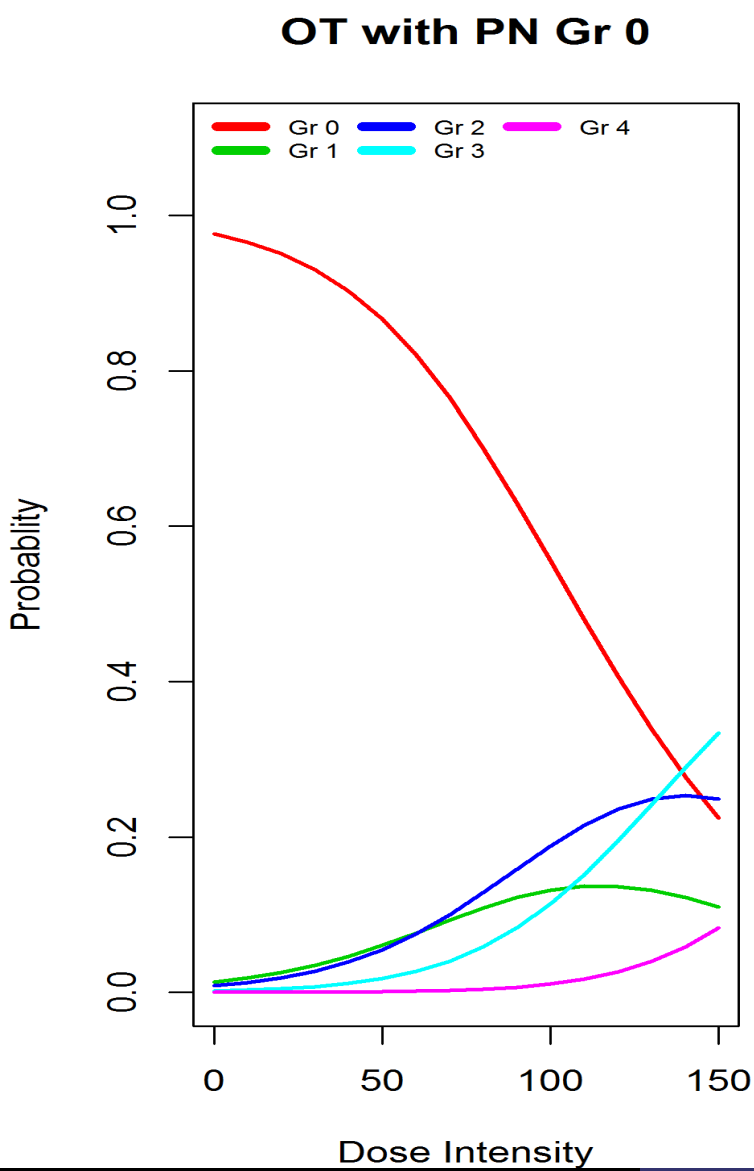

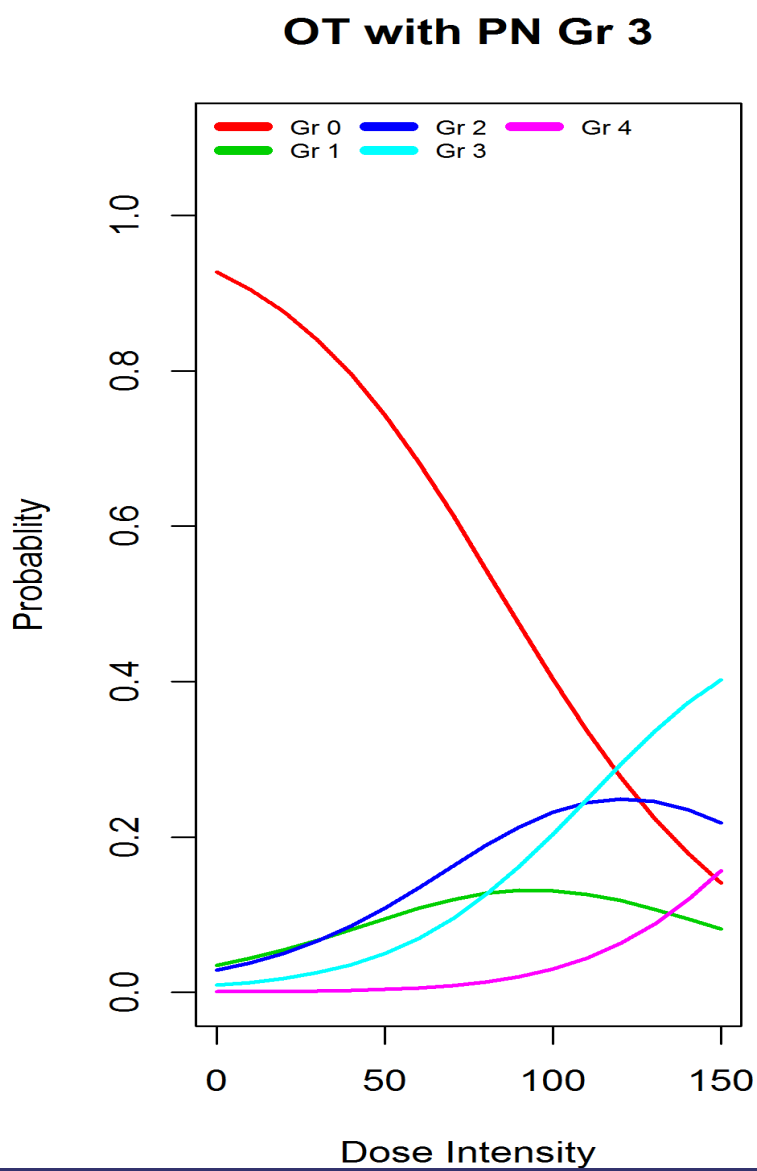

Modèle probit catégoriel temporel univarié

Back

Données

- $X_{i,t}$: Données disponibles pour le patient i au temps t
- $Y_{i,t}$: Données catégorielles observées chez le patient i au temps t

Modèle temporel

$$Y_{i,t} = \sum_{k=1}^{K-1} k \times \mathbb{1}(c_k \leq Z_{i,t} < c_{k+1})$$

$$Z_{i,t} = X_{i,t}\beta + \rho \times \left(Z_{i,t-1} - X_{i,t-1}\beta\right) + \sqrt{1-\rho^2} \times \varepsilon_t$$

$$\varepsilon_t \sim \mathcal{N}(0,1)$$

Cadres d'inférence

- Algorithme de Gibbs
- Algorithme Espérance-Maximisation

Modèle probit catégoriel temporel univarié (Gibbs)

Back

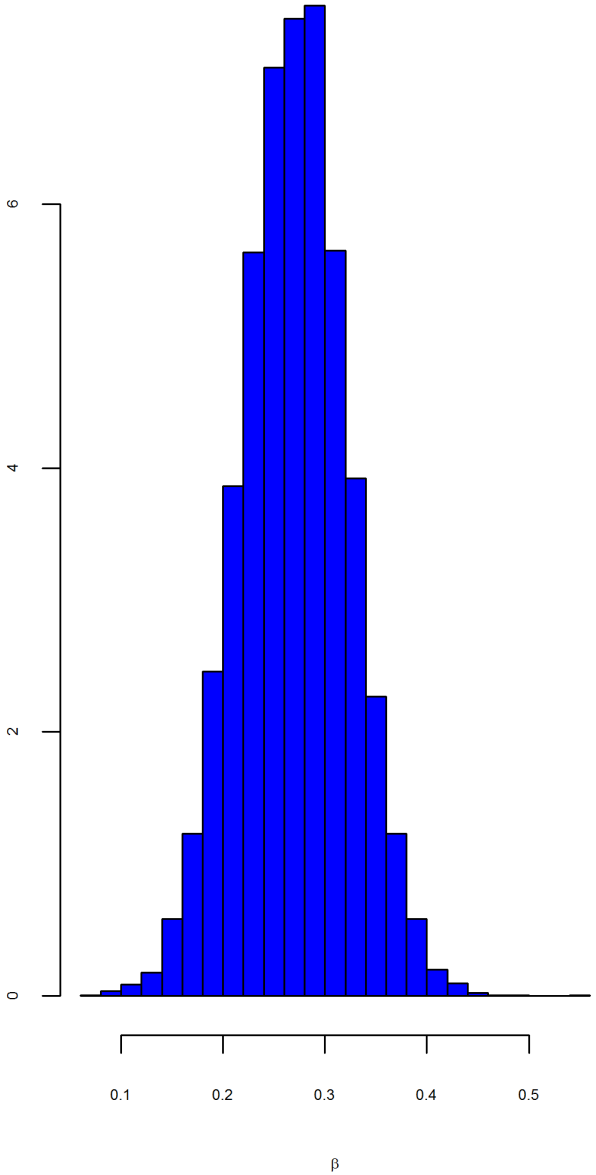

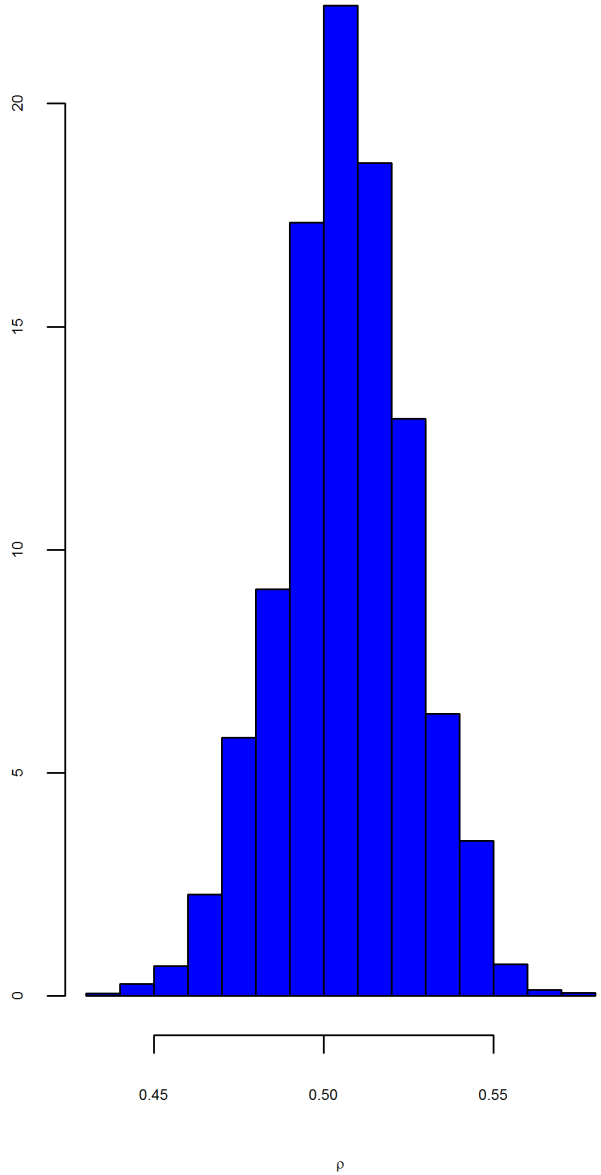

Modèle probit catégoriel temporel univarié (Gibbs)

Back

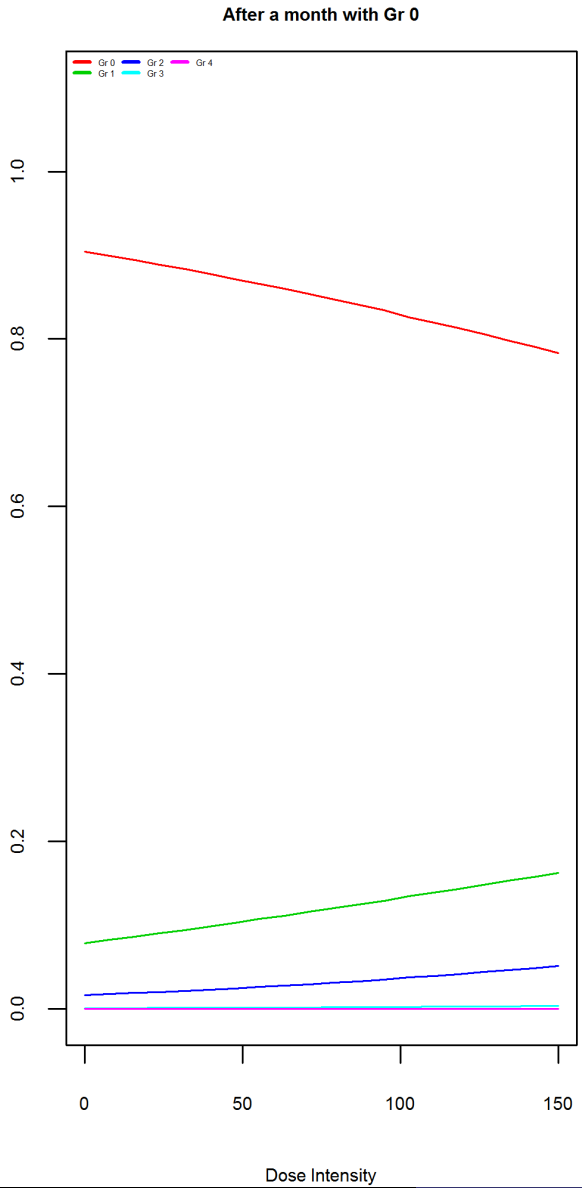

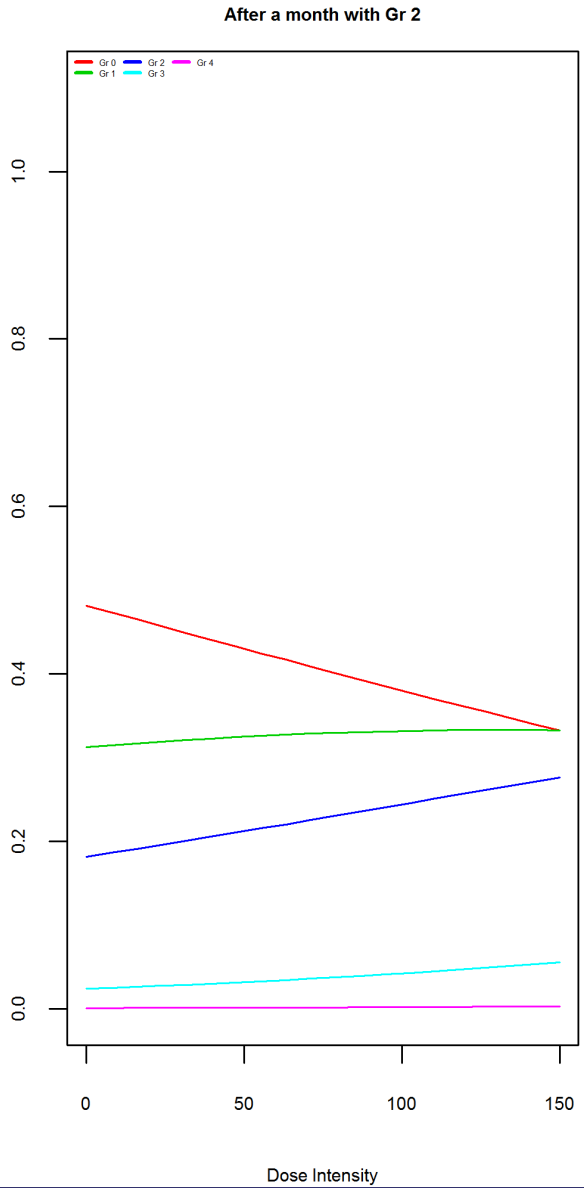

Modèle probit catégoriel temporel multivarié

Back

Données

- $X_{i,j,t}$: Données disponibles pour le patient i au temps t
- $Y_{i,j,t}$: Données catégorielles observées chez le patient i au temps t pour le critère j

