



HAL
open science

Utility-based optimization of phase II / phase III clinical development

Jihane Aouni

► **To cite this version:**

Jihane Aouni. Utility-based optimization of phase II / phase III clinical development. Statistics [math.ST]. Université Montpellier, 2019. English. NNT : 2019MONT032 . tel-02382685

HAL Id: tel-02382685

<https://theses.hal.science/tel-02382685>

Submitted on 27 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE POUR OBTENIR LE GRADE DE DOCTEUR DE L'UNIVERSITÉ DE MONTPELLIER

En Biostatistique

École doctorale Information Structures Systèmes

Utility-based optimization of phase II / phase III clinical development

Présentée par Jihane AOUNI

Le 16 septembre 2019

Sous la direction de Jean-Noël BACRO

Devant le jury composé de

M. Franz KOENIG, Professeur, Medical University of Vienna

Mme Virginie RONDEAU, Directrice de recherche, INSERM

M. François COQUET, Professeur, ENSAI

M. Ali GANNOUN, Professeur, Université de Montpellier

Jean-Christophe LEMARIE, Directeur de statistique, Effi-Stat

Jean-Noël BACRO, Professeur, Université de Montpellier

Gwladys TOULEMONDE, Maître de conférences, Université de Montpellier

Loïc DARCHY, Chef du groupe de méthodologie statistique, Sanofi R&D

Pierre COLIN, Biostatisticien, Sanofi R&D

Rapporteur

Rapporteuse

Examinateur

Président du jury

Invité

Directeur de thèse

Co-encadrante de thèse

Co-encadrant de thèse

Co-encadrant de thèse



Contents

Acknowledgements	14
Preamble: Brief presentation of Sanofi	16
Résumé	18
Summary	24
1 Introduction	30
1.1 Drug Development	30
1.2 Methodologies for dose selection in drug development	31
1.3 Designs for dose-finding studies: fixed and adaptive designs studies	32
1.4 Objectives of the thesis	33
2 A brief literature review	36
2.1 Adaptive designs	36
2.2 Modeling procedures	39
2.2.1 Dose-response modeling	39
2.2.2 Model Averaging approach	40
2.3 Optimization procedures	42
2.3.1 Optimal designs	42
2.3.2 Bayesian optimal designs	43
2.4 Utility functions	44
2.5 Summary and outlook	49
3 Utility functions: how to build them and what contributions for design optimization can we expect?	54
3.1 General notations and main notions	54
3.2 Utility functions: constructions and properties	55
3.3 Seamless design and utility function: some first quite disappointing results	58
3.3.1 Introduction and notations	58
3.3.2 Optimal patient allocation: unattractive results	59
3.4 Sponsor's strategy: Optimal dose and decision rules	59
3.5 Conclusions	63

4	Dose selection in the context of a phase II dose-finding study	66
4.1	Materials and Methods	67
4.1.1	Dose-response modeling	67
4.1.2	Decision-making framework	67
4.1.2.1	Utility function	68
4.1.2.2	Optimal dose and decision rules	68
4.1.2.3	Optimal dose estimation method: Batching approach	70
4.1.3	Simulations	73
4.1.3.1	Simulation protocol	73
4.1.3.2	Simulation scenarios for efficacy and toxicity	75
4.2	Results	76
4.2.1	No activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)	77
4.2.2	No activity with low toxicity scenario	79
4.2.3	No activity scenario with high toxicity scenario	81
4.2.4	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)	82
4.2.5	Sigmoid scenario with acceptable safety scenario (except for highest dose)	84
4.2.6	Sigmoid scenario with low toxicity scenario	86
4.2.7	Sigmoid scenario with high toxicity scenario	87
4.2.8	Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)	89
4.3	Conclusions	90
5	Utility-based dose-finding in practice: some empirical contributions and recommendations	92
5.1	Estimation of the posterior distribution of the utility values	92
5.1.1	Exploring properties of MCMC based posterior distributions: examples of two particular single simulated studies	93
5.1.1.1	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15)	95
5.1.1.2	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)	105
5.1.1.3	Conclusions	114
5.1.2	Properties of Bayesian estimators of utilities - posterior means and medians	116
5.1.2.1	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15): posterior means of utilities	116
5.1.2.2	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15): posterior medians of utilities	120
5.1.2.3	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20): posterior means of utilities	123
5.1.2.4	Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20): posterior medians of utilities	127
5.1.2.5	Conclusion	130
5.2	Sponsor's strategy: Optimal dose and decision rules	131
5.2.1	Comparisons of alternative decision rules	131

5.2.2	Simulation results	132
5.3	Influence of phase II sample size - sample size recommendations	136
5.4	Alternative criteria for interim analysis	141
5.5	Concluding Remarks	146
6	Discussion and Conclusions	148
6.1	Discussion	148
6.2	Perspectives	152
6.3	General conclusion	155
	Bibliography	156
	Appendices	166
A.1.	Complementary developments on utility-based seamless design optimization . .	166
A.2.	Convergence and autocorrelations	188
A.3.	Results for low efficacy threshold, $\text{PoS} > 0.30$	189
A.4.	Results for high efficacy threshold, $\text{PoS} > 0.90$	190
A.5.	Boxplots related to the posterior distribution	191
A.6.	Results for U2	206
A.7.	Sensitivity analysis	207

List of Figures

1	Clinical trials	16
1.1	Summary of clinical trial phases (Source: Sanofi).	30
2.1	Clinical trial steps.	37
2.2	Theoretical curves (see text for details).	51
4.1	Theoretical curves, no activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	78
4.2	Theoretical curves, no activity scenario with low toxicity scenario. The PoS and utility curves are superimposed.	79
4.3	Theoretical curves, no activity scenario with high toxicity scenario. The PoS and utility curves are partially superimposed, especially for low doses.	81
4.4	Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20). The PoS and utility curves are partially superimposed, especially for low doses.	82
4.5	Theoretical curves, Sigmoid scenario with acceptable safety scenario (except for highest dose). The PoS and utility curves are partially superimposed, especially for low doses.	84
4.6	Theoretical curves, Sigmoid scenario with low toxicity scenario. The PoS and utility curves are almost identical.	86
4.7	Theoretical curves, Sigmoid scenario with high toxicity scenario. The PoS and utility curves are partially superimposed, especially for low doses.	87
4.8	Theoretical curves, Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20). The PoS and utility curves are partially superimposed, especially for low doses.	89
5.1	Violin plot example.	94
5.2	Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15). The PoS and utility curves are superimposed for small dose values.	95
5.3	Violin plots for one simulated study, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	96
5.4	Violin plots for one simulated study, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	97
5.5	Violin plots for one simulated study, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	99

5.6	Violin plots for one simulated study, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	100
5.7	Violin plots for one simulated study, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	102
5.8	Violin plots for one simulated study, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	103
5.9	Violin plots for one simulated study, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	106
5.10	Violin plots for one simulated study, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	107
5.11	Violin plots for one simulated study, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	109
5.12	Violin plots for one simulated study, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	110
5.13	Violin plots for one simulated study, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	112
5.14	Violin plots for one simulated study, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	113
5.15	Violin plots of posterior means of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	116
5.16	Violin plots of posterior means of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	117
5.17	Violin plots of posterior means of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	117
5.18	Violin plots of posterior means of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	118
5.19	Violin plots of posterior means of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	118
5.20	Violin plots of posterior means of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	119
5.21	Violin plots of posterior medians of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	120
5.22	Violin plots of posterior medians of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	121
5.23	Violin plots of posterior medians of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	121
5.24	Violin plots of posterior medians of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	122
5.25	Violin plots of posterior medians of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	122
5.26	Violin plots of posterior medians of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).	123
5.27	Violin plots of posterior means of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	124
5.28	Violin plots of posterior means of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	124

5.29	Violin plots of posterior means of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	125
5.30	Violin plots of posterior means of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	125
5.31	Violin plots of posterior means of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	126
5.32	Violin plots of posterior means of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	126
5.33	Violin plots of posterior medians of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	127
5.34	Violin plots of posterior medians of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	128
5.35	Violin plots of posterior medians of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	128
5.36	Violin plots of posterior medians of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	129
5.37	Violin plots of posterior medians of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	129
5.38	Violin plots of posterior medians of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	130
5.39	Relative utility loss graph comparing the five decision rules.	133
5.40	Relative utility loss, Sigmoid scenario - Bad safety (progressive toxicity scenario (and toxicity of highest dose = 0.20)).	137
5.41	Relative utility loss, Plateau scenario - Bad safety (progressive toxicity scenario (and toxicity of highest dose = 0.20)).	137
5.42	Relative utility loss, Sigmoid scenario - Good safety profile (low toxicity scenario).	138
5.43	Relative utility loss, Plateau scenario - Good safety profile (low toxicity scenario).	138
6.1	Iso-quant utility curves, Safety as a function of Efficacy, $k = 2, 0.5$.	150
6.2	Effect size as a function of the dose.	170
6.3	PoS by dose for various values of the safety parameter.	171
6.4	Utility as a function of phase II sample size and safety sa .	172
6.5	Utility as a function of phase II sample size for three dose-response profiles.	173
6.6	Utility as a function of phase II sample size for two dose-response profiles.	174
6.7	Theoretical curves for U5, Sigmoid scenario.	177
6.8	Theoretical utilities as a function of the dose and f .	183
6.9	Theoretical curves for U9, Sigmoid scenario.	185
6.10	Utility expectations as a function of f , Sigmoid scenario.	187
6.11	Utility expectations as a function of f , Plateau scenario.	187
6.12	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.	191
6.13	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.	192
6.14	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.	192

6.15	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$	193
6.16	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$	193
6.17	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$	194
6.18	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$	195
6.19	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$	195
6.20	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$	196
6.21	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$	196
6.22	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$	197
6.23	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$	197
6.24	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$	199
6.25	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$	200
6.26	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$	200
6.27	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$	201
6.28	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$	201
6.29	Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$	202
6.30	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$	203
6.31	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$	203
6.32	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$	204
6.33	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$	204
6.34	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$	205
6.35	Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$	205
6.36	Prior dose-response distribution plots.	207
6.37	Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $t = 0.30$	209

List of Tables

4.1	Simulation results, no activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	78
4.2	Simulation results, no activity scenario with low toxicity scenario.	80
4.3	Simulation results, no activity scenario with high toxicity scenario.	81
4.4	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	83
4.5	Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).	84
4.6	Simulation results, Sigmoid scenario with low toxicity scenario.	86
4.7	Simulation results, Sigmoid scenario with high toxicity scenario.	88
4.8	Simulation results, Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	89
5.1	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.	96
5.2	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.	97
5.3	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.	98
5.4	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.	98
5.5	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.	99
5.6	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.	99
5.7	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.	100
5.8	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.	101
5.9	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.	102
5.10	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.	102
5.11	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$.	103

5.12	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$	104
5.13	Mean and median relative estimation bias per dose, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$	105
5.14	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$	106
5.15	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$	107
5.16	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$	108
5.17	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$	108
5.18	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$	109
5.19	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$	109
5.20	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$	110
5.21	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$	111
5.22	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$	112
5.23	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$	112
5.24	Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$	113
5.25	Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$	113
5.26	Mean and median relative estimation bias per dose, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$	114
5.27	Values of the relative utility loss depending on N_2	133
5.28	Simulation results, $N_2 = 250$, comparison of five decision rules for dose selection.	134
5.29	Simulation results, $N_2 = 500$, comparison of five decision rules for dose selection.	134
5.30	Simulation results, $N_2 = 1000$, comparison of five decision rules for dose selection.	135
5.31	Simulation results, interim at $N'_2 = 100$ and $N'_2 = 250$, stop at interim if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j \text{data}] \geq l, l = 0.80, 0.90$	142
5.32	Simulation results, utility mean differences criterion: Domination criterion 1 with $X = 0.10, 0.20$	144
5.33	Simulation results, utility median differences criterion: Domination criterion 2 with $X = 0.10, 0.20$	145
6.1	Balanced design -patients are equally allocated to doses- for $U5$, Sigmoid and Plateau scenarios.	178
6.2	Optimal design -optimizing the dose allocation ratio- for $U5$, Sigmoid and Plateau scenarios, with $f = 0.20$	178

6.3	Optimal design -optimizing the dose allocation ratio- for $U5$, Sigmoid and Plateau scenarios, with $f = 0.25$.	179
6.4	Simulation results for $U5$ with the balanced design (patients are equally allocated to doses), Sigmoid scenario, by increasing f .	181
6.5	Optimal versus non-optimal design for $U9$, where optimal design consists here in optimizing the global patient allocation between phase II and phase III, Sigmoid and Plateau scenarios.	186
6.6	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	189
6.7	Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).	189
6.8	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	190
6.9	Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).	190
6.10	Optimal versus non-optimal design for $U2$, where optimal design consists here in optimizing the patient allocation to doses and the global patient allocation between phase II and phase III simultaneously, Sigmoid and Plateau scenarios.	206
6.11	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	208
6.12	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).	209
6.13	Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $t = 0.30$.	210

Acknowledgements

At first I would like to thank Mr. Loïc DARCHY, head of the Statistical Methodology Group, for giving me the opportunity to complete my PhD thesis within a huge multinational and very well reputed pharmaceutical industry like Sanofi, and Mr. Jean Noel BACRO for giving me the chance to be a part of the IMAG laboratory and to graduate from the University of Montpellier. I would like to thank all the people that worked in the Research and Development department of Sanofi. With their patience and openness they created an enjoyable working environment. Furthermore I want to thank my four tutors, Mr. Loïc DARCHY, Mr. Pierre COLIN, Mr. Jean Noel BACRO and Mrs. Gwladys TOULEMONDE, for their guidance and support. We experienced great things together and they have shown me a beautiful part of team work. Mr. DARCHY helped me learn new statistical skills and techniques and gave me much confidence in myself. Mr. COLIN, a Bayesian expert, shared his knowledge with me, regularly followed my work and helped me improve my programming skills. Mr. BACRO and Mrs. TOULEMONDE shared their mathematical, statistical and methodological expertise with me, along with their teaching experiences, even when those weren't necessarily related to clinical research domain. They also helped me a lot on the logistic and administrative level, as well as writing my thesis manuscript. Without my tutors, I could not have done the work by myself therefore I am grateful for them.

In addition, I would particularly like to thank all the members of my thesis jury, my four tutors mentioned above, my thesis referees Mr. Franz KOENIG and Mrs. Virginie RONDEAU, my thesis examiner Mr. François COQUET and the jury president Mr. Ali GANNOUN, for the time they have dedicated to read my manuscript, for their participation in my thesis defense, for their advices and for their useful comments. Thanks to their expertise in the field of statistics and their evaluation / appreciation of my thesis work, I managed to get my doctoral degree, which I am very proud of.

I would also like to thank Mr. Christopher JENNISON and Mr. Eric PARENT, members of my thesis committee, for committing themselves to follow the progress of my work, to propose new ideas from an external and objective point of view. In addition, their expertise in clinical trials, Bayesian designs, dose-finding and optimization was absolutely useful for the achievement of my goals.

A special thanks to my parents and family (I hope you are all proud of me), to my friends Joelle Saade, Jinane Alloul, Fadi Jabroun, Joe Raad, Annarosa Quarello, Yann Vasseur, Florent Nacry, etc., to my colleagues Bernard Sebastien, Emma Gérard, Marie-Karelle Rivière, Gabriel Bologna, Wencan Zhu, Emma français, Pascaline Piccard, Lydie Melas-Melt, Cécile

Barkian, Emilie Gérard, Anissa Elfakir, Clémence Rigaux, Françoise Kueny, and to all the people who have always been around.

Lastly and most importantly, to my best friend Rana Jreiche that I met 8 years ago: rarely does a best friend get a chance to hear how special, amazing, and warm she is, and VERY FEW people have the chance to have such dedicated and trustworthy friend like you my dear Rana JREICH, because best friends are just hard to come by. That's why I wanted to let you know how much you mean to me. Thank you for being my closest friend, you were always by my side, in good times and bad. Thank you for being such an inspiration to me.

Preamble: Brief presentation of Sanofi

Sanofi is a multinational pharmaceutical company headquartered in Paris, France, as of 2013 the world's fifth-largest by prescription sales. The company was formed as Sanofi-Aventis in 2004 by the merger of Aventis and Sanofi-Synthélabo, which were each the product of several previous mergers. It changed its name to Sanofi in May 2011. Sanofi engages in the Research and Development (R&D), manufacturing and marketing of pharmaceutical drugs principally in the prescription market, but the firm also develops over-the-counter medication. It is one of the global leading companies in its field, pledged to innovate and develop healthcare solutions to its patients from all around the world. More than 110,000 people at Sanofi are dedicated to make a difference on patients' daily life, wherever they live and enable them to enjoy a healthier life. Sanofi transforms scientific innovation into healthcare solutions from the prevention stage to the final treatment. The company covers 7 major therapeutic areas: : human vaccines, rare diseases, multiple sclerosis, oncology, immunology, infectious diseases, diabetes and cardiovascular solutions and consumer healthcare.



Figure 1: Clinical trials, a very important role in biostatistics.

Medicine has had a massive evolution in the last century. Sanofi was one of the companies who helped in this evolution because of the hard work of its >100,000 employees who committed their efforts to improve patient's everyday life by ensuring them a healthier heart and soul. Innovating in the researches of serious diseases treatment over the past years had a great impact

on human's life. In fact, the average life expectancy of people all over the world has doubled in that time. To be committed to the uninterrupted advancement of health is one of Sanofi's trademarks. Because of the quality of its individuals, Sanofi has proved her character and excellence on the healthcare market by expanding new treatments not only to treat diseases, but also to prevent them in the future for anyone, worldwide.

Sanofi helped the developing world throughout the years with its employees expertise. Thus, her employees invented and developed the first to be known vaccine for dengue virus that infected 390 million of people every year which range is threatened to enlarge due to climate change.

Sleeping sickness treatments have rescued more than 180,000 people. Yet, they started the first vaccinations campaigns in the Philippines and they are close to eradication polio! Furthermore, they developed new innovative treatments for killer diseases like cardiovascular diseases and diabetes and there is more yet to come.

Furthermore, health inequalities is threatening people all around the world (1/3 of world's population have no contact with healthcare support), and Sanofi's task in providing healthcare solution products is just one part of the job to exterminate this crisis.

At Sanofi, they take notice to its stakeholders with a collaborative approach to accommodate the way they operate to meet their commitment to health and bring to light those innovations between people who are in need. To do so, Sanofi partnered with the GAVI Alliance, the World Health Organization, the Bill and Melinda Gates Foundation and many more.

Lastly, the goal at Sanofi goes beyond producing medicine, they are dedicating their work to make a real positive impact in people's lives and as a leader they do not reckon this as a task but as a responsibility.

My thesis work was accomplished (thanks to CIFRE PhD funding), in Sanofi R&D (almost 26% of employees works for R&D), in the Biostatistics and Programming department, supported by the demand and need of Sanofi. Indeed, the idea of the thesis was identified following the problems related to phase II, in particular, dose selection problems, intermediate analyzes to evaluate the sample size, within the framework of saving the recruitment of patients (budgetary saving), without loss of power, etc. These problems are related to different therapeutic areas (Diabetes, oncology, etc.) and may concern several types of drugs.

Résumé

Contexte, Motivation

Les études de recherche de dose constituent une étape majeure du développement de médicaments et doivent être menées avec soin. De façon plus générale, la méthode de détermination de la dose est un facteur clé du succès ou de l'échec de la phase III et des phases ultérieures du développement du médicament. Par exemple, [Sacks et al.2014] affirment qu'environ 16% des échecs de développement de médicaments sont dus à une sélection inadéquate des doses et, plus globalement, une sélection inadéquate des doses a été considérée comme l'une des raisons du déclin observé de la productivité de la R&D dans l'industrie pharmaceutique à la fin des années 90.

Hormis les indications en Oncologie, dans la plupart des cas, les méthodes traditionnelles de détermination de la dose étaient essentiellement fondées sur les procédures de tests multiples: la dose bien tolérée la plus élevée parmi les doses significativement (en utilisant une procédure de test multiple appropriée qui contrôle l'erreur globale de type I) supérieures au placebo, par rapport au critère d'efficacité de l'étude de phase II, a été choisie pour les phases ultérieures. La méthodologie de tests multiples n'étant pas assez robuste et informative, ceci a motivé le développement d'approches plus récentes, fondées sur la modélisation, comme l'approche MCP-Mod [Bretz et al.2005].

Aujourd'hui, la tendance actuelle, dans le même esprit, est clairement de considérer le choix de la dose comme un problème d'estimation statistique et non plus comme un problème d'essais multiples.

Cependant, les méthodologies les plus récentes se sont surtout concentrées sur la caractérisation de la relation dose-réponse liée à l'efficacité (même si la méthodologie MCP-Mod tenait compte d'une fonction dose-réponse non monotone qui fait implicitement référence aux contraintes de limitation de dose liées à la sécurité). Il est maintenant clair que le processus de sélection des doses doit explicitement tenir compte de la toxicité potentielle liée à la dose du médicament.

La motivation de mon travail de thèse était d'examiner le problème de la sélection des doses dans le cadre de la théorie de la décision et d'évaluer les propriétés des procédures de sélection des doses basées sur l'utilité.

Objectifs

L'objectif global de cette thèse était d'aborder le problème de la sélection des doses en développement clinique sous l'angle de la théorie de la décision et des fonctions d'utilité. Les thèmes de la sélection des doses et des études de détermination des doses devaient être étudiés sous leurs différents aspects:

- Règles de décision
- Plans d'étude
- Méthodologie d'analyse statistique

L'objectif de départ de mon travail de thèse était de proposer et d'étudier divers cadres décisionnels fondés sur les fonctions d'utilité et les règles de décision associées. Ces fonctions d'utilité doivent être appropriées et pertinentes par rapport aux décisions que le promoteur doit prendre: concevoir l'essai de phase II (taille totale de l'échantillon de phase II et l'allocation des patients aux doses), définir le moment de l'analyse intermédiaire dans le cas des études "Seamless" de phase II/phase III, décider de poursuivre ou non la phase III et choisir la dose pour la phase III quand cela est pertinent.

Le problème de la conception optimale a déjà été étudié par de nombreux auteurs dans un contexte général ainsi que dans le contexte spécifique de la sélection des doses. Le plus souvent, dans ces travaux, la notion de D-optimalité a été utilisée (basée sur le déterminant de la matrice d'information de Fisher). Ainsi, un des objectifs initiaux de la thèse était de classer les conceptions en fonction de leur capacité à maximiser l'utilité moyenne et d'identifier une conception optimale de cette manière.

L'objectif principal de cette thèse était de proposer une méthodologie statistique spécifique pour analyser les données de l'étude de détermination des doses afin d'éclairer les règles de décision définies dans le cadre décisionnel proposé.

Principaux résultats

Premièrement, dans le contexte d'un plan d'étude "Seamless" de phase II/phase III, nous avons défini un cadre de décision statistique dans lequel le promoteur doit prendre des décisions séquentielles dans le but de maximiser l'utilité future attendue. Pour ce faire, nous avons proposé et discuté diverses formes de fonctions d'utilité: pour chacune d'elles, le calcul de leurs espérances impliquait le calcul de la Probabilité de Succès (PoS) en phase III. En termes de méthodologie statistique, nous avons considéré une approche fréquentiste: le promoteur analyse les données de l'analyse intermédiaire (la partie phase II du Seamless design) à l'aide d'un modèle paramétrique de type Emax via une estimation du maximum de vraisemblance, mais nous avons considéré la possibilité que le promoteur prenne en compte l'incertitude concernant son estimation de la fonction dose-réponse pour prendre ces décisions. Nous nous attendions à ce que ce cadre de travail permette de comparer différents Seamless designs, un plan d'échantillonnage, un 'design' étant défini par le ratio entre la taille de l'échantillon au

moment de l'analyse intermédiaire et la taille totale de l'échantillon, ainsi que par la répartition des patients entre les groupes de dose au début de la phase II. À cette fin, nous avons effectué des simulations d'essai dans le but d'identifier les conceptions "Seamless" optimales, pour certaines des fonctions d'utilité les plus pertinentes parmi l'ensemble des fonctions proposées, mais cet exercice a été abandonné par la suite: les conceptions optimales identifiées étaient, dans la plupart des cas, très proches de la conception équilibrée standard. Mais ce travail a également souligné l'importance cruciale de la taille de la phase II avec, pour certains scénarios, une allocation optimale allouant plus de patients en phase II qu'en phase III, ce qui n'est pas réaliste dans la pratique.

Par conséquent, la partie majeure de la thèse portait sur la situation encore plus fréquente de la sélection des doses dans le contexte de l'étude de phase II de détermination de la dose, avec un échantillon de taille fixe et un plan équilibré. Pour ce faire, nous avons proposé un cadre de décision statistique un peu plus simple que celui mentionné ci-dessus, où les valeurs d'utilité sont attribuées aux doses elles-mêmes, puis indirectement attribuées aux décisions à la fin de l'étude de phase II, et sont égales à la valeur d'utilité de la dose choisie pour la phase III ou une valeur nulle si l'on décide de ne pas poursuivre le développement du médicament après la phase II. Maintenant, le problème du promoteur est de trouver la meilleure dose, c'est-à-dire celle qui a l'utilité la plus élevée.

Nous avons considéré une valeur d'utilité intégrant deux composantes: une composante liée à l'efficacité (la probabilité de succès = la puissance d'une étude de phase III - avec 1000 patients par exemple - de cette dose par rapport au placebo) et une composante liée à la sécurité. Nous avons choisi de caractériser l'efficacité de la dose de cette manière parce que cette définition est directement applicable à tous les types de critères d'efficacité (quantitatifs, binaires, temps jusqu'à l'événement, etc.) et même si elle peut être exprimée de manière équivalente en terme de taille de l'effet, elle a une interprétation plus directe, notamment pour les cliniciens, en termes de puissance pour une étude de référence de phase III. Pour la composante sécurité, nous avons choisi de la caractériser par la probabilité d'observer un taux de toxicité inférieur ou égal à un seuil donné (que nous avons fixé à 0.15 pour nos simulations) en phase III (toujours pour une étude de 1000 patients au total). Cette approche présente l'avantage d'être similaire aux concepts utilisés dans les essais de phase I en Oncologie, qui visent notamment à trouver la dose liée à une toxicité limitante (notion de "Drug limiting Toxicity"). Nous avons élaboré un cadre décisionnel basé sur la fonction d'utilité suivante : $U(d) = PoS(d) \times \mathbb{P}(tox_{obs} \leq 0.15)^2$, mais la même méthodologie serait facilement applicable à d'autres fonctions d'utilité similaires (en considérant par exemple un exposant $k \neq 2$ pour la composante sécurité).

Pour réaliser l'analyse et identifier la dose optimale, nous préconisons l'utilisation d'une méthode bayésienne, au lieu d'une approche fréquentiste de maximum de vraisemblance: elle a l'avantage de fournir un ensemble plus riche de règles de sélection des doses et, par définition de l'approche, de permettre au promoteur d'utiliser des informations externes déjà disponibles.

Pour sélectionner la dose optimale, nous avons proposé une règle de décision originale basée sur la probabilité a posteriori qu'une dose donnée soit la dose optimale. Nos simulations ont montré que cette nouvelle méthode est supérieure à la règle de sélection de dose basée sur le classement des utilités calculées après l'estimation bayésienne des paramètres (moyennes ou

médianes a posteriori) et a des performances similaires à la règle de sélection de dose basée sur le classement des doses par l'utilité espérée a posteriori. De plus, nous pensons que notre méthodologie proposée prend mieux compte de l'incertitude que les dernières approches plus traditionnelles et présente l'avantage d'être facilement adaptée à la définition des règles d'arrêt pour les analyses intermédiaires de l'essai de phase II: nous avons proposé une règle d'arrêt comme "s'arrêter à l'analyse intermédiaire si la probabilité a posteriori que la dose choisie soit la dose optimale est ≥ 0.80 ", qui peut être facilement comprise par une équipe clinique.

En outre, pour améliorer l'efficacité de l'identification de la dose optimale, nous avons proposé une application originale de l'algorithme MCMC de Metropolis-Hastings. Nous avons d'abord formé des moyennes par lots de longueur 150 des valeurs MCMC d'utilité simulées, et nous avons utilisé ces moyennes de lots pour calculer les probabilités a posteriori d'intérêt. Ces versions lissées et plus concentrées de la probabilité a posteriori permettent une différenciation plus facile des doses, et par exemple, une augmentation de la probabilité de s'arrêter aux analyses intermédiaires (sur la base de la règle d'arrêt définie ci-dessus).

Nous avons étudié les propriétés des règles de décision en simulant des essais de phase II de différentes tailles : 250, 500 et 1000 patients. Pour les deux derniers plans (500 et 1000 patients en phase II), nous avons également évalué l'intérêt d'effectuer une analyse intermédiaire lorsque la moitié des patients sont recrutés (c'est-à-dire avec les 250 premiers et les 500 premiers patients respectivement). L'objectif était alors d'évaluer si, pour les essais de phase II de plus grande taille, la possibilité de choisir la dose au milieu de l'étude et de poursuivre l'étude jusqu'à la fin si l'analyse intermédiaire n'est pas concluante, pouvait réduire la taille de l'essai de phase II tout en préservant la pertinence du choix final des doses.

Les simulations montrent que l'estimation d'une dose optimale est une tâche difficile et exigeante. Par exemple, pour la plupart des scénarios présentant un profil d'efficacité satisfaisant, la probabilité de choisir d'aller en phase III suite à une étude de phase II avec 250 patients était le plus souvent égale à moins de 60%. Heureusement, ces probabilités de prendre la bonne décision augmentent avec la taille de l'échantillon. Nous avons constaté, lors de la simulation d'un seul essai (mais avec une taille d'échantillon croissante), que la distribution a posteriori des valeurs de l'utilité se concentrait progressivement autour des vraies valeurs à mesure que la taille de l'échantillon augmentait, mais ce processus peut être lent dans certaines situations. En ce qui concerne le choix de la dose, la probabilité de choisir la bonne dose (conditionnelle à la décision du promoteur d'aller en phase III) augmente également en fonction de la taille de l'échantillon. En termes d'estimation des valeurs de l'utilité elles-mêmes (c'est-à-dire en mettant de côté l'adéquation des règles de décision), nous avons constaté une convergence des estimations bayésiennes (moyenne ou médiane a posteriori) vers la valeur vraie, mais la moyenne a posteriori présentait moins de variabilité que la médiane.

Pour les profils d'efficacité et de sécurité montrant un pic clair de la valeur d'utilité pour une seule dose donnée, une sélection précise de la dose peut être atteinte avec une taille d'échantillon limitée. Si plusieurs doses adjacentes présentent des valeurs d'utilité similaires, l'identification de la dose optimale est plus difficile et nécessite plus de patients.

En effet, quelles que soient les contraintes de temps et de budget, le promoteur a toujours

intérêt à mener de grandes études de phase II pour prendre des décisions précises concernant la fin du programme de développement ou le choix de la dose. Mais, dans la pratique, la taille de l'échantillon de l'étude de phase II est nécessairement limitée par des contraintes budgétaires et des contraintes de temps. Raison pour laquelle nous avons évalué si le fait de permettre une analyse intermédiaire, lorsque la moitié des patients sont recrutés, peut réduire, en moyenne, la taille de l'échantillon de l'essai de phase II tout en préservant les caractéristiques opérationnelles (comme la capacité à prendre les bonnes décisions) du design et des règles de décision. Comme nous l'avons mentionné ci-dessus, nous avons arrêté l'essai à l'analyse intermédiaire lorsque la probabilité a posteriori que la dose choisie soit la dose optimale était ≥ 0.80 . Nous avons également évalué d'autres types de règles d'arrêt, sur la base des valeurs d'utilité estimées lors de l'analyse intermédiaire, et donc moins faciles à interpréter que notre règle d'arrêt retenue, mais celles-ci n'ont apporté aucune amélioration significative par rapport à notre approche. C'est la raison pour laquelle nous ne les avons pas considérées davantage, dans la suite du travail. De plus, nous avons évalué d'autres valeurs seuils, comme 0.70 ou 0.90, mais la valeur de 0.80 a apparemment montré le meilleur compromis entre l'exactitude de la sélection de la dose et la fréquence de terminaison précoce de l'étude à l'examen intermédiaire (le plus souvent l'essai se termine prématurément avec une bonne prise de décision, mieux c'est). Pour certains scénarios, en particulier lorsque la meilleure dose présente un avantage évident en termes d'utilité par rapport aux autres, cette approche présente de bonnes propriétés: avec une probabilité assez importante d'arrêt à l'analyse intermédiaire, elle permet de réduire considérablement la taille de l'échantillon tout en maintenant les propriétés de la grande taille d'échantillon du design fixe. Pour d'autres scénarios, elle est moins utile, car l'étude est rarement interrompue au moment de l'analyse intermédiaire, le promoteur n'étant pas en mesure de déterminer clairement la meilleure dose à ce stade-là. De façon générale, nous avons conclu que ces analyses intermédiaires n'augmentaient que légèrement le risque de prendre à tort la décision d'aller en phase III tout en réduisant la taille moyenne de l'échantillon de l'essai de phase II.

En conclusion, nous avons élaboré un cadre flexible de sélection des doses fondé sur la théorie de la décision et l'utilité pour les études de phase II de détermination de doses, présentant des caractéristiques de fonctionnement satisfaisantes. Ce cadre de travail a également permis de planifier des analyses intermédiaires avec des règles d'arrêt qui peuvent être facilement définies et interprétées par l'équipe clinique dans ce même cadre.

Summary

Context, Motivation

Dose-finding studies are a major milestone in drug development and should be attentively coordinated. More generally, the dose-finding methodology is a key factor of success or failure in the phase III and later phases of the drug development. As an example, it is claimed in [Sacks et al.2014], that approximately 16% of the failed drug development are due to inadequate dose selection and more globally, inadequate dose selection was considered as one of the reason of the observed decline of the productivity of pharmaceutical industry R&D observed in the late 90's.

Apart from the oncology indications, in most cases, traditional dose-finding methodologies were essentially driven by multiple testing procedures: the highest well tolerated dose amongst those significantly (in using a suitable multiple testing procedure that controls the global type I error) superior to the control, with respect to the efficacy criterion in the dose-finding phase II study, was selected for later phases. As the need for more sound and informative approaches became more clear, more recent approaches moved the methodology from the multiple testing methods to the "modeling" based methods with the MCP-Mod [Bretz et al.2005] approach. Now current trend, in this same spirit, is clearly to consider the dose selection as a statistical estimation problem and not anymore as a multiple testing problem.

But still those most recent methodologies mostly focused on characterizing the dose-response relationship related to efficacy (even though the MCP-Mod methodology considered non monotonic dose-response function which implicitly refers to safety related dose limitation constraints). It is now clear that the dose selection process must explicitly consider the potential dose related toxicity of the drug.

The motivation of this work was to consider the dose selection problem in the framework of Decision Theory and assess properties of utility-based dose selection procedures.

Objectives

The global objective of this thesis was to address the problem of dose selection in clinical development with the point of view of Decision Theory and utility functions. The topics of dose selection and also of dose-finding studies were to be studied in their various aspects:

- Decision rules

- Study designs
- Statistical analysis methodology

The first objective of this work was to propose and study various decision-making framework based on utility functions and related decision rules. Those utility functions must be suitable and relevant with respect to the decisions the sponsor has to make: design the phase II trial (total phase II sample size and weights of the doses arm), define the timing of the interim analysis in case of seamless phase II/phase III designs, decide to continue in phase III or not and choose the dose for phase III when relevant.

The problem of optimal design was already studied by many authors in general settings as well as in the specific context of dose selection. Most often, in those works, the notion of D-optimality (based on the determinant of the Fisher information matrix) was used. One of the objectives was to rank designs with respect to their ability to maximize the mean utility and identify an optimal design in this manner.

The last objective of this thesis was to propose a specific statistical methodology to analyze the dose-finding study data in order to inform the decision rules defined within the proposed decision-making framework.

Main Results

First, in the context of seamless phase II/phase III study design, we have defined a Statistical Decision framework in which the sponsor needs to take sequential decisions with the objective of maximizing the expected future utility. For this matter, we proposed and discussed various forms of utility functions: for all of them, the calculation of their expectations involved the calculation of the Probability of Success (PoS) in phase III. In terms of statistical methodology, we considered there a frequentist approach: the sponsor analyzes the data of the intermediate analysis (the phase II part of the seamless design) using a parametric model of the Emax type via maximum likelihood estimation but we considered the possibility that the sponsor takes into account the uncertainty regarding his estimation of the dose-response function to take these decisions. We expected this framework to enable comparisons of different seamless designs, a design being defined by the ratio between the sample size at the interim analysis and the total sample size and also by the distribution of patients among the dose groups at the beginning of phase II. For this purpose, we performed trial simulations with the objectives of identifying the optimal seamless designs, for some of the most relevant utility functions discussed, but this exercise was not fully successful: the optimal designs identified were, in most cases, very close to the standard balanced design. But this work has also highlighted the crucial importance of the size of phase II with, for some scenarios, an optimal allocation allocating more patients in phase II than in phase III, which is not realistic in practice.

Therefore, the major part of the thesis was focused on the even more frequent situation of dose selection in the context of the phase II dose-finding study with a fixed sample size and a balanced design. For this purpose, we have proposed a statistical decision framework a little bit simpler than the one above mentioned: utility values are assigned to the doses themselves,

and then indirectly assigned to the decisions at the end of the phase II study: as being equal to the utility value of the selected dose for phase III or a null value if it is decided not to pursue the drug development after phase II. Now the sponsor's problem is to find the best dose, that is to say, the one having the highest utility.

We considered a utility value integrating two components: an efficacy-related component (the PoS = the power of a phase III trial - with 1000 patients for instance - of this dose versus placebo) and a safety-related component. We have chosen to characterize the efficacy of the dose in this manner because this definition is directly applicable for all types of efficacy criterion (quantitative, binary, time to event, etc.) and even though it can be equivalently expressed in term of effect size, it has a more direct interpretation, especially for clinicians, in terms of power for a reference phase III study with total sample size of 1000 patients for instance. For the safety component, we chose to characterize it by the probability of observing a toxicity rate lower or equal to a given threshold (that we set to 0.15 for our simulations) in phase III (still for a trial of 1000 patients in total). This approach has the advantage of being similar to the concepts used in phase I trials in oncology, which particularly aim to find the dose related to a limiting toxicity (notion of "Drug limiting Toxicity"). We have developed a decision-making framework based on the following utility function: $U(d) = PoS(d) \times \mathbb{P}(tox_{obs} \leq 0.15)^2$, but the same methodology would be easily applicable for other similar utility functions (considering for instance an exponent $k \neq 2$ for the safety component).

For conducting the analysis and identifying the optimal dose, we advocate for the use of a Bayesian method, instead of a frequentist maximum likelihood approach: it has the advantage of providing a richer set of dose selection rules and, by definition of the Bayesian approach, allows the sponsor to use external information already available.

To select the optimal dose we proposed an original decision rule based on the posterior probability for a given dose to be the optimal one. Our simulations have shown that this new method is superior to dose selection rule based on the ranking of utilities computed following Bayesian estimation of the parameters (posterior means or posterior median) and has similar performance than dose selection rule based on the ranking of the doses by the posterior expected utility. However we believe that our proposed methodology better accounts for uncertainty than the latter more traditional approaches, and has the advantage to be easily adapted for defining stopping rules for interim analyses of the phase II trial: we proposed stopping rule such as "stop at interim analysis if the posterior probability for the selected dose to be the optimal dose is ≥ 0.80 ", which can be easily understandable by a clinical team.

In addition, to improve efficiency in the identification of the optimal dose we proposed an original application of Metropolis-Hastings MCMC algorithm. We first formed batch means of length 150 of the mcmc simulated utility values and used those batch means to compute the posterior probabilities of interest. These smoothed and more concentrated version of the posterior probability allow an easier differentiation of the doses and for instance, increase the probability to stop at the the interim analyses (based on the stopping rule defined above).

We studied the properties of decision rules by simulating phase II trials of different sizes: 250, 500 and 1000 patients. For the last two designs (500 and 1000 patients in phase II), we

have also evaluated the interest of performing an interim analysis when half of the patients are enrolled (i.e. with the first 250 and the first 500 patients included respectively). The purpose was then to evaluate whether or not, for larger phase II trials, allowing the possibility of choosing the dose in the middle of the study and continuing the study to the end if the interim analysis is not conclusive, could reduce the size of the phase II trial while preserving the relevance of the final dose choice.

The simulations show that estimating an optimal dose is a difficult and demanding task. For instance, for most of the scenarios with a satisfactory efficacy profile, the probability of making the choice of going to phase III following a phase II study with 250 patients was most often less than 60%. These probabilities of making the good decision increase with the sample size. We noticed, in simulating one single trial (but with sequentially increasing sample size) that the posterior distribution of the utility values progressively concentrated around the true values as the sample size increases, but this process can be slow in some situations. Concerning the dose selection, the probability of selecting the right dose (conditional on sponsor's decision to go to phase III) also increases as the sample size increases. In terms of estimation of utility values themselves (i.e. in putting aside adequacy of decision rules), we noticed convergence of the Bayesian estimates (posterior mean or median) towards the true value, but the posterior mean showed less variability.

For those efficacy and safety profiles that show a clear peak of utility value for one given dose, accurate dose selection can be achieved with limited sample size. In case several adjacent doses show similar utility values, the identification of the optimal dose is more challenging and requires more patients.

In fact, regardless of time and budget constraints, the sponsor has always interest to run large phase II studies to make accurate decisions regarding the termination of the development program or the selection of the dose. But, in practice, the sample size of the phase II study is necessarily limited by budget and time constraints. This is why we assessed whether allowing an interim analysis, when half of the patients are enrolled, can reduce, on average, the sample size of the phase II trial while preserving the operational characteristics (like ability to make the good decisions) of the design and decisions rules. As mentioned above we stopped the trial at the interim analysis when the posterior probability for the selected dose to be the optimal dose was ≥ 0.80 . We also assessed other types of stopping rules, based on the utility values estimated at the interim analysis, and therefore less easily interpretable than our retained stopping rule, but they did not bring any significant improvement compared to our approach, this is why we did not consider them further. Moreover, we assessed other threshold values, 0.70 or 0.90, but the value of 0.80 showed apparently the best trade-off between the accuracy of the dose selection and the frequency of termination of the trial at the interim examination (the more often the trial is terminated prematurely with the good decision-making the better). For some scenarios, in particular when the best dose shows a clear benefit in utility as compared to the others, this approach has good properties: with a quite large probability of study termination at interim analysis, it enables to reduce the sample size while maintaining the properties of the fixed large sample size design. For some other scenarios, it is less useful as the study is rarely terminated at the interim analysis, the sponsor being unable to clearly identify the best dose at interim analysis. In general, we concluded that those interim analyses

only slightly increased the risk of wrongly taking the decision to go to phase III while reducing the average sample size of the phase II trial.