Modèle multivarié

$$Y_{i,j,t} = \sum_{k=1}^{K-1} k \times \mathbb{1}(c_{j,k} \leq Z_{i,j,t} < c_{j,k+1})$$

$$Z_{i,j,t} = X_{i,j,t}\beta_j + \rho_j \times \left(Z_{i,j,t-1} - X_{i,j,t-1}\beta_j\right) + \sqrt{1-\rho_j^2} \times \varepsilon_{j,t}$$

$$\left(\varepsilon_{1,t}, \ldots, \varepsilon_{J,t}\right) \sim \mathcal{N}(0, \Sigma)$$

Cadres d'inférence

- Algorithme de Gibbs
- Algorithme Espérance-Maximisation

Modèle probit catégoriel temporel multivarié

Back

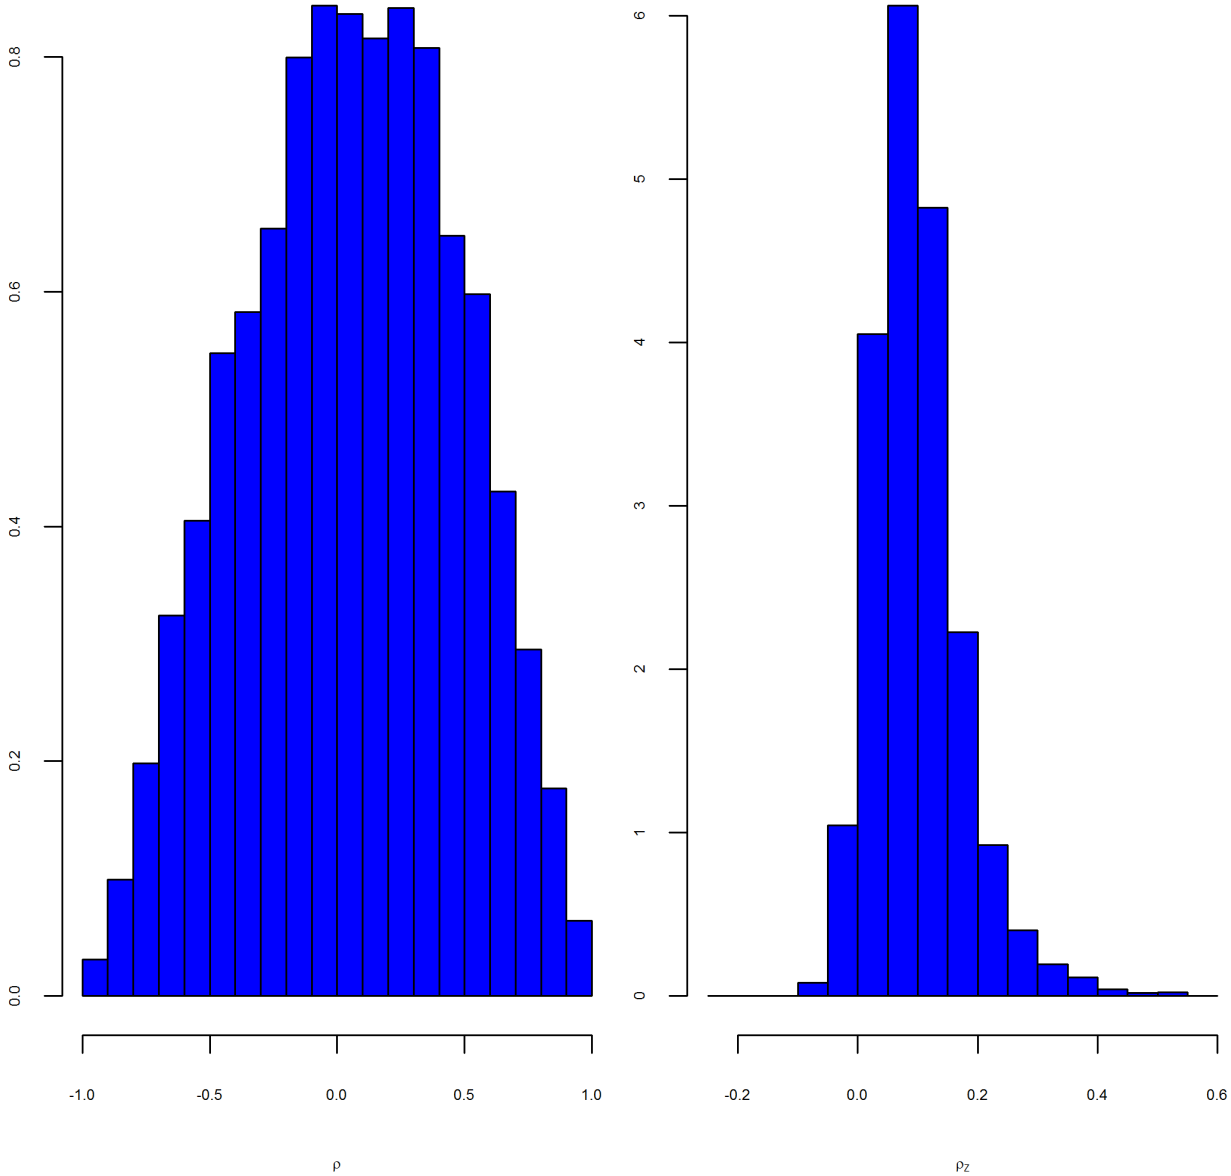

Common writing for GLM and Markov

Back

Notations

- Patient i monitored from week 1 to week T_i
- $Y_i = (Y_i^{(1)}, \ldots, Y^{(T_i)})$: Observed toxic grade of patient i
- $X_i = (X_i^{(1)}, \ldots, X^{(T_i)})$: Predictors and covariates of patient i
- θ: Model parameters

Common writing for GLM and Markov

Back

Notations

- Patient i monitored from week 1 to week T_i
- $Y_i = (Y_i^{(1)}, \ldots, Y^{(T_i)})$: Observed toxic grade of patient i
- $X_i = (X_i^{(1)}, \ldots, X^{(T_i)})$: Predictors and covariates of patient i
- θ: Model parameters

Likelihood

$$\begin{aligned}
[Y_i^{(1)}, \ldots, Y_i^{(T_i)} | \theta, X_i] &= [Y_i^{(1)} | \theta, X_i] \times [Y_i^{(2)}, \ldots, Y_i^{(T_i)} | Y_i^{(1)}, \theta, X_i] \\
&= [Y_i^{(1)} | \theta, X_i] \times [Y_i^{(2)} | Y_i^{(1)}, \theta, X_i] \times \ldots \\
&= \prod_{t=1}^{T_i} [Y_i^{(t)} | Y_i^{(t-1)}, \ldots, Y_i^{(1)}, \theta, X_i^{(t)}]
\end{aligned}$$

GLM with correlation

Back

Model for binary data

$$\theta = (\mu, \beta, \rho, \sigma)$$

$$\mathbb{P}(Y_i^{(t)} = 1|\theta, X_i^{(t)}) = \phi\left(\mu + X_i^{(t)}\beta\right)$$

$$\mathbb{V}(Y_i = 1|\theta, X_i) = \sigma^2 AR1(\rho)$$

$$AR1(\rho) = \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 \\ \rho & 1 & \rho & \rho^2 \\ \rho^2 & \rho & 1 & \rho \\ \rho^3 & \rho^2 & \rho & 1 \end{pmatrix}$$

Estimation performed through a GEE algorithm (Zeger 1986) - R (geepack)

Heterogeneous Markov model - Discrete time

Back

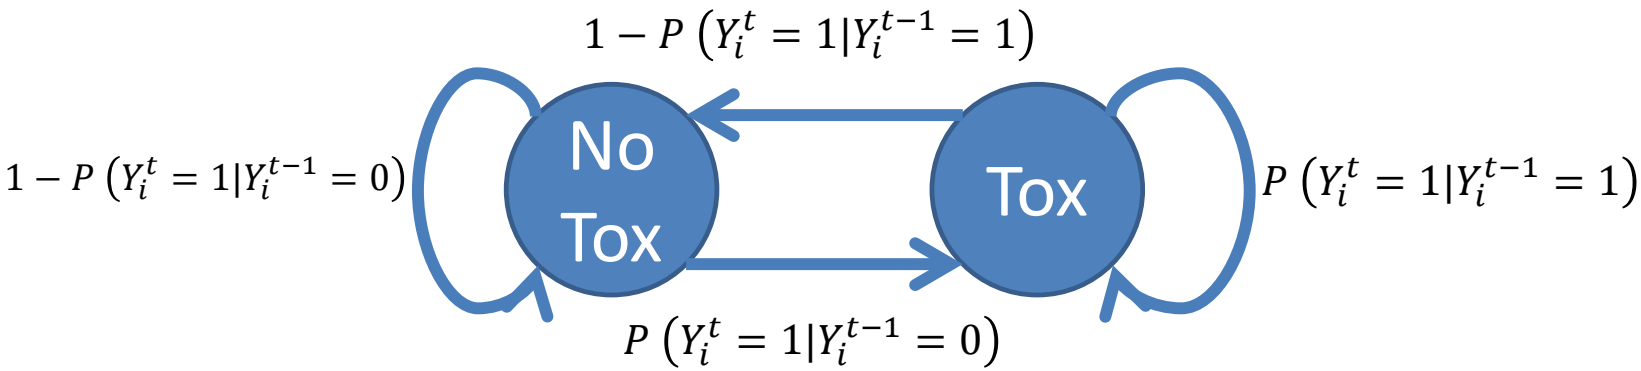

0 = No Tox $(ToxGrade \leq 1)$

1 = Tox $(ToxGrade \geq 2)$

Heterogeneous Markov model - Discrete time

Back

Markov hypothesis

- $[Y_i^{(1)}, \ldots, Y_i^{(T_i)} | \theta, X_i] = \prod_{t=1}^{T_i} [Y_i^{(t)} | Y_i^{(t-1)}, \theta, X_i^{(t)}]$
- $\theta = (\mu, \beta, \beta_{pY})$
- Equal slopes (β) does not depend on $Y_i^{(t-1)}$ (avoid misinterpretation)

Heterogeneous Markov model - Discrete time

Back

Markov hypothesis

- $[Y_i^{(1)}, \ldots, Y_i^{(T_i)} | \theta, X_i] = \prod_{t=1}^{T_i} [Y_i^{(t)} | Y_i^{(t-1)}, \theta, X_i^{(t)}]$
- $\theta = (\mu, \beta, \beta_{pY})$
- Equal slopes (β) does not depend on $Y_i^{(t-1)}$ (avoid misinterpretation)

Probit model for binary data

$$\mathbb{P}(Y_i^{(t)} = 1 | Y_i^{(t-1)} = 0, \theta, X_i^{(t)}) = \phi\left(\mu + X_i^{(t)}\beta\right)$$

$$\mathbb{P}(Y_i^{(t)} = 1 | Y_i^{(t-1)} = 1, \theta, X_i^{(t)}) = \phi\left(\mu + X_i^{(t)}\beta + \beta_{pY}\right)$$

Estimation performed by likelihood maximization - R (glm)

Arche

Back

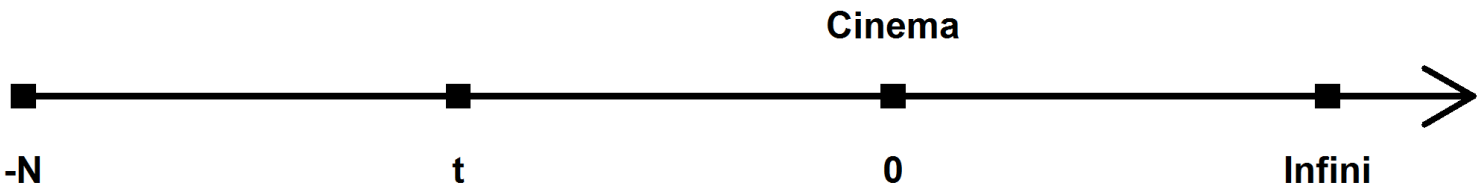

Structure du problème

- Probabilité d'avoir une place libre: loi de Bernoulli
- À chaque étape
  - Places de parking précédentes observées (ou loi *a priori*)
  - Plusieurs observations possibles: place libre ou occupée
  - Plusieurs décisions possibles: se garer ou continuer
  - État des connaissance x_t, Décision u_t

Archetype de décision séquentielle stochastique

Back

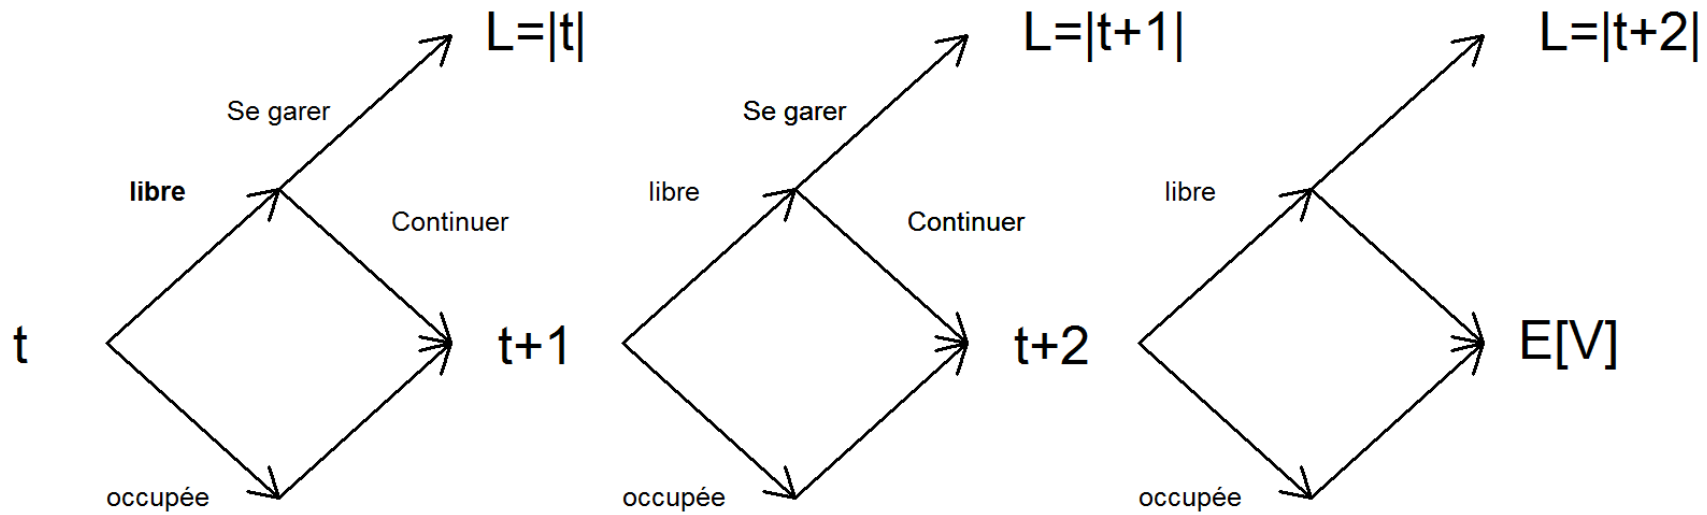

Archetype de décision séquentielle stochastique

Back

Modèle probabiliste

Probabilité d'avoir une place libre $p \sim \mathcal{B}eta(\alpha, \beta)$

Fonction de coût immédiat

$$L(x_t, u_t, t) = \begin{cases} |t| & \text{si } y_t = \text{Place libre } \& \ u_t = \text{Se garer} \\ 0 & \text{sinon} \end{cases}$$

Équation de Bellman

$$V(x_t, t) = \min_{u_t} \left\{ L(x_t, u_t, t) + \frac{\alpha_t}{\alpha_t + \beta_t} V(X^1_{t+1}, t+1) + \frac{\beta_t}{\alpha_t + \beta_t} V(X^0_{t+1}, t+1) \right\}$$

Archetype de décision séquentielle stochastique

Back

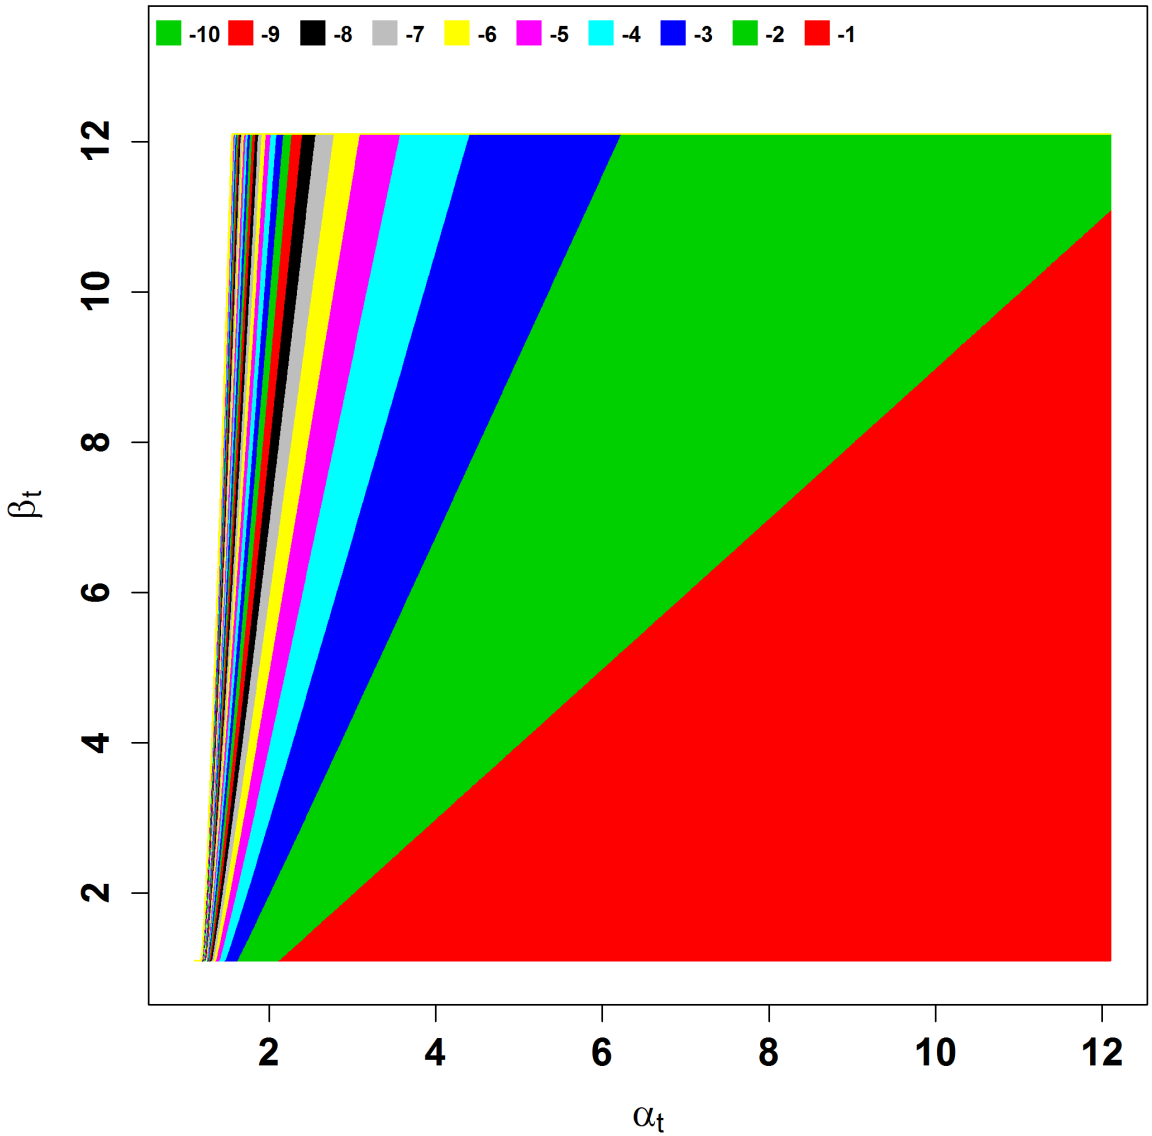

Problème de l'embauche

Back

Source

- Parmigiani G. & Inoue L. Décision Theory - Principles and Approaches
  - Chapter 12 - Dynamic programming[a]

[a] Également appelé "The secretary problem", "Beauty-contest problem" ou "Keynesian beauty contest"

Structure du problème

- Probabilité d'avoir un candidat de rang plus élevé
- À chaque étape
  - Rang relatif par rapport aux candidats précédents
  - Total de N candidats
  - Plusieurs décisions possibles: embauche ou refus définitif

Problème de l'embauche

Back

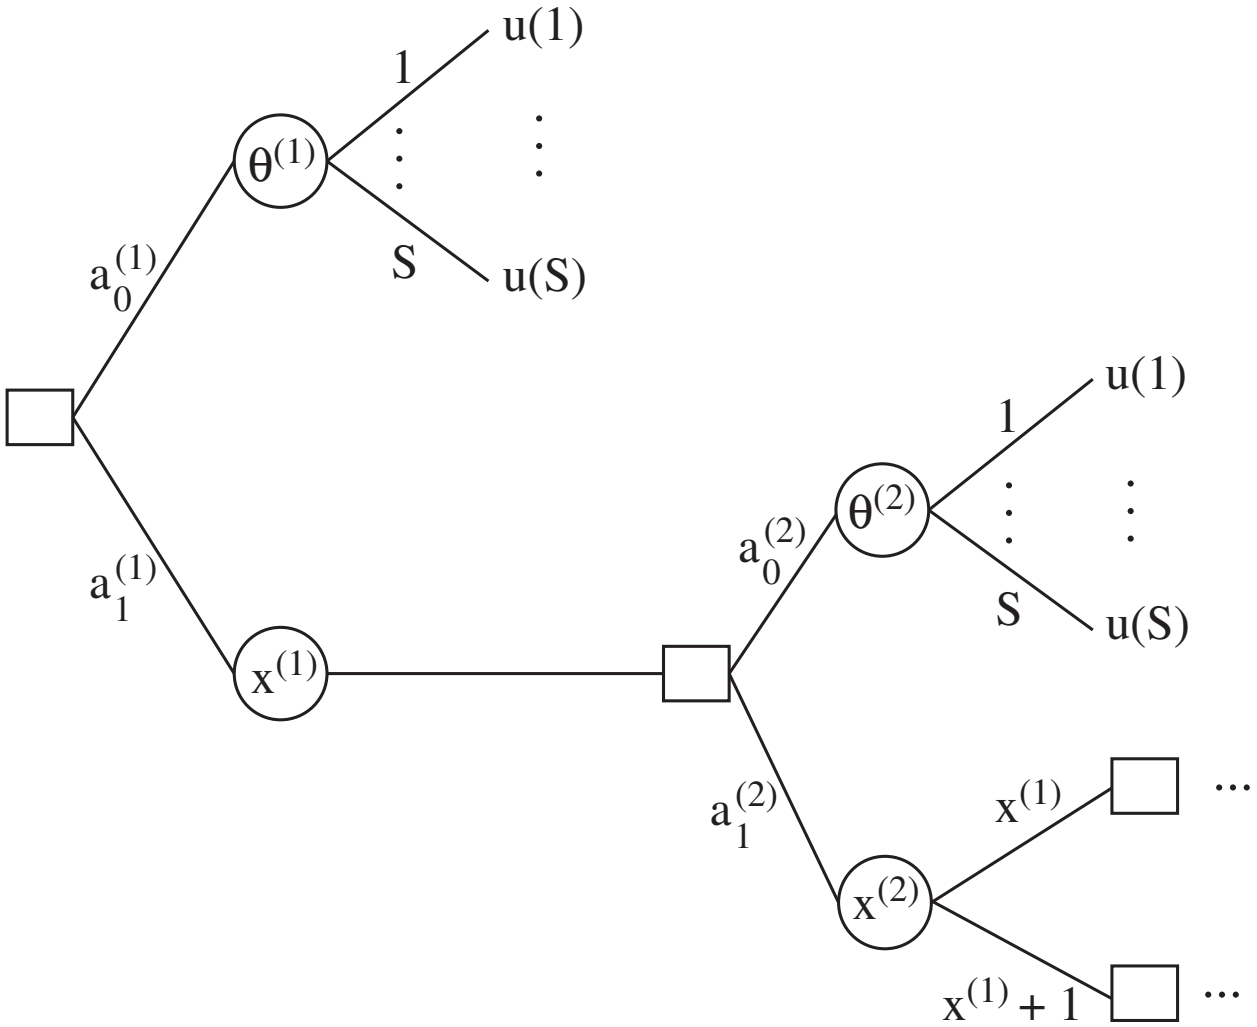

Problème de l'embauche

Back

$$\pi(\theta^{(s)}|x^{(s)}) = \binom{\theta^{(s)}-1}{x^{(s)}-1}\binom{S-\theta^{(s)}}{s-x^{(s)}}\binom{S}{s}$$

$$\mathcal{U}(a_0^{(s)}) = \sum_{\theta=x^{(s)}}^{S-s+x^{(s)}} u(\theta)\binom{\theta-1}{x^{(s)}-1}\binom{S-\theta}{s-x^{(s)}}\binom{S}{s}$$

$$\mathcal{U}^*(a_1^{(s)}) = b(s,x^{(s)})$$

$$= \frac{1}{s+1}\sum_{x=1}^{s+1} b(s+1,x)$$

Problème de l'embauche

Back

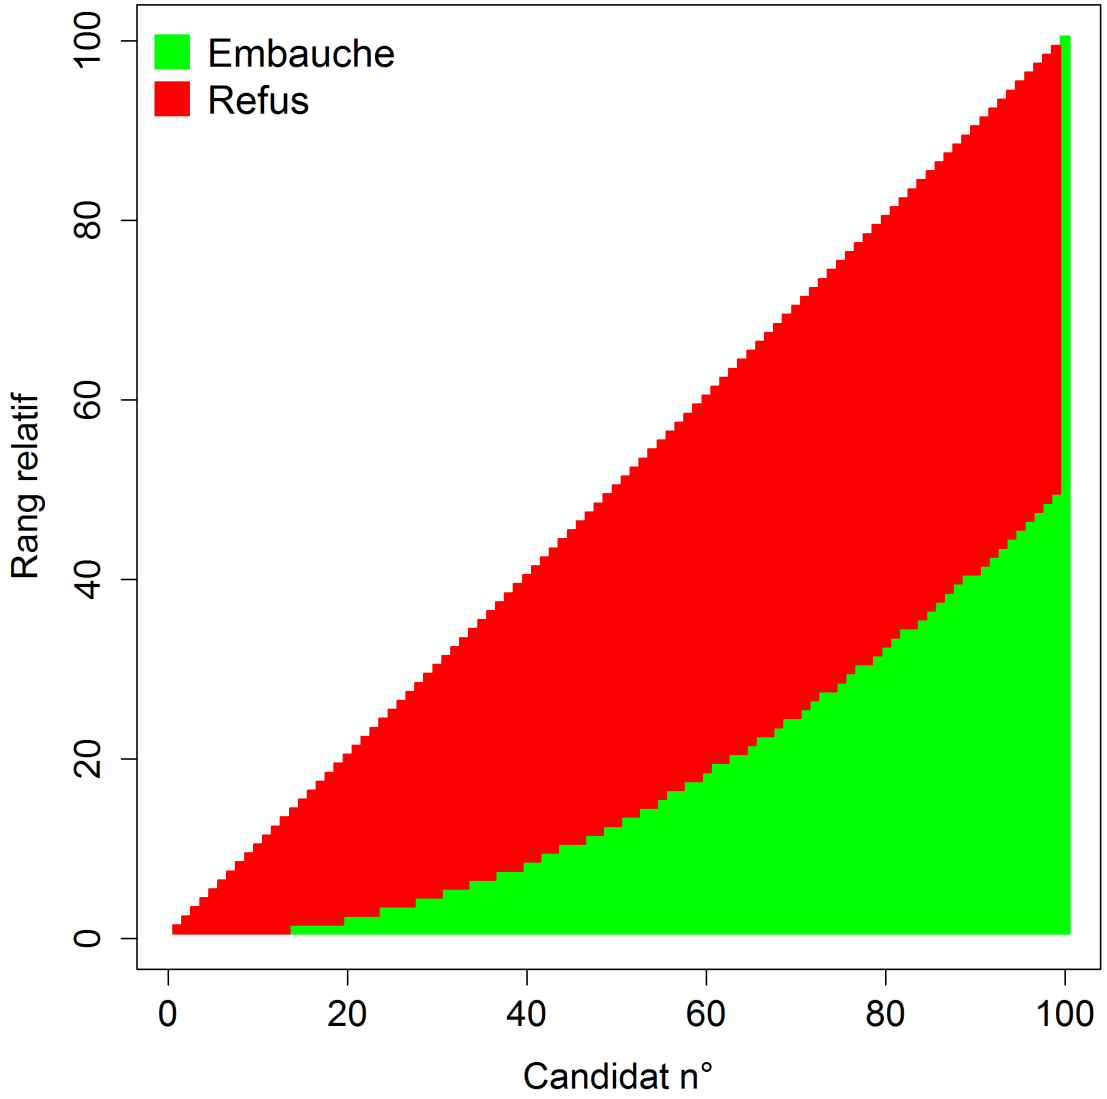

Clinical trial: Adaptive allocation

Back

Source

Parmigiani G. & Inoue L. Décision Theory - Principles and Approaches. Chapter 12 - Dynamic programming

Context

- 2 treatments A and B
- Prior $(\theta_A, \theta_B) \sim \mathcal{B}eta(1,1)$
- 4 experiments remaining
- 96 identical experiments to plan
- Actions at each step $\{a_A^{(k)}, a_B^{(k)}\}$

Decision process

$$\mathcal{U}\left(a_i^{(n)}\right) = \mathbb{E}\Big[\text{Nb of cured patients}\Big]$$

Clinical trial: Adaptive allocation

Back

Available information

- Prior $(\theta_A, \theta_B) \sim \mathcal{B}eta(1,1)$
- Resp. 2 and 1 patients
- Resp. 1 and 0 responses
- Posterior $\theta_A \sim \mathcal{B}eta(2,2)$, $\theta_B \sim \mathcal{B}eta(1,2)$