To conclude, we developed a flexible Decision Theory/ utility-based dose selection framework for phase II dose-finding studies that has satisfactory operating characteristics. This framework allowed also to plan interim analyses with stopping rules that can be easily defined and interpreted by the clinical team within this same framework.

Chapter 1

Introduction

1.1 Drug Development

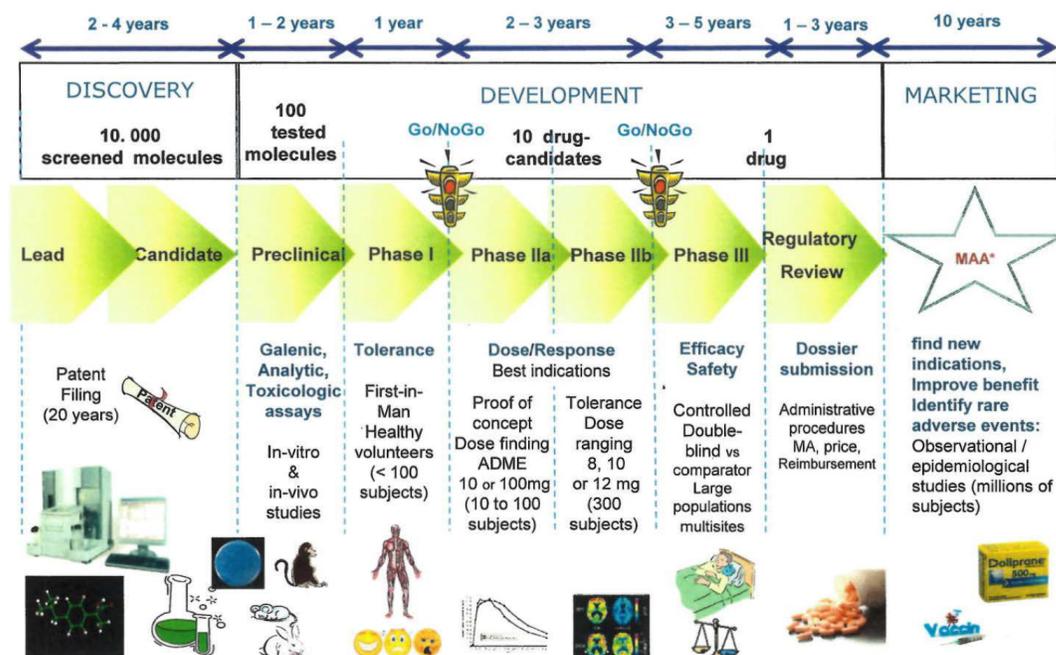


Figure 1.1: Summary of clinical trial phases (Source: Sanofi).

Before digging into the thesis subject, here are the steps for a drug development in clinical research (see Figure 1.1 above).

There are mainly three major axes in the clinical development of a new drug: the research phase (pre-clinical phase including biological and animal research), the development phase (research in human: phase I, phase II and phase III), and the marketing part (phase IV).

After the discovery phase, preclinical and animal testing studies are conducted (1 to 2 years), followed by phase I of tolerance and human volunteer trials (less than 100 subjects for a duration of about one year), then "Go / NoGo" decision step for Phase II, i.e. the decision to perform (or not) the late phase study.

Phase II is often divided into two parts for a duration of 2 to 3 years: phase IIa and phase IIb. The phase IIa consists in a relatively small (from 10 to 100 patients) Proof of Concept (PoC) study where the aim is to confirm the pharmacological or clinical activity of the drug, very often with a high dose. The phase IIb often consists in a dose-finding study, where the objective is to identify the optimal dosage of the product in terms of efficacy and tolerance.

Then, there is the "Go / NoGo" step to decide whether or not to proceed with phase III. In phase III, which usually takes 3 to 5 years, the trials are conducted on several thousand patients representative of the population to which the treatment is intended. These are comparative trials in which the developing drug is compared to an already marketed effective treatment or to a placebo, i.e. a treatment without pharmacological activity.

Finally, there is the phase IV, where trials of this phase are carried out once the drug is marketed, on a often very large number of patients (up to tens of thousands of people) for a duration of about 10 years. These trials allow to deepen the knowledge of the drug under the actual conditions of use and to evaluate its tolerance on a large scale. Pharmacovigilance thus makes it possible to detect very rare undesirable side effects (also called adverse events) which could not be detected before, during other test phases.

1.2 Methodologies for dose selection in drug development

Definition of the dosage and dosage schedule is a key question during clinical development of a new drug, and it is the objective of the so-called dose-finding studies. It is now well documented that poor dose selection is a root cause for failures or delays in drug approval [Sacks et al.2014]. The goal is to satisfy the requirement that patients should be exposed only to the amount of drug that they really need. In other words, the idea is to determine an acceptable dose level and ideally, the characterization of the dose-response relationship. Safety problems will definitely be generated when the dose is too high, although a failed program will result due to a difficulty of authenticating a satisfactory effectiveness, when the selected dose is too low. Here comes the major role of dose-finding studies in drug development, these studies should be rigorously implemented.

Traditionally, apart from oncology indication, the search of the optimal dose resulted from a sequential process: first the set of efficacious doses (it could include only one dose or in the worst case scenario, none) was identified and second, the highest dose considered as "safe" or "well tolerated" was selected for the late development phases. Also, the first step related to the identification of the efficacious doses were driven by multiple-testing procedures: the set of efficacious doses was defined as the set of doses that were significantly different from placebo in the dose finding study after adjustment for multiplicity. Various multiple-testing procedures can be considered: Dunnett's procedure is widely used for the quantitative variables, more

recent general gatekeeping procedures [Dmitrienko and Tamhane2007] are also used.

A more recent approach that requires the assessment of the dose-response (for efficacy) relationship is the Multiple Comparison Procedure and Modeling (MCP-Mod) [Bretz et al.2005]. It uses a predefined set of candidate models for the dose-response relationship. Once the evidence of a drug effect is established at the MCP step using multiple contrast tests, a Mod step is used to estimate the dose meeting the expectations of the sponsor.

It is now becomingly accepted that finding the right dose should be rather considered as an estimation problem than a multiple testing problem [Li et al.2017]. This latter traditional approach, as well as the more recent MCP-Mod procedures generally consider efficacy and safety sequentially: doses associated with statistically significant differences versus the control, for the multiple testing approach, or doses with desired difference versus control, for the MCP-Mod approach, are identified first and then the highest dose amongst them considered as "well tolerated" is generally chosen. An alternative approach should rather rank the doses using efficacy and safety assessments simultaneously, via utility functions.

On the other hand, in many settings the dose selection is mainly driven by efficacy. In absence of safety considerations, one typically searches for the dose which is near the plateau, e.g. the dose reaching 90% or 95% of the maximal efficacy denoted by ED90 and ED95. This holds for monotonic dose-responses. Higher doses will unduly expose the patients to potential toxicity issues while lower doses may represent a substantial loss of efficacy. Another dose of interest is the Minimum Effective Dose (MED), i.e. the smallest dose associated with a statistically significant and clinically relevant effect. The range of doses between the MED and the ED90/ED95 constitutes the interesting dose zone. When serious safety issues arise within this interesting zone, the dose selection becomes more challenging and involves multiple criteria. Additional toxicity may counterbalance a gain of efficacy and one needs to introduce some utility score balancing both efficacy and safety.

This thesis aims to extend the modeling approach based on Decision Theory and utility functions to optimize decision rules (dose choice) and designs. The choice of the utility function approach was driven by Decision Theory [Savage1954] that claims that utility functions are the most natural and consistent way to describe and rank preferences or decisions.

1.3 Designs for dose-finding studies: fixed and adaptive designs studies

The standard dose finding study is parallel arms study with 4 or 5 doses and a control group (in general a placebo group) with a balanced design. One of the objectives of the thesis is to propose utility functions based decision rules in this fixed sample size design but also use utility-based considerations to optimize the design as well. The problem of optimizing the design for the purpose of dose selection has already been addressed (see for example [Bretz et al.2010] where the aim is to establish efficient study designs to estimate relevant target

doses) but the methodology is most often based on C- or D-optimality. Some work dedicated to the optimization of designs based on utility functions exist but they are rather sparse: we can mention [Foo and Duffull2017], in which the design is optimized by minimizing the expectation of a cost. But as of today, one lacks of a unified framework based on Decision Theory aimed at assessing and comparing several dose-finding strategies and designs.

For several years adaptive designs became more and more popular in the pharmaceutical industry and in particular much attention was brought on adaptive seamless designs. In fact, these designs offer the prospect to incorporate different drug development phases into a single trial, allowing the use of accumulated data and information of the trial, and consequently a promising dose-finding strategy. Because of the growing importance of those designs we aim also to develop a utility-based decision-making framework (for dose selection) for those designs also. A corollary objective is to propose some utility-based tools to optimize the designs: first in terms of ratio between phase II and phase III sample sizes, and, second, in patient allocation to doses at the beginning of phase II.

1.4 Objectives of the thesis

The global objective of this thesis is to address the problem of dose selection in clinical development with the point of view of Decision Theory and utility functions. The problem is to be addressed in its various aspects:

- Use Decision theory and utility functions to rationalize and optimize decision-making related to the choice of dose
- Use this same framework to optimize the design of the phase II trial. This work is conducted in both contexts of fixed sample size phase II trials and adaptive seamless phase II/phase III trials

So this thesis starts by a comprehensive review of the bibliography addressing the utility approach, adaptive designs and dose optimality: the literature review is summarized in **Chapter 2**.

The first section of **Chapter 3** is devoted to the common materials and methods applied to most of the work contained within this thesis. It describes all the required denotations, as well as the mathematical formalization of the efficacy dose-response modeling approach, along with the model-based framework, including Probability of Success (PoS) computation in particular.

Because of the growing importance of adaptive designs in the pharmaceutical industry, we consider this framework first. So in the following sections of **Chapter 3**, we study the problem of optimizing the dose selection as well as the design, in the context of adaptive seamless designs: for this latter objective the aim is to identify the best timing for the interim analysis (ratio between the sample size at interim analysis and the total sample size) and the optimal allocation of patients within the doses arms at start of the study. We first propose and discuss several types of utility functions and then we assess, through simulations, their ability to identify an optimal design.

The **Chapter 4** is dedicated to the problem of optimizing the dose selection in the context of phase II trial with a fixed sample size design, this is the most frequent case in practice. For this purpose, a two-component utility-based approach is proposed to optimize the dose selection process in order to maximize the PoS in phase III. The first component is for efficacy and the second component is for safety. So we consider, for a dose d , a utility function of the form: $U(d) = (\text{efficacy term}(d))^h \times (\text{safety term}(d))^k$, and we propose decision rules based on a Bayesian methodology. Properties of these decisions rules are compared using simulations assuming various efficacy and safety profiles of the drug. More precisely, we propose a dose-ranging trial (phase IIb) comparing J doses of a new product versus placebo followed by a pivotal phase III trial with a single dose selected versus placebo. Efficacy is characterized by a unique continuous endpoint which is supposed to be the same in phase II and in phase III. Safety is modelled using a binary endpoint, with "0" denoting no toxicity and "1" denoting the presence of toxicity, mimicking the Drug Limiting Toxicity (DLT) commonly used approach in oncology. In general, definition of toxicity will typically be project dependent but as a generic example one may think of permanent treatment discontinuation due to Adverse Event (AE). Toxicity could also be defined as a composite endpoint involving multiple AEs of interest.

The key question to which this thesis aims to respond is to know, the sample size of phase III being fixed, how the sample size of the phase IIb study impacts the quality of dose selection and the chances of successful development, in order to provide recommendations on this sample size.

Moreover, a sequential design (with futility and efficacy rules at the interim analysis) is compared to a fixed design in order to allow one to hasten the decision to perform the late phase study. Operating characteristics of this approach are extensively assessed by simulations under a wide range of dose-response scenarios.

In **Chapter 5**, we assess the properties of the posterior distributions (posterior means or posterior medians), by simulating trials with sequentially increasing sample size, and we graphically examine (through violin plots) the properties of the estimated posterior distributions of the utility values. We also perform exploratory analyzes, comparing different designs, alternative decision rules for dose selection, relative utility loss criterion to make recommendations on phase II sample size, and several stopping criteria for interim analysis.

Finally, **Chapter 6** summarizes our Bayesian decision-making framework, addressing the proposed method, and discussing the choice of the utility function, thresholds related to decision rules, and decision criteria for interim analysis. Some perspectives are highlighted at the end of the chapter, suggesting, for instance, prior assessment to guide the sponsor and improve the decisions, and advocating a re-evaluation of the choice of dose-response models, in terms of robustness, with the possibility to perform a Model Averaging approach.

Chapter 2

A brief literature review

2.1 Adaptive designs

In general, adaptive designs consist in using information and accumulated data in a trial, in order to modify some aspects of the design without compromising the validity of the study [Pallmann et al.2018]. Thus, the defined changes can then be implemented according to the interim analyses results, carried out at some planned time points in advance. In practice, an adaptive design allows one to modify characteristics such as the number of patients, the possibility of combining two phases in the same trial, the eligibility criteria in the trial, randomization rules, treatment's dose and duration, treatment groups (closure, enrichment, etc.), or even evaluation criteria.

It is worth noting that an adaptive seamless design [Pallmann et al.2018]; [Bretz et al.2006]; [Bauer and Kieser1999] is a particular definition of adaptive designs, and is used to combine two clinical studies into a single study, it could be a seamless phase I/II design, or a seamless phase II/III design. For example, a seamless phase I/II design consists in combining safety and activity assessment into one trial. A seamless phase II/III consists in combining selection and confirmatory stages into one trial, "Seamless phase II/III designs are aimed at interweaving the two phases of full development by combining them into one single, uninterrupted study conducted in two stages" [Bretz et al.2006]: at the interim (after stage 1), one selects a treatment for instance, this treatment is maintained together with the comparator arm (placebo or control arm) in stage 2. The final comparison of the privileged treatment with the comparator arm may incorporate both stages patients (stage 1 + stage 2), and is fulfilled while controlling the multiple type I error rate at a pre-specified level α .

In [Bretz et al.2006], authors discussed flexible test procedure allowing for adaptively selecting hypothesis at interim and allowing combination of learning and confirming in a single seamless trial. In other words, they discussed all necessary tools required for an adaptive design implementation: adaptive tests, hypothesis selection, combination tests, the closure principle and multiple testing in adaptive designs. [Bretz et al.2006] is deeply rooted in the frequentist framework, mentioning that the flexibility of the adaptive approach allows use of Bayesian decision tools in the interim.

Confirmatory adaptive designs with Bayesian decision tools for a targeted therapy in oncology were assessed in [Brannath et al.2009]. Authors used adaptive multiple test design: combination tests for time-to-event data (examples of time-to-event data could be: time from entry

to hospital until discharged, time from being registered into a research study until the disease of interest is occurred, time from an examination and diagnosis of a disease until death, etc.). They used Bayesian decision tools as well, defining decision rules and thresholds and calculating posterior and predictive probabilities. In the same context, a Bayesian adaptive approach was proposed in [Grieve et al.2013], for pulmonary hypertension disease, in order to select a dose for phase III, accounting for both efficacy and safety modeling as well. Authors put more emphasis on such designs, and recommend the improvement of methodological aspects compared to traditional designs, with the purpose of modifying some characteristics such as sample size re-estimation, patient-enrichment, and other components intending to reach the goal, in an optimal and efficient way, to select the final dose.

Adaptive (seamless) designs have always been essentially used in the last two decades, in order to select doses. Recently, these designs were used in a different context, which is the optimization of stage 2 by optimizing patients allocation, see [MacCallum and Bornkamp2015]; [Geiger et al.2012]; [Christen et al.2004]; [Grieve2017]. In their most general definition, seamless designs can be described in Figure 2.1 below.

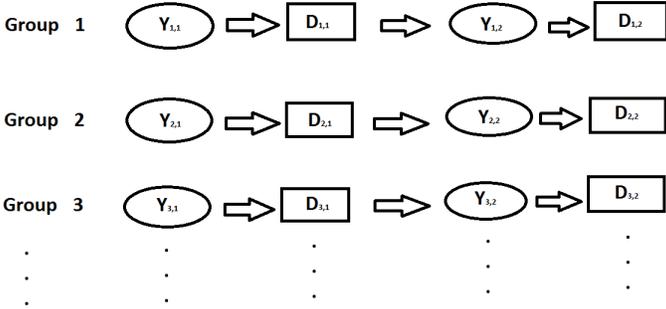


Figure 2.1: Clinical trial steps.

This figure illustrates the different stages of a general seamless design. There are different stages (usually 2 as in the graph), each step ending with a decision impacting the next stage and the study ends with a final decision:

- $Y_{i,j}$ represents the data of the i th arm at step j .
- $D_{i,j}$ is the decision on the i th arm at the end of the j step. The decisions $D_{i,j}$ can impact the data $Y_{i,j+1}$ of the following step, in very different ways:
 - The decision may simply consist of stopping or keeping the i th arm in the study.
 - The decision may consist of changing the sample size of the i th arm (and possibly all the arms).
 - The decision can change the probability of patient allocation in the i th arm for the next step $j + 1$.

This work was explored at the beginning of my thesis (see **Chapter 3** and Appendix A.1): no real gain was noticed when optimizing patient's allocation to doses, and according to several scenarios, the optimal design was almost the balanced design. So, optimizing the dose allocation ratio in stage 2 of the dose-finding study offered very little improvement in regard to significantly increased operational complexity and consequently, this optimization part was removed from the scope of this thesis (see **Chapter 3** and Appendix 1).

Note that the dose allocation optimization for stage 2 may be handled in two ways: based on observed stage 1 data only (by applying adaptive randomization, for example see recent publications such as [Lin et al.2016]; [Geiger et al.2012]; [Christen et al.2004]; [Grieve2017]) or based on observed stage 1 data and expected data in stage 2 (by computing the PoS or the conditional power [Li et al.2002] for example). In [Grieve2017], the author discusses adaptive designs, Bayesian designs and adaptive randomization based on the medical literature, highlighting some specific case studies, basically aiming to select doses, with a particular interest in using both efficacy and toxicity information. However, despite all the highlighted benefits, the author does not hesitate to emphasize the disadvantages that can result from such designs (including different sources of bias, such as time bias, accrual bias, etc.).

Conversely, it is noted that despite some disadvantages of adaptive designs, these designs are more robust than conventional designs, and have less impact on type I error and other sources of bias, see [Bretz et al.2009] for example, where multiple topics were discussed, such as adaptive tests for a single null hypothesis, conditional error function for adaptive combination tests, multiple testing in confirmatory adaptive designs and closure principle.

In [Geiger et al.2012], the work includes highly adaptive designs and dose-response/dose-toxicity detection. The interesting part of this publication was the real application in a particular therapeutic axis, more precisely, type 2 Diabetes mellitus. The aim of the study was to apply a Bayesian adaptive plan for stage 1 to allocate patients to doses, and a fixed plan for stage 2, in order to select two doses of 'Dulaglutide' for phase III. Dose-selection was based on decision rules along with safety and efficacy modeling. Sample size re-estimation was also performed in order to increase the probability of showing superiority to a comparator arm (an effective treatment). Advantages of their implemented adaptive design have been clearly highlighted, compared to a classic fixed design. The gain offered by deviating from the classical balanced randomization scheme was marked by the possibility to optimize the dose allocation ratio for stage 2. In addition, authors demonstrated efficacy using all accumulated data before and after the adaptive interim analysis while preserving a strong control of type I error. Additional trials were conducted to evaluate the performance of the recommended doses of the drug.

A new approach aiming to handle these problematics (safety and efficacy combination, dose selection, etc.) has been proposed in [Christen et al.2004]; [Houede et al.2010]; [Gajewski et al.2015]; [Kirchner et al.2016]. The idea is based on maximising a utility function. Utility functions represent the benefits of the stage 2 and the final dose recommendation. In [Gajewski et al.2015] and [Christen et al.2004], a Bayesian adaptive design is built according to the utility approach (see Section 2.4).

This approach will be discussed in more detail in Section 2.4, where the utility concept is defined, along with some examples of utility functions.

2.2 Modeling procedures

2.2.1 Dose-response modeling

Modeling the dose-concentration-effect relationship is an important aspect in clinical research studies, and can be done using various dose-response models [Pineiro et al.2014]; [Ting2006], linear or nonlinear, such as: linear, logistic, Emax (simple, Sigmoid or standardized form), quadratic, exponential and linear in log-dose models.

The most frequently used model is the Emax model. Its simplest form can be given by $f(d, \theta) = E_0 + E_{max} \times d / (ED_{50} + d)$ where d is the dose, θ represents the model parameters vector to estimate. In other words, $\theta^t = (E_0, E_{max}, ED_{50})$, E_0 is the placebo effect, E_{max} is the maximum effect change of a dose compared to placebo, and ED_{50} is the dose giving half of the maximum effect change. A more complex form of this model is its Sigmoid version, where an additional parameter is included, ensuring greater shape flexibility. It has the following form: $f(d, \theta) = E_0 + E_{max} \times d^g / (ED_{50}^g + d^g)$ where g is called the "Hill" exponent (the slope) reflecting the shape of the dose-effect curve.

Statistical analysis of dose-response curves may also be performed by a linear model, having the following general form: $f(d, \theta) = E_0 + \delta \times d$, where $\theta^t = (E_0, \delta)$. However, the concept of linear dose-response relationship has its disadvantages since it may not apply to non-linear situations, see [Vandenberg et al.2012].

Another possible model to be considered would be the linear in log-dose model, having the following form: $f(d, \theta) = E_0 + \delta \times \log(d + c)$, where $\theta^t = (E_0, \delta)$, E_0 represents the placebo effect, c represents a positive fixed value aiming to avoid placebo problems (when $d = 0$) and δ represents the slope associated with $\log(d + c)$.

Other models may be more appropriate, depending on the circumstances such as the exponential / power model, having the following form: $f(d, \theta) = E_0 \times \exp(d/\delta)$, where δ handles the increase or decrease rate in the effect (δ could be either positive or negative). This model may detect the existence of a convex or sub-linear dose-response relationship.

On the other hand, another dose-response model could be the quadratic one, having the following form: $f(d, \theta) = E_0 + \beta_1 d + \beta_2 d^2$, where $\theta^t = (E_0, \beta_1, \beta_2)$. This model has the advantage of detecting the existence of a non-monotonic dose-response relationship, in a convex/U-shape (when $\beta_2 > 0$), or in a concave/umbrella-shape (when $\beta_2 < 0$).

Finally, there is the logistic model, with the following form: $f(d, \theta) = E_0 + E_{max} / (1 + \exp[(ED_{50} - d)/\delta])$, where $\theta^t = (E_0, E_{max}, ED_{50}, \delta)$. Note that parameters E_0 , ED_{50} and E_{max} have the same definitions as the Emax model previously described. δ controls the change with dose rate in effect. This model has the advantage of allowing both effect increase and decrease (E_{max} parameter could be either positive or negative).

It is worth noting that in general, the maximum likelihood is applied in order to estimate the model's parameters.

The Emax model is the most frequently used one for efficacy in dose-finding framework [Miller et al.2007]; [Comets2010]; [Pineiro et al.2014]. This model assumes a monotonic (either increasing or decreasing) dose-response. It also offers the possibility to estimate the maximal treatment effect, the dose which produces 50% of the maximal effect and the placebo effect. This is the main reason that led us to model our efficacy via an Emax dose-response model in the following chapters of this thesis. However, since it is also known that the Emax model (es-

pecially in its Sigmoid version [Pinheiro et al.2014], i.e. with an exponential parameter known as 'Hill' exponent, which determines steepness) can be very difficult to estimate, with much imprecision, if the design is not rich enough [Dutta et al.1996], we chose a three-parameter Emax function to describe our dose-response model (i.e. exponent=1) in the following work of this thesis (see **Chapters 3, 4, 5**, and Appendix A.1). Speaking of the difficulty of estimating the Emax model, it is specified in [Dutta et al.1996] that parameters estimation of a Sigmoid Emax model definitely leads to an eventual imprecision (i.e. the amount or degree of random error in a calculation, usually represented by the standard deviation, coefficient of variation, or range), and if the concentration-effect relationship extremes are not empirically investigated, this imprecision could be enormous. However, these extremes are not firmly examined in general, due to safety issues. Note that, as specified in [Dutta et al.1996], when parameters of the model are used as metrics, pharmacometric imprecision (as a result of a deficient study design) can be an essential cause of the common considerable variability stated in pharmacodynamics. If one seeks to reduce the variability in pharmacodynamics characterization, using data descriptors based on parameters of the estimated model could be a good solution. In fact, in pharmacodynamics, when testing clinically significant covariates, using data descriptors could be a powerful indicator for these tests.

In the modeling approach, one seeks to validate a given model, which is usually a difficult step in a clinical study. To do this, several methods could be applied. A possibility to choose a given model among a set of candidate ones, is to apply a model-selection approach. The model with minimum criteria value is chosen as the best model to fit the data. The most used criteria to select models are: AIC (Akaike's Information Criterion, [Akaike1975]), BIC (Bayesian Information Criterion, [Schwarz1978]) and AICc (Second-order Akaike Information Criterion, [Snipes and Taylor2014]). These criteria are briefly described below.

The AIC criterion has the following form: $AIC = -2\log L(\hat{\theta}) + 2 \times \dim(\theta)$, where θ is the vector of the model's parameters, $\hat{\theta}$ is the estimate of θ , $L(\hat{\theta})$ is the maximized likelihood function and $\dim(\theta)$ is the number of free parameters in the model.

The BIC criterion is defined as: $-2\log L(\hat{\theta}) + \log(n) \times \dim(\theta)$, where n is the sample size and $\log(n)$ is the penalty for the number of parameters depending on the model selection criterion. The AICc was proposed in order to avoid overfitting issues: there is a large probability that AIC may overfit (by selecting models that have too many parameters) when the sample size is small. So AICc is an AIC with a correction for small sample sizes.

The AICc criterion has the following form: $AIC + \frac{2\dim(\theta)^2 + 2\dim(\theta)}{n - \dim(\theta) - 1}$. AICc is accordingly an AIC with an additional penalty term for the parameters number.

The choice of a particular criterion may be conditioned by the final objective of the analysis. From the literature, AIC seems to perform better than BIC in selecting a predictive model, see [Chakrabarti and Ghosh2011] for example.

Instead of selecting a specific model as described above, one may prefer to consider a weighted combination of the candidate models. This approach is presented in the following subsection.

2.2.2 Model Averaging approach

Since it is possible to select the dose with a bad dose-response model, one might search for more robust models than others, and this is where Model Averaging becomes more interesting. It is shown that Model Averaging is better than obtaining estimates from one model

only [Turkheimer et al.2003] when the true model is unknown.

Model Averaging has the advantage of substituting model selection (even though this might be a heavy approach in some cases, depending on the situation, when one already has too many parameters to estimate) by allocating weights to different pre-specified models of interest. For example, for a given function F of the mean difference versus placebo Δ , with the Model Averaging approach, the sponsor chooses the dose $d^* = d_{j^*} : j^* = \arg \max_{j \leq J} \sum_l w_l F(\hat{\Delta}_{d_j}^{M_l} / s)$ where $\hat{\Delta}_{d_j}^{M_l}$ is the mean difference between the dose d_j and the placebo estimated with the M_l model, and where w_l are model weights to specify or compute.

These weights can be all equal (balanced), or can be calculated with AIC or log-likelihood [Liukko2007].

It is worth noting that there is the R package `MuMIn` [Bartoń2018] for model selection and Model Averaging based on information criteria (AIC, AICc, log-likelihood and alike).

The Bayesian Model Averaging (BMA) [Hoeting et al.1999] is the most used method for weight specification. Let's say one has only one true model from several potential models. One specifies a prior over parameters for each potential model, followed by a prior probability that each one of them is the true model. At that time, the posterior distribution is the weighted average of the individual models meaning that the weights are Bayesian posterior probabilities that the given model is the true model, conditional on the data.

Likewise, the use of the square error loss in the derivation of optimal weights for Model Averaging is investigated in [Nguefack-Tsague2014]. One does not really know if there is an optimal weighting scheme or not. Numerous authors have proven their existence in others methodological frameworks, for example see [Liukko2007]; [Turkheimer et al.2003]; [Giattin et al.2007]; [Forster2000]. In [Nguefack-Tsague2014], the investigations continue its work on the derivation of optimal weights for Model Averaging using square error loss. The presence of these weights can be proved on a theoretical level, however, when one tries to estimate these weights, one can clearly see that they are no longer optimal. An example of linear regression is used to demonstrate that the estimated weights concluded by those Model Averaging estimators probably do not have a better efficiency compared to post-model selection and other averaging estimators [Nguefack-Tsague2014]. A theoretical justification is provided for this phenomenon in their publication [Nguefack-Tsague2014]. One must therefore be wary of the least squares approach, especially for the purpose of deriving optimal weights.

Once dose-response effects are detected, based on model selection or Model Averaging presented in Section 2.2, and depending on specific testing procedures such as MCP-Mod (multiple comparison procedures and modeling, this part is not developed in this thesis) [Thomas2017]; [Jones et al.2011]; [Pineiro et al.2014], one may then proceed with the determination of optimal designs, described by optimal doses d_1, \dots, d_J and their weights w_1, \dots, w_J , by defining some specific algorithms, according to different optimality criteria. Here, d_1, \dots, d_J are assumed to be already determined. So the problem is now to maximize the objective (e.g. D-optimality, defined below) over the $K + 1$ -dimensional vector (w_1, \dots, w_J) which is a constrained optimization problem.

All these possible strategies will be listed in the following section.

2.3 Optimization procedures

2.3.1 Optimal designs

In [Wong1994], robust properties of several optimality criteria such as A, C, D, E, T and G-optimal designs are compared and the relationships among these optimality criteria were studied in [Rady et al.2009]. The most common criteria used are D and C-optimality, where D-optimality does not address the specific issue of identifying at best a pre-defined dose of interest (e.g. MED (Minimum effective dose), ED50 (dose giving half of the maximum change in effect), ED90 (dose giving 90% of the maximum change in effect), etc.) but instead maximizes the determinant of the global Fisher information matrix. Conversely, C-optimality consists in minimizing the variance of the best linear unbiased estimator of a predetermined linear combination of model parameters. The linearity constraint might be a genuine issue when the dose of interest is a rather complex function of model parameters. This is notably the case when e.g. one searches for the dose which achieves a minimum pre-defined efficacy versus placebo within the Emax model.

An algorithm for searching the best two-stage designs under different criteria (for example, designs that minimize the maximum sample size, optimum and "minimax" designs) was proposed in [Tsai2006].

Practical considerations for optimal designs in clinical dose-finding studies were also discussed in [Bretz et al.2010]: in this paper, the aim is to establish efficient study designs to estimate relevant target doses. Authors consider two optimal designs, derived from C-optimality criterion. This criterion is used to minimize a function depending on the minimum effective dose, or to minimize a function depending on the dose achieving a certain percentage ($100p\%$, $0 < p < 1$) of the maximum treatment effect. They also compare these optimal designs with D-optimal designs for two dose-response models: linear and Emax models. Furthermore, they discuss extensions to robust designs accounting for model uncertainty. A particular issue could be the misspecification of the regression model. In fact, when the true model is unknown, the problem of uncertainty could be generated, and therefore, it is better to consider robust designs under uncertainty, where the concept of robustness here is highlighted by its ability to prevent its functioning from being affected by noises, coming from different sources.

In [Miller et al.2007], authors talk about how to construct a "good" design: acquiring valuable designs adopting the best design theory helps us determine a suitable model for describing the dose-effect relationship.

One needs to keep in mind that only one or two doses must be adopted in phase III. Numerous sources will help achieve presumptive knowledge or guesses based on prior information, mentioning the preclinical information, data with similar drug and earlier trials with the same drug [Miller et al.2007]. Sometimes, it is also possible to formulate assumptions for expected effects, which is not an obvious thing to do, even hard to advise, requiring a certain level of expertise, and a consensus with a panel of experts (doctors, clinicians, statisticians).

But a key point here, is to identify the main purpose of dose-finding trial. The purpose could be for instance, the recommendation of one dose, after checking the efficacy level and the toxicity rate, so basically after evaluating the benefit / risk ratio. However, these targets and recommendations may be exposed to unexpected changes due to external factors, or updates of some specific information/data during the development phase [Miller et al.2007].

Another interesting publication [Musuamba et al.2017] presents some guidance and methodologies establishing a bundle of interesting materials and methods for regulators and drug developers. These methods include three methods for study design optimization: clinical trial simulations, Fisher information matrix (FIM)-based methods (D-optimality), and adaptive studies. In addition, there are five advanced methods for data analysis (empirical regression models, MCP-Mod, Model Averaging, quantitative systems pharmacology models and pharmacometrics models). There is also evidence of pairwise comparisons (mainly because of historical reasons). In fact, this paper confers the challenges faced when implementing these methods alongside with their benefits and disadvantages. Obviously, it was confirmed that the selection of the dose for phase III must not be approached by hypothesis testing because it is an estimation issue. The precise choice of what method(s) to adopt is not advised nor suggested as a systematic scheme for attaining some particular object and will rely on lots of features. Therefore, one can state that dose selection for phase III trials should be enlightened by appropriate and well-planned dose-finding studies. Thus, methods should be adopted together to establish a reliable and accurate study of the dosing regimen philosophy as well as dose-finding reasoning, because there are a lot of beneficial and applicable methods which are consistent.

In [Musuamba et al.2017], authors underlined some applications in several therapeutic areas, such as oncology, immunology, etc.

Despite a few limitations, PK / PD modeling (pharmacokinetic / pharmacodynamic modeling [Meibohm and Derendorf1997]) appears to be well justified and valued in these areas, due to its ability to describe dose-response relationships, through mathematical equations such as linear or Emax models, and its possibility to account for the temporal evolution of the intensity of the dose-response effect. Authors also discussed antibiotics and dose-finding in special populations (children and elderly people) in [Musuamba et al.2017]. In such populations, several factors are related, like age, organ degradation or evolution, and such factors should not be neglected because of their potential effect on PK/PD. In fact, these factors could cause the risk of bleeding, or some dysfunctions at certain levels.

2.3.2 Bayesian optimal designs

A large number of dose-finding studies are conducted in a Bayesian framework. In this context, one can apply both non-adaptive or adaptive Bayesian optimal design, but this is where an important question arises: should one prefer an adaptive or a non-adaptive Bayesian optimal design?

A non-adaptive Bayesian optimal design is obtained by implementing the theory of optimal design. It is a design that does not vary according to information acquired from an interim analysis of the available data.

An adaptive Bayesian optimal design is defined by considering the contribution of supplementary information about the dose-response relationship when expanding it upon an interim analysis which is often performed in the dose-finding trials.

As proposed in [Miller et al.2007], optimal designs could be based on the asymptotic variance of the nonlinear least squares estimate of the dose-effect function at dose d compared to placebo, via a non-adaptive Bayesian optimal design or an adaptive Bayesian optimal designs deriving

a posteriori probabilities for each scenario (an adaptive Bayesian optimal design arises here because of the adjustment made on the allocation ratios for the dose arms). Authors seek to compare the adaptive Bayesian optimal design with the efficiencies of the non-adaptive one. Although an adaptive design is more challenging from an operational angle, however, and as reported in [Miller et al.2007], no one has ever given a reasonable and satisfying disadvantage of the non-adaptive Bayesian optimal design regarding the effectiveness compared to the adaptive one. Adopting a non-adaptive design could help us avoid a slight effect on the distributional abilities and quality of the estimates produced by the adaptation. Accordingly, in this situation, it is more likely preferable to adopt a non-adaptive Bayesian optimal design [Miller et al.2007].

Practically, adopting a non-adaptive Bayesian optimal design without changing the design for stage 2 can be for example implemented if an interim analysis is necessary because of ethical reasons or any other reason.

On the other hand, the adaptation causes some deviations from the normal distribution, yet, these deviations are not so crucial. Thus, if the adaptive design shows good efficiency gain then apparently, one needs to welcome this matter positively. Similarly, on a practical level, it is more likely to perform an adaptive design at the expense of other distributional properties of the resulting estimates, if there is a confirmation of these results from an effectiveness perspective [Miller et al.2007].

As elaborated in [Miller et al.2007], if the adaptive design has advantages over a non-adaptive optimal design, do these advantages count principally on the accuracy of the deductive knowledge? If the deductive information is perfectly false, the application of adaptive designs may then have advantages. Hence, either the first part in the trial itself (stage 1) must be adopted or a distinct pilot trial must be performed to derive the final design of the trial. In [Miller et al.2007], authors define these approaches when there is definitely no information concerning the dose-effect curve previously accessible.

Not to forget that there is frequently an attainable prior information and the implementation of non-adaptive optimal designs is confirmed in phase IIb trials. Again, this emphasizes the preference of using non-adaptive Bayesian optimal designs.

In the following section, a more flexible approach compared to the optimal designs will be presented. As highlighted below, this approach, based on a utility function, allows to take into account some specific aspects of interest.

2.4 Utility functions

Utility functions are generally introduced and defined within the decision theory framework. Utility describes the preferences of the "decision maker" and classifies/orders decisions. A decision theory result suggests that, in a risky environment, all decision rules can be compared and classified using the expectation of a certain function, called the utility function.

When conducting a clinical trial, there is an alternation of decisions (choice of doses, allocation of patients, etc.) and observation of random data (as in patients responses). The multitude of these possible cases can be translated into utility functions to be defined according to expectations and goals. Optimization can be formalized thanks to such functions, representing the benefits of the stage 2 and the final dose recommendation.

This utility approach is more flexible compared to the optimal design approach because it enables to account for: safety issues, economical/financial aspects, etc.

An example of utility function is the Bayesian predictive Probability of Success (PoS) in phase III, which can be maximized across all doses tested in stage 2, see [Temple2012]; [Patel et al.2012]; [Patel et al.2013]; [Kirchner et al.2016]; [Antonijevic et al.2013]; [Antonijevic et al.2010]; [Gajewski et al.2015] for example. Such a utility function incorporating a PoS component is also proposed in this thesis report, see **Chapters 3** and **4**. Note that the PoS is the power of a test weighted by the uncertainty on the parameters: by applying a Bayesian approach, one assigns distributions for the parameters, and for each value of the parameter one computes the power of the study to be predicted, the PoS is then the average power, averaged on the distribution of the parameters (see **Chapters 3** and **4**).

Other examples of utility functions could be functions including a safety / penalty term penalizing high doses, rewards, global costs and/or costs per patient, see [Kirchner et al.2016]; [Temple2012]; [Antonijevic et al.2013]; [Patel et al.2012]; [Patel et al.2013]; [Foo and Duffull2017]; [Antonijevic et al.2010]; [Gajewski et al.2015]. Several examples of utility functions defined by economic or financial considerations are proposed and assessed in this thesis, see **Chapter 3**. A new type of utility functions, including both a PoS component related to efficacy, and a probability of controlling over toxicity component related to safety, is proposed in this thesis report, see Section 2.5.

In [Christen et al.2004], the authors assume that it is complicated to generalize a single utility function. A realistic and rational choice of utility function is not so obvious, and may be highly dependent on the treatment or the therapeutic axis, functions may incorporate safety, costs and rewards parameters, which can be difficult to predefine or arbitrarily fix without prior knowledge. Instead, they propose considering a set of possible utility functions, rather than focusing on one particular function: "*we denote with V a set of utility functions, and assume the decision maker is unwilling or unable to further specify a single utility function $u \in V$. The problem has been studied from the perspective of sensitivity analysis to mis-specifications in the utility function and prior distribution*" [Christen et al.2004].

Authors propose a dynamic programming rule which consists in doing backward induction, and a well detailed algorithm is described in particular, alternating sequence of expectation and maximization. Utility-based decisions consist of dose selection, within a Bayesian framework, based on posterior probabilities: at every stage of the trial, the next patient is allocated to the selected dose, the trial may be stopped for futility, with no treatment recommendation, and doses may be dropped during the trial, if they are judged to be less effective than others. However, these doses are not totally excluded from the trial and may be reused in randomization, which consists in allocating patients to doses within the "non-dominated" set. And here, "non-dominated" set refers to the set of superior doses, dominating the others. Indeed, dropped doses may dominate other doses later on, when the posterior probabilities change: a dose may be inferior at a given time point, than superior at another time. Hence an adaptive randomization for dose allocation, carried out from sequential design, based on expected utility, to define the set of "non-dominated" doses. In other words, authors "*describe a Phase II clinical trial for finding optimal dose levels*". They use the algorithm above in order to design a sequential clinical trial (which is a particular type of clinical trials among many others, see <https://www.scientific-european-federation-osteopaths.org/different-types-of-clinical-trials/>), in which observations are evaluated as they

are produced and the total number of participants depends on the accumulated results, "*using backward induction, where non dominated treatments (actions) are randomized to select the treatment for the next patient in the trial*" [Christen et al.2004].

A study example was conferred in the oncology field, in order to find the optimal dose pegylated interferon for platinum resistant ovarian cancer, see [Christen et al.2004].

In the following, several utility functions recently proposed in the literature are presented.

In [Temple2012], a utility function taking into account safety is proposed and defined as follows: costs and rewards are included, the author considers a fixed phase III sample size N_3 . The PoS is defined not only in terms of statistical significance, but also "clinical" significance, and by implicitly taking into account the safety. A utility function is used to select doses; the aim is not to optimize the sample size in phase II nor the allocation ratio, but to find the optimal dose for phase III. As for the choice of the utility function, the author precises that to claim success in phase III, two things need to be observed: two phase III trials with a significant difference from placebo and a reasonable safety profile.

Therefore, for each of the two trials of phase III, the author adopts a two-arm parallel group design while taking into consideration the active dose diagnosed in phase II against the placebo dose. Each arm contains $n_3 = N_3/2$ subjects, selected to insure suitable ability of identifying a clinically significant difference from placebo. The final inference in phase III amounts to a pairwise comparison among both the placebo dose $dose_0$ and the active dose $dose_j$; that's why one can identify this inference as frequentist. Hence, the awaited responses within subjects in phase III are assumed to be normally distributed with variance σ^2 .

In [Temple2012], the type I error is controlled in each phase III trial at the two sided 5% alpha level. Thence, if the inferior bound of the two-sided 95% confidence interval for the difference against placebo is greater than 0, then the author assumes that there is a meaningful difference from placebo, in other words, an actual dose-response. To calculate and evaluate the probability of detecting a treatment effect in one phase III trial, denoted by PDR , the author uses the following equation:

$$PDR(v_j, v_0) = 1 - \Phi\left(-1.96 + \frac{v_j - v_0}{\sqrt{2\sigma^2/n_3}}\right)$$
, where Φ denotes the $N(0, 1)$ distribution function, v_j denotes the true mean response at dose $dose_j$ and v_0 denotes the true mean response at placebo.

Due to safety apprehensions, it is probable for a drug to fail the marketing process although it is effective; hence, the author added a concept of safety to her criteria for overall achievement and success. Here's where the author introduces a quadratic function for the probability of the insufficient drug due to safety manners, denoted by PSF (probability of safety failure):

$$PSF(j) = 0.2\left(\frac{dose_j}{dose_J}\right)^2$$
 with $j = 1, \dots, J$, where J is the number of doses. Note that from now on, $dose_1 < \dots < dose_J$ is assumed.

The value 0.2 is the maximum probability of phase III failing because of safety issues, according to the function above. The reason behind the choice of a quadratic function is to insure that, at higher doses, the increment of safety from a dose to another is bigger than at lower doses. For the sake of balancing the gain in efficiency from enlarging a dose level with the growing risk in the safety profile, the author had to integrate a term for safety purposes in the utility function. If one does not include the safety concept into the study, the dose with the bigger difference in mean efficacy response from placebo would readably be chosen (which is the

highest dose in light of monotonic dose-response profiles), because it leads to the highest utility. One always has to seek an adequate safety profile and also seek success in both of the two phase III trials so that one can go to the market. In this case, the PoS in phase III for dose $dose_j$ is calculated as follows: $\mathbb{P}(v_j, v_0) = (\text{the power of each test})^2(1 - 0.2(dose_j/dose_J)^2) = PDR(v_j, v_0)^2(1 - PSF(j))$.

A utility function is proposed here in order to take into account rewards and costs issues, the PoS in phase III and safety issues. Let N_2 be the total sample size of the phase II dose-finding trial. Each phase III trial has two arms, and to the extent of going to the market, one requires successful phase III trials. The total sample size of every phase III trial is $N_3 = 2n_3$, where n_3 is the number of subjects in each arm. Suppose that c_2 and c_3 are costs per subject in phase II and phase III respectively. If a drug makes it to the market, one designates the plausible reward or benefit for the company as 'R'. If the PoS in phase III is too low, then the development process will be interrupted directly after phase II, which will reveal the cost of the failed phase II trial through the utility for this dose. The author represented this utility function, denoted by $u(v, a)$, as follows (with $j = 1, \dots, J$):

$$u(v, a) = \begin{cases} -c_2N_2 & \text{if } a = 0 \\ \mathbb{P}(v_j, v_0)R - c_2N_2 - 2c_3N_3 & \text{if } a = j \end{cases}$$

In other words, the utility function is defined with the following components: R which is the multiplicative factor of the reward and $\mathbb{P}(., .)$ which is a form of POS, it integrates the power of the test in phase III ($PDR(v_j, v_0)$) and an implicit safety-dependent term ($PSF(j)$). There are two phase III studies and the overall power corresponds to $(\text{the power of each test})^2 = PDR(v_j, v_0)^2$.

In [Temple2012], the author proposed then three different decision rules.

Decision rule 1: "*we select the minimum dose estimated to have a 1.3 change from placebo*", which is done by selecting the minimum dose with a clinically meaningful difference (CMD) from the placebo to be taken into phase III.

Decision rule 2: "*we select the dose that maximizes the PoS in phase III based on a point estimate of the dose-response curve*". Taking the fitted dose-response curve as a base for this decision rule, the author chooses the dose that maximizes the PoS in phase III. Not to forget that the author mentioned a supplemental condition to advance to phase III, which is that the PoS in phase III must be higher than the prespecified threshold x . This value took place in her study so that a feasible chance of success to advance to phase III is guaranteed.

Decision rule 3: "*we select the dose that maximizes the posterior PoS in phase III*". This is a Bayesian decision rule, so here one chooses the dose that maximizes the posterior PoS. Repeatedly, the author added a condition to advance to phase III, that is to say, the PoS in phase III must be higher than the pre-specified threshold x . The author applies a prior distribution to the model parameters in order to take decisions depending on the posterior distribution for each v_j , given the phase II data $Y = y$.

Finally, concerning the choice of reward R and cost per patient, this could usually be based on prior knowledge, already conducted trials, the knowledge of experts, etc.

Note that, in [Temple2012], the author does not justify the choice of cost = 1 and $R = 12000$: "*We assume the cost of a subject in phase II is the same as a subject in phase III and assign $c_2 = c_3 = 1$, where this represents 1 unit of some larger monetary value. One assigns a reward for successfully getting the drug to market of $R = 12000$. Hence the return for a compound to market is 12000 times the cost of a phase II or phase III subject*". These choices are arbitrary

and one could have very different values depending on the study.

In the end, the author does not try to directly maximize her utility but instead, considers three different decision criteria (choice of dose) and detects the one engendering the highest utility.

In her discussion, one of her conclusions is that: "*Using decision rule 2 also improves the expected gain and reduces the gain bias compared with decision rule 1*". In the beginning of my thesis work, I assessed this same utility function proposed by [Temple2012], but with different decision rules, in order to select the optimal dose, and to optimize seamless designs (in terms of patient allocation to doses, and global patient repartition between phase II and phase III, assuming a fixed total sample size phase II + phase III = constant), see **Chapter 3** and Appendix A.1 for further details.

In the same way, there is a publication for Diabetes [Antonijevic et al.2013], and another one for Neuropathic pain [Patel et al.2012], where authors show that the choice of dose based on the maximization of utility is superior, for all examined criteria, to the method which consists in choosing the closest dose to a target one, i.e. the one that gives an average target Δ , where Δ is the mean difference versus the placebo or the comparator arm, identified as relevant.

Different types of utility functions can be used: with or without economic consideration (i.e. costs, phase II & phase III, potential gains if success, etc.). Utilities with Reward are utilities with a systematic cost for Phase II, a cost for Phase III if one decides to 'Go' to this phase, and a reward if Phase III is successful. Apparently, in the literature, this type of utility (called "Expected Net Present Value (ENPV)") is rarely used to select doses. However, these criteria are sometimes used to evaluate and classify various dose selection methods [Temple2012]; [Antonijevic et al.2010].

In [Foo and Duffull2017], an initial exploration of a fully cost-driven design was presented. A design criterion was proposed that represented the minimum expected cost of an early phase clinical study, where costs include resource use as well as study failure. In other words, the design is optimized by simply minimizing $\mathbb{E}(\text{cost})$.

As already briefly mentioned in Section 2.1, authors of [Gajewski et al.2015] carried out a Bayesian adaptive design based on a utility function, along with decision rules to stop the trial for success. This function aims to combine both efficacy and safety, posterior distributions generated by MCMC (Markov Chain Monte Carlo [Ravenswaaij et al.2018]; [Geyer et al.2011]) are used within utility-based decision rules to stop the trial at interim, modify the sample size, etc. The decision rule suggested in [Gajewski et al.2015] is based on maximizing the utility function, accounting for efficacy and safety and comparing several treatment arms. The criterion to stop the trial early for success is defined as follows: $\mathbb{P}_{post}[U_{arm_i} = U_{max}] > 0.90$, where U represents the utility. Authors also proposed an algorithm called SSA (Sweet Spot Accrual rate), where finding the 'sweet spot' refers to three main parameters: the average time to finish the study, the average patients number required and the study cost [Gajewski et al.2015]. This rate algorithm is highly related to the Bayesian adaptive design and could be used to carry out a sensitivity analysis, aiming to assess the cost-benefit of a design, and to judge the status of the trial: a trial could not be achieved if the rate is very slow, whereas acquiring valuable information for adaptation could not be guaranteed if the rate is very fast [Gajewski et al.2015].

Despite the fact that this algorithm is limited by size and duration criteria, it is flexible because it could be applied to different utility-based criteria to stop the trial, different types

and number of endpoints, types of studies, and could even be used in an adaptive version (i.e. changing the accrual rate design and patterns).

In [Kirchner et al.2016], authors mention some interesting points regarding the utility function definition: they work in the context of oncology, more specifically in a time-to-event setting. Note that in clinical trials, patients are recruited over a period and followed up to a fixed date or possible event such as death or recurrence of a tumour (negative), conception or discharge from hospital (usually positive) or cessation of breast feeding (neutral). The time between recruitment and the event is 'time to event' or survival (even when death is not the event in question). Subjects in the trial who survive to the fixed date but for whom the event has not occurred are said to have a censored survival time. So in [Kirchner et al.2016], authors present methods for a program-wise phase II/III planning that aim at determining optimal phase II sample sizes and Go/NoGo decisions in a time-to-event setting. Optimization is based on a utility function that takes into account fixed and variable costs of the drug development program and potential gains after successful launch. They sometimes use the Bayesian approach (but in some cases, they do not consider a prior law on parameters, but particular fixed values), they optimize the utility according to the phase II sample size and the threshold (efficiency) to go to phase III; they have no design concern: there are only 2 arms and the designs are balanced, the sample size of phase III is variable, it is calculated according to the result of phase II with the classical approach. The utility is defined by $U = \text{gain} - \text{cost}$: there is a unit cost per patient (not the same in phase II as in phase III), and also a fixed cost (independent of the sample size). The gain here depends on the results of phase III; the better the result of phase III, the greater the gain; this is why 3 levels of efficiency in phase III are introduced: small, medium, wide; benefits correspond to each level respectively.

Some work extensions are also mentioned: 'multiple arms' for dose selection, case where the gain decreases with the sample size of phase III, and interim analysis.

Safety consideration is not addressed in [Kirchner et al.2016], however this point is very fundamental in oncology and elsewhere, in other therapeutic areas.

2.5 Summary and outlook

Adaptive seamless designs for clinical trials have attracted a lot of attention because they offer the possibility to combine different phases of drug development into a single trial. A very popular adaptation consists of selecting / dropping doses based on accumulating data. Such designs typically start with a learning dose-response curve phase (stage 1). At some interim time-point, one looks at the data to estimate and model the dose-response relationship.

Adaptation may also turn into a modification to the dose allocation ratio putting more emphasis on the most promising doses or even increasing sample size to tentatively rescue a poorly responsive trial.

Consequently, it is proposed to use modeling procedures within two-stage adaptive trials in order to obtain model-based dose-effect estimates at interim to guide early futility stopping and/or re-design stage 2 (e.g. choice of doses, sample size, dose allocation ratio) and analysis (e.g. dropping of inadequate dose-response models). The dose selection and testing strategies will obviously depend on sponsor's objectives and its degree of expectations from the dose-finding trial (e.g. detection of dose-response signal only, type I error control among pairwise comparisons of individual doses against placebo, soundness of selection of dose(s) for phase

III, etc.).

Adaptive combination tests ensuring a strong control of type I error could be developed for confirmatory purposes. More liberal testing procedures (such as those including both stage one data and second stage data in the final testing) may also be studied (these procedures were not developed in this thesis).

One can investigate adaptive designs where both the interim and the final decision are made either on a clinically relevant endpoint and/or appropriate surrogate endpoint.

It is also important to check if it is worth increasing the number of candidate dose-response models compared to some standard set such as Linear, Emax, Exponential, Quadratic, and to possibly consider incorporating model-averaging strategies to achieve better dose selection at interim.

At interim, data may be used to re-design stage 2 by selecting remaining doses according to some pre-defined optimality criteria. As previously stated, the most commonly used optimality criteria are C- and D-optimality.

Another optimality criteria may be the Bayesian posterior probability for a dose of being within a pre-defined optimum range of efficacy (against placebo) that could be used to help the sponsor select doses for stage 2. Bayesian posterior probabilities associated with each dose may also guide the dose allocation ratio to be used for stage 2 putting more weight on those doses with high posterior probabilities. Such criteria are deemed to be more specific than traditional C- or D- optimality criteria.

As an alternative/more recent approach, optimization can also be formalized thanks to utility functions which represent the benefits of the stage 2 and the final dose recommendation, and enable to account for safety issues, economical/financial aspects, etc. Gain should be defined and there are various possibilities to do it, and gain expressed via utility functions should be maximized.

Another important aspect of this approach is how to define the decision rules to "Go" to phase III. Choosing the optimal dose and the optimal phase II sample size could be based on several constraints in the decision rule, depending on prior knowledge and expected goals.

Some main perspectives/extensions are to put uncertainty on the penalty of toxicity, in other words put efficacy and safety at the same level and avoid arbitrary choices for toxicity. More specifically, one could fully characterize subject's safety by using a binary safety outcome mimicking the drug limiting toxicity (DLT) concept commonly applied to oncology. Transposing the DLT concept to phase II / phase III trials this could correspond to the permanent treatment discontinuation due to adverse event coded yes/no. The penalty considered for each subject would depend on the expected probability of DLT / treatment discontinuation at the given dose. Prior distributions could also be used for these probabilities. The optimization framework would consist of optimizing the patient global allocation ratio between phase II and phase III looking for the best trade-off between selecting the best/good dose(s) for phase III and maximizing the PoS in phase III.

One could therefore consider the efficacy component as a PoS involving both the effect size of the dose d and the sample size of the phase III study. For the toxicity, what will generally count is to control it under a maximum admissible threshold with a sufficiently high level of confidence. This threshold can be expressed in several ways: independently from the placebo arm or not. It may just consist of controlling the toxicity of the dose d under an acceptable threshold t , without taking into account the toxicity of the placebo (supposed to be negligible),

or, on the contrary, controlling for example the relative risk versus placebo under an admissible threshold RR_A , or the absolute difference versus placebo under a threshold Δ_A (where Δ_A is an admissible difference versus placebo). The toxicity component of the utility function could then be related to the probability of observing an upper bound of the $1 - \alpha'$ level confidence interval of the relative risk $RR < RR_A$ at the end of phase III, under the hypothesis that the true toxicity rates of placebo and dose d are those estimated by the model that one chooses in advance at the end of phase II.

The α' notation is used to emphasize that the required confidence level is not necessarily $1 - \alpha$ here, where α is the significance level used for efficacy in phase III. If the reference to placebo is deleted, the toxicity component could then just be the probability of observing an upper bound for the estimated rate of toxicity lower than t at the end of phase III, under the assumption that the true toxicity rate of the dose d is the one estimated by the chosen model at the end of phase II. An even simpler alternative could also be the probability of observing a toxicity rate lower than t at the end of phase III. If the estimated RR at the end of phase II is already greater than RR_A for the dose d , or if the toxicity rate of the dose d estimated at the end of phase II is already greater than t , then this should be a reason for not selecting this dose for phase III (in other words, this becomes the toxicity filter).

Consequently, the utility function could have a theoretical form, depending on: the PoS function, the toxicity penalty representing the probability of observing a toxicity rate greater than t in phase III, the dose and tuning parameters reflecting the respective contribution of efficacy and safety to the utility function, as illustrated in Figure 2.2.

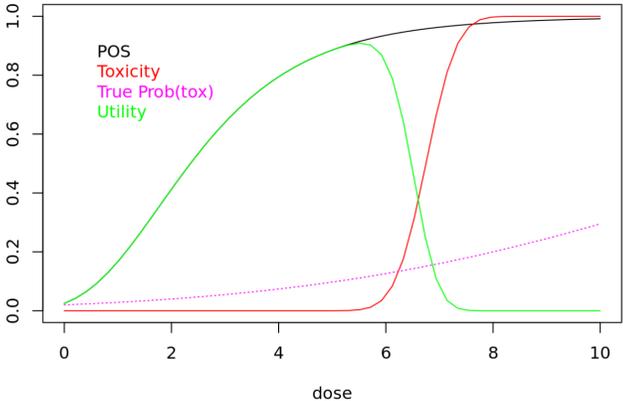


Figure 2.2: Theoretical curves (see text for details).

Figure 2.2 illustrates an example of a feasible utility function aiming to define an optimal dose and an optimal patient distribution between phase II and phase III. For a given dose-response model, four doses and a placebo: the PoS is represented in black, the probability of having more than t of toxicity observed in phase III in red, the theoretical toxicity in purple, and the utility in green. Utility function optimization could be done in a Bayesian context, putting non-informative or informative prior distributions on the model parameters, and computing posterior distributions with an MCMC approach.

In addition, it is also interesting to assess whether or not an interim data inspection strategy

for phase II, with the possibility of re-evaluating the sample size of the phase II study, significantly improves the decision quality of the dose for the phase III (several criteria could be proposed to decide whether to stop or not at interim). One must quantify the loss of utility of the design with an interim analysis compared to the fixed design with the maximum of patients. This should be weighed against the number of patients saved on average.

One could also organize and reflect upon simulations around practical questions such as: in case the interim analysis does not identify "for sure" a better dose, but clearly identifies useless doses (because they are not effective or toxic), is it then possible to abandon one or two doses for the final analysis, without unduly diminishing the final utility?

All these perspectives correspond to the following work developed in this thesis.

Chapter 3

Utility functions: how to build them and what contributions for design optimization can we expect?

The aim of this chapter is to give guidance on an operationally seamless design, in terms of: timing of the interim analysis, design of the first part (known as stage 1) of phase II. The main underlying hypothesis is that the sponsor takes its decisions (decision to continue the trial after interim analyses, choice of dose) in maximizing a utility function that assigns a value to each decision. Several utility functions will be proposed and their properties will be discussed.

3.1 General notations and main notions

In this section, we delineate the common materials and methods applied to most of the work contained within this thesis. We specifically describe the mathematical formalization of a phase II/phase III development program, aiming to define all necessary notations and calculations related to the dose-response modeling of efficacy, and to the PoS.

- Dose values are denoted by d , and dose indices are denoted by d_j , $j = 1, \dots, J$ where J is the maximal dose index.
- $Y_{d,i}$ represents the random efficacy response of patient i in dose d arm, with $i = 1, \dots, n_{d_j}$, where n_{d_j} is the number of patients for the dose d_j in phase II study. It is assumed that $Y_{d,i} \stackrel{iid}{\sim} N(m(d; \theta), \sigma^2)$ where $m(d; \theta)$ is the expected mean effect of dose d , and σ is the residual variability (standard deviation of residual error). The empirical mean responses in dose d and placebo are denoted by \bar{Y}_d and \bar{Y}_0 respectively, and we note $\bar{\Delta}(d)$ the difference of the two.
- N_2 and N_3 denote the phase II and III sample size respectively.
- w is a vector in \mathbb{R}^J , representative of the phase II design: w_{d_j} is the proportion of patients allocated to dose arm d_j ($\sum_{j=1}^J w_{d_j} = 1$).

- It is assumed that the expected mean dose-response for efficacy $m(d; \theta)$ follows an Emax model:

$$m(d; \theta) = \theta_1 + \frac{\theta_2 \times d}{\theta_3 + d}, \theta = (\theta_1, \theta_2, \theta_3)^t$$
 - $\theta_1 = E_0$ is the placebo effect
 - $\theta_2 = E_{max}$ is the maximum effect compared with placebo
 - $\theta_3 = ED_{50}$ is the dose with half of the maximum effect
- The expected mean difference versus placebo is called $\Delta(d) = m(d; \theta) - m(0; \theta)$.

The computation of the Probability of Success (PoS)

The PoS that we consider for efficacy in our utility functions is defined as follows.

Suppose that $\bar{\Delta}(d) = \bar{Y}_d - \bar{Y}_0$ is the difference in observed mean effects between dose d and the placebo in phase III.

The expectation of $\bar{\Delta}(d)$, $\Delta(d)$, is equal to $m(d; \theta) - m(0; \theta)$; this difference $m(d; \theta) - m(0; \theta)$ does not depend on E_0 parameter.

Our null hypothesis H_0 assumes that $m(d; \theta) - m(0; \theta) = 0$. Our statistic of interest is defined as $Z = \frac{\bar{\Delta}(d)}{\sqrt{2SE^2}}$, where $SE^2 = \sigma^2/(N_3/2) = 2\sigma^2/N_3$, and N_3 is divided by 2 because we only have two doses in phase III, the placebo and the chosen dose. Under H_0 , Z follows a standard Normal distribution. The "Success" is defined by a significant test versus placebo in the phase III. In other words, a successful phase III trial means that $\bar{\Delta}(d) \geq z_{1-\alpha} \times \sqrt{2SE^2}$ (assuming without loss of generality that positive values favor the test drug), where $z_{1-\alpha}$ is the $1 - \alpha$ quantile of the standard Normal distribution. A unilateral α level of 0.025 is considered in our calculations: if $Z > 1.96$, H_0 is rejected in favor of $H_1 : m(d; \theta) > m(0; \theta)$.

Assuming a particular alternative hypothesis $H_1 : m(d; \theta) - m(0; \theta) > 0$, the true PoS can then be written as:

$$\begin{aligned} PoS(d, \theta) &= \mathbb{P}_{H_1}(Z \geq 1.96) = \mathbb{P}_{H_1}(\bar{\Delta}(d) \geq 1.96 \times \sqrt{2SE^2}) \\ &= \Phi\left(\frac{m(d; \theta) - m(0; \theta) - 1.96 \times \sqrt{2SE^2}}{\sqrt{2SE^2}}\right), \end{aligned}$$

where Φ denotes the standardized Normal distribution function.

3.2 Utility functions: constructions and properties

In order to define and construct our utility functions, we consider the following assumptions:

- (i) In the context of a seamless design, $N_2 + N_3$ is a fixed constant, N_{tot} .

- (ii) The relative sample size of the phase II study with respect to the total sample size (phase II + phase III) is described with a parameter f , $0 \leq f \leq 1$.
- (iii) The $N_2 (= f \times N_{tot})$ patients are distributed in 4 doses and 1 placebo.
- (iv) The $N_3 (= (1 - f) \times N_{tot})$ patients are distributed in two arms: the selected dose and the placebo, each one with $N_3/2$ patients.

In this chapter we define utility functions that assign numerical values to the sponsor decisions at the end of the phase II part of the study. A typical example is the following utility function [Temple2012], that assign values to a combination of two decisions (these decisions are described in more detail in Section 3.4):

- First, there is the Go/NoGo decision for entering phase III. Let us denote the cost per patient by γ , and the financial reward if the program is successful by R , then:
 - if we choose not to go to phase III, then we have the cost of the phase II trial, equal to $-\gamma N_2$;
 - if we decide to go to phase III, then the value depends on a random event, success or not of the phase III trial: if the phase III trial is successful, then we have the final gain, G , equal to reward-total cost, $G = R - \gamma N_{tot}$; if it is not successful, then there is only a total cost (no reward), and the final gain is negative, $G = -\gamma N_{tot}$. This can also be expressed as a function of an indicator function, $\mathbf{1}('Success')$, as follows: $G = R \times \mathbf{1}('Success') - \gamma N_{tot}$, where 'Success' refers to the Success of phase III.
- Second, if the decision is to go to phase III, then the sponsor must choose the adequate dose within the doses tested in the phase II study.

The utility is in fact random after the phase II stage, as the final gain depends on the success, or not, of the phase III part. Therefore, from the sponsor's point of view, the expectation of final gain, $\mathbb{E}(G|phase\ II, Go) = R \times \mathbb{E}\mathbf{1}('Success') - \gamma N_{tot}$, if the decision is to go to phase III, is the key quantity to assess. This expectation can be represented by a utility function U , that the sponsor will assess at the end of the phase II, $U(d, f) = \mathbb{E}(G|phase\ II, Go)$; it depends on the expectation of the final gain after the phase III, and, as a consequence, depends on the PoS of the phase III part, because $\mathbb{E}\mathbf{1}('Success') = PoS(d, f)$.

In the example of the utility function proposed in [Temple2012], the success of the phase III trial is defined by both simultaneously a statistically significant comparison with the control in the phase III trial and the absence of safety issues in the same phase III trial. As it is assumed that efficacy is independent from the safety, the probability of success is then equal to the product of the probability of a statistically significant difference with the control which corresponds to the power function of the phase III trial, and the probability of absence of safety issues: this latter probability is modeled by a function of the dose d , equal to $1 - sa(d/d_J)^2$, where d_J is the maximal dose, and sa a fixed value related to the safety. We will note this global Probability of Success by $PoS_{adj}(d, f) = PoS(d, f) \times (1 - sa(d/d_J)^2)$, where $PoS(d, f)$

is the standard Probability of Success related to efficacy, i.e. the power function of the phase III trial. The obtained utility function, named $U0$ in the following, has appealing properties, in particular the easy interpretation of the parameters. But the problem with such type of utility functions is that some of the parameters (in particular, the reward R , and the safety parameter sa) are not known with enough confidence or precision at the beginning of the drug clinical development.

More generally, a proper utility function should have the following properties:

- It must depend on success of the phase III study (higher utility in case of success)
- It must be a non-monotonic (preferably concave) function of the dose with a unique maximum value (increasing then decreasing): such a shape reflects the bi-dimensional aspect of the utility function, one increasing with the dose (efficacy component) the other one decreasing with the dose (safety component)

Examples of utility functions capable of verifying these conditions according to the appropriate parameters, are shown below and are discussed thereafter. All the utility functions below include an increasing function of the dose (a function of the PoS, representative of the drug efficacy) multiplied by a decreasing penalty function of the dose (representative of the drug toxicity). Note that for $U0$, $U1$, $U2$, $U3$ and $U7$ defined below, those conditions are verified in case of a 'Go' decision for phase III only. In some of the following utility functions, in order to normalize the effect of the dose (so that it does not depend on the dose unit), the effect is expressed as a function of the relative efficacy, denoted by $\delta = \text{efficacy} / \text{maximum efficacy}$. So with the Emax dose-response model, we have:

$$\delta = \frac{E_{max} \times d / (ED_{50} + d)}{E_{max}} = d / (ED_{50} + d).$$

Therefore δ varies between 0, for a null dose (placebo), and 1, for a very large (or "infinite") dose.

For the sake of simplicity, and in order to facilitate the reading, in the following of the manuscript (next sections of **Chapter 3**, **Chapter 4** and **Chapter 5**), we will drop the parameters in the notations of the quantities of interest when there is no ambiguity. For instance, we will note $PoS(d, f)$ instead of $PoS(d, f, \theta)$, $U(d, f)$ instead of $U(d, f, \theta)$, etc. For all the utility functions below, we assume that $E_{max} > 0$ and $ED_{50} > 0$. Parameter $c > 0$ (introduced in some of the utility functions below) will represent a parameter penalizing high doses, the higher the c , the higher the penalty for safety:

$$\begin{aligned} U0(d, f) &= -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + R \times PoS_{adj}(d, f)) \\ U1(d, f) &= -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + PoS(d, f) \times (R - c(\delta - 0.95)^2)) \\ U2(d, f) &= -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + PoS(d, f) \times R(1 - \delta)) \\ U3(d, f) &= -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + PoS(d, f) \times R(1 - \delta)^2) \\ U4(d, f) &= PoS(d, f) \times (1 - c(\delta - 0.95)^2) \\ U5(d, f) &= PoS(d, f) \times (1 - c \times \delta) \\ U6(d, f) &= PoS(d, f) \times (1 - c \times \delta)^2 \\ U7(d, f) &= -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + PoS(d, f) \times (R - c(\delta - 0.95)^2 \mathbf{1}(\delta > 0.95))) \\ U8(d, f) &= PoS(d, f) \times (1 - c(\delta - 0.95)^2 \mathbf{1}(\delta > 0.95)) \end{aligned}$$

$U9(d, f) = PoS(d, f) \times (1 - c \times (\frac{d}{d_J})^2)$ (where d_J is the highest dose)

$U10(d, f) = PoS(d, f)^h \times (1 - PoT(d))^k$ (where $PoT(d)$ is the probability of toxicity for dose d)

$U11(d) = PoS(d)^h \times \mathbb{P}(tox_{obs}(d) \leq t)^k$ (where tox_{obs} is the observed proportion of patients having a toxicity in phase III, and t is a safety parameter controlling over-toxicity, see **Chapter 4** for further details).

In $U10$ and $U11$, parameters h and k reflect the respective contributions of efficacy and safety to the utility function; for instance, the higher the k , the higher the penalty for safety.

The utility functions $U1$ to $U8$ suggest different efficacy penalties which do not explicitly refer to a safety component. For these utility functions, in order to normalize the dose effect (so that it does not depend on the dose unit), the effect is expressed as a function of δ : the main problem with such a definition is that, from the sponsor's point of view, the utility function depends also on the estimation of the efficacy dose-response model (the penalty depends on δ) which may lead to an increase in uncertainty and bad choices after phase II; on the contrary, $U9$, $U10$ and $U11$ utility functions do not depend on the efficacy dose-response model, this could be a more rational choice intending to avoid estimation problems. But the major difference between these three utility functions is that safety is implicitly considered in $U9$, through the probability of absence of safety issues, $1 - c \times (d/d_J)^2$, whereas the utility functions $U10$ and $U11$ explicitly identify both an efficacy and a safety component. These particular functions will be discussed in Section 3.3.2.

3.3 Seamless design and utility function: some first quite disappointing results

The aim of this section is to give guidance on an operationally seamless design, in terms of: timing of the interim analysis, design of the phase II part (stage 1), and choice of the optimal dose. The main underlying hypothesis is that the sponsor takes its decisions (decision to continue the trial after interim analyses, choice of dose) in maximizing a utility function that assigns a value to each decision. To do that, some of the utility functions defined in the previous section will be discussed according to their related simulation results.

3.3.1 Introduction and notations

When the sponsor takes the decisions regarding the dose or the Go/NoGo choices, it uses the parameter estimates. When computing the global/empirical utility expectation and the PoS of the chosen dose (for all simulated phase II studies), in order to assess the performance of the decision rules, one uses the true parameter values.

Before study starts: sponsor's general strategy is to maximize (in phase II design, w , and N_2/N_{tot} ratio, f) the expected utility. A frequentist approach was used to compute the parameter estimates of dose-response model: sponsor's decisions are driven by maximum likelihood estimations of the model parameters (see Appendix A.1 for further details). Decision rule of the sponsor is only based on point estimate $\hat{\theta}$ of model parameter vector θ . We compared

the efficiency of the decision rules through clinical trial simulations, corresponding results are shown in Appendix A.1. In order to save computational time, simulations were not conducted in simulating individual patients but in simulating directly the maximum likelihood parameter estimates by sampling them with a Normal distribution $N(\theta; \mathcal{I}_\theta^{-1})$, where \mathcal{I} is the Fisher information matrix.

3.3.2 Optimal patient allocation: unattractive results

All utility functions presented in Section 3.2 were assessed through several simulation scenarios, but for sake of simplicity, only some particular functions of interest are presented in Appendix A.1. Regarding the optimisation of patient allocation to doses as well as the global patient allocation between phase II and phase III, we compared in Appendix A.1, through simulations, results obtained in using respectively the utility functions $U5$ (because it is representative of utility functions defined with the relative effect δ) and $U9$ (because it does not depend on the efficacy dose-response model, and implicitly refers to a safety component) only. We do not detail those simulations further in this chapter (we leave the description of the methodology and the discussion for Appendix A.1) as the results obtained were not fully satisfactory: the optimal designs identified were, in most cases, very close to the standard balanced design.

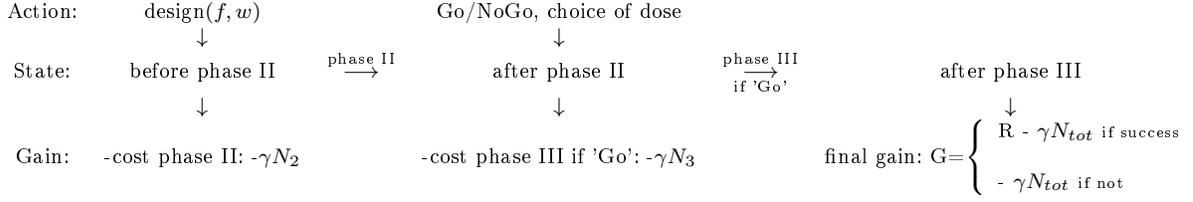
Therefore, for the major part of the thesis we have chosen to abandon the seamless design framework and to focus on the even more frequent situation of dose selection in the context of the phase II dose-finding study with a fixed sample size and a balanced design. For this purpose we considered that utility functions of the form $U10$ or $U11$ were the most appropriate, because they explicitly identify both an efficacy and a safety component. The construction of $U10$ was motivated by [Temple2012]: the idea was to improve the utility function proposed by the author (which has a similar form to $U9$), by explicitly modeling the safety component (instead of considering the absence of safety issues probability through a fixed parameter, c). But, what is striking about $U10$ is that efficacy and toxicity are not treated at the same level. Indeed, the efficacy component, $PoS(d, f)$, involves both the effect size of the dose d and the sample size of the phase III study whereas in comparison, the toxicity component, $(1 - PoT(d))^k$, only involves the toxicity level of the dose d without involving the sample size of the phase III study at all. This asymmetry is troublesome when it comes to optimizing the allocation of patients between phase II and phase III. For this reason for the next chapters and the major part of the thesis, we considered a utility function of the form $U11$, and we focused on the optimization of the dose selection process by maximizing the success probability in phase III.

However, we describe in Section 3.4 a general decisional framework related to the optimization of adaptive designs (in terms of patient allocation to doses and allocation ratio between phase II and phase III) suitable for various utility functions, even though this optimization part is abandoned in the following chapters.

3.4 Sponsor's strategy: Optimal dose and decision rules

Our utility-based decision framework can be described in the context of a Markov Decision Process [Bellman1957]. In particular, the most comprehensive decision framework in our con-

text, the one corresponding to the utility $U0(d, f) = \mathbb{E}(G|phase II, Go)$, can be described by the following graph:



This graph can be understood as follows:

- At start of the study, the sponsor can act on the design of the trial: the timing of the interim analysis (the ratio, f , of the phase II sample size by the total sample size) and the allocation of the phase II patients to the dose arms (vector w). This action has a cost which is proportional to sample size in phase II: γN_2 .
- When the phase II part is completed, the sponsor analyzes the data and takes two decisions: decides to go the phase III or not and, in case of positive answer, chooses the dose for the phase III.
- When the phase III is completed: if it is successful then there is a reward, R , and the final gain is $G = R - \gamma N_2 - \gamma N_3 = R - \gamma N_{tot}$; if it is not successful then there is total cost, and the final gain is negative, $G = -\gamma N_{tot}$.

In an uncertain environment like this one, the sponsor's strategy (the set of actions) is to optimize the final gain expectation (if the decision is to go to phase III) before phase II, $\mathbb{E}(G|Go) = \mathbb{E}^{phase II} \mathbb{E}(G|phase II, Go)$. We remind that $U0(d, f) = \mathbb{E}(G|phase II, Go)$. According to the Bellman Dynamic Programming principle [Bellman1957], which consists in optimizing a decision by breaking it down into a sequence of decision steps over time, this optimisation should be performed backwards:

1. given the phase II trial has been performed, the optimal decisions (Go/NoGo, choice of the optimal dose that we denote by d^*) maximize $U0(d^*, f) = \max_d U0(d, f)$; at this stage only the final gain is random: it depends on the success or not of the phase III trial; therefore: $U0(d^*, f) = \max_d [R \times \mathbb{E}_d(\mathbb{1}('Success')) - \gamma N_{tot}]$, where $\mathbb{E}_d(\mathbb{1}('Success')) = PoS_{adj}(d, f)$ as defined in Section 3.2; in other words, $U0(d^*, f)$ depends on PoS_{adj} , which plays a key role in the calculations, as already discussed in Section 3.2
2. the sponsor chooses 'Go' if the estimated PoS_{adj} associated to the best dose d^* , that we note $\widehat{PoS}_{adj}(d^*, f)$, is large enough, say greater than 0.30 for instance, and if $\widehat{U0}(d^*, f)$ is greater than 0. This is equivalent to $\widehat{PoS}_{adj}(d^*, f) \geq 0.30$ and $R \times \widehat{PoS}_{adj}(d^*, f) > \gamma N_{tot}$, otherwise the sponsor should choose 'NoGo'
3. the optimal design maximizes the expected gain $\mathbb{E}U0(d^*, f)$

For the utility functions $U1, U2, U3$ and $U7$, the same decision process is applied, as these functions also depend on costs and rewards.

For the utility functions $U4, U5, U6, U8$ and $U9$, the decision process has been slightly simplified: because for those utilities there is no reference to economic costs, we have proposed to

base the decision to go to phase III or not on a minimal value of the PoS only, that we have also set to 0.30: the sponsor decides to go in to phase III if the estimated PoS associated to the best dose is ≥ 0.30 .

For $U10$ and $U11$, additional constraints are considered for the safety components: the sponsor decides to go to phase III if the estimated PoS associated to the best dose is ≥ 0.30 and $\widehat{PoT}(d^*) < 0.15$ (for $U10$), or if the estimated PoS associated to the best dose is > 0.60 and $\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq t) > 0.50$ (for $U11$); see **Chapters 4, 5 and 6** for further details regarding $U11$ and the choice of these minimal/maximal thresholds.

In the following, we detail the methodology related to the PoS computation and to the design optimization, for all utility functions $U0, \dots, U11$.

The computation of Probability of Success (PoS)

The Emax model will be used also by the sponsor as "working" model to estimate the mean dose-response relationship. The efficacy PoS computed by the sponsor, for dose selection, uses the point estimate $\hat{\theta}$ of θ .

In the method shown in Section 3.1, the sponsor uses the raw value of the estimate of the model parameters to estimate the PoS as if it was the true parameter value. In a more conservative approach, the sponsor might want to consider the uncertainty in the parameter value: in that case, a semi-Bayesian/hierarchical approach can be used, hierarchical in the sense that it characterizes the law of $\bar{\Delta}(d)$ conditional on phase II and on $m(d; \theta)$, then the law of $m(d; \theta)$ conditional on the estimate $\hat{\theta}$ (i.e. conditional on phase II). In order to introduce uncertainty, the sponsor can consider that θ is random: a classical way to do this (see [Ghosh et al.2007] chapter 4) is to consider that, conditionally on the phase II data, θ follows a Normal distribution, centered at the estimated parameter vector obtained with the maximum likelihood approach, $\hat{\theta}$, and with covariance matrix equal to the inverse of the Fisher matrix of $\hat{\theta}$. In other words it is considered that θ is a random variable following $N(\hat{\theta}, \mathcal{I}_{\hat{\theta}}^{-1})$. Therefore, in the sequel of this paragraph, $\hat{\theta}$ denotes the non-random vector of the maximum likelihood estimates.

With the delta-method, sponsor deduces the distribution of $m(d; \theta)$ given phase II, and then deduces the distribution of $\bar{\Delta}(d)$ given phase II.

Hierarchical model The hierarchical model approach is as follows:

- $\bar{\Delta}(d) \mid \theta \sim N(m(d, \theta) - m(0, \theta), 2SE^2)$, this can be written as: $\bar{\Delta}(d) = m(d, \theta) - m(0, \theta) + e$, with $e \sim N(0, 2SE^2)$.
- It is then assumed that $\theta \mid \hat{\theta} \sim N(\hat{\theta}, \mathcal{I}_{\hat{\theta}}^{-1})$. By application of the delta-method: $m(d, \theta) - m(0, \theta) \mid \hat{\theta} \sim N(m(d, \hat{\theta}) - m(0, \hat{\theta}), \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)^t \mathcal{I}_{\hat{\theta}}^{-1} \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right))$, which can be written as $m(d, \theta) - m(0, \theta) = m(d, \hat{\theta}) - m(0, \hat{\theta}) + \varepsilon$, where $\varepsilon \sim N(0, \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)^t \mathcal{I}_{\hat{\theta}}^{-1} \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right))$, and $\nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)$ denotes the gradient of $m(d, \hat{\theta}) - m(0, \hat{\theta})$.

- Therefore, $\bar{\Delta}(d) = m(d, \hat{\theta}) - m(0, \hat{\theta}) + e + \varepsilon \Leftrightarrow \bar{\Delta}(d) \sim N(m(d, \hat{\theta}) - m(0, \hat{\theta}), 2SE^2 + \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)^t \mathcal{I}_{\hat{\theta}}^{-1} \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right))$.

As a consequence, the estimated PoS, $\widehat{PoS}(d; f)$, is equal to:

$$\Phi \left(\frac{m(d, \hat{\theta}) - m(0, \hat{\theta}) - 1.96 \times \sqrt{2SE^2}}{\sqrt{2SE^2 + \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)^t \mathcal{I}_{\hat{\theta}}^{-1} \nabla \left(m(d, \hat{\theta}) - m(0, \hat{\theta}) \right)}} \right).$$

Accounting for uncertainty induces decrease of the estimated PoS, consequently, sponsor is encouraged to increase the dose to compensate.

We essentially considered the case in which the sponsor only uses the estimate of the parameter to compute the PoS (i.e. without taking into account the uncertainty in the parameter value), but in the next chapter we consider a full Bayesian approach, see Section 4.1.2.1 of **Chapter 4**.

Optimizing the design

We remind that the choice of the optimal dose d^* depends on $\hat{\theta}$, therefore, we will note $d_{\hat{\theta}}^*$ instead of d^* in the sequel, in order to explicitly refer to this dependency for the computation of the overall expectation of the maximized utility function in d . We remind that the maximized utility in d can be written as $U(d^*, f) = \max_d U(d, f)$. Therefore, for a given design (summarized by the quantities f, w), the overall expectation $\mathbb{E}_{w,f}^{phaseII}(U(d^*, f))$ of the maximized utility function in d can also be written as $\mathbb{E}_{w,f}^{\hat{\theta}} U(d_{\hat{\theta}}^*, f)$ where $\hat{\theta}$ has a Normal distribution $N(\theta, \mathcal{I}_{\theta}^{-1})$, see Section 3.3.1.

The expectation $\mathbb{E}_{w,f}^{\hat{\theta}} U(d_{\hat{\theta}}^*, f)$ can be computed by numerical integration, or by simulations (if there is a large number of parameters or if the computation time of the expectations with numerical integration routines is very long).

Regarding $U0$, the expectation was computed via numerical integration:

$$\mathbb{E}_{w,f}^{\hat{\theta}} U0(d_{\hat{\theta}}^*, f) = \mathbb{E}_{w,f}^{phaseII}(U0(d^*, f)) = \int (((-\gamma \times N_2) \times \mathbf{1}(NoGo(\hat{\theta}))) + (-\gamma \times N_{tot} + R \times PoS_{adj}(d_{\hat{\theta}}^*, f) \times \mathbf{1}(Go(\hat{\theta})))) p(\hat{\theta}) d\hat{\theta}, \text{ where } \mathbf{1}(Go(\hat{\theta})) \text{ and } \mathbf{1}(NoGo(\hat{\theta})) \text{ are the random decisions to go or not to phase III, depending on } \hat{\theta}, \text{ and } p(\hat{\theta}) \text{ is the density of the Gaussian distribution } N(\theta, \mathcal{I}_{\theta}^{-1}). \text{ With this approach, computation time was very long.}$$

For $U1, \dots, U10$, as we anticipated very long computation time with numerical integration routines (like the case of $U0$), we used simulations instead, and we computed the expectation via a Monte Carlo approach. The expectation $\mathbb{E}_{w,f}^{\hat{\theta}} U(d_{\hat{\theta}}^*, f)$ can therefore be estimated by:

$$\frac{1}{N_{sim}} \sum_{r=1}^{N_{sim}} U(d_{\hat{\theta}_r}^*, f)$$

where the $\hat{\theta}_r$ are sampled from $N(\theta, \mathcal{I}_\theta^{-1})$, and N_{sim} designates the number of simulations.