Clinical trial: Adaptive allocation

Back

Available information

- Prior $(\theta_A, \theta_B) \sim \mathcal{B}eta(1,1)$
- Resp. 2 and 1 patients
- Resp. 1 and 0 responses
- Posterior $\theta_A \sim \mathcal{B}eta(2,2)$, $\theta_B \sim \mathcal{B}eta(1,2)$

If A is chosen

$$x = 1 \Rightarrow u(A) = 1 + 96 \times \frac{3}{5} = 58.6$$

$$x = 0 \Rightarrow u(A) = 0 + 96 \times \frac{2}{5} = 38.4$$

$$U(A) = 58.6 \times \frac{1}{2} + 38.4 \times \frac{1}{2} = 48.5$$

Clinical trial: Adaptive allocation

Back

Available information

- Prior $(\theta_A, \theta_B) \sim \mathcal{B}eta(1,1)$
- Resp. 2 and 1 patients
- Resp. 1 and 0 responses
- Posterior $\theta_A \sim \mathcal{B}eta(2,2)$, $\theta_B \sim \mathcal{B}eta(1,2)$

If B is chosen

$$x = 1 \Rightarrow u(B) = 1 + 96 \times \frac{1}{2} = 49$$

$$x = 0 \Rightarrow u(B) = 0 + 96 \times \frac{1}{2} = 48$$

$$U(B) = 48.5 \times \frac{1}{3} + 48 \times \frac{2}{3} = 48.33$$

Clinical trial: Adaptive allocation

Back

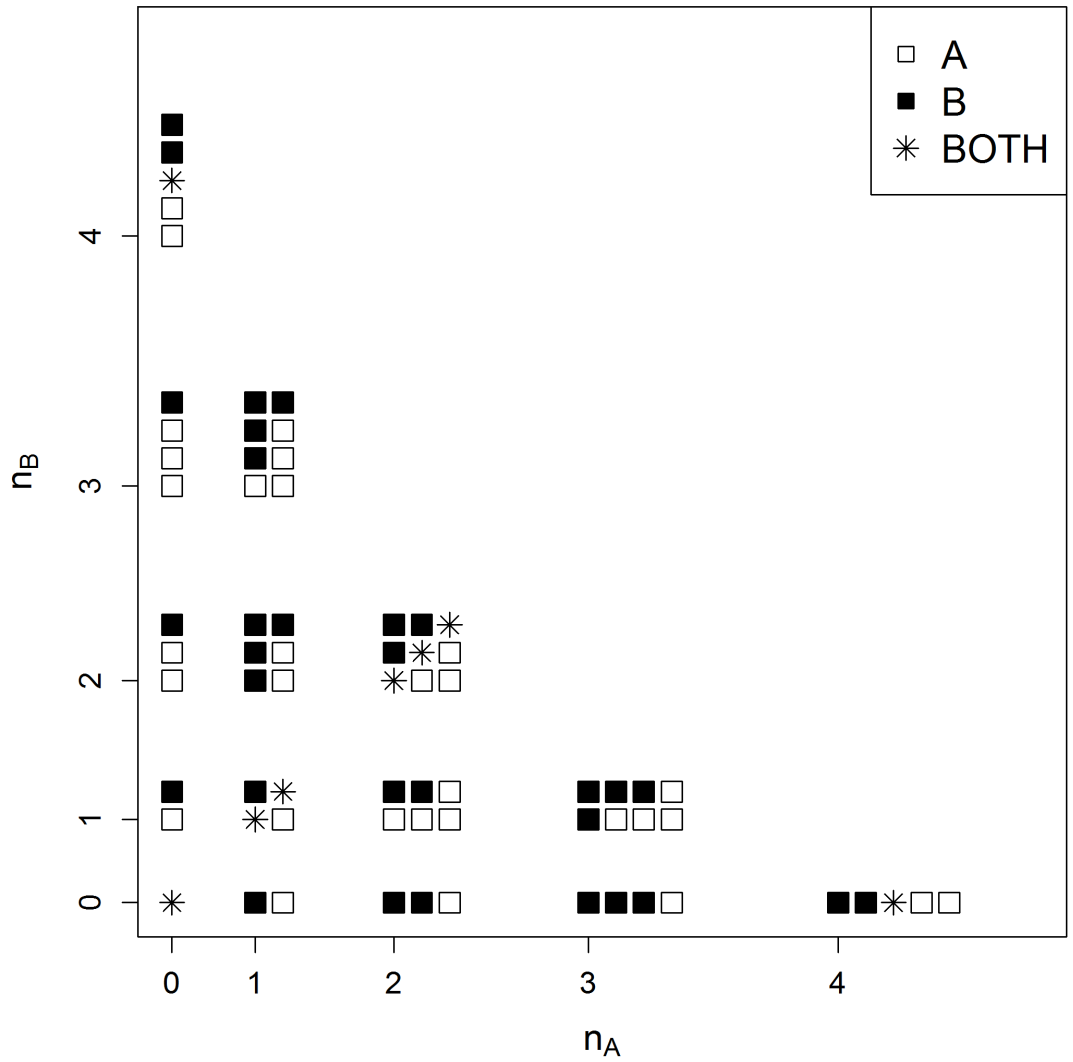

Clinical trial: Adaptive allocation

Back

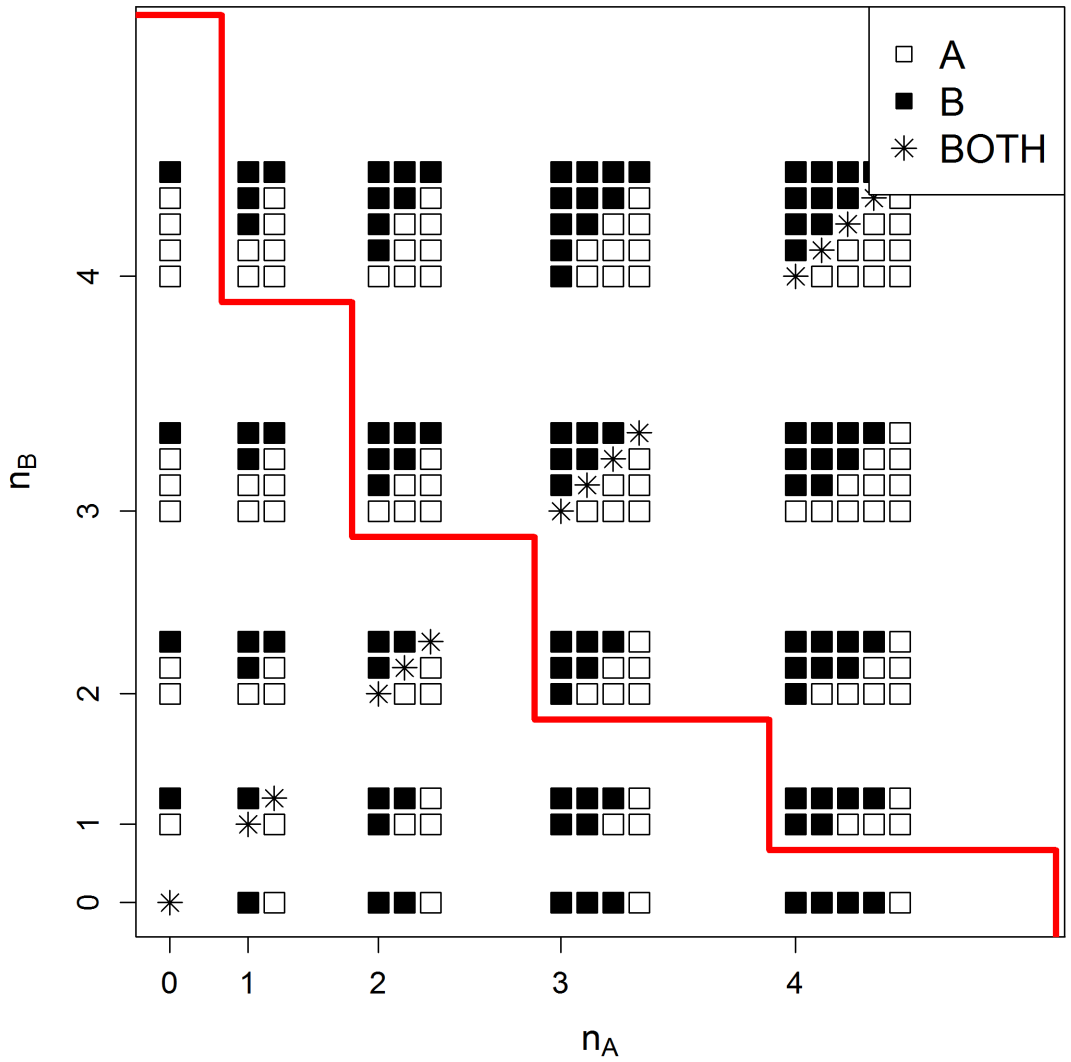

Clinical trial: Adaptive allocation

Back

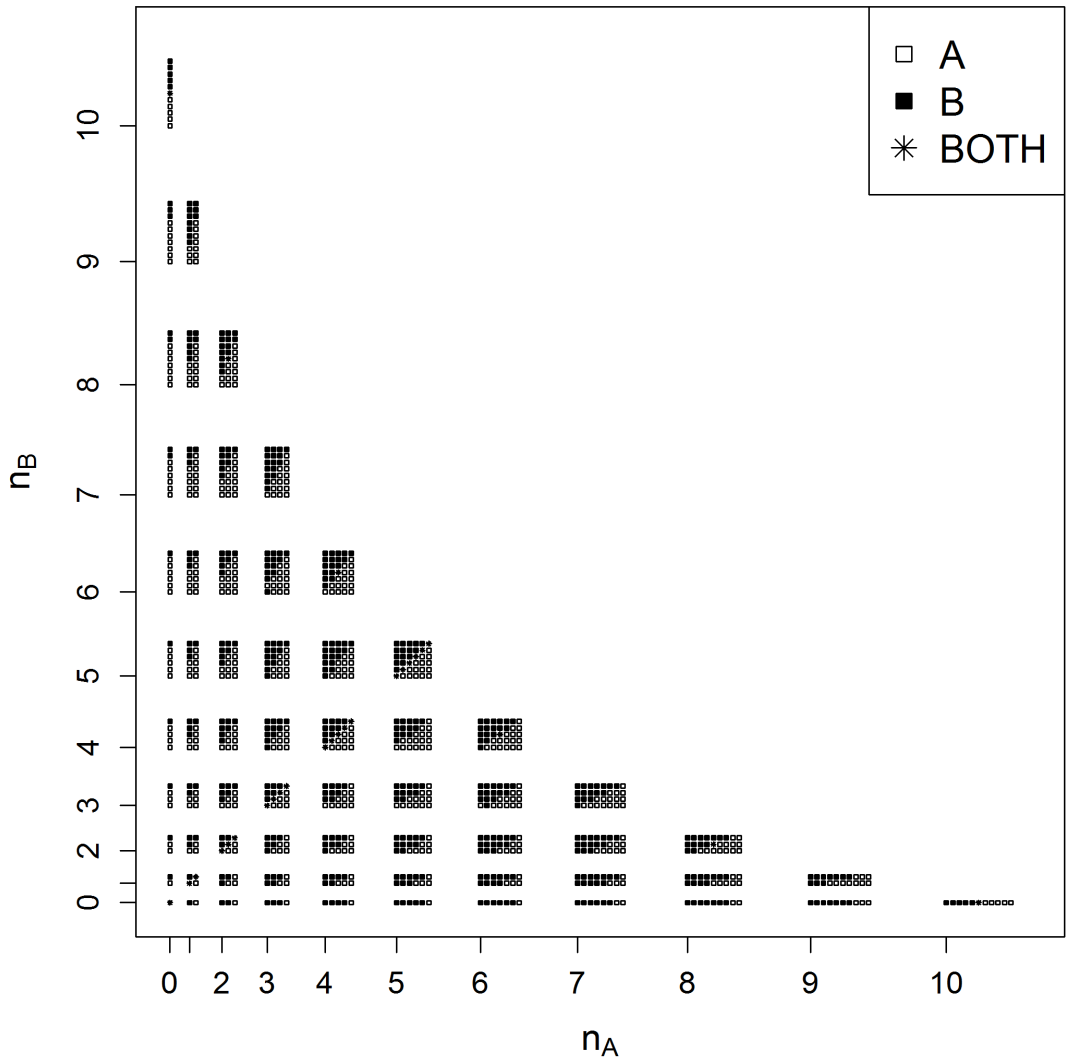

Curse of dimension

Back

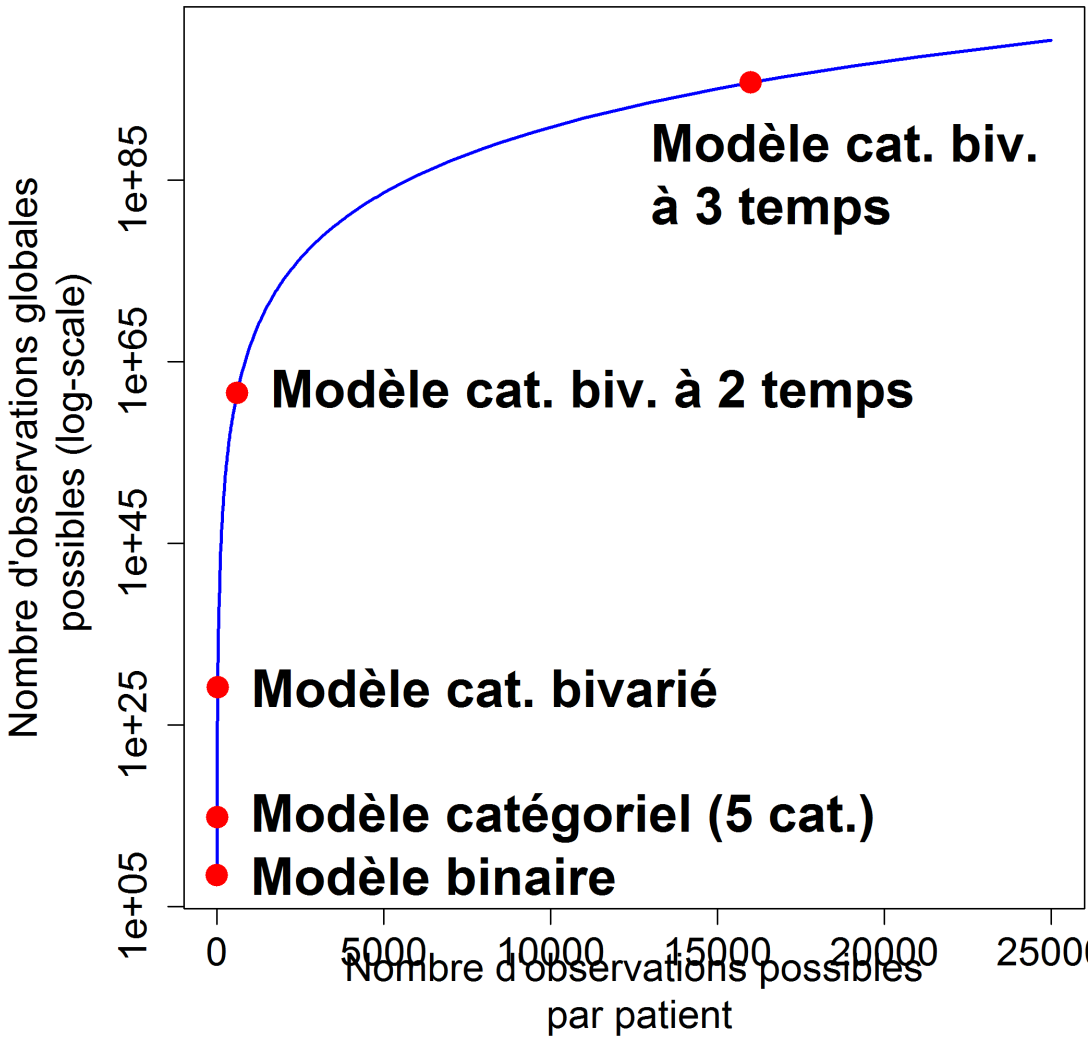

Groupe de lecture AgroParisTech - Parmigiani & Inoue

Back

Reference

- Parmigiani, G. & Inoue, L. Decision theory: principles and approaches. Wiley Series in Probability and Statistics, 2009