For $U11$, the estimation approach is different (the reason is explained in **Chapter 4**) and is applied in a different context (see Sections 3.3.2 and 3.5); it is based on MCMC iterations, and the posterior utility distribution is computed using a particular method called the batching method, see **Chapter 4** for further details.

For more details regarding the computation of the Fisher information matrix and other computational/programming aspects, see Appendix A.1.

We remind that the sponsor's strategy is to optimize the final gain expectation (if the decision is to go to phase III) before phase II, $\mathbb{E}(G|Go) = \mathbb{E}^{phaseII}(\mathbb{E}(G|phase\ II, Go))$, and this optimization is performed backwards (see previous discussion regarding the Bellman Dynamic Programming). In the end, the strategy can be summarized as follows:

- After phase II, for a given utility $U(d, f)$:
 The sponsor computes $U(d, f)$ for each dose d , then computes the optimal dose $d^* = d_\theta^* = \arg \max_d U(d, f)$, then decides to go or not to phase III (see pages 53 and 54 for details)
 - for $U0, U1, U2, U3$ and $U7$, decides if worth going to phase III if $\widehat{PoS}(d^*, f) \geq 0.30$ and if $\widehat{U}(d^*, f) > 0$
 - for $U4, U5, U6, U8$ and $U9$, decides if worth going to phase III if $\widehat{PoS}(d^*, f) \geq 0.30$ only
 - for $U10$ and $U11$, uses additional rules and different efficacy/safety thresholds previously discussed (page 54)
- Before phase II:
 The sponsor's strategy before the phase II consists in optimizing the timing of the interim analysis (the ratio, f , of the phase II sample size divided by the total sample size) as well as the allocation of the phase II patients to the dose arms (vector w). Mathematically this can be written as: $(w^*, f^*) = \arg \max_{w, f} \mathbb{E}_{w, f}^{phaseII}(U(d^*, f))$. In practice, $\mathbb{E}_{w, f}^{phaseII}(U(d^*, f))$ is computed via numerical integration or estimated through simulations, as already explained. The optimisation is conducted using Nelder-Mead algorithm (after logistic transforms to ensure that $0 < f < 1$ and $\sum_d w_d = 1$) with the R 'optim' function. This optimisation could be conducted either separately (optimize in f value while w is fixed at a standard value like the one for the balanced design, or in w value while f is fixed at a desired value corresponding to sample size of interest, see Appendix A.1 for simulation results examples, related to $U5$ and $U9$) or simultaneously (optimize f and w at the same time, see Appendix A.6 for simulation result example related to $U2$).

3.5 Conclusions

In this chapter, we have proposed a general decision-making framework, suitable for comparing and optimizing seamless phase II/ phase III designs, based on utility functions. We have re-

viewed and discussed various forms of utility functions, that either were previously mentioned in the literature or appeared reasonable for us.

Because we think that utility functions defined by economic or financial considerations (such as the cost of phase III, expected financial reward in case of successful launch of the drug) are difficult to specify with enough confidence or precision at the beginning of the drug clinical development, we preferred to focus on utility functions only defined by efficacy and safety (explicitly or implicitly) considerations. We performed then a simulation study with those utility functions that appeared the most appropriate to us, in particular the $U5$ and $U9$ utility functions. Unfortunately, the results obtained (shown in Appendix A.1) were not fully satisfactory, as the optimal designs identified were surprisingly, in most cases, very close to the standard balanced design.

For this reason, we preferred, for the major part of this thesis that will be described in the next chapters, to abandon the seamless design framework and to focus on the even more frequent situation of dose selection in the context of the phase II dose-finding study with a fixed sample size and a balanced design (in this case, PoS and utility U do not depend on f anymore, i.e. f is now removed from all previous formulas, we will denote $PoS(d)$ instead of $PoS(d, f)$, $U(d)$ instead of $U(d, f)$, etc.). For this purpose, we used a utility function of the following family only: $U11(d) = PoS(d)^h \times \mathbb{P}(tox_{obs}(d) \leq t)^k$, where PoS is the power of the comparison versus placebo in a reference phase III trial with $N_3/2$ patients per arm (N_3 is now fixed) and $\mathbb{P}(tox_{obs}(d) \leq t)$ is the probability of observing a toxicity rate lower than t in the dose arm of this same phase III reference trial. Further details and explanations will be provided in the next chapter.

Chapter 4

Dose selection in the context of a phase II dose-finding study

In this chapter, we propose a phase II design within a decision-making framework, based on our final utility function. Unlike the previous chapter (seamless design) where we aimed to optimize the allocation ratio and patients distribution between phase II and phase III, while searching for the best dose for phase III, the idea here is to drive sponsor's decision with respect to the continuation, or not, of the drug development as well as the selection of the best dose for the phase III.

In fact, for the specific analysis of phase II, we thought it would be interesting to consider a utility function of the following form $U(d, \theta, \lambda) = (\textit{efficacy}(d))^h \times (\textit{safety}(d))^k$, where d represents the dose, and, θ and λ are the vectors of efficacy and safety dose-response model parameters respectively. For the sake of simplicity, and in order to facilitate the reading, we will drop the parameters in the notations of the quantities of interest when there is no ambiguity (as done in **Chapter 3**); we will write $U(d)$ instead of $U(d, \theta, \lambda)$. Efficacy is an increasing component of the dose, this term depends on the effectiveness of doses, particularly on effect sizes. We chose to express it depending on the PoS, i.e. the power of a phase III trial, of N_3 patients: this has the advantage of "normalizing" this component by varying it between 0 and 1. Defining efficacy in terms of PoS can be used for any type of efficacy criterion: quantitative, binary, time to event criterion, etc. Contrarily, safety is a decreasing term depending on the dose. This term depends on the toxicity of the doses. We chose to express it according to the probability of observing a toxicity rate lower than t in the dose arm, during a phase III trial of N_3 patients in total: this also has the advantage of "normalizing" this component by varying it between 0 and 1.

If the choice of dose was left to an expert panel, the question it would have to face would be to note and prioritize a treatment according to its PoS in phase III and the probability of controlling the over-toxicity in phase III in accordance with a pre-specified threshold that depends mostly on the therapeutic area. The question would be "with which treatment would you prefer to go to phase III?". Clearly, a strong PoS would be mitigated by a low probability of controlling the over-toxicity. So "best" doses do not necessarily correspond to highest PoS since toxicity is also taken into account. This can be translated by our final utility function

proposal, $U11(d) = PoS(d)^h \times \mathbb{P}(tox_{obs}(d) \leq t)^k$, where h and k are parameters reflecting the respective contributions of efficacy and safety to the utility function; the higher the k , the higher the penalty for safety, and t is a safety parameter controlling over toxicity in phase III. PoS reflects the dose efficacy, and $\mathbb{P}(tox_{obs}(d) \leq t)$ is the probability of controlling over-toxicity, i.e. the probability of observing a toxicity rate less than or equal to t in phase III. The detailed calculation of each of these two components will be presented in Section 4.1.2.1.

A specific characteristic of the proposed utility function is that both its efficacy and safety components depend on the sample size of the phase III study. This choice is intended to reflect real life conditions where Go/NoGo decisions and dose selection at the end of phase II always relate to the sample size the sponsor can afford for a superiority phase III trial. This can be viewed as a pragmatic choice.

4.1 Materials and Methods

4.1.1 Dose-response modeling

We chose to model efficacy via an Emax model (defined in **Chapter 3**), and safety via a Probit model. Note that the probability of toxicity is usually given by a logistic regression, but in practice, it is known that Probit and logistic models are almost indistinguishable (both link functions are symmetric). So we expect similar results with a logistic or probit type safety model. However, logistic model may be better suited if one seeks to interpret parameters and outputs in terms of Odds Ratio. Contrariwise, the Probit model has some advantages if one has to model several events simultaneously: using a multivariate normal distribution and subsequently a Probit link to each of its components, makes it possible to model occurrences of several correlated events. It is less straightforward to model occurrence of correlated events using a logistic link function.

Here is the mathematical formalization of our modeling approach with all the necessary notations and calculations:

- (i) The N_3 (the sample size of the phase III study) is now assumed to be constant.
- (ii) For safety, we used the following Probit model: $\pi(d) = \mathbb{P}(W = 1|d) = \Phi(\lambda_1 + \lambda_2 \times d)$, $\lambda = (\lambda_1, \lambda_2)^t$, where $\lambda_1 = a$ is the intercept parameter, $\lambda_2 = b$ is the dose effect, W is the binary toxicity outcome for one patient, 1 for toxicity and 0 if no toxicity, and Φ is the Cumulative Distribution Function (CDF) of the standard normal distribution.
- (iii) For efficacy, the Emax model defined in Section 3.1 is used.

4.1.2 Decision-making framework

In the following, we will discuss our proposed utility function, as well as computations of each of its components (efficacy and toxicity-related components).

4.1.2.1 Utility function

As already mentioned in **Chapter 3** and at the beginning of this chapter, we chose the following utility function $U(d) = U11(d) = PoS(d)^h \times \mathbb{P}(tox_{obs}(d) \leq t)^k$. Each of its components calculation will be presented below. In all the following of this chapter, the random vectors of efficacy and safety model parameters, specified in a Bayesian framework by the sponsor, will be denoted by θ and λ respectively.

The PoS (which is also the phase III power) that we consider in our utility function is defined in **Chapter 3**. The formulas are given for the case of balanced treatment groups.

We use a MCMC approach, [Ravenswaaij et al.2018]; [Geyer et al.2011], particularly a Metropolis-Hastings algorithm to capture the posterior of the model parameters and key quantities of interest: utility, PoS, etc. For instance, samples from the posterior of the PoS can be obtained from MCMC iterations:

$\widehat{POS}_i(d) = \Phi\left(\frac{m(d; \theta^{(i)}) - m(0; \theta^{(i)}) - 1.96 \times \sqrt{2SE^2}}{\sqrt{2SE^2}}\right)$, where $\theta^{(i)}$ is the vector of efficacy model parameters θ simulated at iteration i . The advantage of Bayesian framework over a purely frequentist approach lies in its ability to account for the uncertainty in parameter values in the decisional process and also, in allowing greater flexibility in the definition of the decision rules (defined in Section 4.1.2.2).

On the other hand, the number of patients having a toxicity is a binomial distribution of parameters $N_3/2$ and $\pi(d)$, where $\pi(d)$ represents the probability of toxicity corresponding to dose d as defined in Section 4.1.1. The observed proportion of patients having a toxicity in phase III, $tox_{obs}(d) = \#\text{patients with toxicity}/(N_3/2)$, gives an estimation of $\pi(d)$. The safety component $\mathbb{P}(tox_{obs}(d) \leq t)$ is then the estimated probability of controlling over-toxicity, i.e. the probability of observing a toxicity rate less than or equal to t in phase III. Likewise, a posterior distribution of toxicity model parameters is obtained using a MCMC approach, where $\lambda^{(i)}$ is the simulated value of the toxicity model parameter vector λ obtained at iteration i .

Note that efficacy and toxicity are modelled and simulated as independent random variables.

4.1.2.2 Optimal dose and decision rules

For each study, the sponsor makes two decisions:

- (i) Identification of the recommended dose: at each MCMC iteration, one identifies the best dose as the dose with the highest utility score: for all doses d_j , we compute an estimation (MCMC) of $\mathbb{P}_{post}(d_j \text{ has the highest utility})$, written as $\widehat{\mathbb{P}}_{post}(d_j = \text{optimal dose} | \text{data})$ in all the following of the manuscript; the recommended dose d^* for phase III is the dose for which this probability is the highest one, i.e. the dose being the most often identified as the best one among all MCMC iterations. In (the unlikely) case two doses have exactly the same probability of being the best dose, the lower dose is chosen and recommended for phase III. We call this Decision rule 1 (we compared alternative decision rules, Decision rule 1*, Decision rule 2, Decision rule 3 and Decision rule 4, see **Chapter 5**).

- (ii) Go / NoGo choice decision: the sponsor computes the average of estimated PoSs and the average of the estimated toxicity probabilities for the recommended dose d^* among all MCMC iterations denoted by $mean_{MCMC}(\widehat{PoS}(d^*))$ and $mean_{MCMC}(\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq t))$ respectively. The 'Go' for phase III is then decided if these averages pass prefixed efficacy and toxicity thresholds denoted by `threshold.eff` and `threshold.safe` respectively. In other words, the sponsor chooses 'Go' if $mean_{MCMC}(\widehat{PoS}(d^*)) > \text{threshold.eff}$ and $mean_{MCMC}(\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq t)) > \text{threshold.safe}$. These thresholds are at the study level, they depend on the therapeutic area and the objectives of the study.

Sequential design We consider also the case of a sequential design and propose an adapted utility-based decisional framework. The sequential design consists in performing an interim analysis when a fraction (for instance half, as in the simulations we performed) of the total sample size has been enrolled: following the interim analysis, the sponsor might decide to terminate the study or to continue until the total planned sample size is enrolled. Regarding the interim analysis, we propose a simple and intuitive method: one stops at the interim analysis if the posterior probability for the selected dose to be the optimal dose is higher than or equal to a threshold l , i.e. if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | data] \geq l$, where $l \in [0, 1]$. The details related to the computational aspects (estimation) of this interim analysis criterion are given in the following Section 4.1.2.3. Threshold l should be wisely predefined by the sponsor, it should be high enough to guarantee accuracy of the dose choice, but not too high, otherwise frequency of early termination will be decreased and studies will be rarely terminated at interim. The design of the study is the following: the sponsor plans a phase II study with N_2 patients (for example $N_2 = 500$ patients), and decides to do an interim analysis with N'_2 patients ($N'_2 < N_2$, for example $N'_2 = 250$). Then, the strategy is as follows: if the dose choice is obvious with N'_2 patients, one stops at the interim analysis (and then chooses whether to go or not to phase III with the chosen dose, see following discussion), otherwise one goes to the end of the phase II study with N_2 patients in total. To decide to stop the phase II study, we took advantage of the Bayesian framework of data analysis: we stop the phase II study with N'_2 patients if the identified dose d^* is very likely to be (conditional to data) the best one, i.e. we stop if " $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | data] \geq l$ ". Note that an early termination of the trial at the interim analysis is not necessarily a positive outcome: we can also stop the analysis for futility, i.e. we stop at interim and we do not Go to phase III, with the same decision criteria as the ones for the fixed design (if $mean_{MCMC}(\widehat{PoS}(d^*)) < \text{threshold.eff}$ and $mean_{MCMC}(\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq t)) < \text{threshold.safe}$ at interim). The values of these thresholds could be different for the interim analysis compared to those predefined for the final analysis. One could imagine more restrictive thresholds for the interim, and less constraining ones for the final analysis. In our work, the interim and final analyses are conducted the same way. But in fact, according to sponsor's objectives related to the interim analysis, they could be conducted completely differently. For instance, if the only aim of the interim analysis is to assess if the drug shows some efficacy or not (with no further objective to identify the optimal dose), then a specific decision rule could be built in relation to the efficacy of the largest dose only (for instance the decision rule could be defined as a minimal PoS in phase III for the largest dose; studies would be stopped if efficacy of the largest dose is insufficient). In this example, studies would be stopped only for futility (we only stop for failure, never for success).

4.1.2.3 Optimal dose estimation method: Batching approach

Posterior distributions could be computed with a MCMC approach. More precisely, it could be done via subsampling method (also known as thinning), [Geyer et al.2011], to remove autocorrelations, or via batching method, [Alexopoulos and Seila1996]; [Fishman and Yarberry1997]; [Schmeiser and Song1996]; [Geyer et al.2011], to avoid both information loss and autocorrelations, and to possibly estimate the variance of a MCMC estimator (the latter issue was not particularly the purpose of implementing this method in this chapter). In this chapter, a batching method is implemented for a different purpose: govern the dose selection process in refining the dose selection rule mentioned in the previous section. Indeed, the main idea is to select the dose $d^* = d_j$ such that $\hat{\mathbb{P}}_{post}(d_j = \textit{optimal dose} | \textit{data})$ has the highest value amongst all doses. Instead of simply computing standard MCMC estimates of those posterior probabilities (i.e. $\hat{\mathbb{P}}_{post}(d_j = \textit{optimal dose} | \textit{data})$ for all doses d_j), we will apply a batching method that will be described thereafter: it consists in computing, first, partial sums of the MCMC utilities iterates over batches of a sufficiently large length. With the latter method, smoother and more concentrated posterior distributions are obtained and therefore, two aspects are ensured: reducing variability and avoiding information loss within the chain. In the following, we delineate the statistical process of this method, and we give a detailed explanation of why this method should not be applied directly to model parameters, but only to estimated utilities (based on parameter estimates).

In the following, we define the Markov chain (X) by $X_i = (X_{1_i}, X_{2_i})^t$, where X_{1_i} and X_{2_i} represent the MCMC Markov chains related to efficacy and safety parameters respectively. The i th element of our Markov chain (X) represents candidate vectors $\theta^{(i)}$ and $\lambda^{(i)}$ of θ and λ respectively, at each iteration i . Let n_{iter} be the total number of MCMC iterations.

Subsampling the original Markov chain at spacing k , is the classical process that takes every k th element of the Markov chain X_1, X_2, \dots , forming a new Markov chain $X_1, X_{k+1}, X_{2k+1}, \dots$. This method is mainly used to reduce autocorrelations. The spacing of iterations that contributes to the thinning process of the Markov chain is usually denoted by $nspac$. Subsampling consists then in keeping one iteration every ' $nspac$ ' iterations.

Batching the Markov chain is a different process. In our context, we are not primarily interested by $\mathbb{E}_{\pi_{post_{\theta, \lambda}}}(\theta, \lambda)$, where $\pi_{post_{\theta, \lambda}}$ is the posterior distribution of θ and λ , but rather by the mean of our utility function $U(d, \theta, \lambda)$, $\mathbb{E}_{\pi_{post_{\theta, \lambda}}}(U(d, \theta, \lambda))$. However, if a classic MCMC batch form is applied to model parameters and defined by partial averages of length B , $(\theta^{(i)B}, \lambda^{(i)B})^t = \frac{1}{B} \sum_{s=0}^{B-1} X_{(i-1)B+s}$ approximates $\mathbb{E}_{\pi_{post_{\theta, \lambda}}}(\theta, \lambda)$, for B large enough (B is chosen so that the partial averages $\theta^{(i)B}$ and $\lambda^{(i)B}$ have the lowest possible autocorrelation). The total number of iterations of a Markov chain, n_{iter} , corresponds to $n \times B$ where n is the number of iterations in the final output. In other words, $U(d, \theta^{(i)B}, \lambda^{(i)B}) \simeq U(d, \mathbb{E}_{\pi_{post_{\theta, \lambda}}}(\theta, \lambda))$ and partial averages $(\theta^{(i)B}, \lambda^{(i)B})^t$ will lead to an estimation of $U(d, \mathbb{E}_{\pi_{post_{\theta, \lambda}}}(\theta, \lambda))$ instead of $\mathbb{E}_{\pi_{post_{\theta, \lambda}}}(U(d, \theta, \lambda))$.

So with this classic batching method (applied directly to model parameters estimates) and with a large B , we estimate the utility of the posterior mean of the parameters instead of

estimating the posterior mean of the utility, this means that results would be different (see Section 5.2.1 for instance).

In order to correctly estimate our desired posterior distributions, we can do the same discussion as above but with the following changes: based on the initial Markov chain representing estimates of our model parameters, we apply the batching method to the estimated utilities (calculated for each $(\theta^{(i)}, \lambda^{(i)})^t$), in order to compute posterior probabilities. The MCMC estimator of interest (described thereafter) is then the empirical mean of the batched utility partial means.

So this is what happens in our actual context, we do not particularly want to calculate the posterior average of the parameters of the efficacy and toxicity models, but to choose the dose, and for that, we must estimate the four posterior probabilities for each of the four active doses, based on Bayesian estimates of $U(d, \theta, \lambda)$. In the following, we remind that d_1, d_2, d_3 and d_4 represent the active dose indices, and correspond to dose values $d = 2, 4, 6, 8$ respectively.

For a given iteration i , and for $j = 1, \dots, J$, we form partial sums $\gamma = \frac{1}{B} \sum_{s=0}^{B-1} \begin{pmatrix} U(d_1; X_{(i-1)B+s}) \\ \vdots \\ U(d_J; X_{(i-1)B+s}) \end{pmatrix}$;

this vector of partial sums approximately follows a normal distribution $N(m_U, \Sigma_U/B)$, centered on the posterior utility mean m_U , with a covariance matrix equal to Σ_U/B , which is the asymptotic covariance matrix divided by the length of the batch B . In fact, by definition, the asymptotic covariance is the matrix Σ such that the covariance matrix of any MCMC mean of length L , $\frac{1}{L} \sum_{l=0}^{L-1} U_l$, is approximately equal to Σ/L , see [Geyer et al.2011] chapter 1.

The desired batching is then directly applied to the utility posterior distribution. Consequently, the posterior probability that the batched utility is largest at a certain dose is:

$$\eta = \frac{1}{n} \sum_{i=1}^n \begin{pmatrix} \mathbb{1}(u_i^B(d_1) > u_i^B(d_j), \forall j \neq 1) \\ \vdots \\ \mathbb{1}(u_i^B(d_J) > u_i^B(d_j), \forall j \neq J) \end{pmatrix}, \text{ where } u_i^B = \frac{1}{B} \sum_{s=0}^{B-1} U(d_j, X_{(i-1)B+s}).$$

Consequently, an estimation of the posterior probabilities computed to choose the optimal dose can be written as follows:

$$\eta = \begin{pmatrix} \widehat{\mathbb{P}}(u^B(d_1) > u^B(d_j), \forall j \neq 1) \\ \vdots \\ \widehat{\mathbb{P}}(u^B(d_J) > u^B(d_j), \forall j \neq J) \end{pmatrix}, \text{ where } \begin{pmatrix} u^B(d_1) \\ \vdots \\ u^B(d_J) \end{pmatrix} \sim N(m_U, \Sigma_U/B).$$

Those estimated probabilities correspond to $\widehat{\mathbb{P}}_{post}(d_j = \text{optimal dose} | \text{data})$ for the four doses d_j , defined in Section 4.1.2.2. Based on these quantities, we recommend dose d_* for phase III for which this probability is the highest one compared to the others (i.e. the dose being the most often identified as the best one among all batches). The interpretation of the dose selection rule after application of the batching method is as follows: we select the dose that has the highest probability $\widehat{\mathbb{P}}(u^B(d) > u^B(d_j), \forall d_j \neq d)$ where the vector u^B is normally distributed, centered

at the vector of the posterior mean utility values and covariance matrix equal to the asymptotic covariance matrix of the MCMC mean estimates of utilities divided by the length of the batch. By using this Gaussian distribution to drive the dose selection, we expect more robust dose selection rules (that can be monitored by the choice of the batch length B) than in simply using the raw MCMC estimates of the posterior distribution of the utility values. Regarding the interim analysis, we stop at interim if this latter probability (of the chosen dose d^*) is higher than a threshold l , in other words, if $\widehat{\mathbb{P}}[u^B(d^*) > u^B(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l$. This is an estimation of the interim analysis criterion defined in Section 4.1.2.2.

Convergence properties of the estimates

Posterior means of utility values The posterior means of the utility values are the quantities $\mathbb{E}_{\pi_{post_{\theta,\lambda}}}(U(d, \theta, \lambda))$. The MCMC estimates of those values are the empirical means: $\frac{1}{n_{iter}} \sum_{i \leq n_{iter}} U(d, \theta^{(i)}, \lambda^{(i)}) = \frac{1}{n_{iter}} \sum_{i \leq n_{iter}} U(d, X_i)$ where X_i is the value at iteration i of the Markov chain generated by the Metropolis-Hastings algorithm ($X_i = (\theta^{(i)}, \lambda^{(i)})^t$ are the estimates of the θ and λ parameters respectively at iteration i of the algorithm) whose stationary distribution, $\pi_{post_{\theta,\lambda}}$, is the posterior distribution of the vector model parameters θ and λ conditional on phase II data, and n_{iter} is the total number of iterations. The convergence of the empirical means $\frac{1}{n_{iter}} \sum_{i \leq n_{iter}} U(d, X_i)$ towards $\mathbb{E}_{\pi_{post_{\theta,\lambda}}} U(d, \theta, \lambda)$ when $n_{iter} \rightarrow \infty$ is justified by the application of the ergodic theorem (see Theorem 6.63 in [Robert and Casella2004]): the necessary condition of irreducibility of the Markov chain (X) is guaranteed by the application of Lemma 7.6 in [Robert and Casella2004]. Indeed, we are using a 'random-walk' version of Metropolis-Hastings algorithm with a normal working conditional density of the form $q(y|x) = \frac{1}{\sqrt{2\pi\tau^2}} \exp(-\frac{1}{2\tau^2}(y-x)^2)$, and then one can easily find a value δ such that $|x-y| < \delta \Rightarrow q(y|x) > \frac{1}{2} \frac{1}{\sqrt{2\pi\tau^2}}$ for instance; therefore Lemma 7.6 in [Robert and Casella2004] can be used, showing the irreducibility of the Markov Chain, then combined with Lemma 7.3 in [Robert and Casella2004], proving the validity of the ergodic theorem.

Batching method Concerning the partial sums $u_i^B = \frac{1}{B} \sum_{s=0}^{B-1} U(d, X_{B(i-1)+s})$, we can notice the two following points:

- The justification of approximate normality of the partial sums u_i^B relies on the validity of the central limit theorems for random-walk Metropolis-Hastings algorithms; a result on this validity can be found in [Roberts and Tweedie1996]: it requires some additional technical assumptions as compared to the ergodic theorem.
- The empirical means $\frac{1}{n} \sum_{i=1}^n u_i^B$ converge also towards $\mathbb{E}_{\pi_{post_{\theta,\lambda}}} U(d, \theta, \lambda)$ when $n_{iter} \rightarrow \infty$, since the empirical mean of the partial sums is in fact the overall mean over the MCMC iterations: indeed $\frac{1}{n} \sum_{i=1}^n u_i^B = \frac{1}{n} \sum_{i=1}^n \frac{1}{B} \sum_{s=0}^{B-1} U(d, X_{B(i-1)+s})$
 $= \frac{1}{n} \frac{1}{B} \sum_{r=0}^{n_{iter}-1} U(d, X_r) = \frac{1}{n_{iter}} \sum_{r=0}^{n_{iter}-1} U(d, X_r)$ (since $n_{iter} = n \times B$);

the convergence of $\frac{1}{n_{iter}} \sum_{r \leq n_{iter}} U(d, X_r)$ towards $\mathbb{E}_{\pi_{post_{\theta, \lambda}}} U(d, \theta, \lambda)$ is justified in the paragraph above. Therefore, even though the use of the batching method does not decorrelate consecutive MCMC iterations (when not applied directly on the model parameter estimates but on the utility estimates instead), the obtained utility estimates are still consistent despite potential autocorrelation within iterations: those estimates as well as the estimates for $PoS(d)$ and $\mathbb{P}(tox_{obs}(d) \leq t)$ are used for the Go/NoGo decision rules.

4.1.3 Simulations

In the following, we describe our simulation protocol and our chosen efficacy/safety dose-response scenarios. Note that all the chosen values stated thereafter are applied for the analysis of each simulated phase II trial and are the same in each scenario.

4.1.3.1 Simulation protocol

We simulated 1000 phase II studies in total (stability of the results was checked, by simulating 5000 phase II studies for a given scenario for instance). For each simulated trial, we made the following assumptions.

Indeed, our models can be applied to different numbers of doses (or even different dosages) but for our simulations, we consider four active doses with the following values, $d = 2, 4, 6, 8$, and one placebo with the following value, $d = 0$.

We consider informative priors for E_0 and ED_{50} , and non-informative prior for E_{max} : $E_{max} \sim N(0, 100)$, $ED_{50} \sim U[1, 10]$ and $E_0 \sim N(0, 1)$. Regarding ED_{50} , we considered this prior as it is consistent with the fact that at this stage of drug development, phase II or phase IIb, the sponsor has quantitative information (based on pre-clinical or phase I/pharmacodynamic studies) about the relevant dose range and that this reflects in design doses. Regarding the prior for E_0 , we assume that, similarly, the sponsor has some information on the range of placebo effect.

The following informative prior distributions for the parameters of the Probit model are considered: intercept $a \sim N(q_{0.05}, 0.10^2)$, where $q_{0.05} \simeq -1.65$ is the normal distribution quantile which corresponds to 5% of adverse event in placebo arm, and dose effect $b \sim U[0, 1]$. The sponsor is considered here to have information on the percentage of toxicity in the placebo group (from epidemiological data, for instance), so the Probit model parameter a is centred around its true value and with limited variability: a coefficient of variation (i.e. ratio between standard deviation (=0.10) and mean ($=|q_{0.05}| \simeq 1.65$) roughly equal to 6%). In real life, these choices are never completely non-informative, we often have an idea on the incidence of adverse event in the placebo arm. Concerning the slope, b , the choice of this prior was motivated by a conservative approach, assuming that the incidence of toxicity was necessarily increasing with the dose.

Sensitivity analyses were conducted in order to examine the performance of the designs with respect to different priors (by considering non informative priors for all model parameters for instance). Results were promising as they were consistent with the ones obtained with the chosen priors in this section, but needed more patients to reach similar properties and decision

rule qualities (see Appendix A.7). Some additional guidelines for prior elicitations are given in **Chapter 6**. Density plots of our prior dose-response distributions are given in Appendix A.7.

We consider $N_3 = 1000$. In practice, phase III sample size is usually set to achieve a statistical power between 80% and 95%. It should be defined based on our understanding of the endpoint, relevant effect and what the drug might achieve. In case overwhelming efficacy is expected by the project team, a smaller phase III sample size can be envisaged as well.

Efficacy and toxicity are modelled and simulated as independent random variables to limit autocorrelation problem. We simulate $n_{iter} = 150000$ safety and $n_{iter} = 150000$ efficacy parameters separately, and then we combine both datasets in order to build the utility score for each dose / iteration. Among these iterations, we discard an initial portion of the Markov chain sample so that the effect of initial values on the posterior inference is minimized: burn-in= $150000/2=75000$ first iterations.

PoS, toxicity component and utility are computed at each MCMC iteration level. Once utilities are estimated based on each $\theta^{(i)}$ and $\lambda^{(i)}$ (75000 estimated utilities after burn-in process, see Appendix A.2 for further details), we implement the batching method to compute posterior probabilities based on the estimated utilities. We consider a batch length, $B = 150$ (the choice of this value is also discussed in Appendix A.2). In the final output, we will then have: $n = 75000/150 = 500$ batches, each batch representing the posterior partial mean of the utility for each dose. Sponsor will use these partial means to rank doses according to utility scores and choose the optimal one as explained in Sections 4.1.2.2 and 4.1.2.3. Tables summarizing simulation results of the 1000 simulated studies are presented in Section 4.2, each result is an average value calculated over all phase II studies.

Regarding the Go/NoGo decision (for the fixed design), we have proposed the decision criteria based on threshold values for the PoS and for the probability of observing a toxicity rate lower than t in phase III. These values will depend on the therapeutic area and the objectives of the study; for efficacy, it could be equal to 0.30 in Oncology for instance; we tested `threshold.eff = 30%`, and it turned out to be too weak and not strict enough (see simulation results in Appendix A.3), we also tested `threshold.eff = 90%` which, as expected, was too restrictive and with this threshold we do not go often enough to phase III (see simulation results in Appendix A.4); we finally kept an intermediate threshold (moderate and reasonable) between the two (`threshold.eff = 60%`). So in simulations, we finally retained an efficacy decision criterion for the PoS, with 60% set as lower bound, and a safety decision criterion for the probability of observing a toxicity rate lower than t , with 50% set as lower bound. For simplicity purposes, the same threshold values are retained for the interim analysis.

We choose $l = 0.80$ for the interim analysis criterion.

We stop at interim if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq 0.80$.

The choice of this threshold is discussed in **Chapter 6** ($l = 0.90$ is tested in **Chapter 5**). We compare fixed designs ($N_2 = 250$, $N_2 = 500$ and $N_2 = 1000$ patients) with sequential designs with an interim analysis when half of the patients are enrolled ($N_2 = 500$ and $N_2 = 1000$ patients with an interim analysis at $N'_2 = 250$ and $N'_2 = 500$ patients respectively). Note

that we also examined the performance of the designs with 50 and 100 patients, but with such small sample sizes, the posterior distributions did not permit to rank the utilities consistently with their true values. Therefore, results related to those designs are not given in this chapter (see **Chapter 5**).

In this work, we consider $t = 0.15$. Note that this is an arbitrary choice, and usually depends on the therapeutic area. For instance, a threshold of 0.30 (or 0.40) is more common in oncology and may vary in other areas; see Appendix A.7 for sensitivity analysis related to the choice of this threshold.

In the following, we consider $h = 1$ and $k = 2$ (respective weights given to the contribution of efficacy and safety to the utility function). The choice of these parameters is discussed in **Chapter 6**. Sensitivity analysis related to the choice of these parameter values were also conducted, see Appendix A.7.

The residual variability σ is assumed to be known and set to the value of 0.5 in the simulations. This value has been chosen in order to have, for one of our most important scenarios, named "Sigmoid" (defined in the following Section 4.1.3.2), a standardized effect of 0.25 for the highest dose ($d = 8$) of our design. This effect size is in the range/order of magnitude of effect size generally targeted in drug development (it is admitted that the standard effect size of clinical importance observed from most clinical trials is within the range of 0.25 and 0.5, see [Cohen1988]). According to simulation results, $\sigma = 0.5$ seemed to be a reasonable choice in terms of estimation quality and dose choice.

The R package `mcmc` was used for sampling the posterior distribution. Computations were parallelized using R package `parallel` to speed up computation.

4.1.3.2 Simulation scenarios for efficacy and toxicity

We remind that $E_0 = \theta_1$, $E_{max} = \theta_2$, $ED_{50} = \theta_3$ for the efficacy model parameters, and $a = \lambda_1$, $b = \lambda_2$ for the toxicity model parameters.

We consider three main efficacy scenarios assumed to be the true ones reflecting the real dose-response:

- (i) Sigmoid scenario: this scenario corresponds to a smooth increase of the effect over the dose range of the design: plateau effect barely reached for the highest design dose (see Figures 4.4, 4.5, 4.6 and 4.7); for this scenario, the true efficacy model parameters values are: $(E_{max}, ED_{50}, E_0) = (0.22, 6, 0)$.
- (ii) Plateau scenario: this scenario begins with an almost linear growth, followed by an inflection, and then stabilizes at the end, which means that the last two doses have the same efficacy (see Figure 4.8); for this scenario, the true efficacy model parameters values are: $(E_{max}, ED_{50}, E_0) = (0.14, 0.9, 0)$.
- (iii) No activity scenario: it is considered to evaluate the type I error (see Figures 4.1, 4.2 and 4.3): the idea is to verify that the clinical trial stops for lack of activity, and not

because of excessive toxicity; for this scenario, the true efficacy model parameters values are: $(E_{max}, ED_{50}, E_0)=(0, 6, 0)$.

We also consider five main toxicity scenarios (each of the efficacy scenario was combined with each of the following toxicity scenarios; in total, 15 scenarios were assessed through simulations and for sake of simplicity, only some particular scenarios of interest are presented in this chapter):

- (i) Scenario with a progressive toxicity; for this scenario, the true toxicity model parameters values are: $(a, b)=(-1.645, 0.076)$, and the theoretical toxicities for $d = 0, d = 2, d = 4, d = 6$ and $d = 8$ doses are: 0.05, 0.07, 0.09, 0.12, 0.15 respectively, so the toxicity probability for the highest dose ($d = 8$) is equal to the fixed over-toxicity threshold ($=0.15$, see Section 4.1.3.1), see Figure 5.2.
- (ii) Scenario with a progressive toxicity; for this scenario, the true toxicity model parameters values are: $(a, b)=(-1.645, 0.100)$, and the theoretical toxicities for $d = 0, d = 2, d = 4, d = 6$ and $d = 8$ doses are: 0.05, 0.07, 0.11, 0.15, 0.20 respectively, so the toxicity probability is strictly higher than 0.15 for the highest dose ($d = 8$), see Figures 4.1, 4.4 and 4.8.
- (iii) Scenario where the safety of all doses is acceptable, except for the highest dose $d = 8$ (in this case $d = 4$ and $d = 6$ doses became quite close in terms of utility); for this scenario, the true toxicity model parameters values are: $(a, b)=(-2.054, 0.152)$, and the theoretical toxicities for $d = 0, d = 2, d = 4, d = 6$ and $d = 8$ doses are: 0.02, 0.04, 0.07, 0.13, 0.20 respectively, see Figure 4.5.
- (iv) Scenario with low toxicity; for this scenario, the true toxicity model parameters values are: $(a, b)=(-1.645, 0.045)$, and the theoretical toxicities for $d = 0, d = 2, d = 4, d = 6$ and $d = 8$ doses are: 0.05, 0.06, 0.07, 0.08, 0.10 respectively, so the toxicity is at 10% even for the highest dose, see Figures 4.2 and 4.6.
- (v) Scenario with high toxicity; for this scenario, the true toxicity model parameters values are: $(a, b)=(-1.645, 0.152)$, and the theoretical toxicities for $d = 0, d = 2, d = 4, d = 6$ and $d = 8$ doses are: 0.05, 0.09, 0.15, 0.23, 0.33 respectively, so the toxicity is already at 15% from the second active dose, see Figures 4.3 and 4.7.

4.2 Results

Denoting a simulation scenario by the product of the associated efficacy and toxicity scenarios, eight particular simulation scenarios are considered: no activity scenario \times scenario with progressive toxicity (and toxicity of highest dose = 0.20), no activity scenario \times scenario with low toxicity, no activity scenario \times scenario with high toxicity, Sigmoid scenario \times scenario with progressive toxicity (and toxicity of highest dose = 0.20), Sigmoid scenario \times scenario with acceptable safety except for highest dose, Sigmoid scenario \times scenario with low toxicity, Sigmoid scenario \times scenario with high toxicity, and Plateau scenario \times scenario with progressive toxicity (and toxicity of highest dose = 0.20). For each simulation scenario, a graph highlighting the corresponding theoretical curves is drawn (see Figures 4.1 to 4.8), where 'Toxicity penalty' red curve represents the probability of observing more than 15% of toxicity in

phase III, and a table summarizing all the results is given (see Tables 4.1 to 4.8). This table contains the following:

- (i) 'E(U)' is the empirical utility expectation of the chosen dose for the 1000 simulated phase II studies among 'Go' and 'NoGo' decisions (utility is set to 0 when it is a 'NoGo' decision)
- (ii) 'Prob(choose(Go))' is the empirical probability of going to phase III with the chosen dose
- (iii) 'Distribution selected doses (Conditional to 'Go')' represents the empirical probabilities of choosing the $d = 2, 4, 6$ and 8 dose respectively among the 'Go'
- (iv) 'Distribution selected doses (Conditional to 'Go') at interim analysis' is the empirical distribution of the chosen doses if we choose 'Go' for the interim analysis
- (v) 'Distribution selected doses (Conditional to 'Go') at final analysis' is the empirical distribution of the chosen doses if we continue to the final analysis and we choose 'Go'
- (vi) 'POS(conditional to 'Go')' is the empirical PoSs mean among the 'Go' with the chosen dose
- (vii) 'Prob(Stop at interim)' is the empirical probability of stopping at the interim analysis
- (viii) '% Stop for futility' is the empirical probability of stopping for futility at interim (so this percentage is included in (vii))
- (ix) 'Mean(N2)' is the mean sample size of the sequential plan
- (x) 'Power' is the global power of the combined phase II / phase III program, defined as the product (ii) \times (vi)

4.2.1 No activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

We started by considering a scenario with no activity to evaluate the type I error: the idea is to verify that the clinical trial stops for lack of activity, and not because of excessive toxicity. The utility function is illustrated in Figure 4.1 and results are given in Table 4.1.

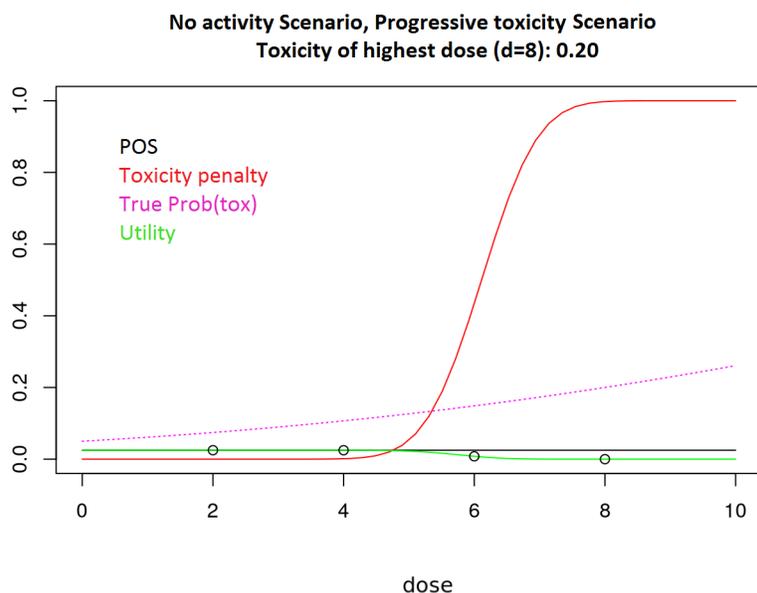


Figure 4.1: Theoretical curves, no activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

No activity scenario Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.003	0.002	0.001	0.001	0.000
Prob(choose(Go))	0.108	0.079	0.059	0.048	0.018
Distribution selected doses (Conditional to 'Go')	0.050 0.860 0.070	0.010 0.900 0.090	0.000 0.900 0.100	0.000 0.960 0.040	0.000 0.720 0.280
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.020 0.980 0.000	-	0.1 0.0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.780 0.220	-	0.000 0.670 0.330	-
POS (conditional to 'Go')	0.025	0.025	0.025	0.025	0.025
Prob(Stop at interim)	-	0.249	-	0.290	-
% Stop for futility	-	0.202	-	0.248	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	438	-	855	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.003	0.002	0.001	0.001	0.000

Table 4.1: Simulation results, no activity scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

In this scenario, the sponsor should not decide to go to phase III since no dose is efficacious as compared to placebo. In terms of probability of wrong decision (decide to go to phase III), it is quite high ($\simeq 11\%$) with a phase II study with $N_2 = 250$ patients. But, as expected, the probability of wrong decision decreases as the sample size increases, reaching the value of approximately 2% for the largest phase II study ($N_2 = 1000$ patients). In the unfavourable

case of wrong decision to go into phase III, the chosen dose is most often $d = 4$. This is due to the fact that the analysis conducted by the sponsor identifies the second dose as the highest "well tolerated" dose (based on the probability of observing more than 15% of toxicity in the phase III study). In such a scenario, the usefulness of conducting an interim analysis when half of the patients are enrolled is debatable. Indeed, the probability of stopping at interim analysis is not negligible (it is around 25% and 29% for sample size of 250 and 500 at interim, respectively) this leads to a decrease of the mean sample size of the phase II study of around 12% and 14% as compared to a fixed sample size design of 500 and 1000 patients respectively. But at the same time, even though the probability of interrupting the study and choose to go directly in phase III is small, conducting an interim analysis inflates the risk of wrongly choosing to go in phase III as compared to the fixed sample size design (risk increases from 6% to 8% with the phase II study with 500 patients and the risk increases from 2% to 5% with the phase II study with 1000 patients).

4.2.2 No activity with low toxicity scenario

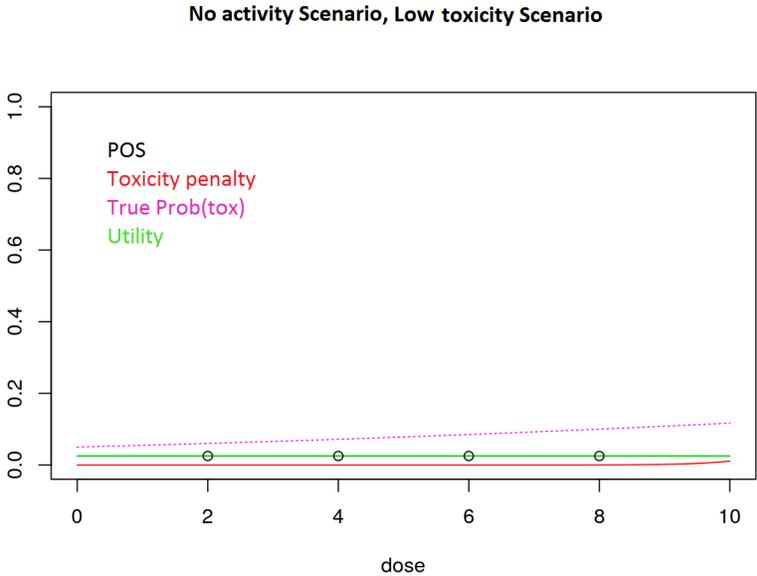


Figure 4.2: Theoretical curves, no activity scenario with low toxicity scenario. The PoS and utility curves are superimposed.

No activity scenario Low toxicity Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.004	0.002	0.002	0.001	0.001
Prob(choose(Go))	0.156	0.098	0.100	0.059	0.052
Distribution selected doses (Conditional to 'Go')	0.000 0.170 0.330	0.000 0.040 0.210	0.000 0.020 0.350	0.000 0.000 0.050	0.000 0.000 0.100
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.000 0.170 0.000 0.830	-	0 0 0 1	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.000 0.280 0.720	-	0.000 0.000 0.120 0.880	-
POS (conditional to 'Go')	0.025	0.025	0.025	0.025	0.025
Prob(Stop at interim)	-	0.146	-	0.244	-
% Stop for futility	-	0.122	-	0.210	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	464	-	878	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.004	0.002	0.002	0.001	0.001

Table 4.2: Simulation results, no activity scenario with low toxicity scenario.

In this scenario in which, again, the sponsor should not decide to go to phase III, the conclusions are roughly similar as for the previous scenario.

- The probability of wrong decision are even slightly larger than in the previous scenario. This is due to the fact that, because of the good safety profile of the drug, the wrong identification of a satisfactory efficacy (as assessed by the estimated POS) is not counter-balanced by toxicity warnings. Because of that, the probability of wrong decision is a little bit high with a phase II study of 250 patients sample size. It reaches the more acceptable levels of 10% and 5% with phase II studies of size 500 and 1000 patients respectively.
- Because of the absence of toxicity, the model on which are based the sponsor's decisions most often identifies the largest dose as the optimal one: especially for the largest phase II studies.
- The impact of the interim analysis is similar as in the previous scenario, but with a lower magnitude because the probability to stop at the interim analysis was slightly lower as in the previous scenario.

4.2.3 No activity scenario with high toxicity scenario

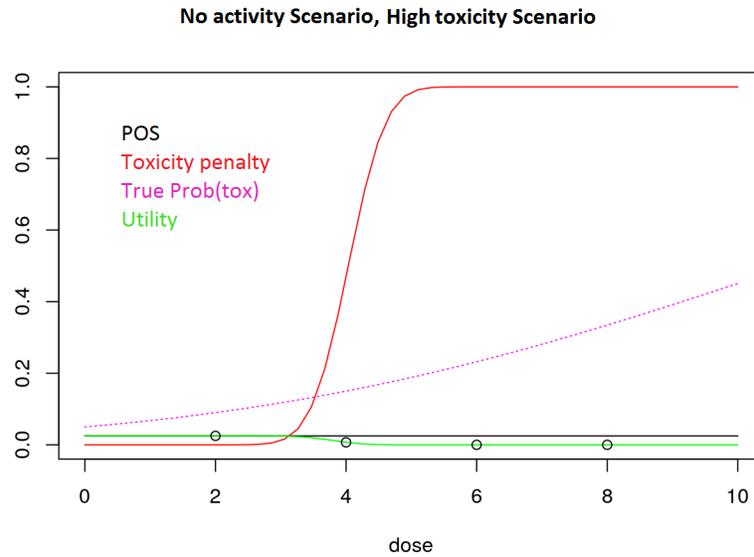


Figure 4.3: Theoretical curves, no activity scenario with high toxicity scenario. The PoS and utility curves are partially superimposed, especially for low doses.

No activity scenario High toxicity Threshold.eff2=0.60 Threshold.safe2=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.001	0.001	0.000	0.000	0.000
Prob(choose(Go))	0.049	0.038	0.014	0.009	0.001
Distribution selected doses (Conditional to 'Go')	0.730 0.270 0.000 0.000	0.790 0.210 0.000 0.000	0.500 0.500 0.000 0.000	0.330 0.670 0.000 0.000	0 1 0 0
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.910 0.090 0.000 0.000	-	0.430 0.570 0.000 0.000	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.170 0.830 0.000 0.000	-	0 1 0 0	-
POS (conditional to 'Go')	0.025	0.025	0.025	0.025	0.025
Prob(Stop at interim)	-	0.652	-	0.745	-
% Stop for futility	-	0.620	-	0.738	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	337	-	628	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.001	0.001	0.000	0.000	0.000

Table 4.3: Simulation results, no activity scenario with high toxicity scenario.

This is the worst case scenario: the drug does not show any efficacy but at the same time the safety is poor with a quasi linear steep dose-toxicity curve. The same conclusions as for the previous scenarios apply:

- The probability of making the good decision (stop after phase II) increases with the sample size of the phase II study. But the profile of the drug is so "bad" that decision making is easy: with only 250 patients in the phase II study, the probability of taking the wrong decision is as low as 5% and is equal to 0.1% for the largest phase II study.
- In this scenario, the interim analyses are very useful since, because taking the right decision (stop for futility because of the high toxicity and the lack of efficacy) does not require many patients, the probability to stop at the interim analysis is extremely high: equal to 65% (equal to 62% for futility) with 250 patients at the interim analysis and equal to 75% (equal to 74% for futility) with 500 patients at the interim analysis.

4.2.4 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

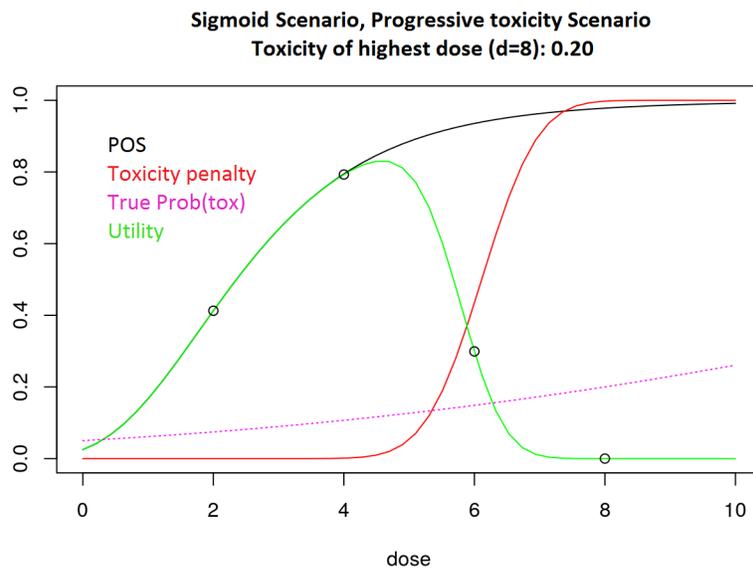


Figure 4.4: Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20). The PoS and utility curves are partially superimposed, especially for low doses.

Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.401	0.473	0.478	0.522	0.547
Prob(choose(Go))	0.555	0.628	0.626	0.672	0.704
Distribution selected doses (Conditional to 'Go')	0.120 0.830 0.050 0.000	0.040 0.910 0.050 0.000	0.020 0.940 0.040 0.000	0.010 0.970 0.030 0.000	0.000 0.970 0.030 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.060 0.940 0.000 0.000	-	0 1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.030 0.890 0.080 0.000	-	0.010 0.900 0.090 0.000	-
POS (conditional to 'Go')	0.757	0.785	0.791	0.796	0.798
Prob(Stop at interim)	-	0.404	-	0.655	-
% Stop for futility	-	0.130	-	0.195	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	399	-	673	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.420	0.493	0.495	0.535	0.562

Table 4.4: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

In this scenario, the optimal dose is the second one ($d = 4$) and the true associated PoS and utility are both approximately equal to 0.8 (see Figure 4.4). With this scenario we can see that the probability of making the good decision (go to phase III with the second dose) is clearly dependent on the sample size of the phase II study, the probability of good decision increasing significantly with the phase II sample size. When it is equal to 250, the sponsor decides to go to phase III with a probability approximately equal to 56%, whereas the global power is equal to 42%: this sample size does not seem large enough for a sufficiently accurate estimation of both efficacy and safety models to allow good decisions. With 1000 patients in phase II, i.e. the largest phase II study, the probability of choosing to go to phase III significantly increases and reaches 70% with the largest phase II study: concomitantly, when the sample size is increased from 250 to 1000 patients, the mean utility and the global power relative increase is approximately equal to 15% .

In terms of choice of dose, the best dose ($d = 4$) is the selected one in most of the cases, even with only 250 patients (chosen with probability equal to 83%). But again, increasing the sample size significantly improves further the dose selection: with 500 patients in the phase II study, the best dose is selected for phase III with a probability approaching 95%.

In this scenario, performing an interim analysis when half of the patients are recruited has some interest: the probability of stopping at interim analysis is quite high (but the probability of wrong stop for futility is not negligible, equal to 13% and 20% for the interim analyses at 250 and 500 patients respectively) which leads to relative decrease of the mean sample size of 20% and 33% as compared to the fixed sample size design with 500 and 1000 patients respectively. This is interesting because this reduction of sample size does not degrade the properties of the design: considering either the probability of going to phase III, the mean utility, the selected doses or the global power, the design with an interim analysis with half patients has very similar properties as the fixed sample size design.

4.2.5 Sigmoid scenario with acceptable safety scenario (except for highest dose)

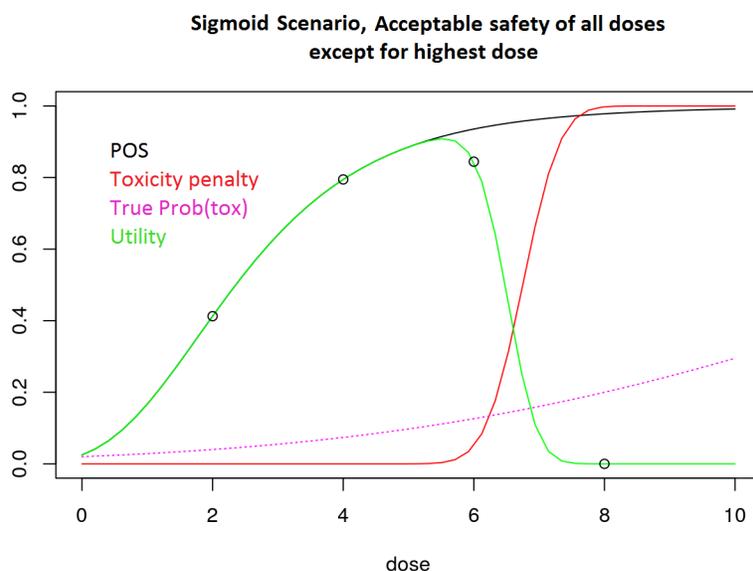


Figure 4.5: Theoretical curves, Sigmoid scenario with acceptable safety scenario (except for highest dose). The PoS and utility curves are partially superimposed, especially for low doses.

Sigmoid Acceptable safety of all doses Except for d=8 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.453	0.523	0.526	0.609	0.607
Prob(choose(Go))	0.577	0.654	0.654	0.753	0.750
Distribution selected doses (Conditional to 'Go')	0.030 0.820 0.140	0.000 0.800 0.200	0.000 0.770 0.230	0.000 0.710 0.290	0.000 0.700 0.300
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.010 0.990 0.000	-	0.1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.660 0.330	-	0.000 0.470 0.530	-
POS (conditional to 'Go')	0.805	0.822	0.827	0.835	0.836
Prob(Stop at interim)	-	0.373	-	0.438	-
% Stop for futility	-	0.106	-	0.099	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	407	-	781	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.465	0.538	0.541	0.629	0.627

Table 4.5: Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).

In this scenario, two doses are very similar in terms of utility: the best dose is the third one ($d = 6$) with a utility approximately equal to 0.84, but the second dose ($d = 4$) is very close

with a true utility approximately equal to 0.79 (see Figure 4.5).

Most of the conclusions related to the previous scenario apply. The probability of choosing to go to phase III increases with the sample size of the phase II study. For the two largest phase II studies (500 or 1000 patients), planning an interim analysis when half of the patients are enrolled, enables to reduce the mean sample size while maintaining the properties of the design. In this scenario the gain in mean sample size is less than in the previous one; this is due to the fact that the probability of stopping at interim analysis is reduced as compared to the previous scenario: this is explained by the similarity in true utility values of the doses 4 and 6, the interim analysis is less likely to identify a dose with posterior probability of being the optimal dose ≥ 0.80 .

As the sample size increases, the probability of selecting the true optimal dose increases as well, but is only equal to 30% when the sample size of the phase II is equal to 1000 patients.

Under this scenario, doses $d = 4$ and $d = 6$ have very similar utilities according to Figure 4.5. Yet dose $d = 4$ is very predominantly and consistently selected from $N_2 = 250$ to $N_2 = 1000$. The reason is that the toxicity penalty is dramatically growing just beyond dose $d = 6$. The uncertainty in the toxicity rate estimation in phase II will result in a significant proportion of simulated trials with a quite high toxicity penalty for dose $d = 6$ disqualifying that dose. Thus, our approach tends to disadvantage doses with a true toxicity rate close to the threshold of 0.15. This can be viewed as a conservative property to protect the patient.

This scenario illustrates the difficulty and the requirements in terms of necessary sample size for the models to provide sufficiently accurate estimations, for both efficacy and toxicity models simultaneously, to enable a proper ranking of doses with very similar utility values.

4.2.6 Sigmoid scenario with low toxicity scenario

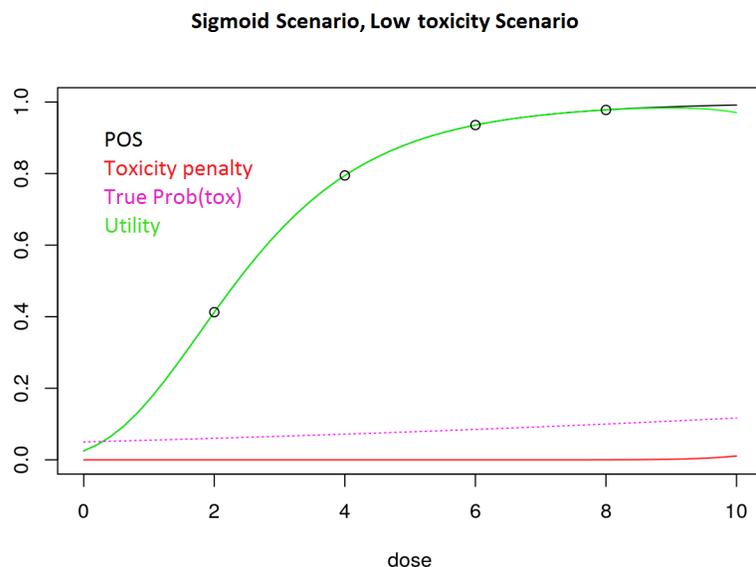


Figure 4.6: Theoretical curves, Sigmoid scenario with low toxicity scenario. The PoS and utility curves are almost identical.

Sigmoid Low toxicity Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.596	0.706	0.712	0.794	0.830
Prob(choose(Go))	0.648	0.746	0.751	0.824	0.861
Distribution selected doses (Conditional to 'Go')	0.000 0.230 0.350 0.410	0.000 0.080 0.380 0.540	0.000 0.070 0.400 0.530	0.000 0.010 0.280 0.710	0.000 0.010 0.290 0.710
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.000 0.350 0.000 0.650	-	0.000 0.030 0.000 0.970	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.050 0.430 0.520	-	0.000 0.000 0.360 0.640	-
POS (conditional to 'Go')	0.921	0.947	0.949	0.964	0.965
Prob(Stop at interim)	-	0.137	-	0.251	-
% Stop for futility	-	0.045	-	0.074	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	466	-	875	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.597	0.706	0.713	0.794	0.831

Table 4.6: Simulation results, Sigmoid scenario with low toxicity scenario.

In this scenario, the drug has very low toxicity, therefore the highest dose ($d = 8$) is the optimal one, with a utility value close to 0.98. The utility value of the preceding dose ($d = 6$) is equal to 0.94, it is quite close to the optimal value.

The same conclusions related to the previous scenarios apply as well. This scenario points out the following facts:

- For a drug profile with several adjacent doses showing similar utility values, identifying the optimal dose with sufficient "assurance" requires a phase II study with quite a large sample size. Indeed with only 250 patients in phase II, the optimal dose (the highest dose) is properly identified with a probability as low as 41%. With 1000 patients, this probability is raised to 71%.
- In this scenario, the interim analysis conducted when half of the planned 500 patients are enrolled is not very useful as the probability to stop the trial is low: around 14%. This is due to the fact that 250 patients are not sufficient to clearly identify the best dose (the distribution of the chosen dose with the fixed design with 250 patients is an illustration of this). With the interim analysis conducted when the first half of the 1000 patients are enrolled the probability to stop the trial raises to approximately 25%: this corresponds to an approximate reduction of the mean sample size of 12%: this is not negligible but it is much lower than in some of the previous scenarios.

4.2.7 Sigmoid scenario with high toxicity scenario

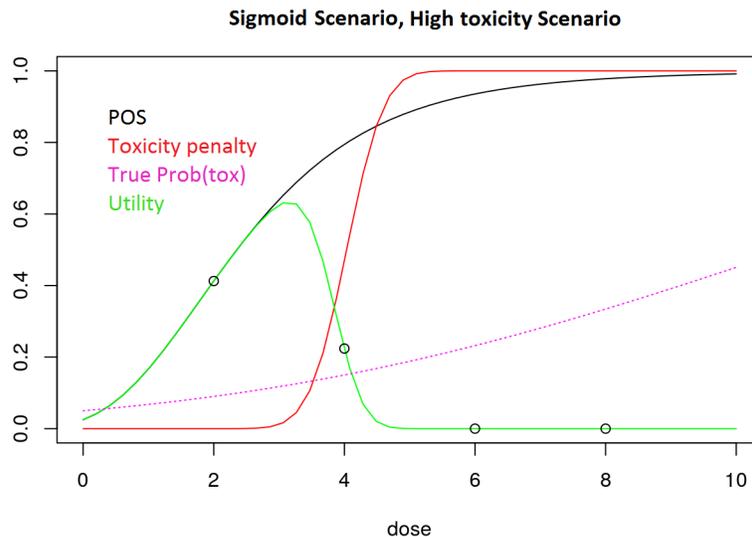


Figure 4.7: Theoretical curves, Sigmoid scenario with high toxicity scenario. The PoS and utility curves are partially superimposed, especially for low doses.

Sigmoid High toxicity Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.140	0.144	0.129	0.131	0.105
Prob(choose(Go))	0.361	0.367	0.337	0.337	0.277
Distribution selected doses (Conditional to 'Go')	0.860 0.140 0.000 0.000	0.890 0.110 0.000 0.000	0.840 0.160 0.000 0.000	0.870 0.130 0.000 0.000	0.830 0.170 0.000 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.970 0.030 0.000 0.000	-	0.980 0.020 0.000 0.000	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.670 0.330 0.000 0.000	-	0.530 0.470 0.000 0.000	-
POS (conditional to 'Go')	0.466	0.453	0.475	0.464	0.477
Prob(Stop at interim)	-	0.669	-	0.684	-
% Stop for futility	-	0.393	-	0.435	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	333	-	658	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.168	0.166	0.160	0.156	0.132

Table 4.7: Simulation results, Sigmoid scenario with high toxicity scenario.

This scenario is interesting in the sense that due to very high toxicity, the best dose is the smallest one ($d = 2$), see Figure 4.7. But at the same time, the PoS associated with this dose is too low, according to the sponsor's criteria (PoS $\simeq 0.4$ which is much lower than the sponsor's threshold of 0.60). Therefore, in this scenario, the good decision is not to go to phase III because the only "safe" dose is not effective enough.

Like for the other scenarios the probability of taking the good decision increases with the sample size: it ranges from 64% for the study with the smallest sample size to 72% with the largest one.

Just like the scenario with no efficacy of the drug (for which the best decision is also not to go to phase III) planning an interim analysis reduce the mean sample size but, at the same time, slightly inflates the probability of wrong decision.

4.2.8 Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

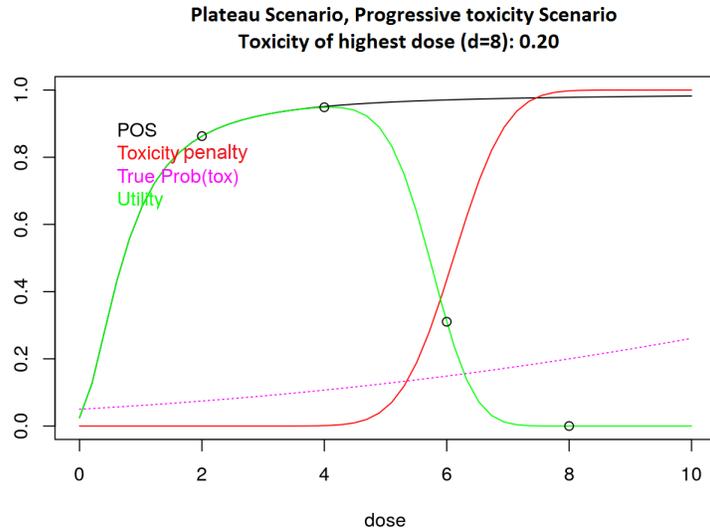


Figure 4.8: Theoretical curves, Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20). The PoS and utility curves are partially superimposed, especially for low doses.

Plateau Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.508	0.572	0.584	0.639	0.666
Prob(choose(Go))	0.558	0.624	0.633	0.686	0.716
Distribution selected doses (Conditional to 'Go')	0.120 0.840 0.040 0.000	0.050 0.910 0.040 0.000	0.030 0.930 0.040 0.000	0.010 0.970 0.030 0.000	0.010 0.970 0.030 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.070 0.930 0.000 0.000	-	0.1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.030 0.900 0.080 0.000	-	0.010 0.920 0.070 0.000	-
POS (conditional to 'Go')	0.941	0.948	0.949	0.951	0.951
Prob(Stop at interim)	-	0.395	-	0.641	-
% Stop for futility	-	0.135	-	0.198	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	401	-	680	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.525	0.591	0.601	0.652	0.681

Table 4.8: Simulation results, Plateau scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

As illustrated in Figure 4.8, in this scenario, the dose-response function reaches a plateau from the first dose: the PoS of the lowest dose ($d = 2$) is approximately equal to 0.86 and is superior to 0.95 for all other doses. Therefore, with very small differences in efficacy between the doses, the best dose will be the highest dose with toxicity lower than 15%: this corresponds to the second dose, $d = 4$. In fact, our approach is able to differentiate between doses $d = 2$ and $d = 4$ which nevertheless have close utilities, favoring dose $d = 4$. This is due to the fact that both doses have a negligible toxicity penalty and their respective true toxicity rates are rather far below the threshold of 0.15. Despite the uncertainty in the toxicity rate estimation in phase II, the risk of having a large toxicity penalty is low for both doses. Thus the dose selection is mainly driven by the PoS, i.e. favoring dose $d = 4$.

The results observed with the previous scenarios apply here as well: the probability of making the good decision increases with the sample size; also, the probability of selecting the optimal dose increases with the phase II sample size. Performing an interim analysis reduces the mean sample size by approximately 20% and 32% as compared to the design with fixed sample size of 500 and 1000 patients respectively.

We can also notice a quite large percentage of studies stopped for futility. This is probably due to a quite high toxicity of the doses $d = 6$ and $d = 8$ that induces, for some occurrences of the simulations, a 'NoGo' decision after the interim analysis, due to the wrongly high estimated toxicity rate of the chosen dose.

4.3 Conclusions

In this chapter, we aimed to propose a decision-making framework, based on the utility function selected amongst those studied at the previous chapter (**Chapter 3**). This decision-making framework enables the sponsor to make the two following decisions after completion of the phase II dose-finding study: decide to go to phase III or not and, in case of positive answer, choose the dose for the phase III trial. The properties of these decision rules were assessed, through simulations, under various safety and efficacy profiles and several sample sizes of the phase II study (250, 500 and 1000 patients). Those simulations confirmed that estimating an optimal dose, optimal with respect to a utility function, is a demanding task: they illustrated the difficulty of simultaneously estimating two complex dose-response models with enough accuracy to properly rank doses using an utility function combining the two. This is due to the fact that, in particular for the smallest sample sizes, the posterior distribution of the utilities, for each of the doses, is not sufficiently concentrated around the true values of the utilities: this often leads to imprecise estimations of the posterior probabilities of the dose with the highest utility score (computed for all doses d_j), which are the quantities used for dose selection, and then to wrong selection of the optimal dose. Without surprises, it appeared that the probability of making the good decision increases with the sample size, but even with the largest sample size, the probability of making the good decision with a large phase II study of 1000 patients only reaches 80% when the drug does not show any toxicity. The simulations clearly show that, regardless of time and budget constraints, the sponsor has always interest in running large phase II studies to make accurate decisions regarding the termination of the development program or the selection of the dose.

We assessed the possibility to conduct an interim analysis, when half of the patients are enrolled, and try to make the selection at this stage: we proposed to stop the trial at the interim analysis if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq 0.80$. For some scenarios, in particular when the best dose shows a clear benefit in the utility as compared to the others, this approach has good properties: with a quite large probability of study termination at interim analysis, it enables to reduce the sample size while maintaining the properties of the fixed large sample size design.

Chapter 5

Utility-based dose-finding in practice: some empirical contributions and recommendations

In this chapter, the aim is to highlight some contributions and recommendations related to the utility function proposed in **Chapter 4** with the same parameter values for h , k and t : $U(d) = PoS(d) \times \mathbb{P}(tox_{obs}(d) \leq 0.15)^2$. In Section 5.1, we assess the properties of the posterior distributions (posterior means or posterior medians), by simulating trials with sequentially increasing sample size, and we graphically examine (through violin plots) the properties of the estimated posterior distributions of the utility values. In Section 5.2, we perform exploratory analyzes, comparing alternative decision rules for dose selection. Section 5.3 is dedicated to the relative utility loss criterion that we suggest in order to make recommendations on phase II sample size. In Section 5.4, we compare several stopping criteria for interim analysis, by comparing different fixed/sequential designs. Finally, Section 5.5 summarizes the main findings of this chapter, and underlines some recommendations regarding specific decision rules for the dose selection and for the interim data inspection. As a reminder, the efficacy and safety models are:

- Efficacy: Emax model defined in **Chapter 2** page 33 and in **Chapter 3** page 48.
- Toxicity: Probit model defined in **Chapter 4** page 59.

We used the same prior distributions as defined in **Chapter 4** page 65, as well as the optimal dose estimation method (batching approach) described in Section 4.1.2.3.

5.1 Estimation of the posterior distribution of the utility values

The aim of this section is only to describe, through two examples (Sigmoid scenario combined with a progressive toxicity scenario where the toxicity of the highest dose is equal to 0.15, and Sigmoid scenario combined with a progressive toxicity scenario where the toxicity of the highest dose is equal to 0.20), the behaviour of the MCMC estimates of the utility values and how their properties are modified by the sample size of the study.

5.1.1 Exploring properties of MCMC based posterior distributions: examples of two particular single simulated studies

To visualize the posterior distribution of the utility values (using the MCMC simulations), we used the violin plots. They are much more convenient and informative than boxplots (see Appendix A.5) to characterize multimodal distributions. Distribution plots are an efficient tool to evaluate the quality of the MCMC utility estimates. The boxplot is a classic approach for visualizing basic distributions. It is convenient for comparing usual summary statistics (such as range and quartiles) for unimodal distributions, but is of poor interest for multimodal distributions. The violin plot is a hybrid of a boxplot and a kernel density plot, with the ability to show multiple peaks of the data distribution. Note that kernel density estimation [Chen2017] is a non-parametric method of estimating the Probability Density Function (PDF) of a continuous random variable. Essentially, at every datum, a symmetric kernel function is created with the datum at its centre. The PDF is then estimated by adding all of these kernel functions and dividing by the number of data to ensure that it satisfies the two properties of a PDF: every possible value of the PDF is non-negative, and the definite integral of the PDF over its support set equals to 1.

More precisely, violin plots are built as follows:

- A kernel estimator of the density function (non-parametric estimator of the density curve) is computed as follows:
 - starting from a kernel v , i.e. a symmetric PDF, in general chosen such that the variance is equal to 1.
 - the density estimate at point y based on sample $(Y_i)_{1 \leq i \leq n}$ is defined by:
$$\hat{f}(y) = \frac{1}{nz} \sum_{i=1}^n v\left(\frac{y - Y_i}{z}\right),$$
 where z is a key tuning parameter, the bandwidth; in general, the magnitude of the chosen bandwidth is of order $n^{-1/5}$
- On each side of the black line, at the middle of the violin plot, the kernel density estimation is displayed.

Violin plots have many of the same summary statistics as boxplots (see Figure 5.1):

- The colored point represents the mean
- The black line represents the median
- The thick white bar in the center represents the interquartile range. Boundaries represent the Q1 and Q3 values
- The thin black line represents the 95% confidence interval under normality assumption

On each side, the black curve is a kernel density estimation to show the distribution shape of the data. Wider sections of the violin plot represent a higher probability that members of the population will take on the associated area (see **Zone A** in Figure 5.1); the skinnier sections

represent a lower probability (see **Zone B** in Figure 5.1).

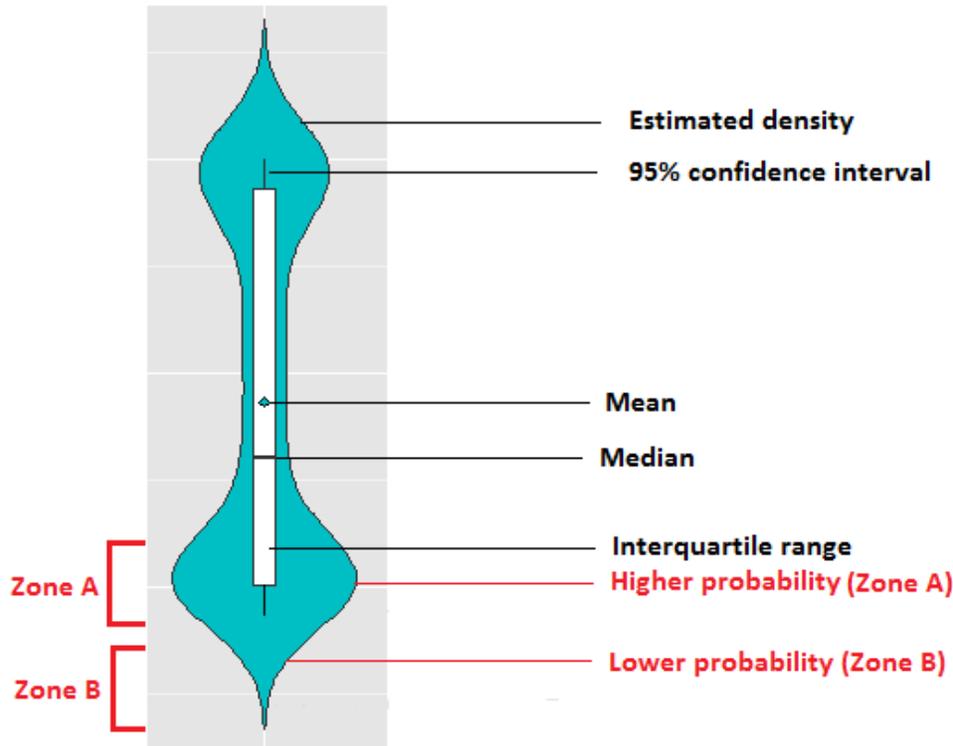


Figure 5.1: Violin plot example.

As already mentioned in the beginning of this subsection, violin plots are more convenient and informative than boxplots to understand the distribution of a dataset. Therefore, in order to highlight the challenges that might occur under real life operating conditions, we simulate one single trial (but with sequentially increasing sample size) and we graphically examine (through violin plots) the properties of the estimated posterior distributions of the utility values. The aim is to find out potential and/or atypical patterns that might occur in the distribution of the MCMC simulated data. Indeed, in real life, we only have one study of interest; the idea here is to inform the sponsor on the complexities that one may encounter in practice, especially with small phase II sample sizes, and to underline difficulties that might arise in the worst case scenarios. Note that the first 50 observations for $N_2 = 100$ are identical to those

generated with $N_2 = 50$ and similarly, the first 100 observations for $N_2 = 250$ are identical to those generated with $N_2 = 100$, etc. In other words, we do not start over from scratch each time we change N_2 . This amounts to generating 1000 observations in one go, and then considering the first 50, the first 100, the first 250, the first 500, etc. We have performed this exercise for two scenario combinations: first with the Sigmoid scenario with progressive toxicity scenario and the second one with a slightly more increased toxicity (toxicity of the highest dose equal to 0.20 instead of 0.15 for the first scenario). Each violin plot is followed by its corresponding summary table of parameter estimates for the Emax and Probit models respectively (summarized by the parameter posterior means over all MCMC iterations), as well as the table of the mean estimates (over all MCMC iterations of the relative study) of the efficacy effects versus placebo and toxicity dose-response respectively. Each mean estimate is computed by $mean_{MCMC}(m(d; \theta) - m(0; \theta))$ for the efficacy versus placebo effect, and by $mean_{MCMC}(\widehat{\mathbb{P}}(W = 1|d, \lambda))$ for the toxicity dose-response.

5.1.1.1 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15)

In the following, theoretical curves of the utility function and each of its components are represented in Figure 5.2, for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), followed by the violin plots of the corresponding estimated utilities for each dose, for one simulated phase II trial, and for $N_2 = 50, 100, 250, 500$ and 1000 respectively.

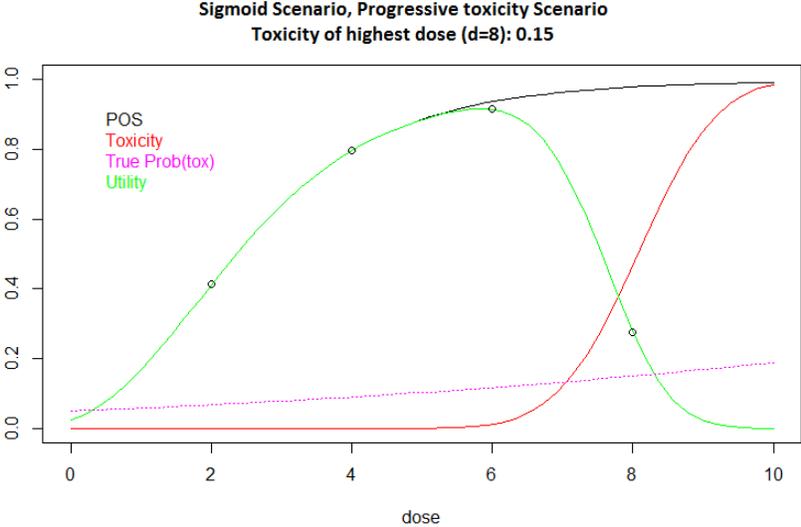


Figure 5.2: Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15). The PoS and utility curves are superimposed for small dose values.

Based on Figure 5.2, we can clearly see that this chosen scenario combination is a very challenging one: two doses have very similar utility values, the toxicity rate is low for all

doses but the last one, and consequently, this could be an unfavorable framework to assess the quality of the estimation, which becomes more challenging in this kind of situation.

In the sequel, we focus on results from a particular simulation of the considered model, to illustrate the type of behaviour that can be obtained according to the values of N_2 .

In all the following violin plots, the true range of the data varies between 0 and 1 (utilities are actually bounded between 0 and 1), but the tails of the violins are not trimmed to the range of the data because of the kernel estimation smoothing. The star symbol represents the true (theoretical) utility value.

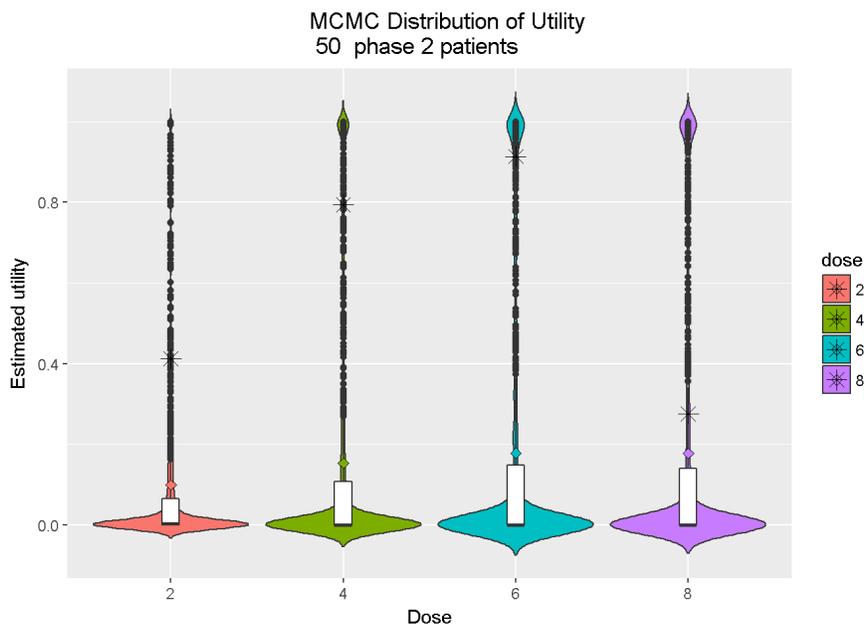


Figure 5.3: Violin plots for one simulated study, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	-0.10	6.00	0.06	-1.68	0.05

Table 5.1: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	-0.02	-0.04	-0.05	-0.06
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.06	0.07	0.09	0.10

Table 5.2: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.

The mean and median values are quite far from their true values in this example of study with only 50 patients in phase II. The posterior distribution do not permit to rank the utilities consistently with their true values.

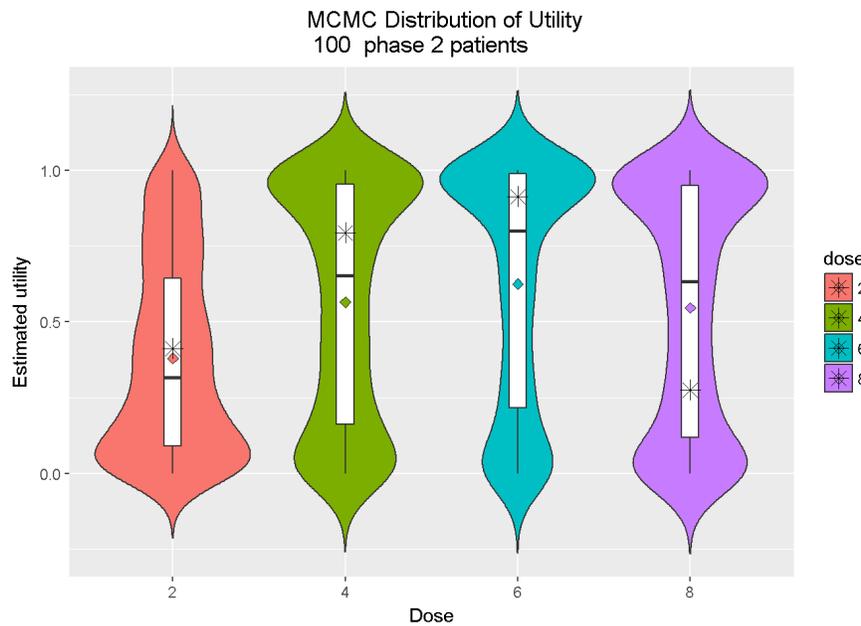


Figure 5.4: Violin plots for one simulated study, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

–	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	0.18	6.04	-0.00	-1.65	0.06

Table 5.3: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.05	0.07	0.09	0.10
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.06	0.08	0.10	0.12

Table 5.4: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.

With 50 more patients and a total of 100 patients, doses are now ranked according to their true ranking, except for $d = 8$ (according to Figure 5.4, the estimated utility of $d = 8$ is higher than the estimated one of $d = 2$, whereas the true utility of $d = 8$ is lower than the true one of $d = 2$). But posterior median and mean values are quite far from their true values.

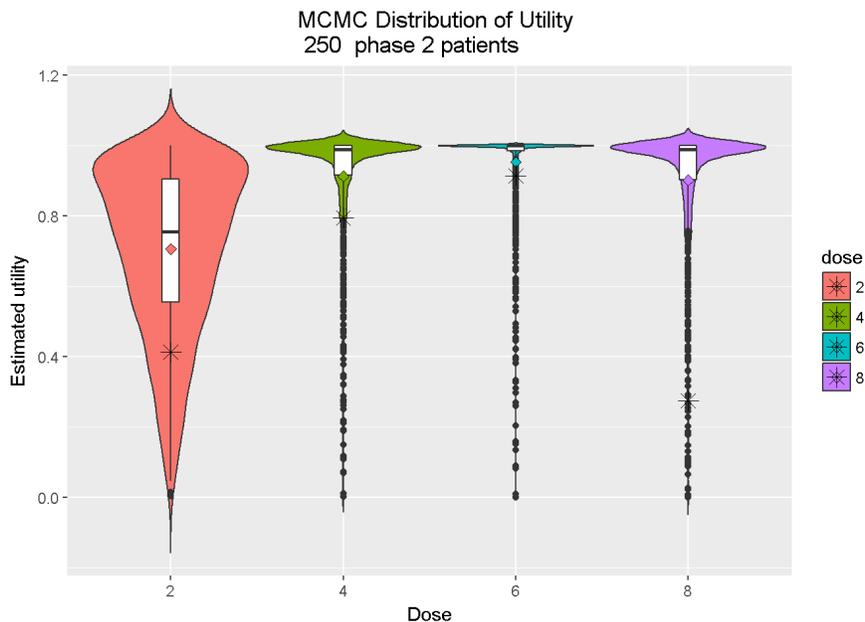


Figure 5.5: Violin plots for one simulated study, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	0.34	6.03	-0.04	-1.67	0.06

Table 5.5: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.