Lectures

- Utility in action - P. Colin
- Changes in utility as information - P. Colin

Fonction d'utilité Activité-Toxicité

Back

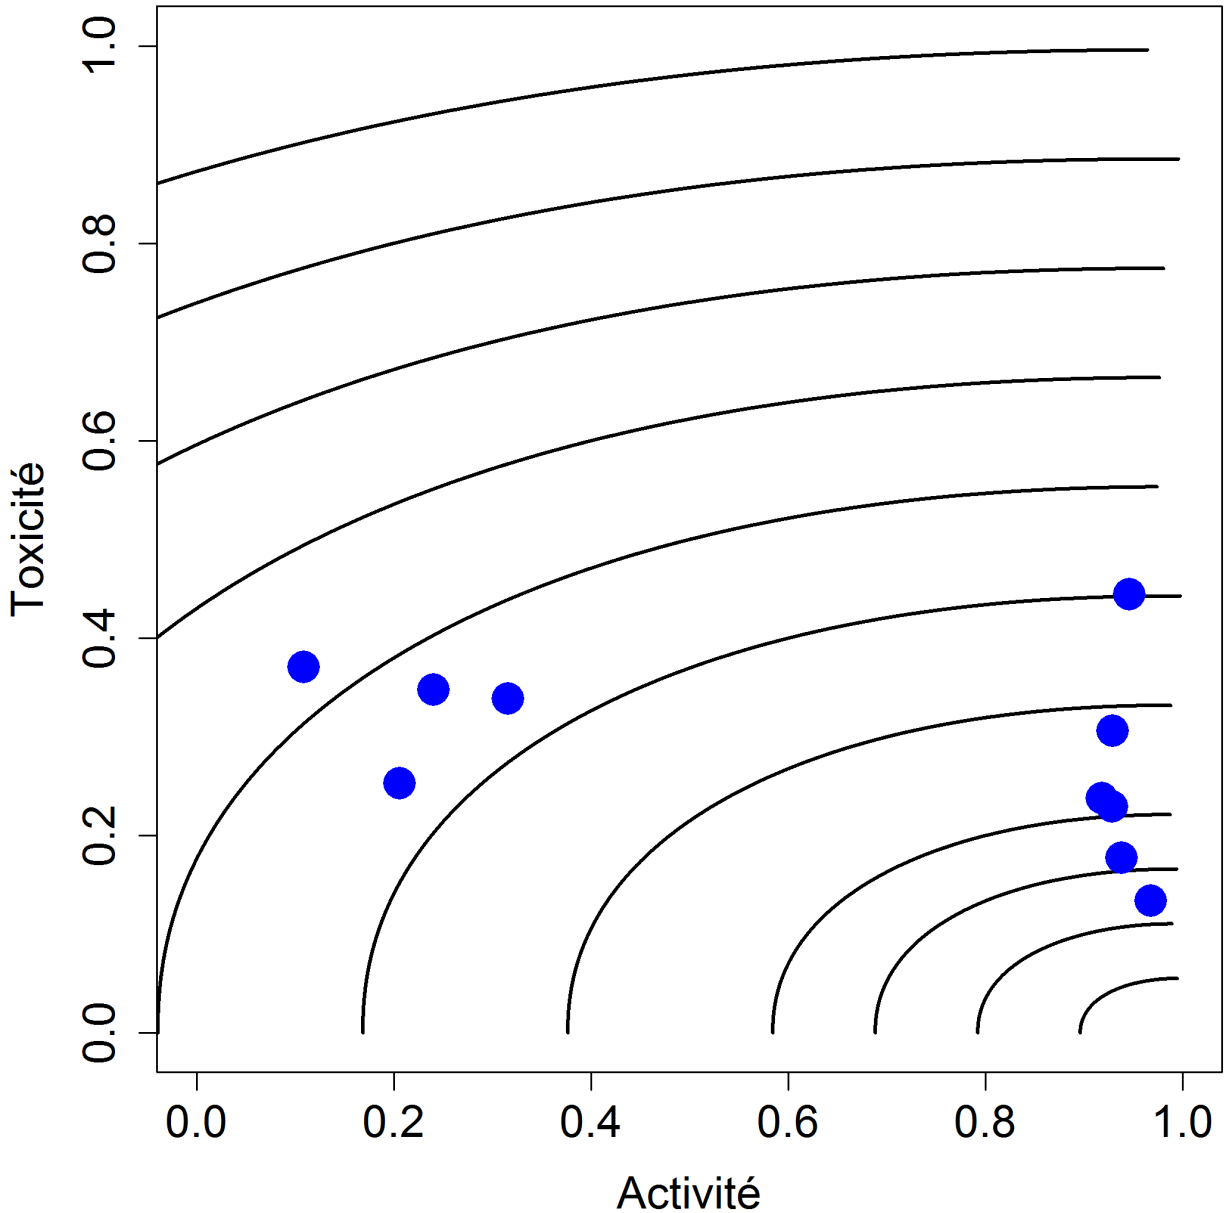

Fonction d'utilité Activité-Toxicité

Back

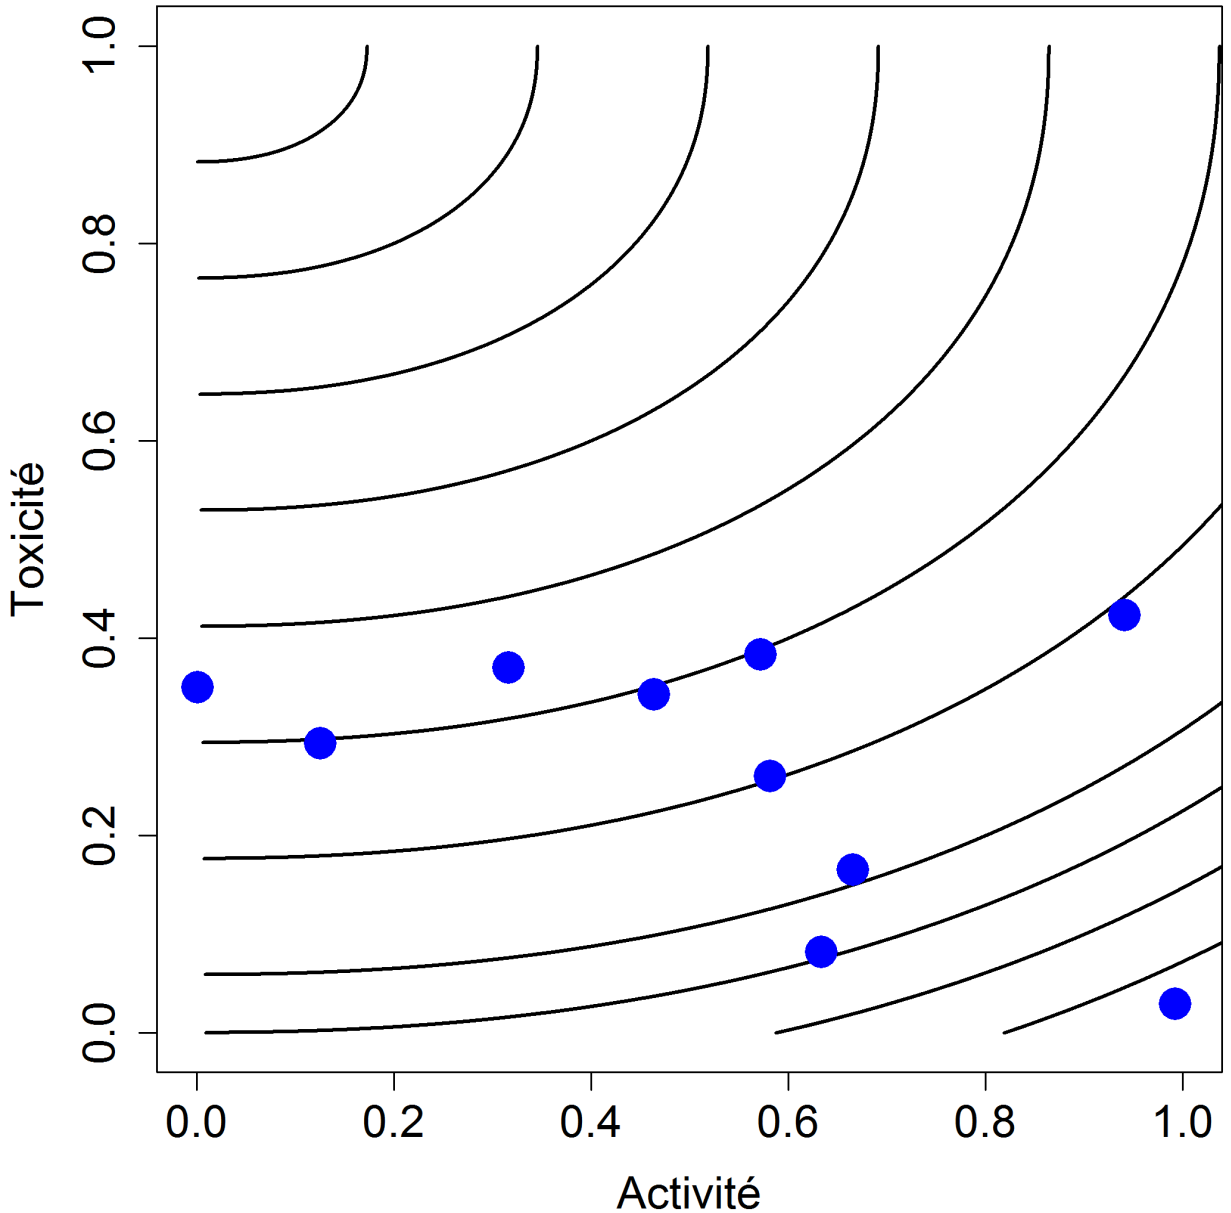

Particle filter - Principle

Back

$$w_i = \frac{\mathcal{L}(\mathbf{y}_1, \theta^{(i)}) \times f(\theta^{(i)})}{g(\theta^{(i)})}$$

où $\mathcal{L}(\mathbf{y}_1, \theta^{(i)})$ est la vraisemblance de $\mathbf{y}_1$ sachant $\theta = \theta^{(i)}$. f est la fonction de densité de la loi *a priori* de θ.

$$w_i := w_i \times \mathcal{L}(\mathbf{y}_2, \theta^{(i)})$$

- Si $\frac{1}{\sum_{i=1}^{N_e} w_i^2}$ est inférieur à $\frac{N_e}{4}$, on applique les étapes suivantes:
- on rééchantillonne les particules $(\theta^{(i)}, \ldots, \theta^{(N_e)})$ selon un échantillonnage résiduel
- on génère $(\tilde{\theta}^{(1)}, \ldots, \tilde{\theta}^{(N_e)})$ à partir d'une loi h.
- $\theta^{(i)}$ est remplacée par $\tilde{\theta}^{(i)}$ avec la probabilité $\min\left\{\frac{\mathcal{L}(\mathbf{y},\tilde{\theta}^{(i)}) \times f(\tilde{\theta}^{(i)}) \times h(\theta^{(i)})}{h(\tilde{\theta}^{(i)}) \times \mathcal{L}(\mathbf{y},\theta^{(i)}) \times f(\theta^{(i)})}, 1\right\}$.

Example

Back

Table : Réponse au traitement de 6 patients pour l'illustration de l'inférence Bayésien par filtre particulaire.

Patient	$\log(d)$	Toxicité	Activité
1	0	0	0
2	0	0	0
3	0	0	0
4	1	0	0
5	1	0	0
6	1	0	1

On commence par simuler N_e particules. Pour cet exemple, $N_e = 10$. On utilise comme fonction g la fonction de densité *a priori* des paramètres. On représente les particules du paramètres α_1.

```
## Error in library(pbivnorm):  there is no package
called 'pbivnorm'
```

Example

Back

```
## Error in library(pbivnorm):  there is no package
called 'pbivnorm'
```

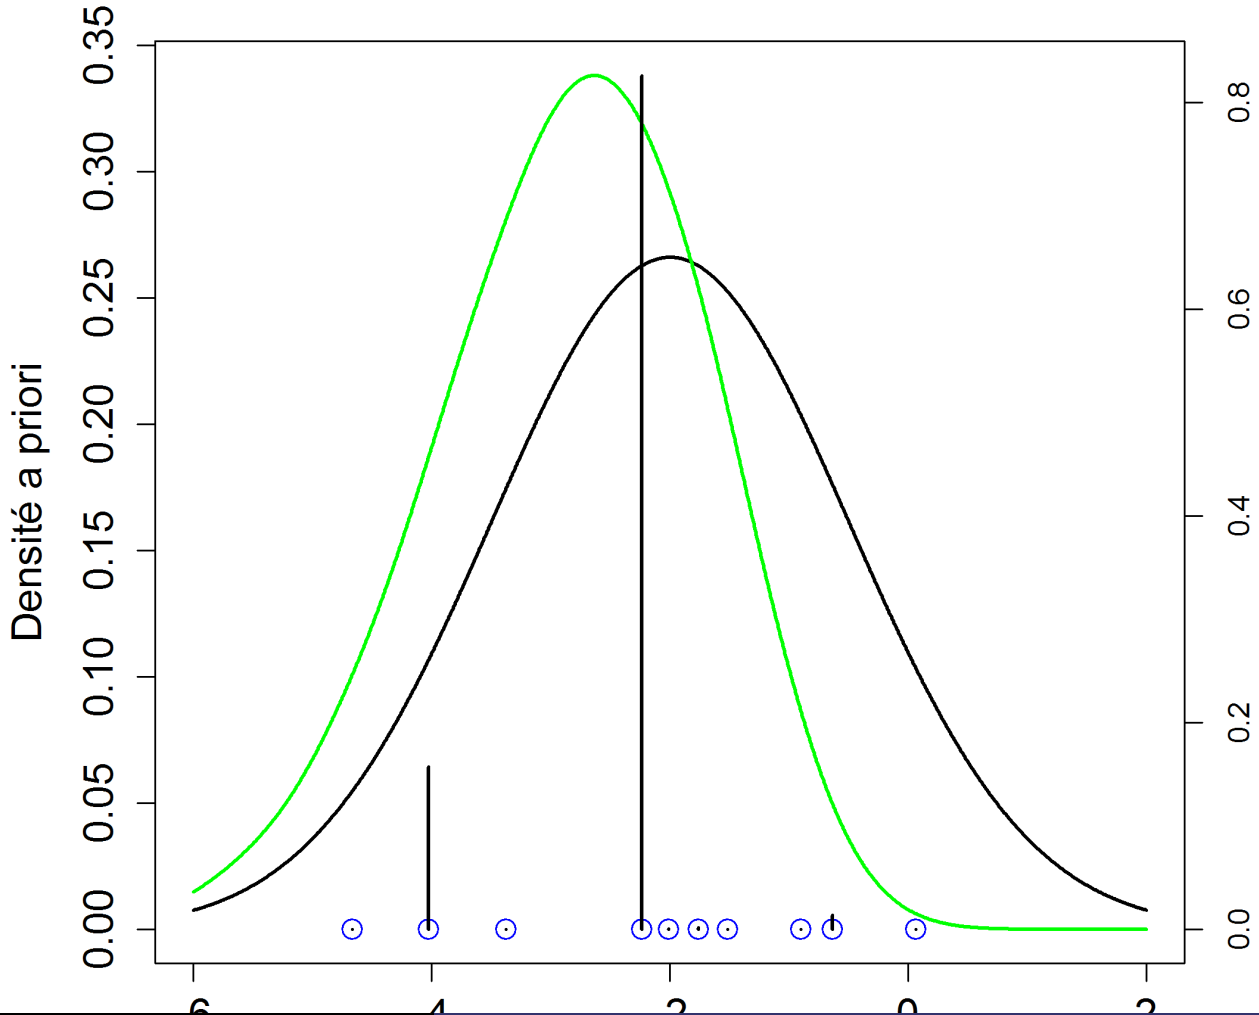

Example

Back

```
## Error in library(pbivnorm):  there is no package
called 'pbivnorm'
```