—	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.09	0.13	0.17	0.19
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.06	0.07	0.09	0.11

Table 5.6: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.

With 150 more patients and a total of 250 patients, the properties of the posterior distribution worsened a little bit, with a lack of differentiation of the doses and a global over-estimation of the utility values. This is due to the fact that the posterior distribution is not stable yet with 250 patients, and this result is related to one particular study.

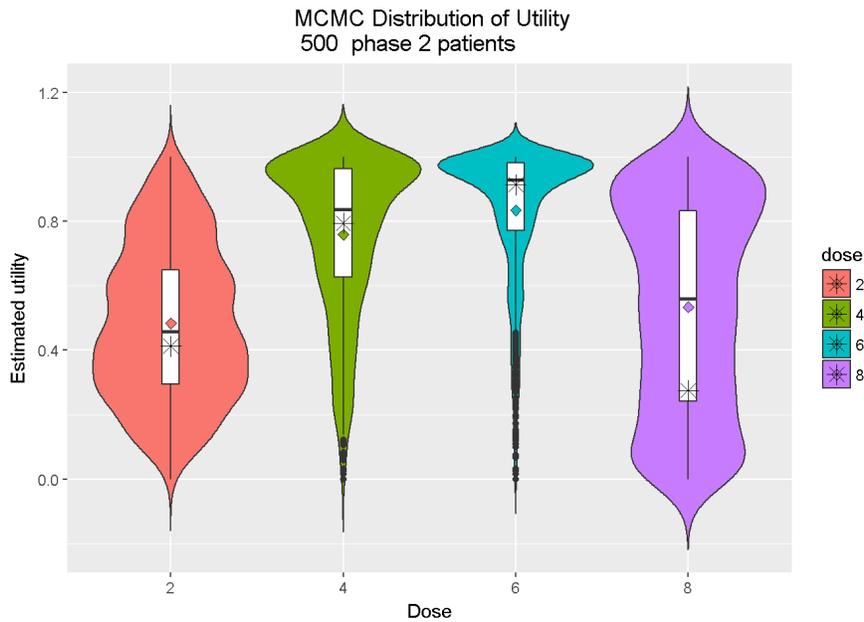


Figure 5.6: Violin plots for one simulated study, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

–	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	0.23	5.79	0.03	-1.65	0.07

Table 5.7: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.06	0.10	0.12	0.13
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.07	0.09	0.11	0.14

Table 5.8: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.

Violin plots above show us that with this chosen toxicity scenario (progressive toxicity scenario and toxicity of highest dose = 0.15) with $N_2 = 500$, estimated utilities are spread out over the whole possible area, with a slight concentration around the lower quartile (Q1) and the upper quartile (Q3) for $d = 8$. We can also see that the estimated utilities are highly concentrated on the biggest values for doses $d = 4$ and $d = 6$. Accordingly, the estimate is consistent with the true value. Globally, with 500 patients, the ranking is close to the correct one: in particular the doses $d = 4$ and $d = 6$ are clearly identified as better than the other two doses, but the utility of the largest dose seems overestimated. This is probably due to uncertainty on safety parameters of this dose, especially that its toxicity is equal to the predefined threshold, 0.15, in our utility function. This uncertainty is also study-related (we remind that this is only one simulated phase II trial), we need to simulate several studies in order to make a general assessment of the estimation quality.

The major difference observed with 500 patients, compared with 250 patients, lies in the distribution of the estimated utilities which became much more dispersed, and less concentrated (compared to 250 patients) around high values. And since this behavior is related to the study, according to the estimations of the model parameters and the dose-response functions (Tables 5.5 and 5.6), it is clear that with 250 patients, the E_{max} parameter is overestimated, and consequently the efficacy is overestimated too, but the toxicity is underestimated, which explains the global overestimation of the utility distribution showing that all the doses are "wrongly" good, with "wrongly" good efficacy and low toxicity. With 250 additional patients (i.e. a total of 500 patients), see Tables 5.7 and 5.8, this phenomenon is adjusted by the inclusion of this new patients cohort: according to the related tables, a better estimate of the E_{max} parameter is clearly noticed (and is close to the true value), which implies a better estimation of the efficacy dose-response function; in addition, a significant improvement of the toxicity estimate is well detected also, and is much closer to the true values.

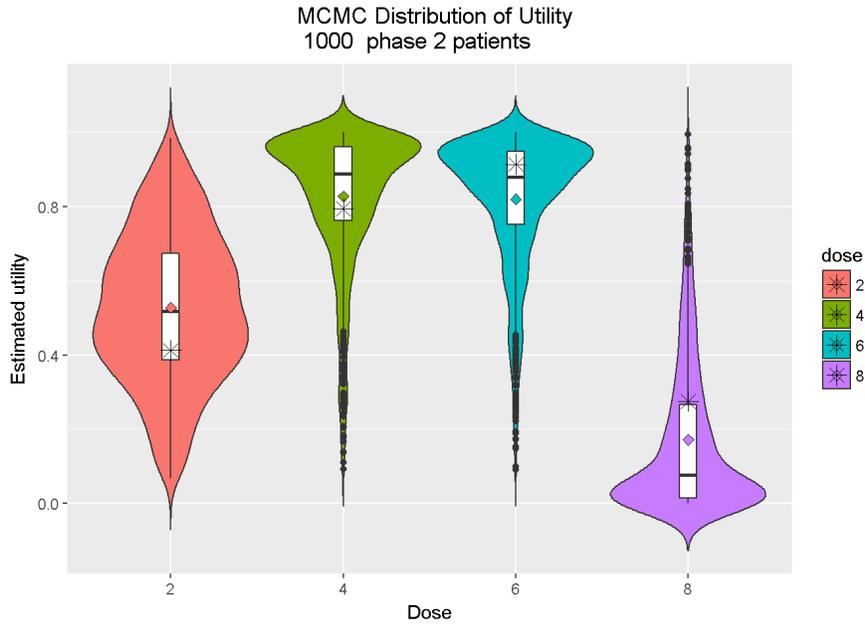


Figure 5.7: Violin plots for one simulated study, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	0.25	6.02	-0.00	-1.66	0.08

Table 5.9: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.

—	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.06	0.10	0.13	0.14
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.07	0.09	0.12	0.16

Table 5.10: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.

With $N_2 = 1000$, no significant changes are noticed compared to $N_2 = 500$, except for the dose $d = 8$, where the additional information provided allows the concentration of the obtained utilities in an area close to the true value of utility, with an underestimation leading to an overconcentration around Q1. As shown in Figure 5.8, previous results are reinforced as the sample size of the phase II study grows. With 1000 patients, the posterior distribution of the utilities clearly identify the doses $d = 4$ and $d = 6$ as the two best doses and the highest dose ($d = 8$) as the worst one. But the posterior distribution of the doses $d = 4$ and $d = 6$ are very close to each other even though, for those two doses the mean and the median of the posterior distribution are close to the true values.

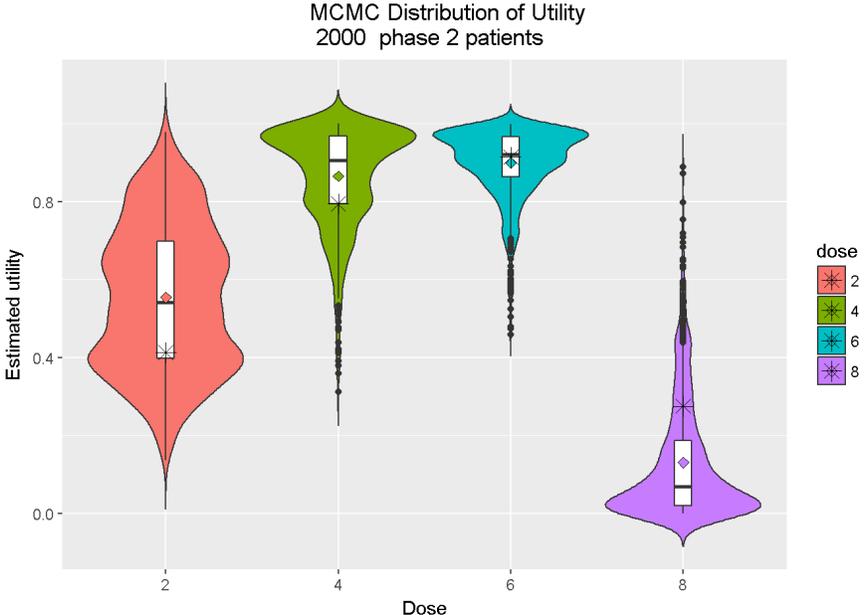


Figure 5.8: Violin plots for one simulated study, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.08
Estimates	0.31	6.17	-0.02	-1.75	0.10

Table 5.11: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.08	0.12	0.15	0.17
True toxicity	0.07	0.09	0.12	0.15
Toxicity Estimate	0.06	0.09	0.12	0.16

Table 5.12: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$.

We conclude, based on this simulated study, with sequentially increasing sample size, that:

- The posterior distribution of the utility values quite quickly identify the two middle doses as superior to the two others
- Globally, as the sample size increases, the distributions get more concentrated around the true values (see Figure 5.8) as we could naturally hope, but the shape of the distribution is very "unstable" for the lowest sample sizes
- A large number of patient is required (2000 patients), for this particular example, to differentiate the two middle-doses and clearly identify $d = 6$ as the optimal one: with this example of study, even with 1000 patients, the two doses $d = 4$ and $d = 6$ could not be distinguished by their posterior distribution of the utility values. This can be explained by the proximity of the two utility values: the utility value of $d = 6$ being only slightly larger than the one of $d = 4$.

We noticed also that the posterior median is generally different from the posterior mean indicating a possible non-symmetric posterior distribution.

On the other hand, we observed in this specific example, a particular chronology in the patterns that occurred in the posterior distributions, as the sample size increased: the posterior distribution was first concentrated on low values (with 50 patients, see Figure 5.3), then dispersed (with 100 patients, see Figure 5.4), then concentrated on high values (with 250 patients, see Figure 5.5), and finally stabilized for large phase II studies ($N_2 \geq 500$, see Figures 5.6, 5.7 and 5.8).

In the following, we will check the bias related to posterior estimates for this same simulated trial, where bias refers to the difference between the estimated utility mean or median of the posterior distribution and the true utility value, for each dose d , i.e. $\text{Bias}(d) = \text{Estimated utility}(d) - \text{theoretical utility}(d)$. We will rather express the bias in percent of the true value and then report the relative bias. We preferred to report the relative bias in order to permit a more straightforward comparison of the bias across the doses [Morris et al.2019]. Let's consider $N_2 = 500$, that corresponds to the sample size of a typical large phase IIb study. We denote the Relative Estimation Bias by REB, and we compute it for the estimated median/mean, for each dose, as follows: $\text{REB}(d) = \text{Bias}(d) / \text{theoretical utility}(d)$, over all batches for the Sigmoid scenario, with progressive toxicity scenario (and toxicity of highest dose = 0.15):

–	$\text{REB}(d = 2)$	$\text{REB}(d = 4)$	$\text{REB}(d = 6)$	$\text{REB}(d = 8)$
REB of posterior mean	$(0.50-0.41)/0.41=0.22$	$(0.79-0.80)/0.80=-0.01$	$(0.85-0.92)/0.92=-0.08$	$(0.57-0.29)/0.29=0.97$
REB of posterior median	$(0.49-0.41)/0.41=0.20$	$(0.82-0.80)/0.80=0.03$	$(0.93-0.92)/0.92=0.01$	$(0.58-0.29)/0.29=1$

Table 5.13: Mean and median relative estimation bias per dose, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.

We can see that the relative bias is rather low for most of the doses, except for $d = 8$, where the relative bias is much more pronounced. Again, this is due to a greater uncertainty for the safety of this dose.

In this example, we simulated a particular trial and we studied the posterior distribution of the utility values, sampled using a Metropolis-Hastings algorithm. Due to the closeness of the true utility values of the two middle doses in this example, we pointed-out that the posterior distribution is very sensitive to the study sample size: small values of N_2 can have dramatic effects. For intermediate values ($N_2 = 250, 500$ patients), the estimated posterior distribution quickly identified the two middle doses as superior to the others, but only with $N_2 > 1000$ patients, the posterior median and mean values of the utility values got really close to the true values, and consequently, the optimal dose was distinctly identified.

This example confirms the importance of the sample size of the phase II trial for the identification of the optimal dose. In the following, we will represent the same violin plots, but this time with a less challenging scenario combination, where the two middle doses are clearly distinguished, as compared to their relative utilities, and consequently, the optimal dose is clearly identified.

5.1.1.2 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

In order to assess how a change in safety profile can modify the properties of the MCMC estimated posterior distribution, we simulated, in the following, the same study (same efficacy

data), only modifying the toxicity data (simulated with another safety parameter value): we show the violin plots of the estimated utilities for the Sigmoid scenario (same data) with progressive toxicity scenario (and toxicity of highest dose = 0.20), for each dose, for one simulated phase II trial, and for $N_2 = 50, 100, 250, 500$ and 1000 respectively. Theoretical curves of the utility function and each of its components corresponding to this scenario combination are represented in Figure 4.4 (see **Chapter 4** page 74).

Based on Figure 4.4, we can clearly see that $d = 4$ has the highest utility value, and thus, $d = 4$ is the optimal dose.

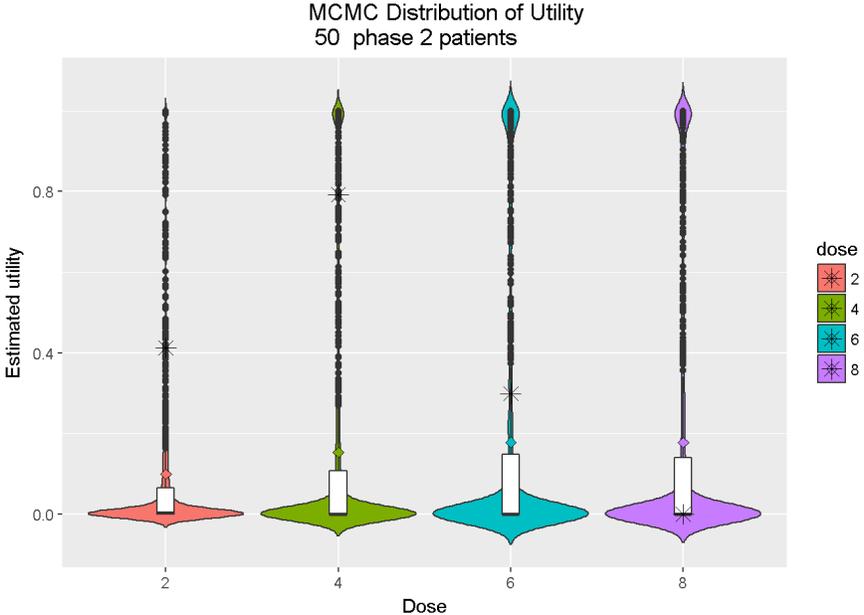


Figure 5.9: Violin plots for one simulated study, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	-0.10	6.00	0.06	-1.68	0.05

Table 5.14: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	-0.02	-0.04	-0.05	-0.06
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.06	0.07	0.09	0.10

Table 5.15: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$.

The mean and median values are quite far from their true values in this example of study with only 50 patients in phase II. The posterior distribution do not permit to rank the utilities consistently with their true values.

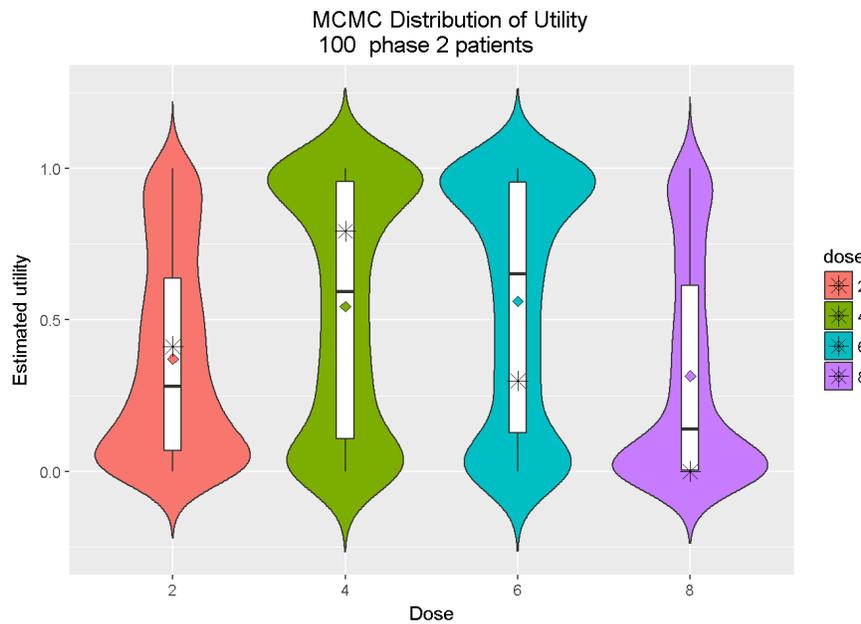


Figure 5.10: Violin plots for one simulated study, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

–	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	0.18	6.02	-0.00	-1.65	0.07

Table 5.16: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.04	0.07	0.09	0.10
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.07	0.09	0.11	0.14

Table 5.17: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$.

With 50 more patients and a total of 100 patients, doses are now ranked according to their true ranking. But posterior median and mean values are quite far from their true values.

The same behavior (according to the value of N_2) is noticed in this scenario combination (see Figures 5.9 to 5.14), compared to the previous one (see Figures 5.3 to 5.8). This is precisely due to the fact that both simulated studies share the same efficacy data, but differ in their safety data.

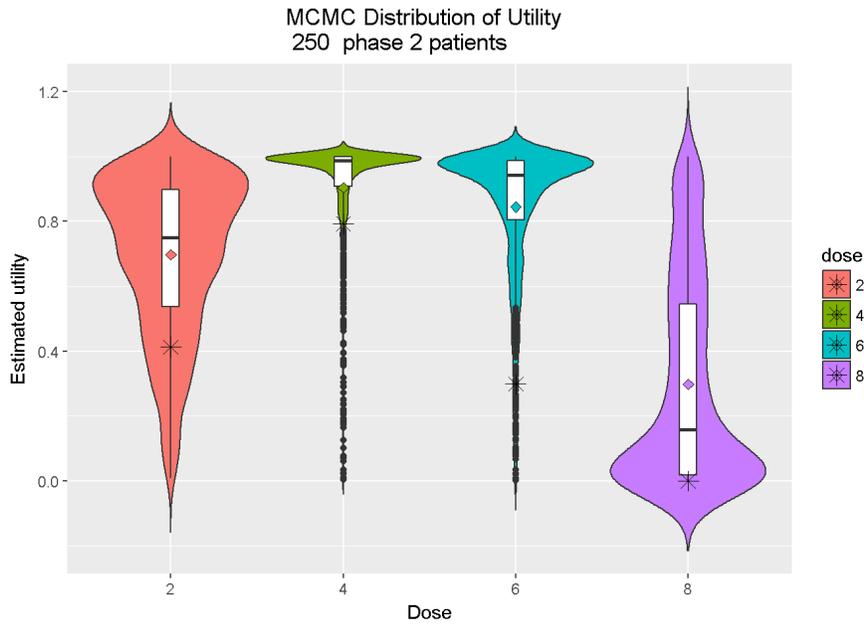


Figure 5.11: Violin plots for one simulated study, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	0.33	6.04	-0.04	-1.67	0.08

Table 5.18: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$.

—	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.08	0.13	0.17	0.19
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.07	0.09	0.12	0.16

Table 5.19: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$.

With 150 more patients and a total of 250 patients, the properties of the posterior distribution slightly improved, with the ability to differentiate the doses but with a global over-estimation of the utility values.

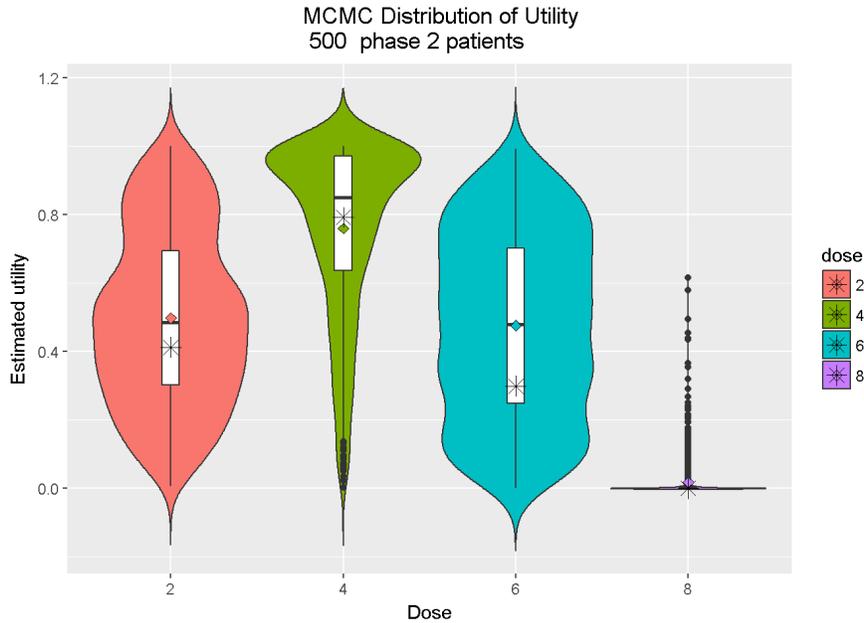


Figure 5.12: Violin plots for one simulated study, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

–	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	0.23	5.81	0.02	-1.65	0.10

Table 5.20: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$.

–	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.06	0.10	0.12	0.14
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.07	0.10	0.14	0.19

Table 5.21: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$.

Violin plots above show us that with this chosen toxicity scenario (progressive toxicity scenario and toxicity of highest dose = 0.20) with $N_2 = 500$, estimated utilities are highly concentrated around the true values for most of the doses. The dose $d = 4$ is clearly identified as the optimal one.

With 250 patients, and according to the estimations of the model parameters and the dose-response functions (Tables 5.18 and 5.19), it is clear that the E_{max} parameter is overestimated, and consequently the efficacy is overestimated too, but the toxicity is underestimated, which explains the global overestimation of the utility distribution showing that almost all the doses are "wrongly" good (except for $d = 8$), with "wrongly" good efficacy and low toxicity. With 250 additional patients (i.e. a total of 500 patients), see Tables 5.20 and 5.21, this phenomenon is adjusted by the inclusion of this new patients cohort: according to the related tables, a better estimate of the E_{max} parameter is clearly noticed (and is close to the true value), which implies a better estimation of the efficacy dose-response function; in addition, a significant improvement of the toxicity estimate is well detected also, and is much closer to the true values. Unlike the previous scenario combination, the toxicity of the highest dose (=0.20) is higher than the 0.15 threshold predefined in the safety component of our utility function; so with 500 patients, and because of the decrease in uncertainty due to the inclusion of the new patients cohort (more information is brought), the toxicity probability of $d = 8$, equal to 0.19 according to Table 5.21, is correctly estimated (close to its true toxicity probability, 0.20), with the following tight 95% confidence interval (CI)¹: CI=[0.15; 0.23], covering the true value of 0.20 with a quite good accuracy/precision. Consequently, $d = 8$ is highly penalized (because of the toxicity penalty imposed when the toxicity rate is higher than 0.15), and its related utility estimate dropped to zero (compared with 250 patients only). Again, this highlights a good property of our utility function, which tends to disadvantage doses with a toxicity probability close/equal or higher than the threshold of 0.15.

For $N_2 = 1000$ and $N_2 = 2000$ patients, conclusions are similar, even slightly better.

¹note that not all confidence intervals are presented in the estimation tables because this is not the primary objective of this subsection; this is only an example of one particular simulated phase II study

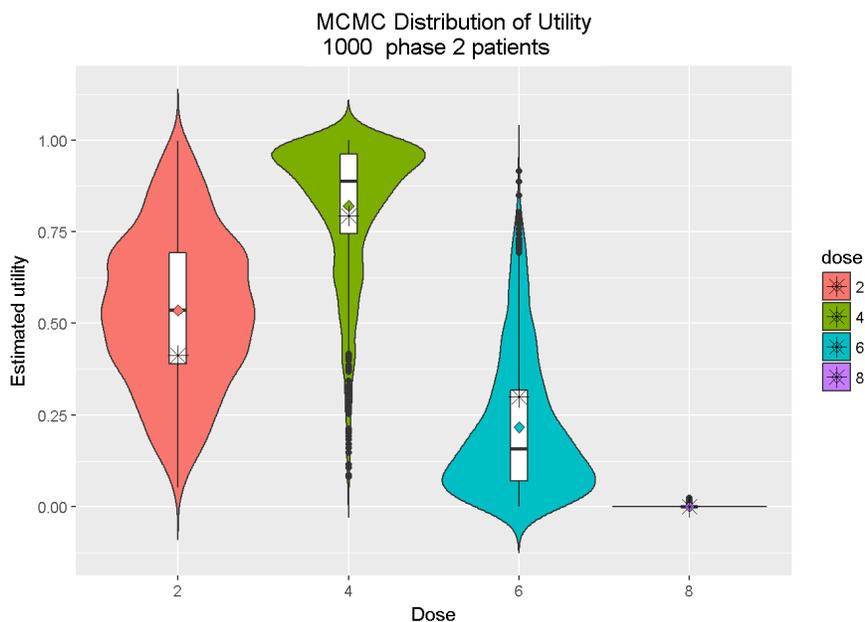


Figure 5.13: Violin plots for one simulated study, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	0.25	5.90	-0.00	-1.67	0.11

Table 5.22: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$.

—	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.07	0.10	0.13	0.14
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.07	0.11	0.15	0.21

Table 5.23: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$.

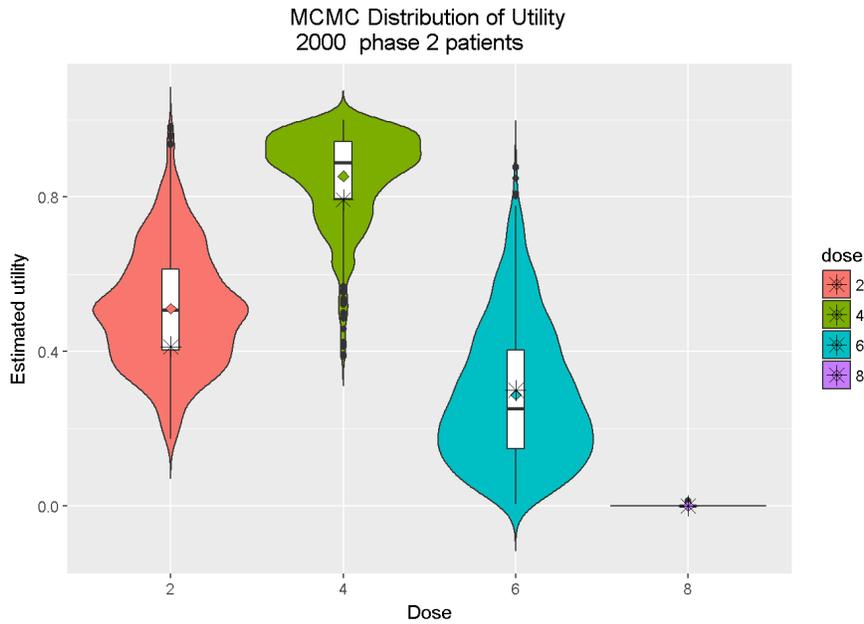


Figure 5.14: Violin plots for one simulated study, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

—	E_{max}	ED_{50}	E_0	a	b
True values	0.22	6	0	-1.65	0.10
Estimates	0.31	6.24	-0.02	-1.75	0.12

Table 5.24: Parameter estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$.

—	$d = 2$	$d = 4$	$d = 6$	$d = 8$
True efficacy	0.06	0.09	0.11	0.13
Efficacy Estimate	0.08	0.12	0.15	0.17
True toxicity	0.07	0.11	0.15	0.20
Toxicity Estimate	0.07	0.10	0.15	0.21

Table 5.25: Dose-response estimates, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$.

We conclude, based on this simulated study, with new safety data (compared to the previous example) and with sequentially increasing sample size, that:

- The posterior distribution of the utility values quite quickly identifies $d = 4$ as superior to the others
- Globally, as the sample size increases, the distributions get more concentrated around the true values as we could naturally hope
- For this particular example, there is no need for a very large sample size to differentiate the doses and clearly identify $d = 4$ as the optimal one, and the highest dose is well estimated
- With N_2 higher than or equal to 1000 patients, estimations have remarkably improved, and the posterior distributions are highly concentrated around the true values
- Posterior median and posterior means showed similar values

In the following, we will check the bias related to posterior estimates regarding this scenario combination. By considering $N_2 = 500$ as in the previous example, we get the following results:

–	REB($d = 2$)	REB($d = 4$)	REB($d = 6$)	REB($d = 8$)
REB of posterior mean	$(0.50-0.41)/0.41=0.22$	$(0.78-0.80)/0.80=-0.03$	$(0.49-0.30)/0.30=0.63$	NA ²
REB of posterior median	$(0.49-0.41)/0.41=0.20$	$(0.83-0.80)/0.80=0.04$	$(0.49-0.30)/0.30=0.63$	NA

Table 5.26: Mean and median relative estimation bias per dose, Sigmoid, with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$.

We can see that the relative bias is low for almost all doses (except for $d = 6$), and contrary to the previous scenario, the highest dose $d = 8$ is very well estimated, because its related toxicity is clearly larger than the 0.15 threshold.

5.1.1.3 Conclusions

In order to highlight the challenges that might occur under real life operating conditions, we simulated one single trial (but with sequentially increasing sample size) and we graphically examined the properties of the estimated posterior distributions of the utility values, through violin plots, for two different scenario combinations. The aim was to find out potential and/or atypical patterns that might occur in the distribution of the MCMC simulated data.

²true value too close to 0 to compute relative bias

Regarding the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), leading to a non-obvious true optimal dose:

- With the first 50 patients, the distribution appeared quite diffuse, apart from peaks at very lowest values.
- With 50 more patients, we start to observe the true values getting closer to the high probability regions. But at this stage, the posterior distribution shape is still very sensitive to the adding of new data, as with 250 patients, unexpectedly, the distribution appears very concentrated around the highest values.
- But starting from 500 patients, the true values of the utility are clearly located within the regions with highest probability, except for the highest dose, for which the convergence appears slower.

Regarding the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), leading to an obvious true optimal dose:

- The optimal dose is identified with only 250 patients, and the distribution is less sensitive to the adding of new data, compared to the previous scenario combination
- With more than 500 patients, the posterior distribution appears more and more concentrated around the true values, especially for the highest dose, because in this scenario, the toxicity of this dose is clearly higher than the 0.15 threshold, and dose $d = 4$ is clearly identified as the optimal dose

We noticed based on these single-trial simulations in two different, but similar, scenario combinations (that differ in their safety data but share the same efficacy data) that, globally, there is some concentration of the posterior distribution around the true values as the sample size increases. But in the first scenario combination, the posterior distribution could not properly capture the utility in the highest dose, unlike in the second scenario combination for which for $N_2 > 500$, predicted values were close to the true values for all doses.

Those findings illustrate how the ability to properly rank doses can depend on slight differences in the true dose toxicity relationship for instance. However, those findings concern only one single trial. In order to confirm and generalize them, we simulated, in the following, 1000 trials (for sample size varying from 50 to 2000 patients) and we assessed the properties of the two Bayesian estimates: the posterior mean and the posterior median, for both scenario combinations.

5.1.2 Properties of Bayesian estimators of utilities - posterior means and medians

5.1.2.1 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15): posterior means of utilities

In the following, we represent the violin plots of posterior **means** of utilities, for the 1000 simulated phase II studies, for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), and for several values of N_2 .

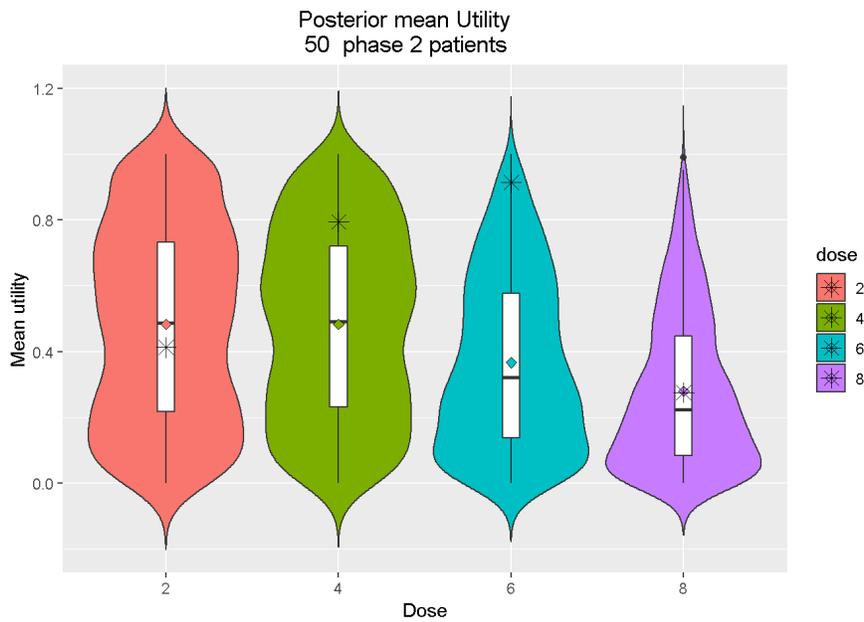


Figure 5.15: Violin plots of posterior means of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

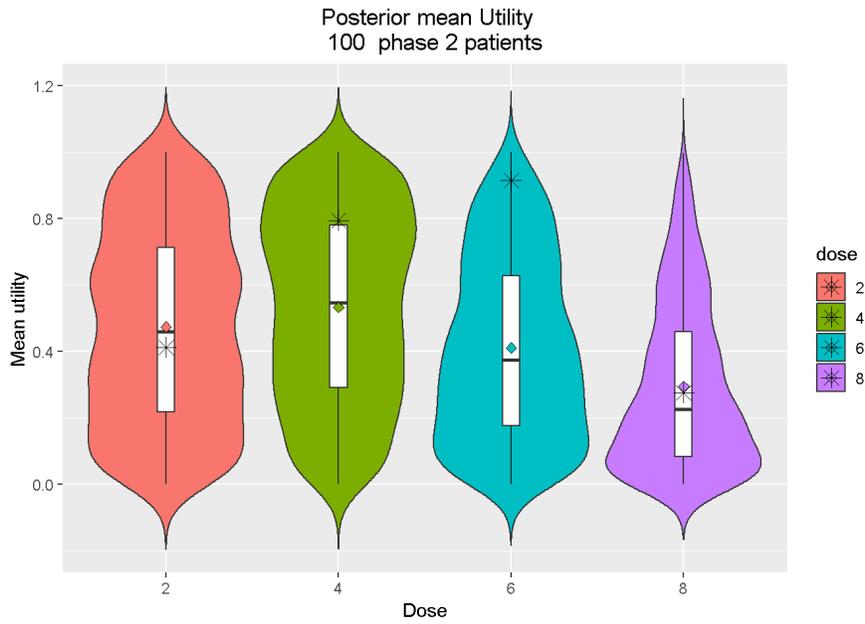


Figure 5.16: Violin plots of posterior means of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

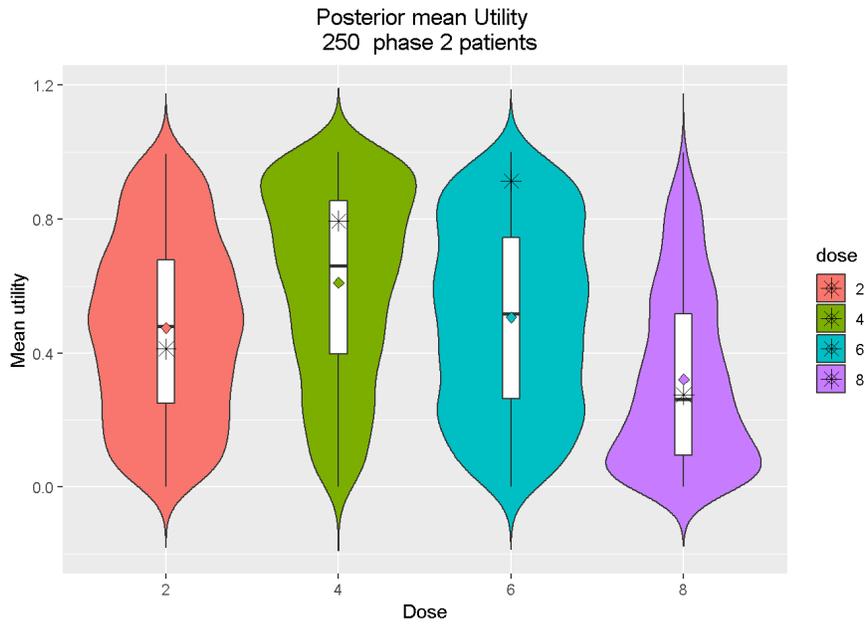


Figure 5.17: Violin plots of posterior means of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

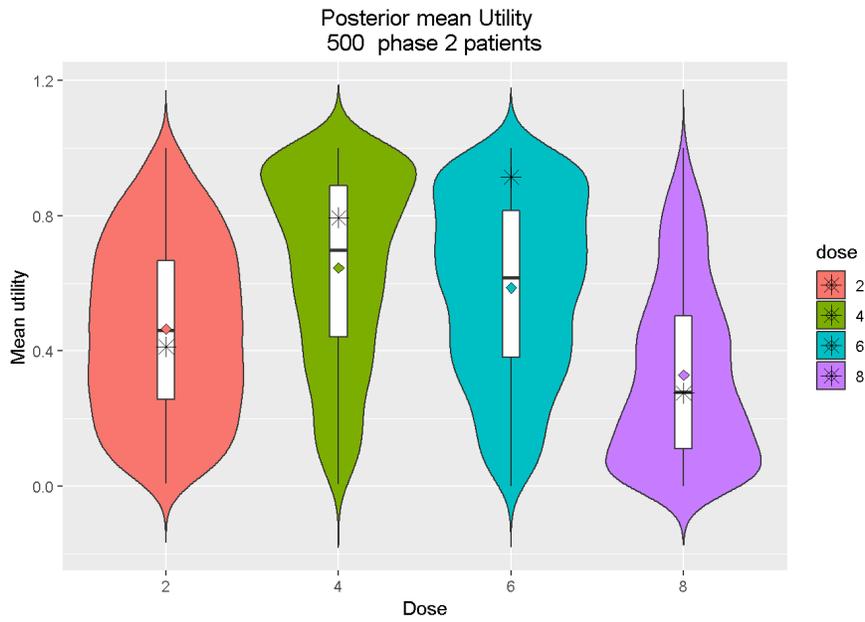


Figure 5.18: Violin plots of posterior means of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

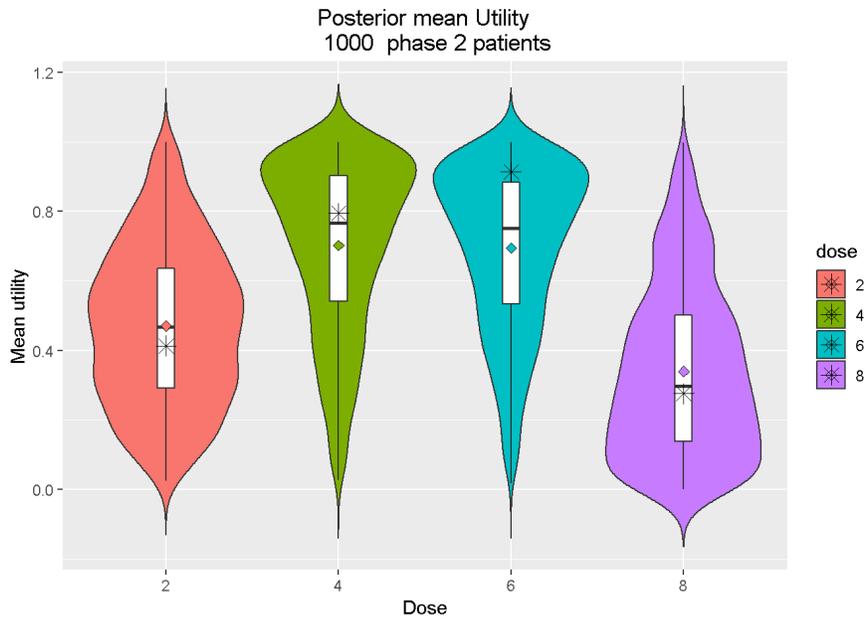


Figure 5.19: Violin plots of posterior means of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

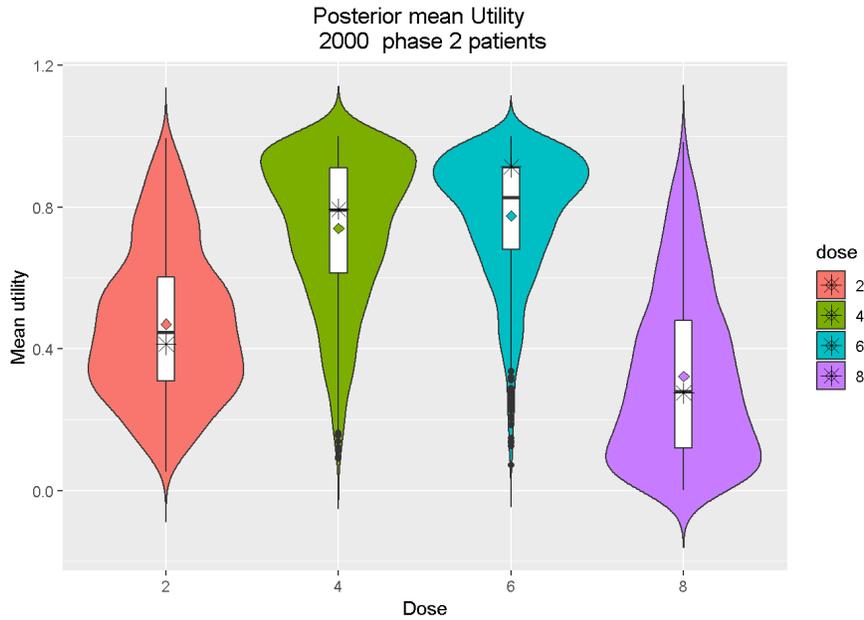


Figure 5.20: Violin plots of posterior means of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

With 1000 simulated phase II studies, even with low values of N_2 , the particular behaviors visualized for one single trial are no longer perceived. Concerning the posterior mean of the utilities, and for the smallest sample size, $N_2 = 50$ patients, the utility values of the two middle doses appear quite underestimated. As expected, the quality of the estimators improves as the sample size grows. But even for a phase II study with 1000 patients, the distribution of the estimates of the middle doses are still quite close, whereas the true value of the utility of $d = 6$ is slightly larger than for $d = 4$. This underestimation of the utility value for $d = 6$ can be due to a less stable estimation of the safety profile, which can be explained by a lower probability of toxicity for the doses < 6 , leading to an inaccurate estimation of the slope. More precisely, the toxicity penalty is dramatically growing just beyond dose $d = 6$. The uncertainty in the toxicity rate estimation in phase II will result in a significant proportion of simulated trials with a quite high toxicity penalty for dose $d = 6$, penalizing that dose. Thus, our approach tends to disadvantage doses with a true toxicity rate close or equal to the threshold of 0.15. This can be viewed as a conservative property to protect the patient. A very large phase II sample size (2000 patients, see Figure 5.20) is needed to distinguish $d = 6$ from $d = 4$ and to detect $d = 6$ as the optimal one. However, still with this very large sample size, $d = 6$ is identified as the optimal dose with a very slight difference as compared to the true difference between the doses.

5.1.2.2 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15): posterior medians of utilities

In the following, we represent the violin plots of posterior **medians** of utilities, for the 1000 simulated phase II studies and for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

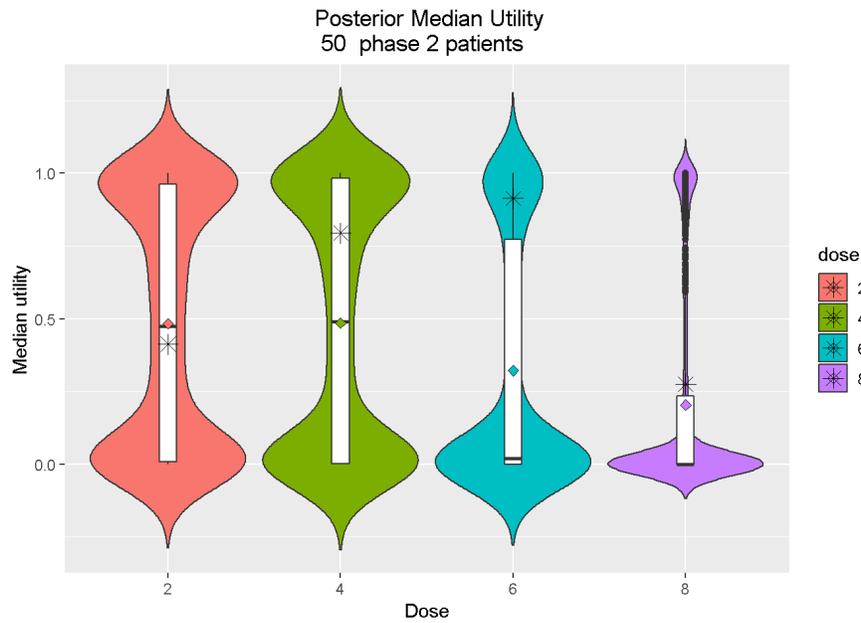


Figure 5.21: Violin plots of posterior medians of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

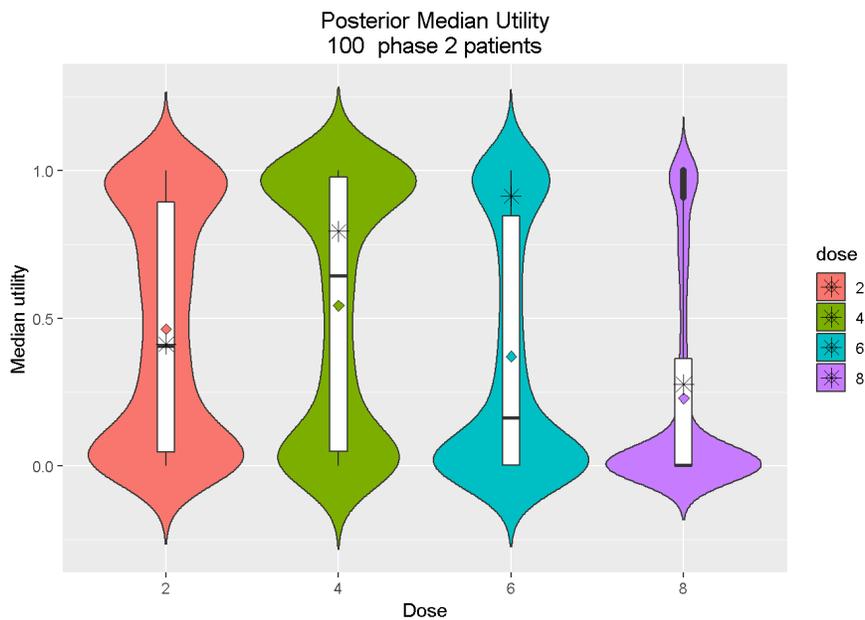


Figure 5.22: Violin plots of posterior medians of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

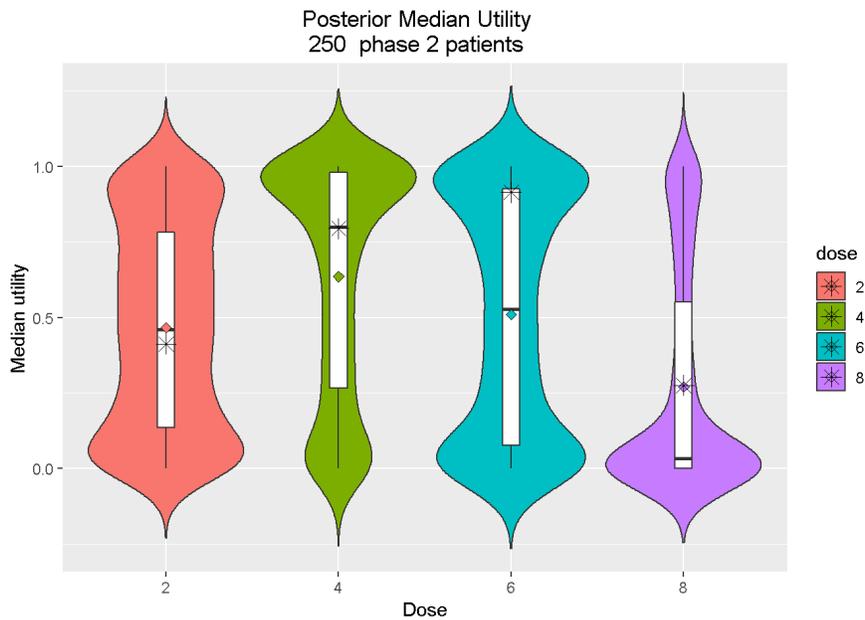


Figure 5.23: Violin plots of posterior medians of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

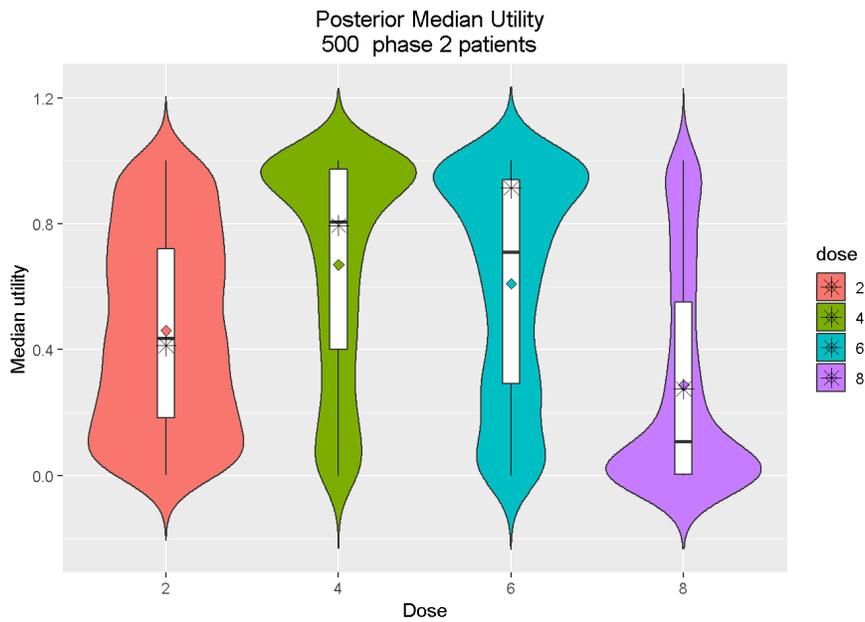


Figure 5.24: Violin plots of posterior medians of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

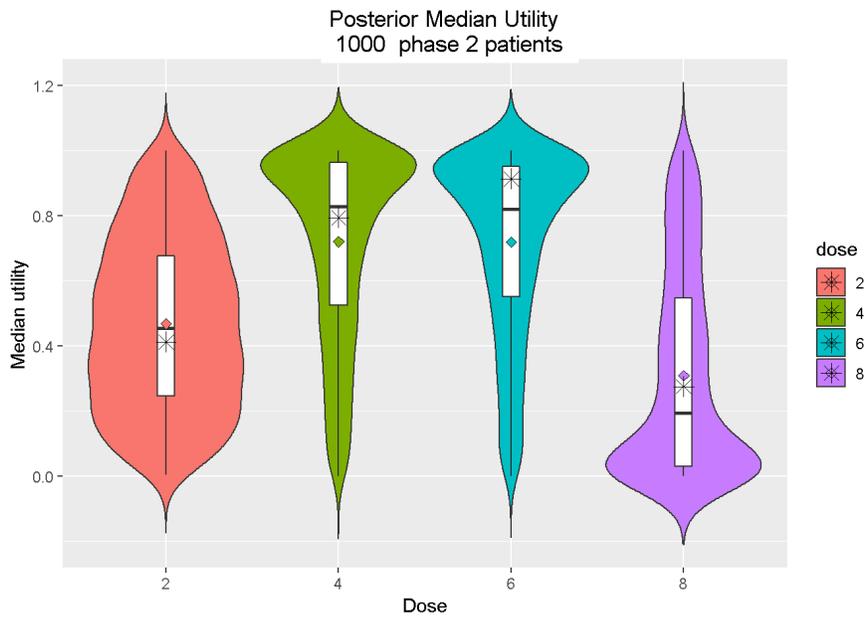


Figure 5.25: Violin plots of posterior medians of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

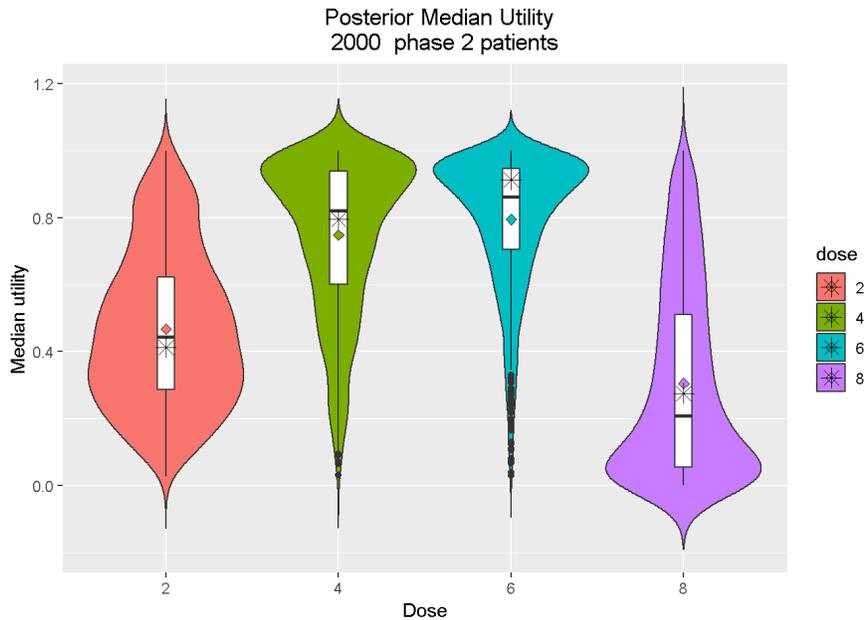


Figure 5.26: Violin plots of posterior medians of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

For the distribution of the posterior median, it appears that it has a larger range, with some bimodality for the low sample sizes (see Figures 5.21 and 5.22), as compared to the posterior mean distribution: the posterior median showing clearly more variability than the posterior mean. But in terms of ability to correctly rank doses, the posterior mean and the posterior median distributions look similar, both estimators have difficulty to capture the utility value of the $d = 6$: the posterior median appears wrongly more accurate due to its increased variability that makes the true value appear within the interquartile interval of the distribution, but in terms of median values, both estimators (posterior mean and posterior median) show a similar bias, even for the phase II study of 1000 patients. According to the violin plots, it takes 2000 patients to properly estimate the dose $d = 6$ with both the posterior mean and median, which is not very realistic according to real life conditions.

5.1.2.3 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20): posterior means of utilities

In the following, we represent the violin plots of posterior **means** of utilities, for the 1000 simulated phase II studies, for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), and for several values of N_2 .

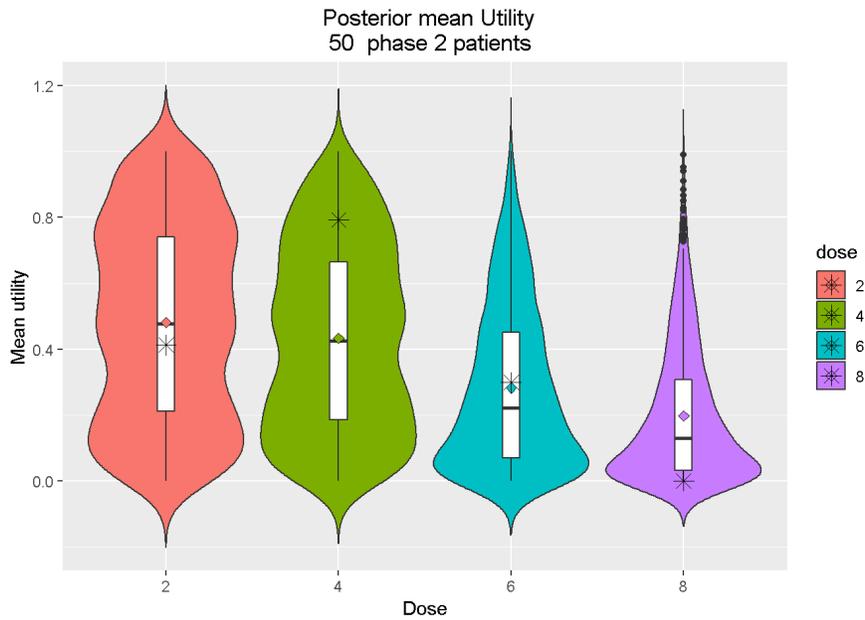


Figure 5.27: Violin plots of posterior means of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

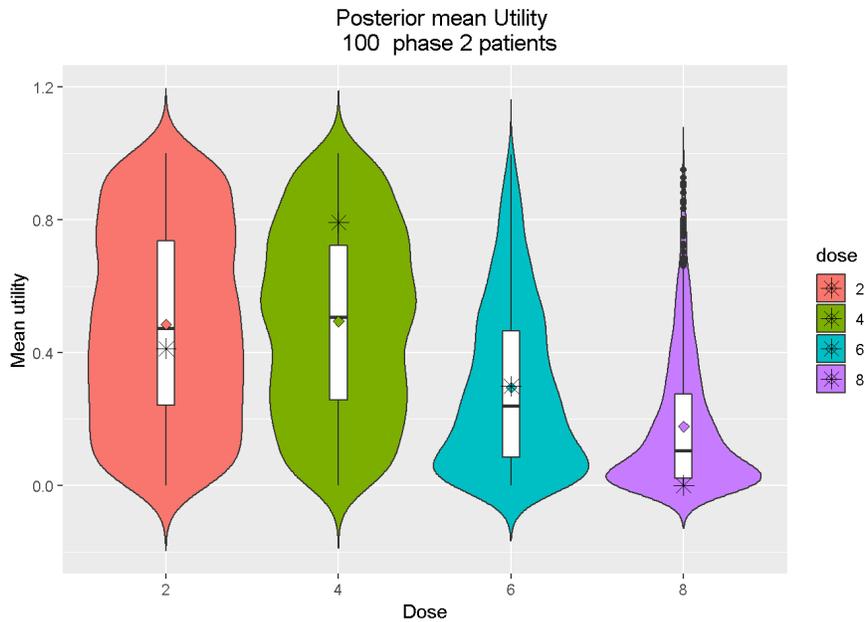


Figure 5.28: Violin plots of posterior means of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

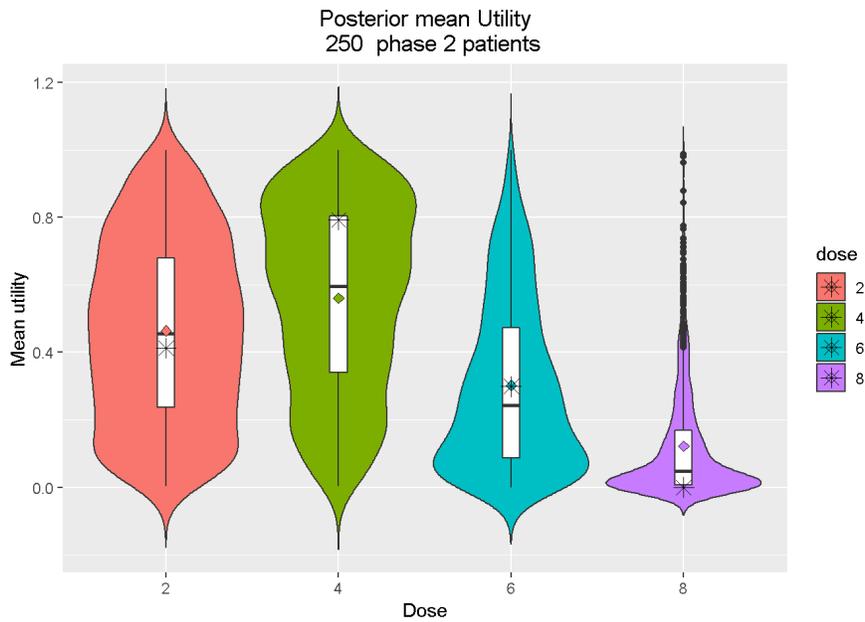


Figure 5.29: Violin plots of posterior means of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

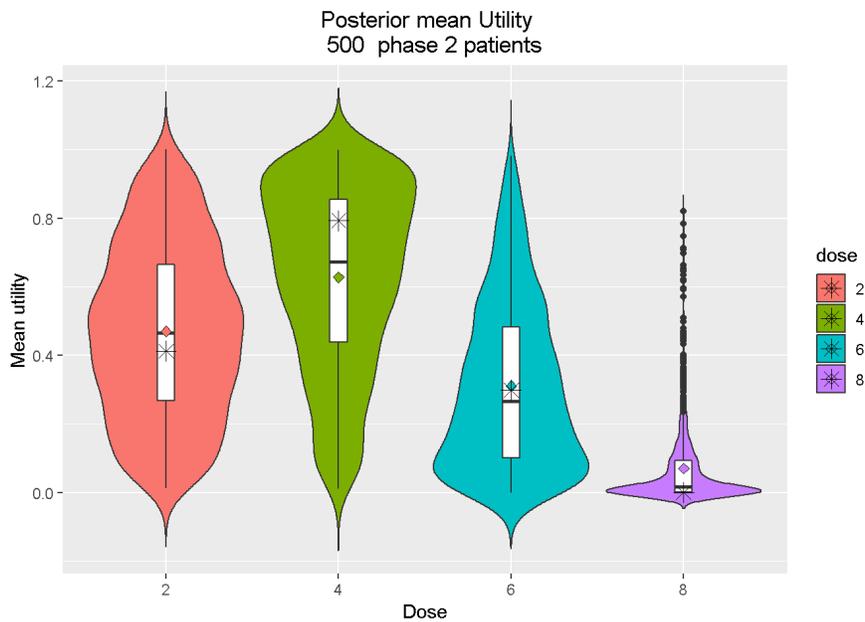


Figure 5.30: Violin plots of posterior means of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

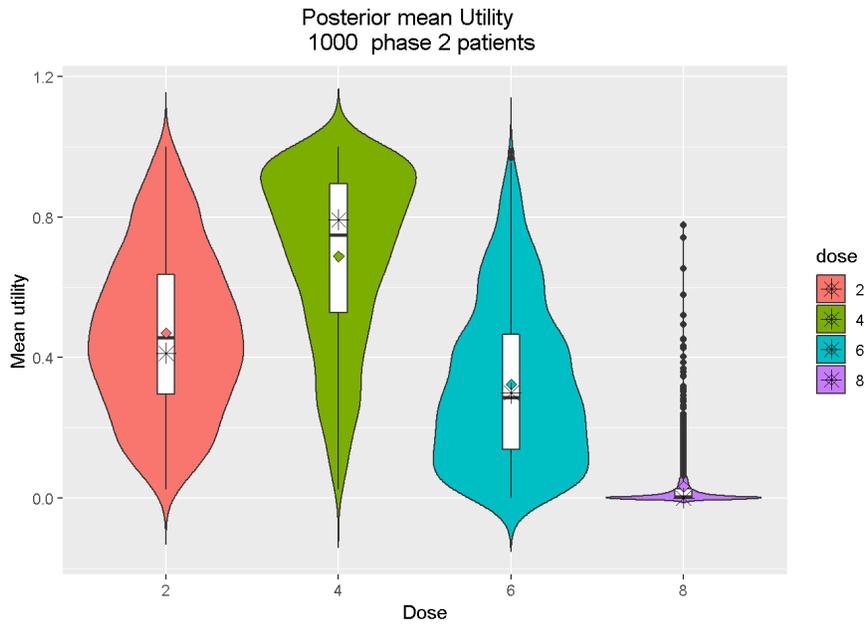


Figure 5.31: Violin plots of posterior means of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

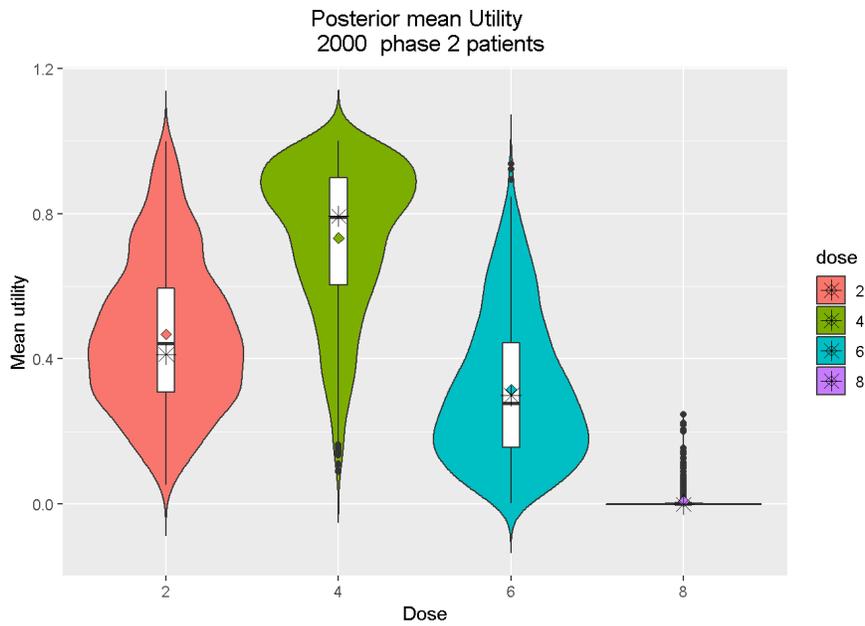


Figure 5.32: Violin plots of posterior means of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

Concerning the posterior mean of the utilities, and for the smallest sample size, $N_2 = 50$ patients, estimated values are not so far from the true values (except maybe for $d = 4$), despite the fact that the posterior distributions appear quite diffuse. As expected, the quality of the estimators improves as the sample size grows. For a phase II study with only 250 patients, the distributions of the estimates of the doses become more concentrated around the true values of the utility, and doses are already well ranked. With more than 250 patients, estimates and dose rankings continue to improve in a remarkable way, reaching their best performances with 2000 patients.

5.1.2.4 Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20): posterior medians of utilities

In the following, we represent the Violin plots of posterior **medians** of utilities, for the 1000 simulated phase II studies and for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

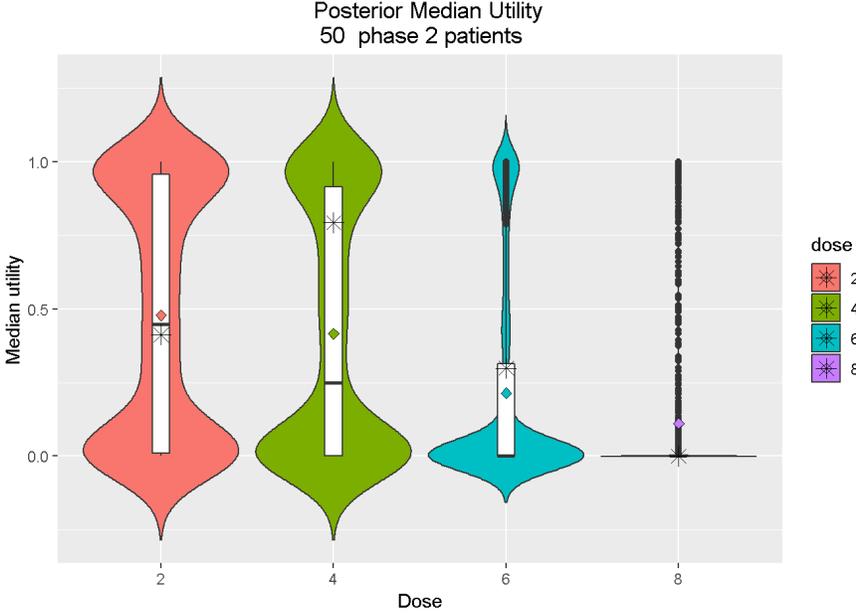


Figure 5.33: Violin plots of posterior medians of utilities, $N_2 = 50$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

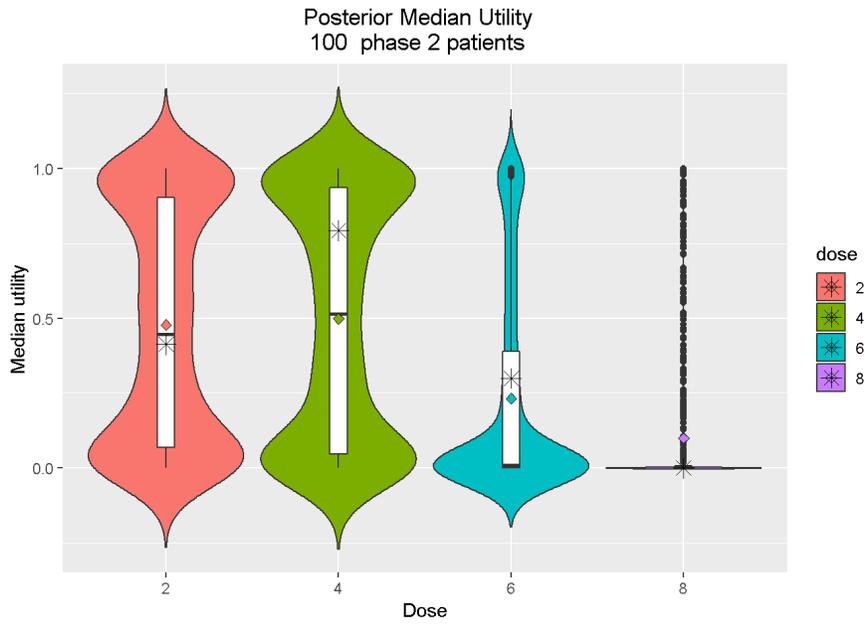


Figure 5.34: Violin plots of posterior medians of utilities, $N_2 = 100$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

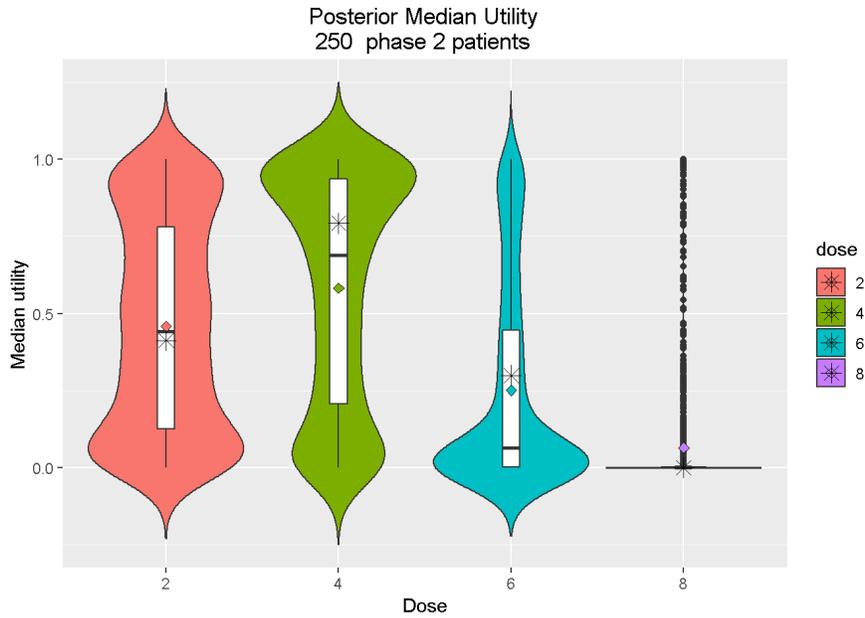


Figure 5.35: Violin plots of posterior medians of utilities, $N_2 = 250$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

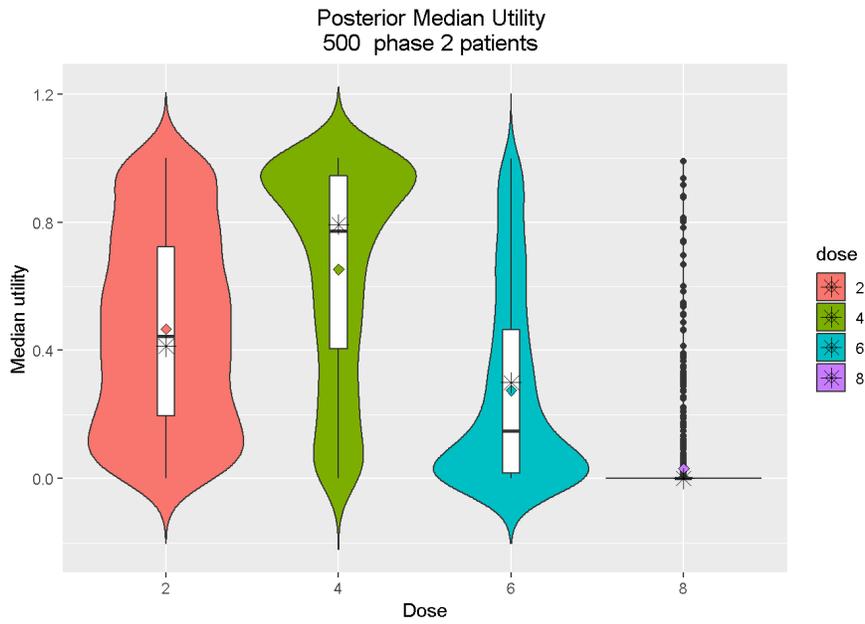


Figure 5.36: Violin plots of posterior medians of utilities, $N_2 = 500$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

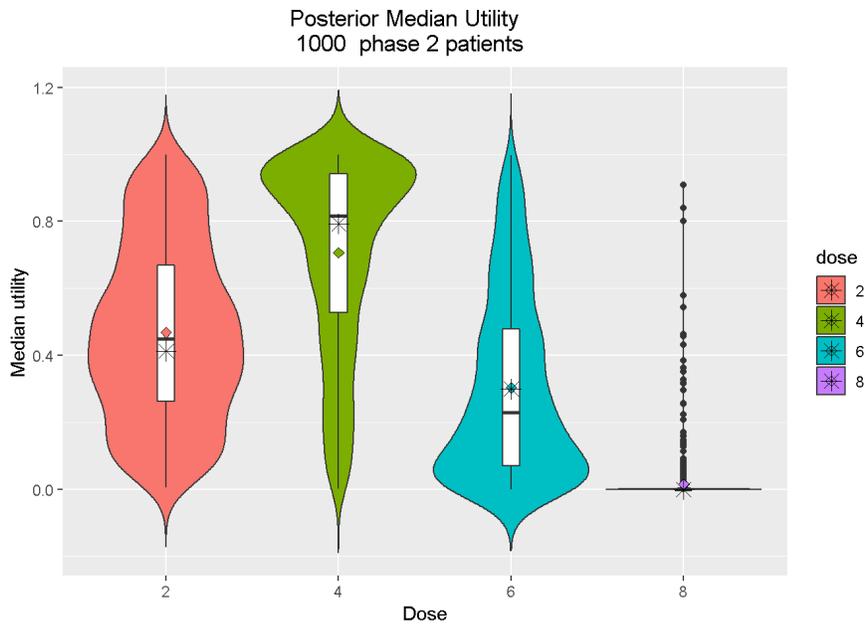


Figure 5.37: Violin plots of posterior medians of utilities, $N_2 = 1000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

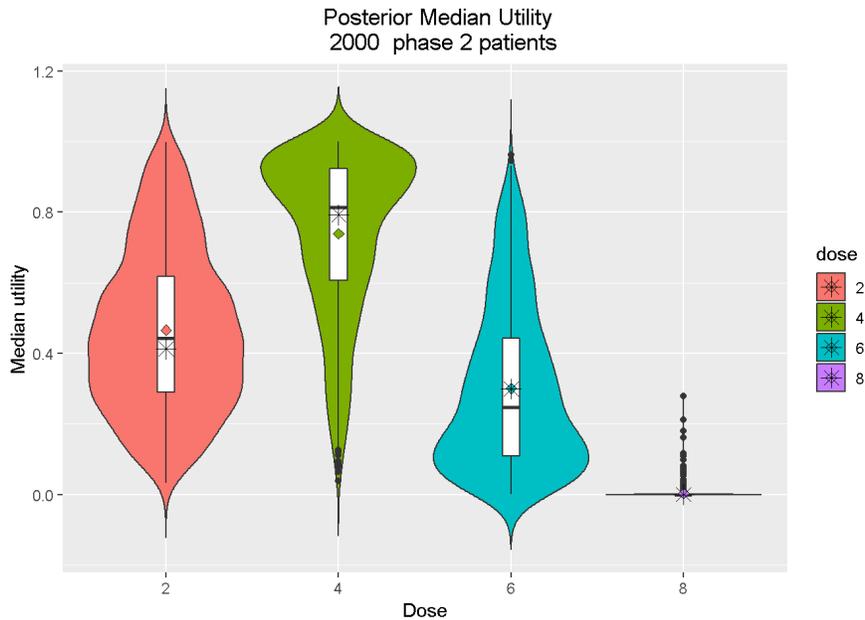


Figure 5.38: Violin plots of posterior medians of utilities, $N_2 = 2000$, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

If we consider now the distribution of the posterior median, and similarly to the previous scenario combination, it appears that it has a larger range as compared to the posterior mean distribution: the posterior median showing clearly more variability than the posterior mean. But in terms of ability to correctly rank doses, the posterior mean and the posterior distribution median look similar.

5.1.2.5 Conclusion

We conclude from these simulations that an accurate estimation of the utility values and a correct ranking of the doses can be very demanding in terms of sample size, especially in the case of a challenging scenario combination where two doses have similar utilities. Estimating the utility values is not the primary goal of the sponsor's analysis (the primary objective is to rank the doses with respect to the utility function), but even though both posterior means and medians do not show any apparent bias for large phase II sample size, the posterior mean appears to be a better estimate, the posterior median showing an inflated variability. So globally, distribution of posterior mean and median are quite different, even if they are both centered around the true value (for the largest sample size), the posterior median seems to exhibit a larger range (with some bimodality for the lowest sample size).

5.2 Sponsor’s strategy: Optimal dose and decision rules

The aim of this section is to compare various decision rules related to the choice of dose by the sponsor. In addition to the decision rule that we have used up to now (Decision rule 1: the sponsor chooses the dose that has the highest probability to be the optimal one, see **Chapter 4**), we will consider four other decision rules that are described thereafter. We remind that 1000 phase II studies are simulated, each study contains 500 batches after burn-in and batching process (applied on 75000 MCMC iterations, see Appendix A.2 for further details).

For each study, the sponsor makes two decisions:

- Identification of the recommended dose:
the sponsor chooses the optimal dose d^* according to one of the decision rules discussed in the following Section 5.2.1: Decision rule 1, Decision rule 1*, Decision rule 2, Decision rule 3 and Decision rule 4.
- ‘Go / NoGo’ decision:
the same rule as in Section 4.1.2.2 is applied: the sponsor computes the average of the estimated PoSs and the average of the estimated toxicity probabilities for the recommended dose d^* among all MCMC iterations denoted by $mean_{MCMC}(\widehat{PoS}(d^*))$ and $mean_{MCMC}(\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq 0.15))$ respectively. The ‘Go’ for phase III is then decided if these averages pass prefixed efficacy and toxicity thresholds denoted by `threshold.eff2` and `threshold.safe2` respectively (these thresholds are simply denoted by `threshold.eff` and `threshold.safe` in **Chapter 4**, since no numbering was needed in the absence of thresholds at the MCMC level, see following discussion). In other words, the sponsor chooses ‘Go’ if:

$$\begin{aligned} mean_{MCMC}(\widehat{PoS}(d^*)) &> \text{threshold.eff2} \\ &\text{and} \\ mean_{MCMC}(\widehat{\mathbb{P}}(tox_{obs}(d^*) \leq 0.15)) &> \text{threshold.safe2.} \end{aligned}$$

5.2.1 Comparisons of alternative decision rules

The aim of this subsection is to compare various decision rules, through simulations, and visualize their performances through the relative utility loss (defined below) graph (see Figure 5.39 and Table 5.27). To do so, we will work on the decision rule of the sponsor (choice of dose), by comparing simulation results with different possible alternatives:

- Decision rule 1: dose that has the greatest probability of being the best, defined in **Chapter 4** page 60
- Decision rule 1*: dose that has the greatest probability of being the best, with additional constraints at the MCMC level (modified version of Decision rule 1, see following discussion), is selected
- Decision rule 2: dose that maximizes $E_{MCMC}(U)$ is selected

- Decision rule 3: dose that maximizes $U(E_{MCMC}(parameter))$ is selected (the batching here is applied directly to the model parameters, see Section 4.1.2.3)
- Decision rule 4: dose that maximizes $U(Median_{MCMC}(parameter))$ is selected (the batching here is also applied directly to the model parameters, see Section 4.1.2.3)

Decision rule 1* is a slightly different version of Decision rule 1; it is defined as follows: the main idea is the same, selecting the dose that is the most likely (according to posterior distribution) to be the optimal dose, but the implementation is slightly different. In order to avoid that the selected dose although having likely the highest utility has at the same time either a too low PoS or a too high probability to have an observed toxicity rate > 0.15 , we modified the dose selection algorithm. We selected the dose with highest probability to be the optimal one with respect to a modified utility denoted by $U'(d)$. This latter utility $U'(d)$ has the same values than $U(d)$, but is set to 0 when the estimated PoS is either lower to a given threshold (threshold.eff1) or when the estimated probability of having an observed toxicity ≤ 0.15 is lower than another threshold (threshold.safe1), these thresholds are applied at the MCMC level. In other words, this modified utility function can be defined, at each MCMC iteration, as follows:

$$U'(d) = \begin{cases} U(d) & \text{if } \widehat{PoS}(d) > \text{threshold.eff1 and } \widehat{\mathbb{P}}(tox_{obs}(d) \leq 0.15) > \text{threshold.safe1} \\ 0 & \text{if not.} \end{cases}$$

By applying efficacy and safety rules (at both MCMC and study levels), the sponsor is more restrictive regarding the dose choices.

We have chosen the relative utility loss as a metric to rank these five decision rules proposed above; it is defined as the difference between the expectation of the utility induced by the considered decision rule and the maximum true utility value within the four doses, named U_{max} , divided by this same maximal value U_{max} . It enables to characterize the quality of the decision rule, through its relative proximity to the ideal best decision rule (always select the optimal dose). It can be defined as follows: $\frac{U_{max} - E(U)}{U_{max}}$, where $E(U)$ is the empirical utility expectation of the chosen dose d^* for the 1000 simulated phase II trials.

5.2.2 Simulation results

In the following, the graph and the simulation results are given for Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), for each of the decision rules defined above, for 1000 simulated phase II studies, by considering threshold.eff1=0.30, threshold.safe1=0.30, threshold.eff2=0.30 and threshold.safe2=0.50 (these thresholds values are also used in the next sections of this chapter).

The relative utility loss functions corresponding to the five decision rules we considered are presented in Figure 5.39.

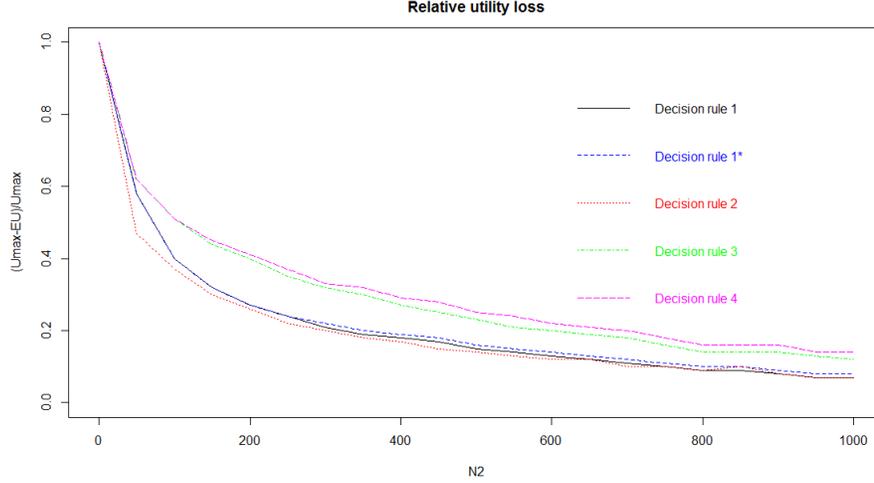


Figure 5.39: Relative utility loss graph comparing the five decision rules.