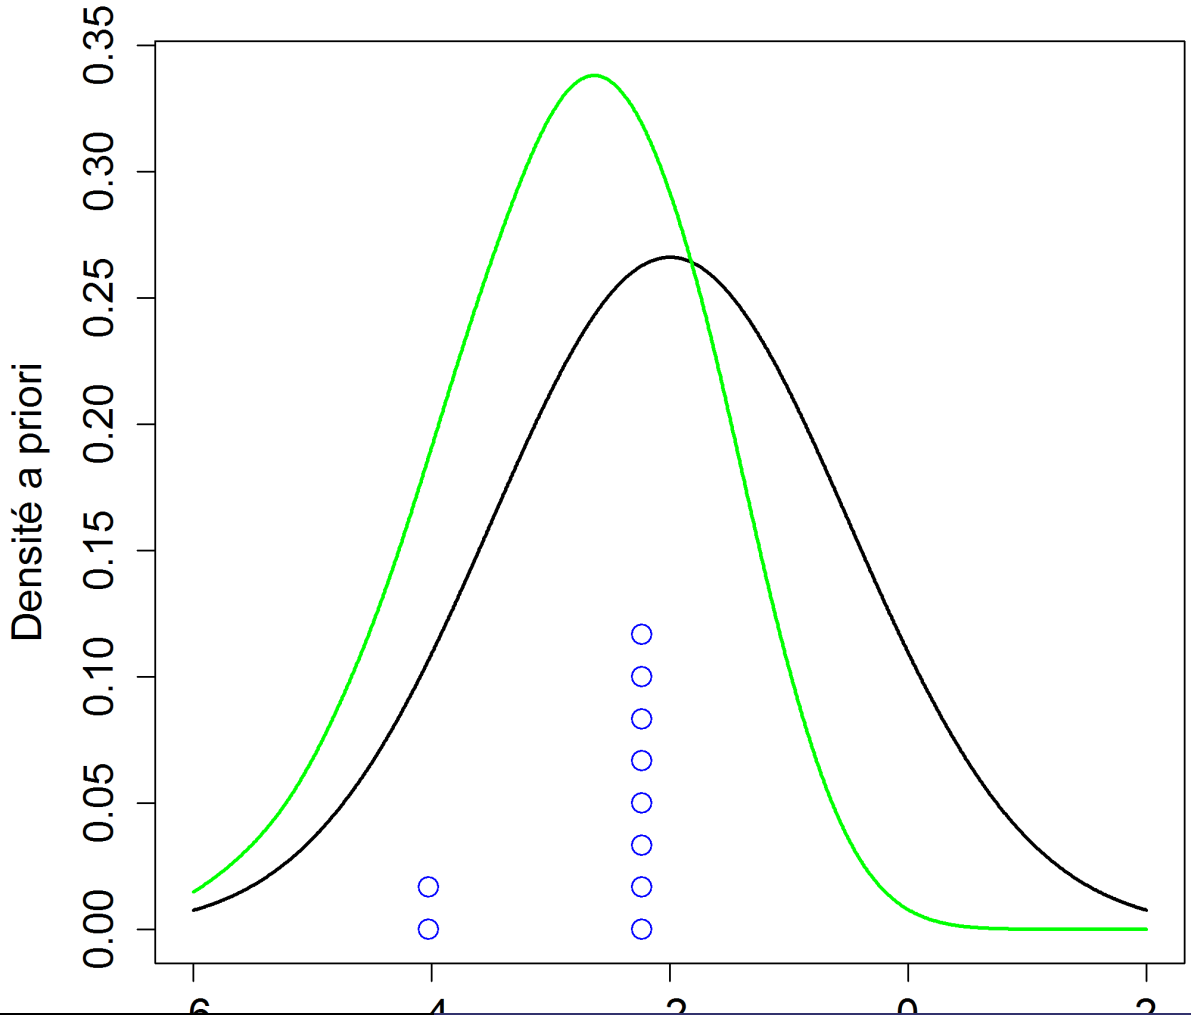

Example

Back

```
## Error in library(pbivnorm):  there is no package
called 'pbivnorm'
```

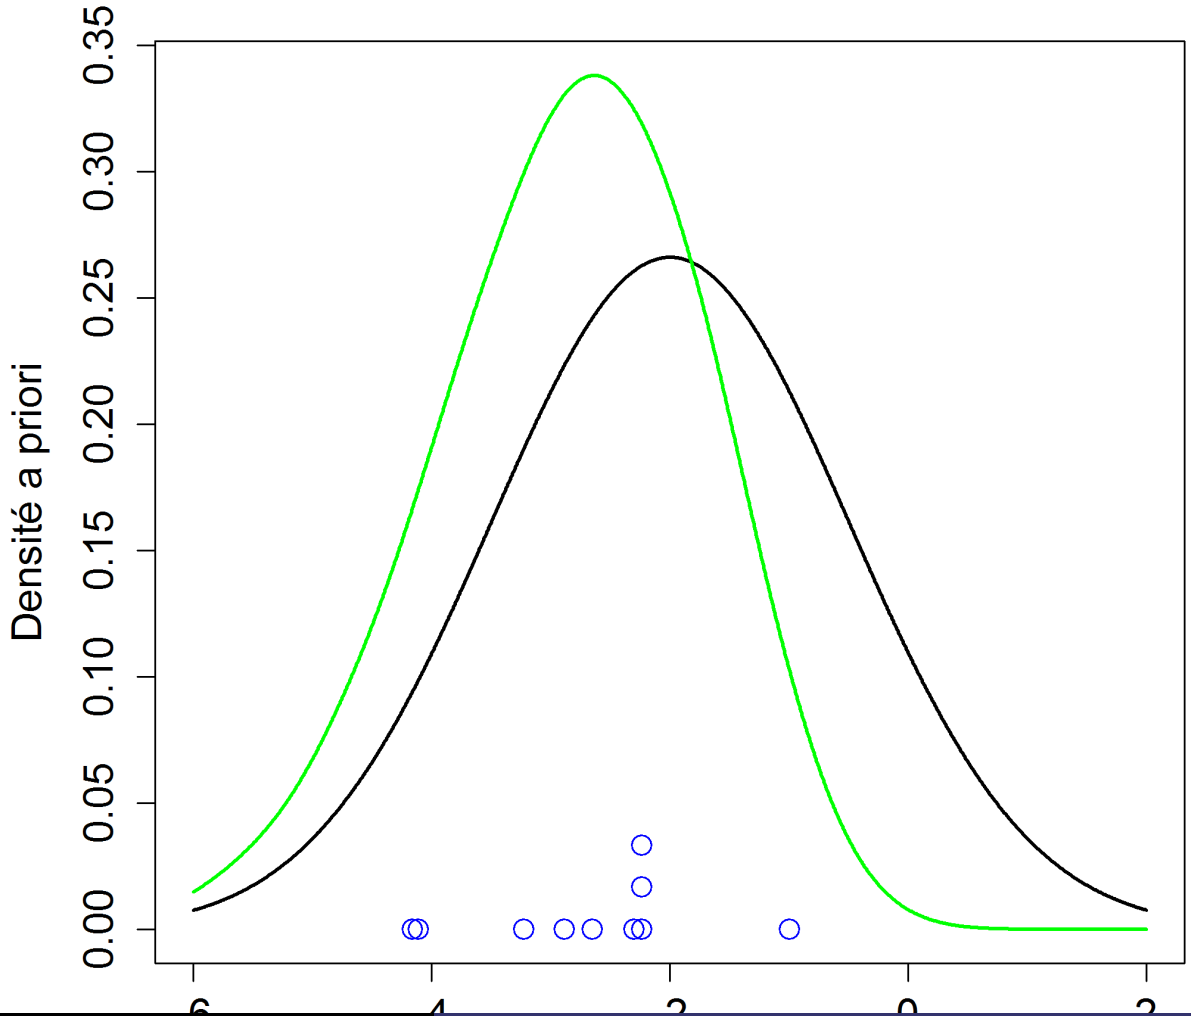

Plans d'expériences algorithmiques

Back

Classification & Exemples

- Plans d'expériences algorithmiques (pas de modélisation)
- Plans d'expériences *aveugles* (aucune anticipation)
- Exemples:
  - 3+3 (A+B)
  - ATD
  -

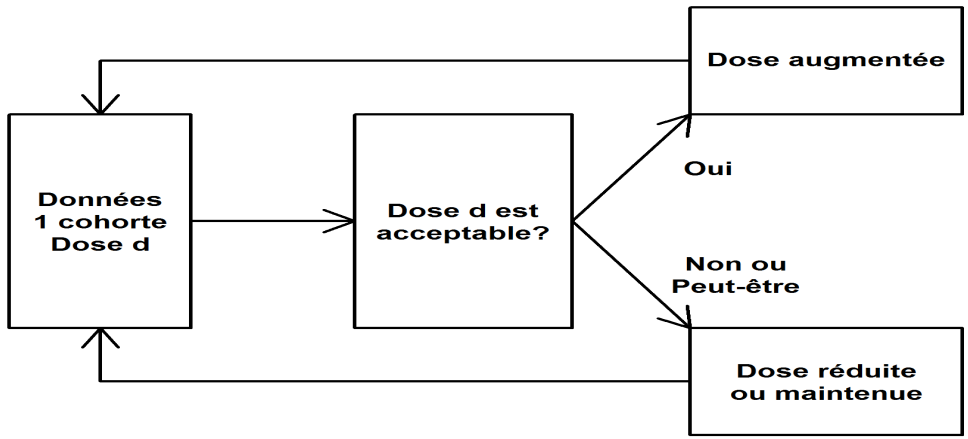

Plans d'expériences *model-based* (1/2)

Back

Classification & Exemples

- Plans d'expériences *model-based* (modélisation de la relation dose-réponse)
- Plans d'expériences *myopes* (anticipation d'ordre 1)
- Anticipent uniquement les observations suivantes et non la déci
- CRM) (Roth 200

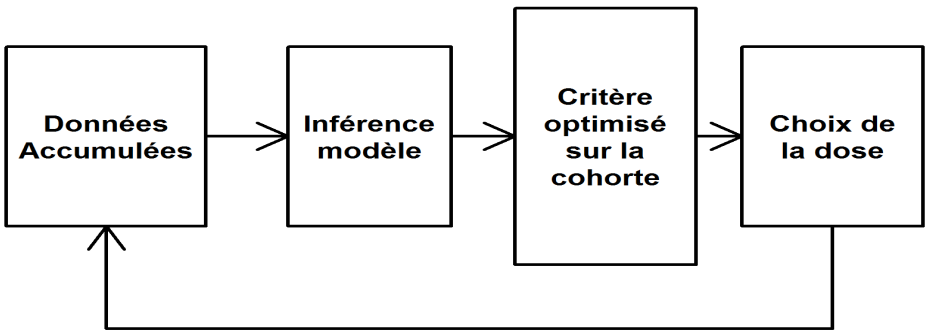

Plans d'expériences *model-based* (1/2)

Back

Classification & Exemples

- Plans d'expériences *model-based* (modélisation de la relation dose-réponse)
- Plans d'expériences *myopes* (anticipation d'ordre 1)
- Anticipent uniquement les observations suivantes et non la décision qui en découlera
- CRM (O'Quigley 1990), EWOC (Babb 1998) et SLOD (Roth 2009)

Plans séquentiels myopes proposés

- K. Roth: Modèle bivarié. Critères D-optimal, L-optimal
- F. Bretz: Modèles univariés d'activité. Optimisation d'une caractéristique de la matice d'information (det, tr...)
- C. Drovandi: Modèle non linéaire univarié. Minimisation de la variance de la MTD estimée

Plans séquentiels "optimaux" en Phase I

Back

SLOD (Sequentially Locally Optimal Design)

- Objectif: Décider de la dose à utiliser pour la cohorte suivante
- Considère chaque nouvelle cohorte comme la dernière
- Optimisation: Ne tient compte que de la cohorte suivante
- Critère à optimiser: Fonction de la matrice de Fisher, Variance de la dose recommandée

Plans séquentiels myopes proposés

- K. Roth: Modèle bivarié. Critères D-optimal, L-optimal
- C. Drovandi: Modèle non linéaire univarié. Minimisation de la variance de la MTD estimée
- F. Bretz: Modèles univariés d'activité. Optimisation d'une caractéristique de la matice d'information (det, tr...)

Plans d'expériences *model-based* (2/2)

Back

Classification & Exemples

- Plans d'expériences *model-based* (modélisation de la relation dose-réponse)
- Plans d'expériences *adaptatifs* (anticipation d'ordre $n \geq 2$)
- Contrôle séquentiel de systèmes dynamiques
- Déte
  situa tions)
- Exer

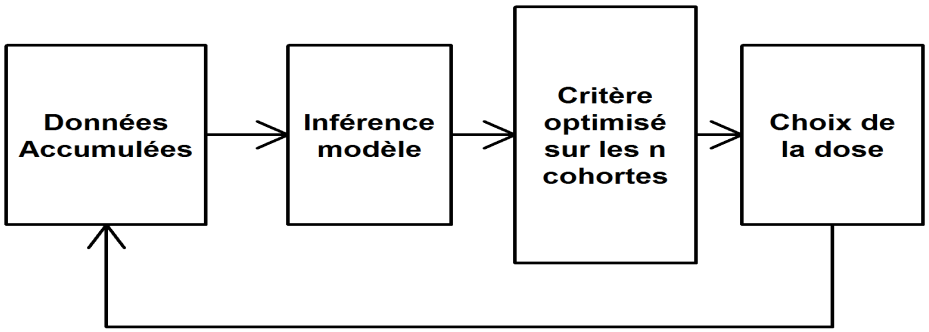

Plans d'expériences *model-based* (2/2)

Back

$doses = \{0, 1, 2\}$ Critère = Var(MTD|Données)

$a_1 = \sum_k r_k$ $a_2 = \sum_k r_k \times d_k$

n_1	n_2	n_3	a_1	a_2	
1	0	0	0	0	2.136(2)
			1	0	4.040(1)
1	1	0	0	0	2.657(3)
			1	0	3.476(1)
				1	1.320(1)
			2	1	4.652(1)
2	0	0	0	0	2.326(2)
			1	0	2.592(1)
			2	0	6.092(1)

n_1	n_2	n_3	a_1	a_2	
1	1	1	0	0	4.568
			1	0	4.659
				1	2.711
				2	1.101
			2	1	5.141
				2	2.990
				3	1.018
			3	3	5.290
2	1	0	0	0	3.158
			1	0	3.035
				1	1.077
			2	0	4.684
				1	2.289
			3	1	7.084
1	2	0	0	0	3.547
			1	0	3.559
				1	1.511
			2	1	3.462
				2	1.184
			3	2	5.677
3	0	0	0	0	2.657
			1	0	2.299
			2	0	3.281
			3	0	8.154

Azriel, D. (2013). Optimal sequential designs in phase I studies

Plans adaptatifs en Phase I

Back

Classification par modélisation

- Plans d'expériences algorithmiques
  - 3+3 (A+B), ATD, Hunsberger, *mTPI* ...
- Plans d'expériences *model-based*
  - CRM, EWOC, SLOD...

Classification par anticipation

- Plans d'expériences *aveugles*
  - 3+3 (A+B), ATD, Hunsberger, mTPI...
- Plans d'expériences *myopes*
  - CRM, EWOC, SLOD...
- Plans d'expériences *adaptatifs*
  - Contrôle séquentiel de systèmes dynamiques

mTPI as algorithmic design

Back

- Proposed as model-based design by authors
- Bayesian inference
- Based on independant binomial models
- Dose recommendation based on current dose level
- No prediction of further patients

mTPI details

Back

Independant binomial models

$$X_i = \text{Number of toxicities at dose i}$$

$$[X_i|\pi_i] \sim \mathcal{B}inom(N_i, \pi_i) \ \forall \ i$$

Bayesian inference

$$[\pi_i] \sim \mathcal{B}eta(a, b) \Rightarrow [\pi_i|X_i] \sim \mathcal{B}eta(a + X_i, b + N_i - X_i)$$

Dose recommendation

$$\underbrace{\frac{\int_0^{p_T-\epsilon_1} [\pi_i|X_i]\mathrm{d}\pi_i}{p_T - \epsilon_1}}_{E} \quad \underbrace{\frac{\int_{p_T-\epsilon_1}^{p_T+\epsilon_2} [\pi_i|X_i]\mathrm{d}\pi_i}{\epsilon_1 + \epsilon_2}}_{S} \quad \underbrace{\frac{\int_{p_T+\epsilon_2}^{1} [\pi_i|X_i]\mathrm{d}\pi_i}{1 - p_T - \epsilon_2}}_{D}$$

3+3 as .... design

Back

3+3 rules

- 0/3 tox. $\Rightarrow$ Escalate
- 1/3 tox. $\Rightarrow$ Stay
- $\geq$ 2/3 tox. $\Rightarrow$ De-escalate

3+3 as .... design

Back

3+3 rules

- 0/3 tox. $\Rightarrow$ Escalate
- 1/3 tox. $\Rightarrow$ Stay
- $\geq$ 2/3 tox. $\Rightarrow$ De-escalate

MLE for binomial model

- $X_i \sim \mathcal{B}inom(N_i, \pi_i) \Rightarrow \hat{\pi}_i^{MLE} = x_i/N_i$
- $\hat{\pi}_i^{MLE} = 0/3 \Rightarrow$ Escalate
- $\hat{\pi}_i^{MLE} = 1/3 \Rightarrow$ Stay
- $\hat{\pi}_i^{MLE} \geq 2/3 \Rightarrow$ De-escalate

3+3 as .... design

Back

3+3 rules

- 0/3 tox. $\Rightarrow$ Escalate
- 1/3 tox. $\Rightarrow$ Stay
- $\geq$ 2/3 tox. $\Rightarrow$ De-escalate

MLE for binomial model

- $X_i \sim \mathcal{B}inom(N_i, \pi_i) \Rightarrow \hat{\pi}_i^{MLE} = x_i/N_i$
- $\hat{\pi}_i^{MLE} = 0/3 \Rightarrow$ Escalate
- $\hat{\pi}_i^{MLE} = 1/3 \Rightarrow$ Stay
- $\hat{\pi}_i^{MLE} \geq 2/3 \Rightarrow$ De-escalate

Conclusion: 3+3 is a **model-based** design

3 + 3 design (algorithmic)

Back

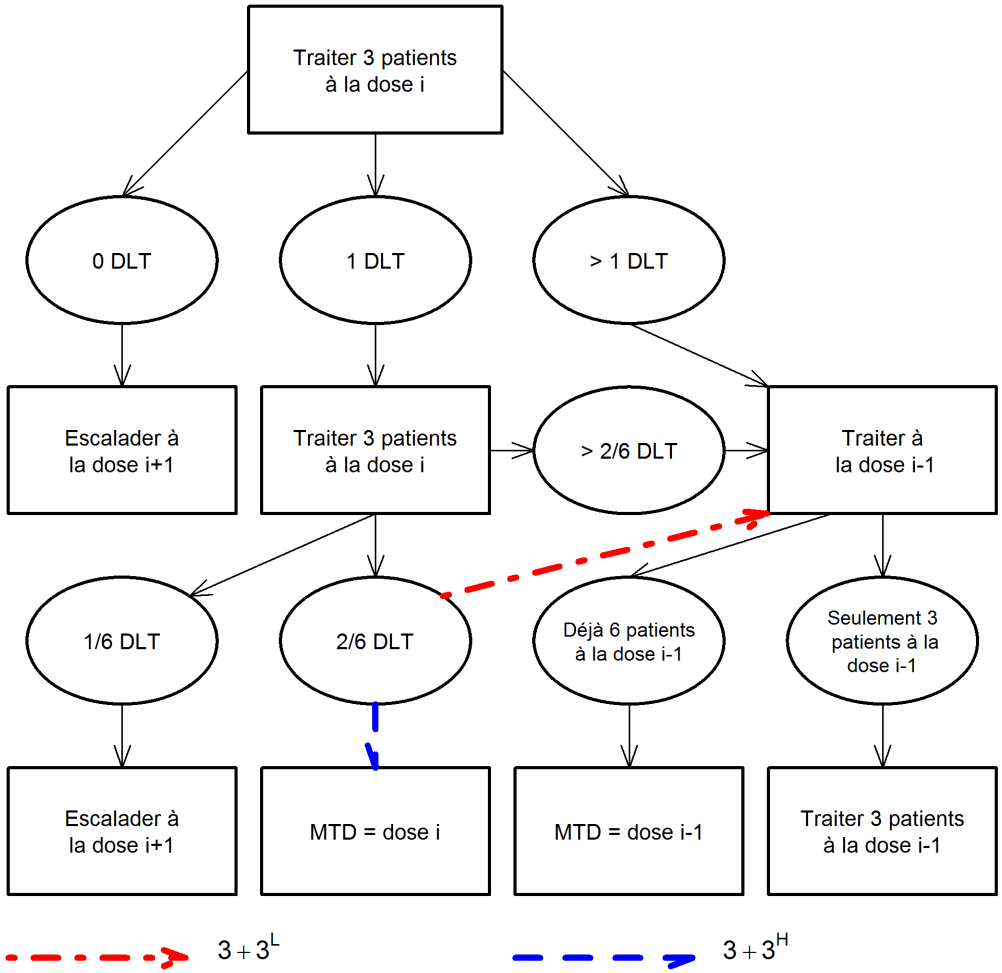

ATD design (algorithmic)

Back

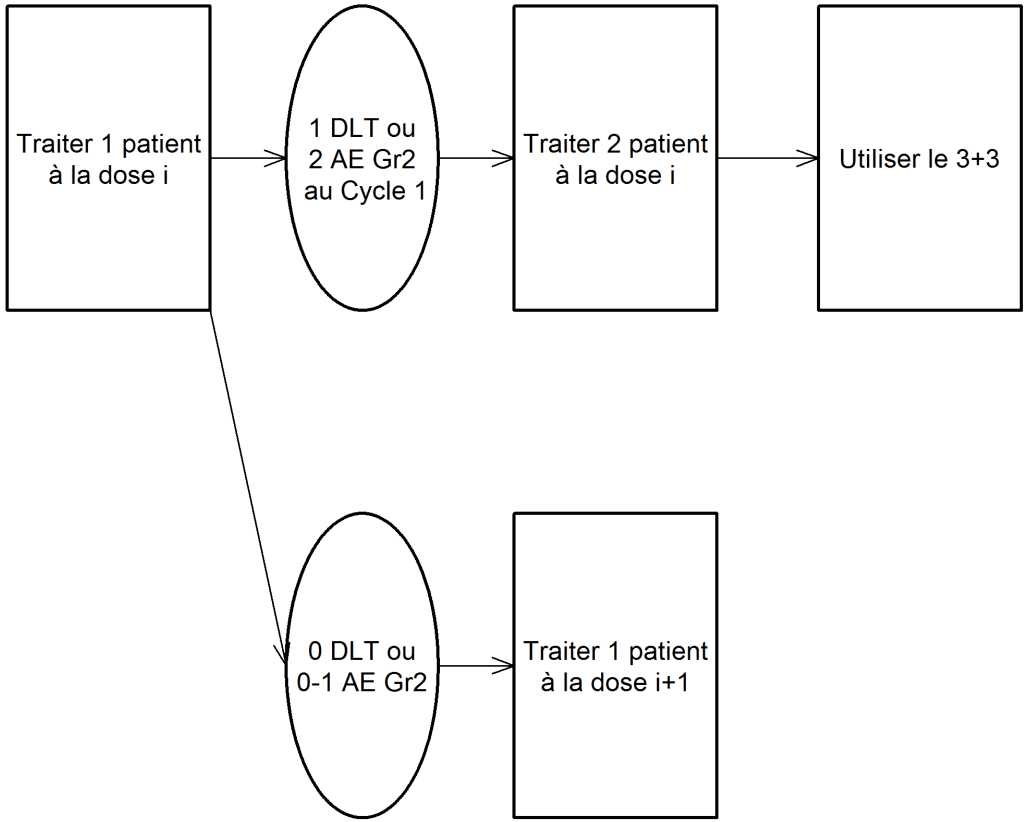

mTPI design (algorithmic)

Back

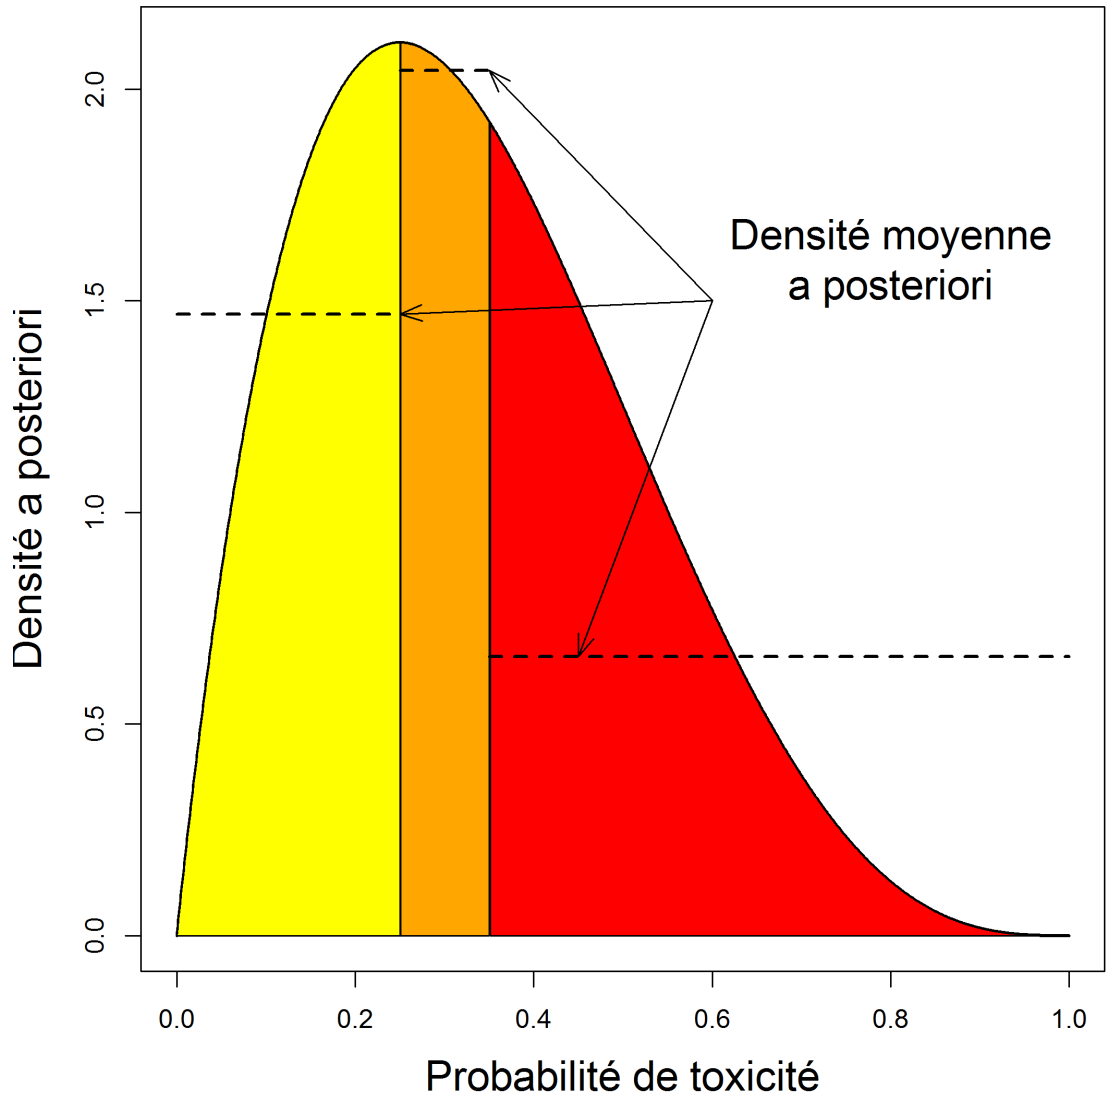

CRM design (model-based)

Back

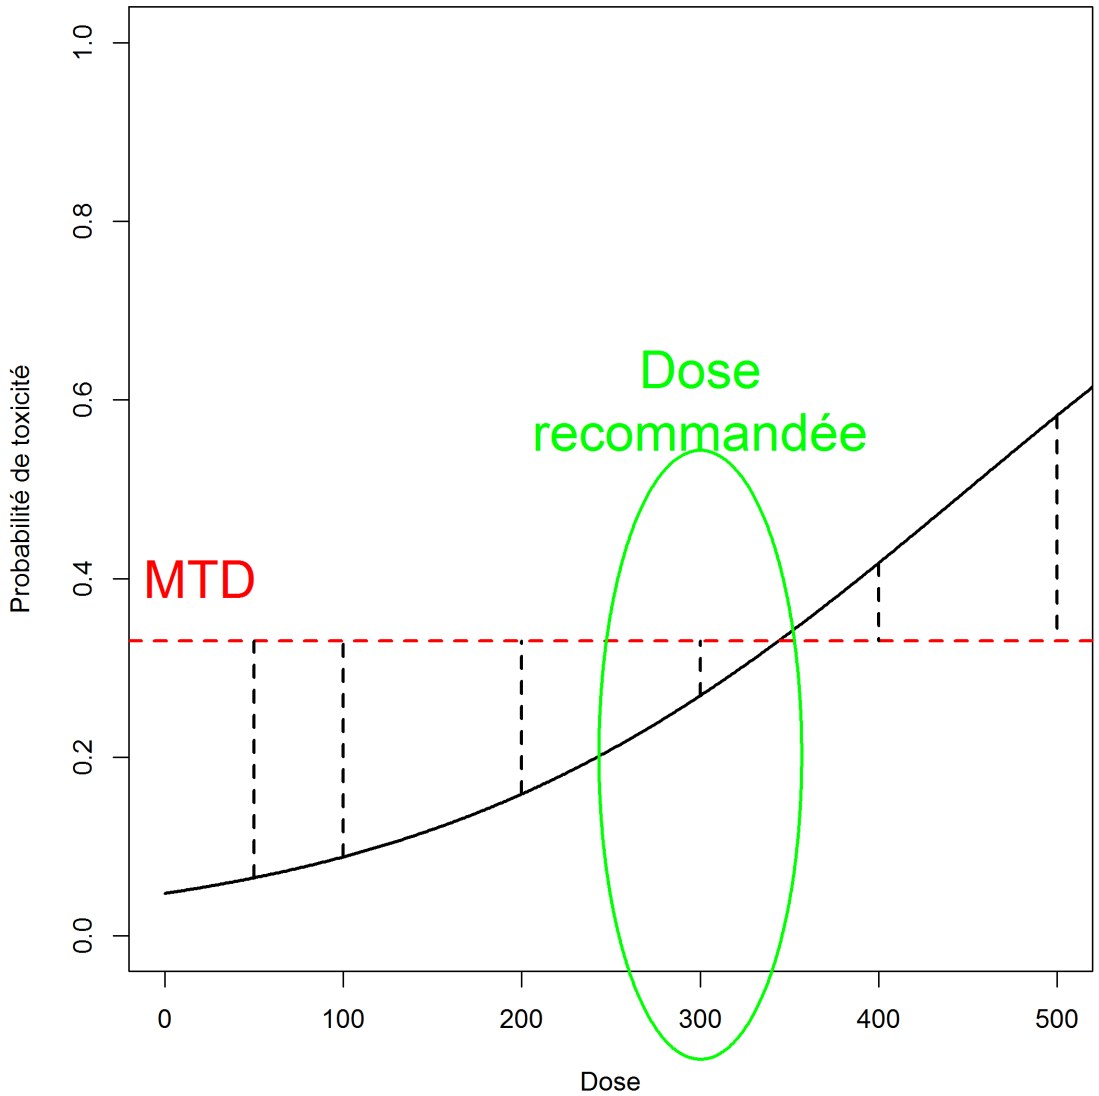

EWOC design (model-based)

Back

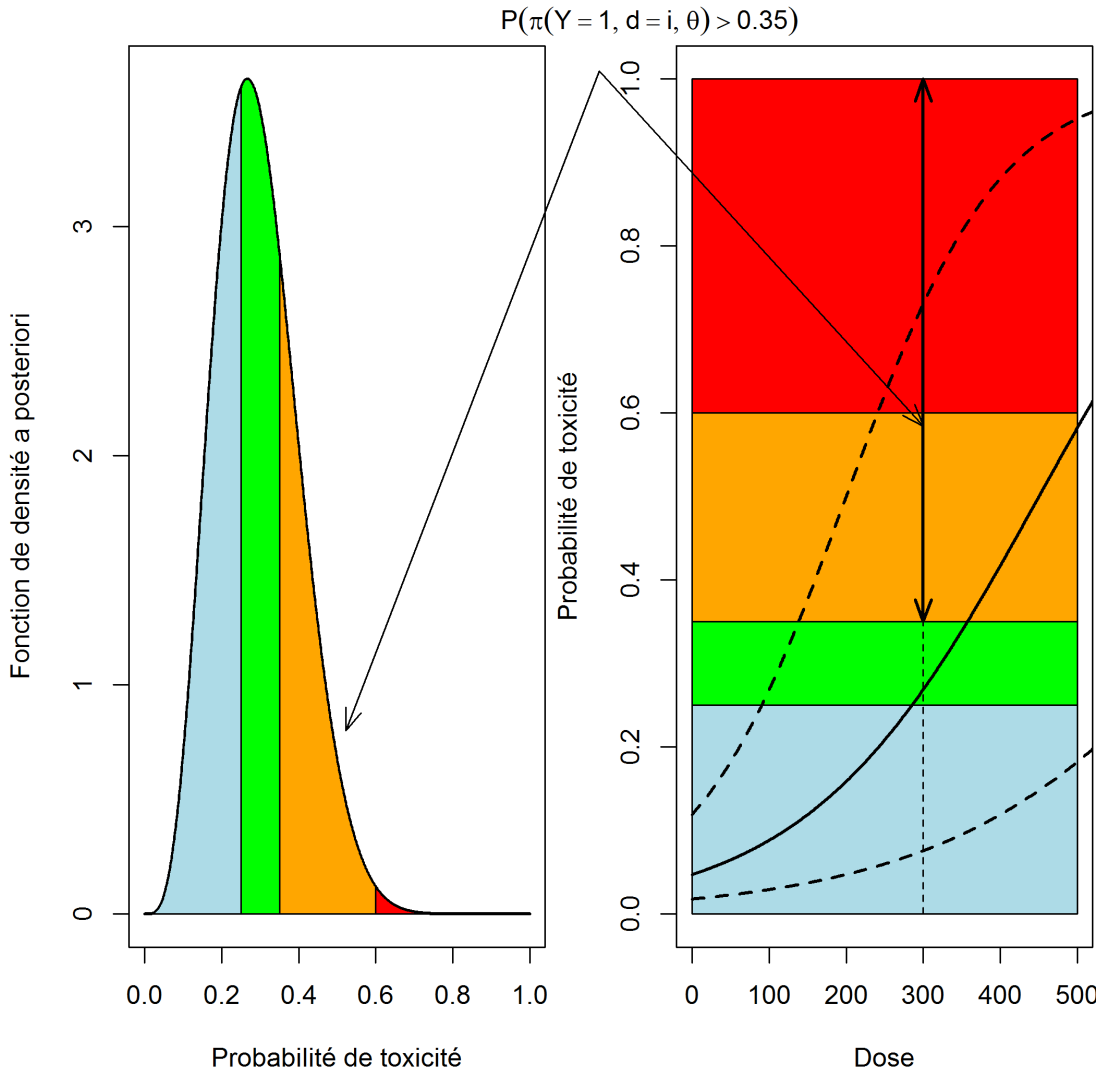

Cost functions

Norm function

- $L(\theta, d_t, y_t, t) = \left(d_t - MTD(\theta)\right)^2$
- $L(\theta, d_t, y_t, t) = \left|d_t - MTD(\theta)\right|$
- $L(\theta, d_t, y_t, t) = \left(\sum_{k=1}^{K} w_k f_k(\theta)^p\right)^{1/p}$

Cost functions

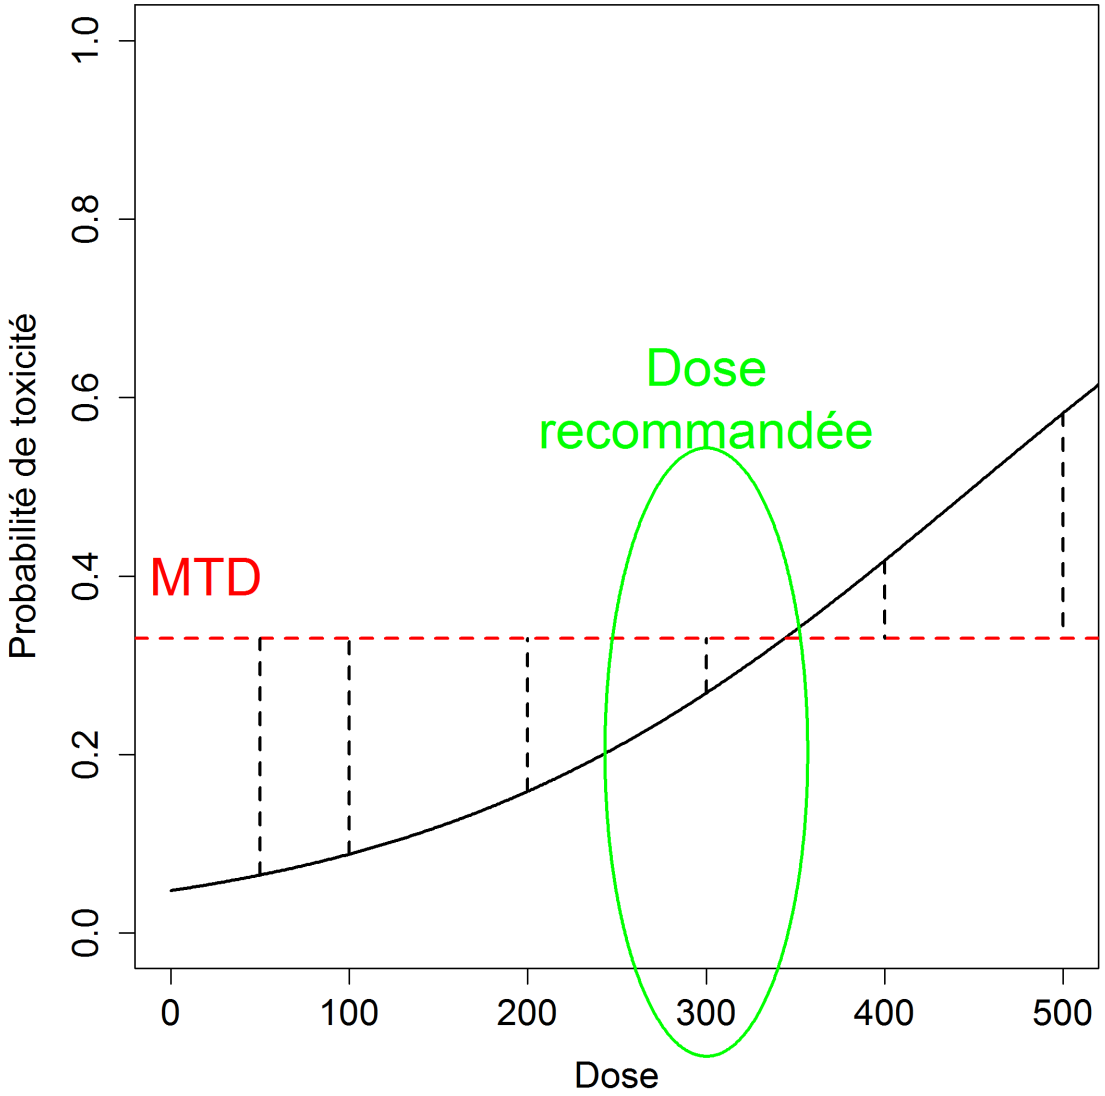

Cost functions

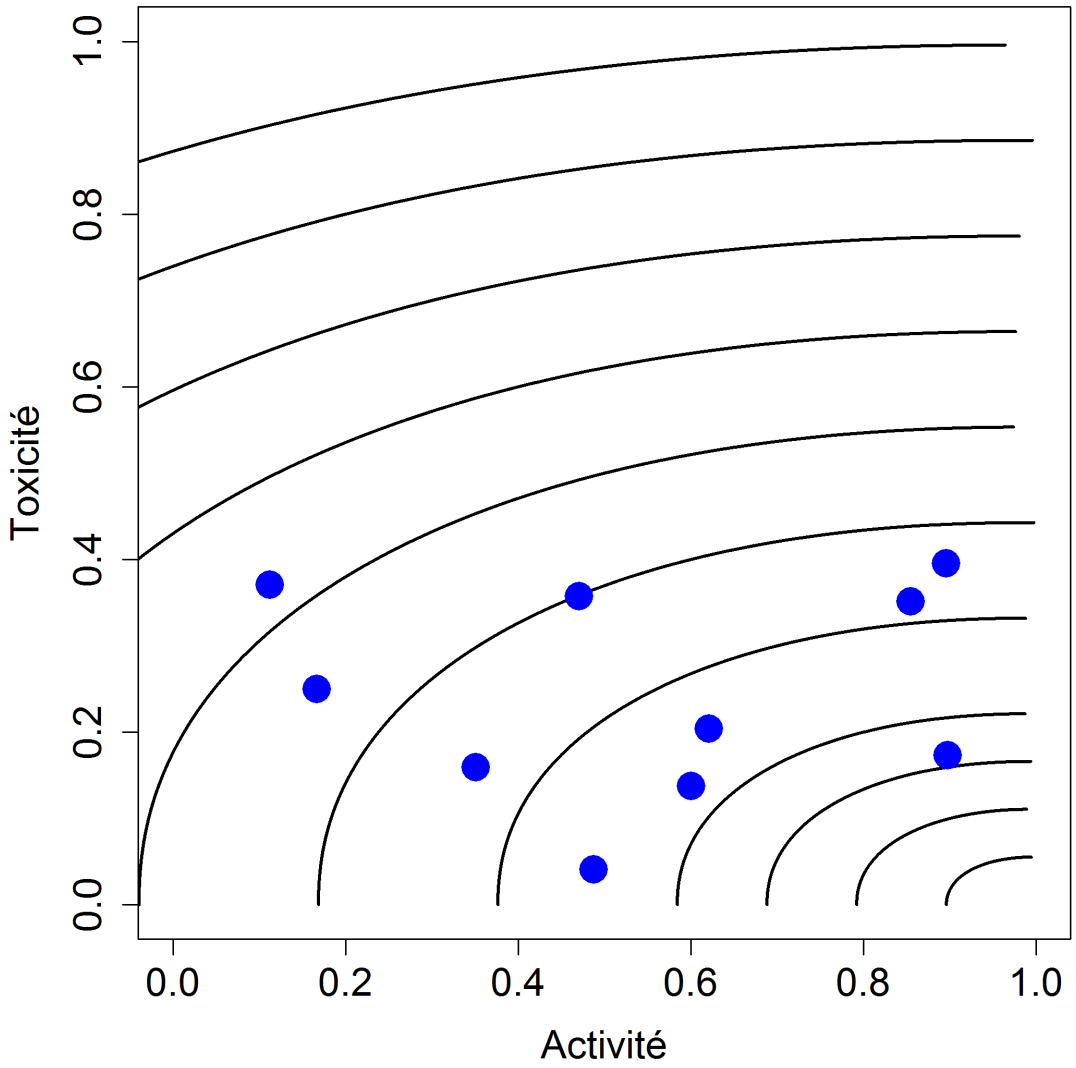

Cost functions

Norm function

- $L(\theta, d_t, y_t, t) = \Big(d_t - MTD(\theta)\Big)^2$
- $L(\theta, d_t, y_t, t) = \Big|d_t - MTD(\theta)\Big|$
- $L(\theta, d_t, y_t, t) = \left(\sum_{k=1}^{K} w_k f_k(\theta)^p\right)^{1/p}$

Information function

- $L(\theta, d_t, y_t, t) = \mathbb{E}\left(\frac{\partial^2 \mathcal{L}(y_t, d_t, \theta)}{\partial \theta^2}\right)$