N_2	Decision rule 1	Decision rule 1*	Decision rule 2	Decision rule 3	Decision rule 4
50	0.58	0.58	0.47	0.62	0.62
100	0.40	0.40	0.37	0.51	0.51
150	0.32	0.32	0.30	0.44	0.45
200	0.27	0.27	0.26	0.40	0.41
250	0.24	0.24	0.22	0.35	0.37
300	0.21	0.22	0.20	0.32	0.33
350	0.19	0.20	0.18	0.30	0.32
400	0.18	0.19	0.17	0.27	0.29
450	0.17	0.18	0.15	0.25	0.28
500	0.15	0.16	0.14	0.23	0.25
550	0.14	0.15	0.13	0.21	0.24
600	0.13	0.14	0.12	0.20	0.22
650	0.12	0.13	0.12	0.19	0.21
700	0.11	0.12	0.10	0.18	0.20
750	0.10	0.11	0.10	0.16	0.18
800	0.09	0.10	0.09	0.14	0.16
850	0.09	0.10	0.10	0.14	0.16
900	0.08	0.09	0.08	0.14	0.16
950	0.07	0.08	0.07	0.13	0.14
1000	0.07	0.08	0.07	0.12	0.14

Table 5.27: Values of the relative utility loss depending on N_2 .

In Figure 5.39, and its related table (Table 5.27), we can see that Decision rule 1, Decision rule 1* and Decision rule 2 are consistently better than Decision rule 3 and Decision rule 4 for all values of the sample size, but the difference between Decision rule 1, Decision rule 1* and Decision rule 2 versus Decision rule 3 and Decision rule 4 is particularly important for the sample size between 200 and 400 patients.

The Decision rule 2 is consistently better than Decision rule 1* and Decision rule 1, but

the difference is small. However, Decision rule 2 (as well as Decision rule 3 and 4) does not take into account uncertainty. Here, the uncertainty is the same for all doses due to the balanced treatment groups, but in case of an unbalanced design or patients drop-outs, this property will no longer be valid for Decision rule 2. Decision rule 1 or Decision rule 1* will then be more robust and more effective. Based on the graph and the table, Decision rule 1 is slightly better than Decision rule 1*.

In order to confirm these results, we made, in the following, a more advanced comparison of these five decision rules, based on the 1000 simulated phase II studies, and we discussed, in Section 5.5, the decision rule that we chose for our final/main results of this thesis.

Simulation results

We remind that in the following tables, 'E(U)' is defined as in Section 4.2 of **Chapter 4**: it is the empirical utility expectation of the chosen dose for the 1000 simulated phase II studies among 'Go' and 'NoGo' decisions (utility is set to 0 when it is a 'NoGo' decision).

For $N_2 = 250$	Decision rule 1 ³	Decision rule 1*	Decision rule 2	Decision rule 3	Decision rule 4
E(U)	0.61	0.61	0.62	0.51	0.50
Prob(choose 'Go')	0.84	0.85	0.85	0.78	0.77
Distribution of selected doses (conditional to 'Go')	0.09 0.84 0.07 0.00	0.07 0.84 0.06 0.04	0.08 0.85 0.06 0.00	0.05 0.73 0.18 0.04	0.05 0.72 0.19 0.05
POS (conditional to 'Go')	0.77	0.78	0.78	0.81	0.81
Power	0.65	0.66	0.66	0.63	0.62

Table 5.28: Simulation results, $N_2 = 250$, comparison of five decision rules for dose selection.

For $N_2 = 500$	Decision rule 1	Decision rule 1*	Decision rule 2	Decision rule 3	Decision rule 4
E(U)	0.68	0.67	0.68	0.61	0.60
Prob(choose 'Go')	0.90	0.90	0.90	0.86	0.85
Distribution of selected doses (conditional to 'Go')	0.02 0.92 0.06 0.00	0.01 0.89 0.09 0.00	0.01 0.93 0.06 0.00	0.01 0.84 0.15 0.00	0.01 0.81 0.17 0.01
POS (conditional to 'Go')	0.80	0.80	0.80	0.81	0.82
Power	0.71	0.72	0.72	0.70	0.70

Table 5.29: Simulation results, $N_2 = 500$, comparison of five decision rules for dose selection.

³full results corresponding to Decision rule 1, for threshold.eff2=0.30, are given in Appendix A.3

For $N_2 = 1000$	Decision rule 1	Decision rule 1*	Decision rule 2	Decision rule 3	Decision rule 4
E(U)	0.74	0.73	0.74	0.70	0.68
Prob(choose 'Go')	0.95	0.96	0.95	0.93	0.92
Distribution of selected doses (conditional to 'Go')	0.00 0.96 0.04 0.00	0.00 0.94 0.06 0.00	0.00 0.96 0.04 0.00	0.00 0.91 0.09 0.00	0.00 0.90 0.10 0.00
POS (conditional to 'Go')	0.80	0.80	0.80	0.81	0.81
Power	0.76	0.77	0.76	0.75	0.75

Table 5.30: Simulation results, $N_2 = 1000$, comparison of five decision rules for dose selection.

Based on these tables, we can see that with Decision rule 2, the dose choice is slightly better compared to Decision rule 1 and Decision rule 1*: one chooses $d = 4$ a bit more often (which is the true optimal dose according to theoretical curves), but the difference is small. However, results are quite similar in terms of expected utilities, probabilities of going to phase III, PoSs and global powers. Decision rule 1 is slightly better than Decision rule 1* in terms of expected utilities and dose choice. Decision rule 3 is clearly worse than Decision rule 1, Decision rule 1* and Decision rule 2. A possible explanation could be that the extreme values of the parameter estimates have an impact on the mean values used and accentuate an estimation bias. However, when considering the estimates median rather than estimates mean (Decision rule 4), we can see that results are very close to those obtained with the mean, even slightly worse.

Globally, Decision rule 1, Decision rule 1* and Decision rule 2 lead to almost similar results and are consistently better than Decision rule 3 and Decision rule 4.

This result is also visually established by Figure 5.39.

We retained Decision rule 1 for our final/main results in **Chapter 4** (without efficacy/safety rules at the MCMC level) and for the next sections of this chapter: the difference is small between Decision rule 1, Decision rule 1* and Decision rule 2, but, the Decision rule 1 (or Decision rule 1*) is easily understandable/interpretable by a clinical team, it better accounts for uncertainty in parameter values, and it fits well to a suitable rule for interim analysis. The reason why Decision rule 1 was retained in preference to Decision rule 1* is discussed in more detail in Section 5.5.

5.3 Influence of phase II sample size - sample size recommendations

In this section, I propose a criterion allowing to choose the sample size of phase II: I will discuss the necessary sample size according to a utility criteria (and not according to power criteria as usually done). For instance, the sample size would be defined as follows: "if the profiles of efficacy and safety are of such type then X patients are required in phase II to have 90% of the maximal utility. If the efficacy and safety profiles are of another type then it takes Y patients in phase II to have 90% of the maximal utility".

This necessary sample size obviously depends on the profile that is unknown, but this is common practice; in order to have more robustness for the power calculations, several alternative scenarios are assessed or, more recently (see [Chuang-Stein2006]), prior distributions are considered for Δ , the true mean difference versus placebo (defined in Section 3.1), and σ , in order to have some Bayesian-averaged sample size calculation. A very similar approach can be used with the new methodology we propose: in the simulation-based sample size calculation, instead of simulating studies with the same hypothesized fixed value of the efficacy and safety parameters, those structural parameters could be sampled from prior distribution representative of the sponsor's expectation related to the new drug.

Our recommendations will be of the same type, even richer (assumptions on both efficacy and safety) and the sample size will not guarantee a power of 90%, if the assumptions are true, but 90% of the maximum utility: this could be a new approach consistent with the true objective of the phase II study (recommending a safe and active dose for phase III study).

In the following, we apply the above criterion allowing to judge whether the phase II sample size is sufficient or not. This criterion is based on the importance of the relative loss of utility defined in Section 5.2: one can say that the size N_2 is sufficient if the global empirical expectation $E(U)$ of the chosen dose (over all 1000 simulations) reaches 90% (or maybe less, 80% for example) of the maximum utility. This gives an idea of necessary sample size of phase II. Then we can vary the efficacy and safety scenarios and decide which is the necessary sample size of phase II according to the efficacy and safety profiles. For each profile, the determination of this sample size will be based on the relative loss of utility graphs, which will be plotted for each profile.

We also compare the phase II recommendations based on this relative utility loss criterion to their corresponding phase II powers using a Student test, denoted by $Power_{t.test}(d)$. These powers can be considered as reference values allowing to judge the interest of this new approach, i.e. of choosing between 80% or 90% of the maximum utility. The phase II power is calculated for each dose separately versus placebo, for a unilateral test at 5% level, which corresponds to the standards of phase II studies. No adjustment for multiplicity is performed. These powers increase as the dose increases, because they depend on $\Delta(d)$, which necessarily increases with the dose.

Below are the results of the sample size determination based on the utility criterion (relative utility loss) defined in Section 5.2, with Decision rule 1. In Figures 5.40, 5.41, 5.42 and 5.43, the values of 0.2 and 0.1 (pointed out by red lines) correspond to a relative utility loss of 20% and 10%, respectively. So we will assess in the following, the number of patients required to reach 80% and 90% respectively of the maximal utility.

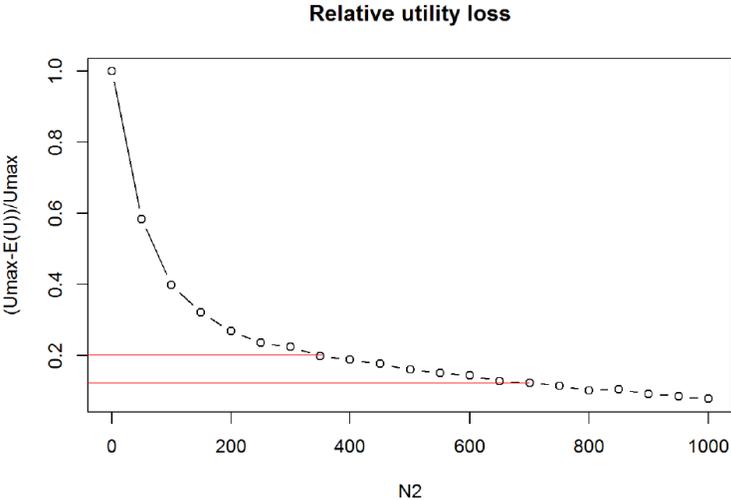


Figure 5.40: Relative utility loss, Sigmoid scenario - Bad safety (progressive toxicity scenario (and toxicity of highest dose = 0.20)).

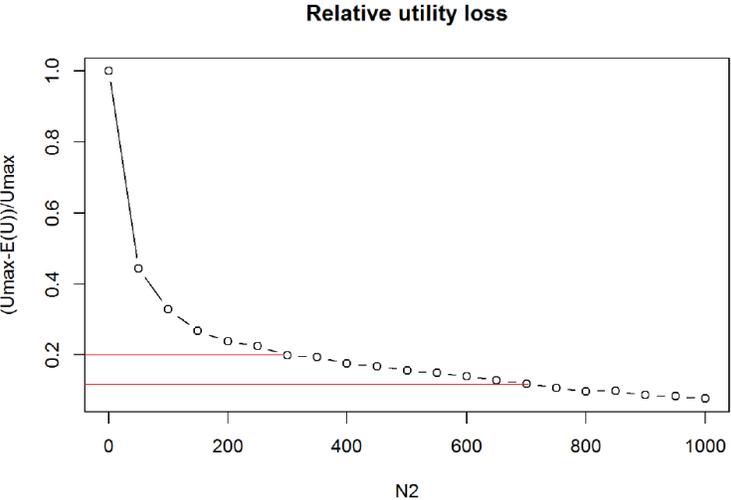


Figure 5.41: Relative utility loss, Plateau scenario - Bad safety (progressive toxicity scenario (and toxicity of highest dose = 0.20)).

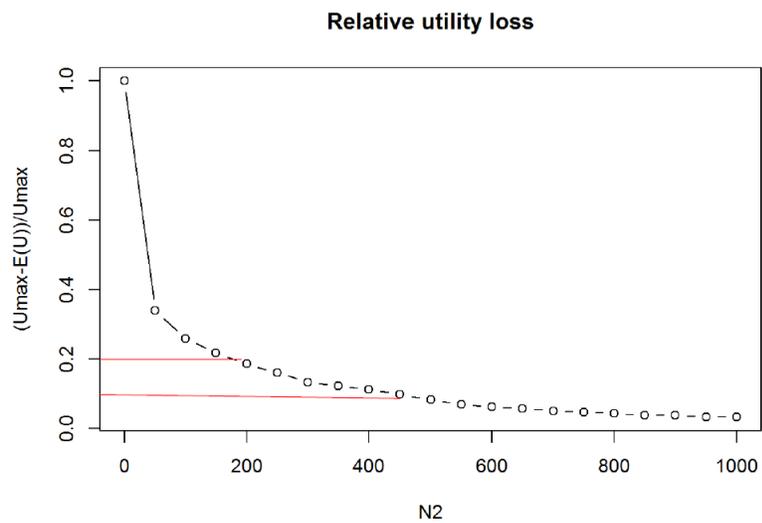


Figure 5.42: Relative utility loss, Sigmoid scenario - Good safety profile (low toxicity scenario).

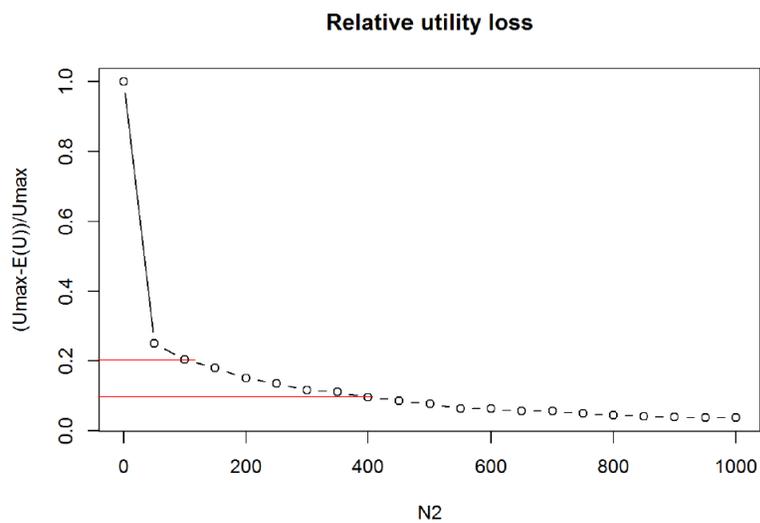


Figure 5.43: Relative utility loss, Plateau scenario - Good safety profile (low toxicity scenario).

One could of course develop or imagine other scenarios (like scenarios presented in **Chapter 4** for instance).

Results:

Regarding the bad Safety profile (progressive toxicity scenario, and toxicity of highest dose = 0.20), we have the following simulation results.

Sigmoid scenario (Figure 5.40):

- $\simeq 350$ patients in phase II to reach 80% of maximal utility
 $Power_{t.test}(2) = 0.16$, $Power_{t.test}(4) = 0.27$, $Power_{t.test}(6) = 0.37$, and
 $Power_{t.test}(8) = 0.44$
- $\simeq 700$ patients in phase II to reach 90% of maximal utility
 $Power_{t.test}(2) = 0.23$, $Power_{t.test}(4) = 0.43$, $Power_{t.test}(6) = 0.58$, and
 $Power_{t.test}(8) = 0.68$

Plateau scenario (Figure 5.41):

- $\simeq 280$ patients in phase II to reach 80% of maximal utility
 $Power_{t.test}(2) = 0.27$, $Power_{t.test}(4) = 0.33$, $Power_{t.test}(6) = 0.36$, and
 $Power_{t.test}(8) = 0.38$
- $\simeq 700$ patients in phase II reach 90% of maximal utility
 $Power_{t.test}(2) = 0.49$, $Power_{t.test}(4) = 0.61$, $Power_{t.test}(6) = 0.65$, and
 $Power_{t.test}(8) = 0.68$

Regarding the good safety profile (low toxicity scenario), we have the following simulation results.

Sigmoid scenario (Figure 5.42):

- $\simeq 200$ patients in phase II to reach 80% of maximal utility
 $Power_{t.test}(2) = 0.12$, $Power_{t.test}(4) = 0.20$, $Power_{t.test}(6) = 0.25$, and
 $Power_{t.test}(8) = 0.30$
- $\simeq 420$ patients in phase II to reach 90% of maximal utility
 $Power_{t.test}(2) = 0.18$, $Power_{t.test}(4) = 0.31$, $Power_{t.test}(6) = 0.41$, and
 $Power_{t.test}(8) = 0.49$

Plateau scenario (Figure 5.43):

- $\simeq 140$ patients in phase II to reach 80% of maximal utility
 $Power_{t.test}(2) = 0.11$, $Power_{t.test}(4) = 0.16$, $Power_{t.test}(6) = 0.21$, and
 $Power_{t.test}(8) = 0.24$

- $\simeq 400$ patients in phase II reach 90% of maximal utility
 $Power_{t.test}(2) = 0.17$, $Power_{t.test}(4) = 0.30$, $Power_{t.test}(6) = 0.40$, and
 $Power_{t.test}(8) = 0.48$

Conclusions:

- In those examples, the sample sizes proposed appear quite small, as compared to those necessary to reach the standard 80% or 90% of the classic phase II power. But one should keep in mind that selecting a dose for phase III (showing a favorable trade-off between efficacy and safety) is a totally different objective from that of searching for statistical significance of the difference of the mean response of the doses versus placebo. In fact reaching statistical significance is not the main objective of phase II (it is an objective for phase III), it is rather to propose the most appropriate dose for phase III.
- To go from 80% of the maximal utility to 90%, it is quite demanding in terms of sample size: we should almost double the number of patients, regardless of efficacy/safety profiles.
- If the dose choice is more difficult (when the safety of the high dose is not very good, i.e. bad safety profile), it is more demanding in terms of number of patients to make good choices.
- So globally, and as for classic phase II power calculation, each incremental probability to achieve the study goal is more and more expensive (in terms of sample size). But we would recommend, in this case, the smaller sample sizes, as they are sufficient to reach 80% of the maximal utility, whereas reaching 90% would require to double the sample size and probably is not worth the investment. Therefore, we would recommend for instance 350 patients for phase II (rather than 700 patients) for the Sigmoid scenario combined with a bad safety profile. This new approach is we think more consistent with the true objective of the phase II study (recommending a safe and active dose for phase III study).

5.4 Alternative criteria for interim analysis

A sub-issue of this thesis was to perceive if an interim data inspection strategy for phase II, when $N'_2 < N_2$ patients are enrolled, can significantly reduce the mean sample size (consequently, budget and time as well) while maintaining the properties of the design (good decision quality of the dose for the phase III). To do so, a sequential design (with futility and efficacy rules at the interim analysis) is compared to a fixed design in order to check the usefulness of interim analysis.

In this section, we considered alternative stopping rules criteria for the interim analysis in addition to the one proposed in **Chapter 4** page 61, as well as several threshold values to stop at interim.

In the following, we reconsider the same interim analysis criterion proposed in **Chapter 4** page 61, by testing this time different l values, $0 < l < 1$, and different designs (smaller phase II sample sizes compared to the ones considered in **Chapter 4**). We recall that the value of threshold l should be predefined by the sponsor, it should be high enough to guarantee accuracy of the dose choice, but not too high, otherwise frequency of early termination will be decreased and studies will be rarely terminated at interim. We compared simulation results for $l = 0.80$ and $l = 0.90$. The choice of these thresholds is also discussed in **Chapter 6**.

We started by comparing the four following fixed/sequential designs via simulation (1000 simulated studies):

- Design 1: 100 patients in phase II.
- Design 2: 500 patients with interim analysis at 100 patients \rightarrow at interim (with 100 patients), one determines the dose d^* :
 - if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l \rightarrow$ stop the study, and choose the optimal dose d^*
 - otherwise continue to the final analysis with 500 patients.
- Design 3: 500 patients with interim analysis at 250 patients \rightarrow at interim (with 250 patients), one determines the dose d^* :
 - if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l \rightarrow$ stop the study, and choose the optimal dose d^*
 - otherwise continue to the final analysis with 500 patients.
- Design 4: 500 patients in phase II.

Our aim here is to obtain a simulation-based comparison between the four designs: Design 1, Design 2, Design 3 and Design 4. The empirical utility expectation of the chosen dose will be the metric used to compare the designs. We will denote these four empirical expectations by $E(U_{D1})$, $E(U_{D2})$, $E(U_{D3})$ and $E(U_{D4})$ respectively.

Normally, as we know that the utility increases with the sample size, $E(U_{D1}) < E(U_{D2}) < E(U_{D3}) < E(U_{D4})$ is expected. But we hope to illustrate that Design 2 or Design 3 leads to quite similar expectation as Design 4, even if they are slightly less efficient (in terms of expected utility) than the latter one. Moreover, for situations where there is a dose that is clearly different/distinguished from others, we hope to stop often at 100 patients or at 250 patients, and these designs will be more economical than Design 4.

In this section, we consider the following combination of efficacy and safety scenarios: Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

Here also, the choice of the dose is governed by Decision rule 1, with values `threshold.eff2=0.30` and `threshold.safe2=0.50`.

	Design 1	Design 2 l=0.80	Design 2 l=0.90	Design 3 l=0.80	Design 3 l=0.90	Design 4
Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff2=0.30 Threshold.safe2=0.50	N2=100	N2=500 interim at N2'=100; stop if : P(best dose data) ≥0.80	N2=500 interim at N2'=100; stop if : P(best dose data) ≥0.90	N2=500 interim at N2'=250; stop if : P(best dose data) ≥0.80	N2=500 interim at N2'=250; stop if : P(best dose data) ≥0.90	N2=500
E(U)	0.476	0.532	0.607	0.619	0.640	0.666
Prob(choose(Go))	0.789	0.757	0.824	0.831	0.857	0.900
Distribution selected doses (Conditional to 'Go')	0.240 0.640 0.000 0.120	0.090 0.820 0.070 0.030	0.030 0.880 0.080 0.000	0.020 0.900 0.070 0.010	0.010 0.910 0.090 0.000	0.010 0.890 0.090 0.000
POS (conditional to 'Go')	0.725	0.776	0.794	0.800	0.804	0.803
Prob(Stop at interim)	-	0.313	0.129	0.339	0.111	-
Mean(N2)=N2x(1- prob(interim)) +N2'xprob(interim)	-	375	448	415	472	-
Power=Prob(choose(Go)) xPOS(conditional to 'Go')	0.572	0.587	0.654	0.664	0.689	0.723

Table 5.31: Simulation results, interim at $N'_2 = 100$ and $N'_2 = 250$, stop at interim if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l$, $l = 0.80, 0.90$.

Based on Table 5.31, we can see that the utilities are well ordered (the larger the N_2 , the larger the empirical expectation $E(U)$ of the chosen dose). The interim analysis with $l = 0.80$ is effective, with a significant probability of stopping at interim (around 30% for both sequential designs with 100 and 250 patients at interim analysis respectively). However, the interim analysis with $l = 0.90$ appears too restrictive even though the global power is slightly larger than for $l = 0.80$. But for instance, for the interim analysis with $N'_2 = 250$ patients, the increase in power appears very small as compared to the expected reduction in the sample size obtained with $l = 0.80$. This also implies that there is not much more stopping for futility with $l = 0.80$ than with $l = 0.90$, otherwise the global power would be much lower.

On the other hand, we can clearly see that $E(U_{D1}) < E(U_{D2}) < E(U_{D3}) < E(U_{D4})$, for

both values of l . But $E(U_{D3})$ is much closer to $E(U_{D4})$ than $E(U_{D2})$, and the best dose is clearly different/distinguished from others with this design (the probability of choosing $d = 4$ dose, which is the optimal dose according to theory, is higher with Design 3 compared to Design 2), which makes it more economical and more beneficial than Design 4.

So we have concluded that:

- $l = 0.90$ may appear too restrictive even though the global power is slightly larger than for $l = 0.80$. But for instance, for the interim analysis with $N'_2 = 250$ patients, the increase in power appears very small as compared to the expected reduction in the sample size obtained with $l = 0.80$
- $l = 0.80$ is economically more interesting: we stop more often at interim analysis \rightarrow we save patients (the average sample size of the phase II trial is reduced)
- In this example, Design 2 (design with $N'_2 = 100$), underperforms Design 3 (design with $N'_2 = 250$), for both values of l , 100 patients is not enough and yet we stop too often, which consequently leads to bad decisions
- Design 3 with $l = 0.80$ is beneficial because we stop quite often at interim (so we save patients) while maintaining the properties of Design 4 (fixed design with $N_2 = 500$).

Another possible criterion for the interim analysis, inspired by [Christen et al.2004], could be based on the difference of the means or medians of the utilities. The idea is to estimate the posterior median or mean of the utility for each dose, then to compute the mean or median differences between each dose d_j and d_k (for $j \neq k$) and to check if these differences are at least equal to 0.15 or 0.20 for example, in favor of a given dose. In other words, a dose d_j would dominate another dose d_k and could be preferred, if the difference between the two relative utility medians or means is at least X ($X = 0.10, 0.15$ or 0.20 for instance) in favor of d_j . So we could define these criteria as follows: stop at interim if $E_{MCMC}(U(d^*) - U(\text{other doses})) \geq 0.10$ or 0.20 , or, stop at interim if $Median_{MCMC}(U(d^*)) - Median(U(\text{other doses})) \geq 0.10$ or 0.20 .

In other words, to justify a dose choice, one could imagine the following two domination criteria:

- Domination criterion 1: a dose d_j would dominate another dose d_k and could therefore be preferred if the difference between the two means of the utility is at least X (for example $X = 0.10$ or 0.15 or 0.20) in favor of d_j
- Domination criterion 2: a dose d_j would dominate another dose d_k and could therefore be preferred if the difference between the two medians of the utility is at least X (for example $X = 0.10$ or 0.15 or 0.20) in favor of d_j

The decision for the interim analysis, based on the selected dose d^* (with Decision rule 1) could then be defined as follows: if d^* dominates all other doses for one of the two domination criteria mentioned above, then stop and go to phase III with d^* .

In Table 5.32, we used the utility mean differences criterion (Domination criterion 1), with

$X = 0.10, 0.20$, and we compared the same fixed/sequential designs previously described: Design 1, Design 2 (with $X = 0.10$ and $X = 0.20$), Design 3 (with $X = 0.10$ and $X = 0.20$) and Design 4.

	Design 1	Design 2 $X=0.10$	Design 2 $X=0.20$	Design 3 $X=0.10$	Design 3 $X=0.20$	Design 4
Sigmoid Progressive toxicity scenario $Tox(d=8)=0.20$ Threshold.eff2=0.30 Threshold.safe2=0.50	$N2=100$	$N2=500$ interim at $N2'=100$; stop if : $E(U(d^*)-$ $U(\text{other}$ $\text{doses}))$ ≥ 0.10	$N2=500$ interim at $N2'=100$; stop if : $E(U(d^*)-$ $U(\text{other}$ $\text{doses}))$ ≥ 0.20	$N2=500$ interim at $N2'=250$; stop if : $E(U(d^*)-$ $U(\text{other}$ $\text{doses}))$ ≥ 0.10	$N2=500$ interim at $N2'=250$; stop if : $E(U(d^*)-$ $U(\text{other}$ $\text{doses}))$ ≥ 0.20	$N2=500$
$E(U)$	0.476	0.639	0.647	0.680	0.668	0.666
$Prob(\text{choose}(\text{Go}))$	0.789	0.919	0.909	0.918	0.905	0.900
Distribution selected doses (Conditional to 'Go')	0.240 0.640 0.000 0.120	0.140 0.770 0.080 0.000	0.080 0.820 0.090 0.000	0.040 0.890 0.070 0.000	0.020 0.890 0.090 0.000	0.010 0.890 0.090 0.000
POS (conditional to 'Go')	0.725	0.753	0.776	0.790	0.799	0.803
$Prob(\text{Stop at interim})$	-	0.214	0.064	0.499	0.034	-
$Mean(N2)=N2x(1-$ $prob(\text{interim}))$ $+N2'xprob(\text{interim})$	-	414	474	375	492	-
$Power=Prob(\text{choose}(\text{Go}))$ $xPOS(\text{conditional to 'Go'})$	0.572	0.692	0.705	0.725	0.723	0.723

Table 5.32: Simulation results, utility mean differences criterion: Domination criterion 1 with $X = 0.10, 0.20$.

Results with the utility median differences criterion (Domination criterion 2), with $X = 0.10, 0.20$, are given in Table 5.33, for the same fixed/sequential designs Design 1, Design 2 (with $X = 0.10$ and $X = 0.20$), Design 3 (with $X = 0.10$ and $X = 0.20$) and Design 4.

	Design 1	Design 2 X=0.10	Design 2 X=0.20	Design 3 X=0.10	Design 3 X=0.20	Design 4
Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff2=0.30 Threshold.safe2=0.50	N2=100	N2=500 interim at N2'=100 ; stop if : med(U(d*))- med(U(other doses)) ≥0.10	N2=500 interim at N2'=100 ; stop if : med(U(d*))- med(U(other doses)) ≥0.20	N2=500 interim at N2'=250 ; stop if : med(U(d*))- med(U(other doses)) ≥0.10	N2=500 interim at N2'=250 ; stop if : med(U(d*))- med(U(other doses)) ≥0.20	N2=500
E(U)	0.476	0.654	0.660	0.686	0.684	0.666
Prob(choose(Go))	0.789	0.928	0.924	0.923	0.920	0.900
Distribution selected doses (Conditional to 'Go')	0.240 0.640 0.000 0.120	0.130 0.790 0.080 0.000	0.100 0.820 0.080 0.000	0.040 0.890 0.060 0.000	0.030 0.900 0.070 0.000	0.010 0.890 0.090 0.000
POS (conditional to 'Go')	0.725	0.755	0.767	0.790	0.795	0.803
Prob(Stop at interim)	-	0.310	0.205	0.489	0.325	-
Mean(N2)=N2x(1- prob(interim)) +N2'xprob(interim)	-	376	418	378	419	-
Power=Prob(choose(Go)) xPOS(conditional to 'Go')	0.572	0.701	0.709	0.729	0.731	0.723

Table 5.33: Simulation results, utility median differences criterion: Domination criterion 2 with $X = 0.10, 0.20$.

The ' $E_{MCMC}(U(d^*) - U(\text{other doses})) \geq 0.10$ ' criterion is quite effective based on simulation results given in Table 5.32, but considering a higher threshold (0.20) becomes too restrictive, and consequently, not enough stops are recorded at the interim analysis. On the contrary, results with the utility median differences criterion, given in Table 5.33, are satisfactory when considering either the smaller threshold (0.10) or the higher one (0.20). The frequency of stopping at interim analysis decreased with the threshold of 0.20, but the difference is quite small as compared to the one with the first domination criterion (Table 5.32).

However, these domination criteria proposed above are complicated and not too intuitive: the "domination" definition is very arbitrary (difference of the medians of the utilities > 0.20 or 0.30 is very difficult to justify given the abstract nature of the utility). But such definition of domination criteria was motivated/inspired by [Christen et al.2004], where authors describe a Phase II clinical trial for finding optimal dose levels, in a different context: patients allocation to doses. They use the following algorithm in order to design a sequential clinical trial: they propose a dynamic programming rule which consists in doing backward induction, and a well detailed algorithm is described in particular, alternating sequence of expectation and maximization. Utility-based decisions consist of dose selection, within a Bayesian framework, based on posterior probabilities: at every stage of the trial, the next patient is allocated to the selected dose, the trial may be stopped for futility, with no treatment recommendation, and doses may be dropped during the trial, if they are judged to be less effective than others. However, these doses are not totally excluded from the trial and may be reused in randomization, which consists in allocating patients to doses within the "non-dominated" set. Here, "non-dominated" set refers to the set of superior doses, dominating the others. Indeed, dropped doses may dominate other doses later on, when the posterior probabilities change: a dose may

be inferior at a given time point, than superior at another time. Hence an adaptive randomization for dose allocation, carried out from sequential design, based on expected utility, to define the set of "non-dominated" doses.

The proposal based on stopping at the interim analysis if $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l$ has similar performance (in terms of mean utility, global power) than the stopping rules based on dose domination criteria Domination criterion 1 and Domination criterion 2 (considering differences in the utility numerical values), but is much more intuitive and simpler to explain to a clinical team.

5.5 Concluding Remarks

We have proposed a sponsor's decision rule based on the posterior probabilities of the doses to be the optimal one (Decision rule 1 or Decision rule 1*): the chosen dose being the one that maximizes this posterior probability; we think that such a rule better accounts for the uncertainty in the parameter values than criteria based on the ordering of numerical estimates of the utilities (like the posterior mean of the utilities for instance, as in Decision rule 2, or the utility of the posterior mean and median of the parameter estimates, as in Decision rule 3 and Decision rule 4). In addition, it is an intuitive and understandable rule, that can be used as the basis to define a stopping rule for the interim analysis (rule based on a lower bound of probability of the chosen dose to be the optimal one).

Regarding sponsor's strategy to choose the optimal dose with Decision rule 1*, by applying efficacy/safety rules at both MCMC and study levels, the sponsor is more restrictive regarding the dose choices. When comparing results between putting efficacy/safety constraints at both MCMC and study levels, and putting efficacy/safety constraints at the study level only (i.e. at the Go / NoGo decision level) as with Decision rule 1, it seemed more reasonable to keep this second approach (Decision rule 1). In fact, this approach is preferred not only because it showed slightly better results compared to Decision rule 1*, but also because thresholds are already arbitrarily predefined with Decision rule 1*, and it becomes harder to justify the choice of these thresholds values. In addition, with the first approach (Decision rule 1*), distributions are underestimated because at each MCMC iteration level, the utility is set to zero if $\widehat{PoS}(d)$ or $\widehat{\mathbb{P}}(tox_{obs}(d) \leq 0.15)$ do not pass the filters (in other words, in the case of utilities with constraints at the MCMC level for efficacy and toxicity, the cluster in 0 is more marked). So the distributions are multimodal with a peak in 0. This is due to our Decision rule 1*, where we seek to "eliminate", for the selection of the dose, the iterations with too low efficacy or too high toxicity.

On the other hand, we conclude that the alternative approaches for the interim analysis addressed in this chapter do not bring significant improvement compared to our proposed stopping criterion in **Chapter 4** which is more intuitive and simpler to explain to a clinician, than criteria based on the numerical value of utility, which remains an abstract quantity.

Chapter 6

Discussion and Conclusions

6.1 Discussion

In the major part of this work, we have attached a utility value to each dose, using the utility function defined as the product of a measure of the dose efficacy (the PoS = the power of the comparison versus placebo in a reference phase III trial with $N_3/2$ patients per arm), and a measure of the dose toxicity (the squared probability of observing a toxicity rate lower than t in a typical phase III trial of $N_3/2$ patients per arm).

Phase III sample size

The size of phase III that we considered ($N_3 = 1000$) is arbitrary. What is important is that the PoS is an increasing function of the dose just as the safety component is a decreasing function of the dose and that these PoS and probabilities of toxicity are representative of the effect size and toxicity profiles of each dose. Nevertheless, a size of 1000 patients seems to be a good compromise, it is approximately in the Phase III range (although 1000 is higher than the average). The considered sample size must not be too high (10000 for instance), otherwise it would be difficult to discriminate doses for efficacy (all PoS will be close to 1). Similarly phase III should not be too small (e.g. 50) for strictly opposite reasons. Sample size for phase III is usually set to achieve a statistical power between 80% and 95%. For the safety, increasing the sample size is equivalent to being more and more radical regarding dose selection: when N_3 is very large (e.g. 10000), if the true percentage of toxicity is higher than t , the safety part of utility will be almost equal to zero; on the contrary, if the true percentage of toxicity is lower than t , the safety part of the utility will be almost equal to 1. However, this could be compensated by exponents h and k assigned to each component of the utility function (see next paragraph). In practice, if 1000 patients are in the sample size range of phase III trials, we keep 1000 to define the utility (even if, once the dose is chosen, one can make a calculation of sample size and choose another sample size for phase III). If 1000 patients are not at all in the range, it is then better to make adaptations (for example, we would consider between 300 and 500 patients in oncology).

Choice of the utility function

We do not claim that this utility function is necessarily the best one, but in addition to the necessary properties that should have utility functions (increase when efficacy increases while safety is fixed and decrease when toxicity increases while efficacy is fixed), it has some desirable properties: it is a smooth and a concave function (at least around the maximum of utility), this guarantees the existence of an optimal dose. Therefore, in practice, choosing a utility function $U(d)$ of the form $U(d) = (\textit{efficacy term}(d))^h \times (\textit{safety term}(d))^k$, with both efficacy and safety terms ranging from 0 to 1, is a pragmatic option. The choice of exponents h and k , enables to give more or less weight to the efficacy and safety terms: large values of the exponents put more constraint to the corresponding term (for instance, for a large value of h , an optimal dose should show a very high efficacy). For instance for a rare disease indication for which there is a clear unmet medical need, there should be less constraint on safety: therefore low values of k should be chosen. On the contrary, for a very competitive therapeutic area, more constraint should be put on the safety side, therefore large values of k should be chosen. In principle, a good option for the sponsor for choosing the utility function, could be to gather some experts that would rank some typical efficacy/safety profiles, those reference rankings being then used by the sponsor to choose a consistent utility function. A possible way to calibrate these values is to adopt the Delphi method [Verhagen et al.1998], which is a forecasting process framework based on the results of several rounds of questionnaires sent to a panel of experts. Several rounds of questionnaires are sent out, and the anonymous responses are aggregated and shared with the group after each round. Concretely, statisticians would propose to experts, depending on the project and the therapeutic area, various efficacy and toxicity profiles. Experts should rank these profiles, and parameters h and k would then be determined according to these rankings, once all reviews are collected. An alternative question to ask clinicians/experts could be, for instance, of the following type: how much is it necessary to improve the safety so that $[PoS = 0.7/\mathbb{P}(tox_{obs} \leq 0.15) = ?]$ has the same benefit risk ratio as $[PoS = 0.8/\mathbb{P}(tox_{obs} \leq 0.15) = 0.65]$? If we plot Iso-quant or Iso-product¹ utility curves for $k = 2$ and $k = 0.5$ for example (see Figure 6.1), with $k = 2$ we get $[PoS = 0.7/\mathbb{P}(tox_{obs} \leq 0.15) = \mathbf{0.7}]$, i.e. a small safety improvement of 5% compensates a loss of efficacy of 10%, and with $k = 0.5$ we get $[PoS = 0.7/\mathbb{P}(tox_{obs} \leq 0.15) = \mathbf{1}]$, i.e. a big safety improvement of 35% is necessary to compensate a loss of efficacy of 10%.

¹curve that shows the different combinations of two factors yielding the same total product

**Iso-quant utility curves, Safety as a function of Efficacy
For $k=2$ and $k=0.5$**

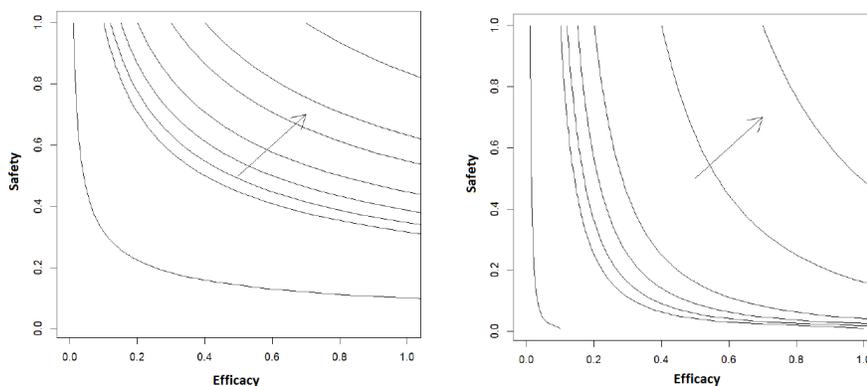


Figure 6.1: Iso-quant utility curves, Safety as a function of Efficacy, $k = 2, 0.5$.

Bayesian approach

To make the necessary decisions, the sponsor must estimate the parameters of the two models, once the data of the phase II study are available. In this work, we have considered the Bayesian framework for the statistical analysis of the phase II data. We advocate for a Bayesian approach as we think it is a more flexible framework for specifying the decision rules. We have proposed a sponsor's decision rule based on the posterior probabilities of the doses to be the optimal one: the chosen dose being the one that maximizes this posterior probability; we think that such a rule better accounts for the uncertainty in the parameter values than criteria based on the ordering of numerical "estimates" of the utilities (like the posterior mean of the utilities for instance). From a technical point of view, in the simulations we have conducted, we have used informative and non-informative priors and computed the posterior probabilities using a MCMC algorithm (Metropolis-Hastings). Informative priors could also be used in order to improve decisions, however, the choice of these priors should be rigorously done, this is discussed in Section 6.2.

Choice of efficacy and toxicity thresholds

Apart from the identification of the best dose, the choice to continue to phase III is a key decision. We have proposed criteria based on threshold values for the PoS, with 60% set as lower bound, and for the probability of observing a toxicity rate lower than t in phase III, with 50% set as lower bound, for the main results of this thesis (see **Chapter 4**). These thresholds have to be determined by the sponsor: for the proposed scenarios, they appeared as a good compromise between the probability of stopping in case of non interesting profile and the probability of going to phase III in case of favourable profile. In practice, to apply the methodology, the sponsor should conduct some simulations to identify the most relevant efficacy and safety thresholds for the targeted, or expected, drug profile.

Impact of sample size

In order to assess the properties of the sponsor's decision-making process mentioned above, we have conducted some simulations (1000 study replicates) under various safety and efficacy profiles and several sample sizes of the phase II study (250, 500 and 1000 patients). The quality of the decision rules were assessed in the light of the frequency, amongst the 1000 study replicates, of the good decisions either for the Go/NoGo decision or the choice of the dose for the phase III. The simulations show that estimating an optimal dose is a difficult and demanding task. For instance, for most of the scenarios with a satisfactory efficacy profile, the probability of making the choice of going to phase III following a phase II study with 250 patients was always less than 60%, except in the scenario in which the drug shows almost no toxicity. This is due to the fact that, with this sample size, the posterior distributions of the utilities, for each of the doses, are not sufficiently concentrated around the true utility values. This leads, often, to imprecise estimations of the posterior probabilities of the dose with the highest utility score (computed for all doses d_j), which are the quantities used for dose selection, and then to wrong selection of the optimal dose. As expected, these probabilities of making the good decision increase with the sample size, but even with the largest sample size, the probability of making the good decision with a large phase II study of 1000 patients only reaches 80% when the drug does not show any toxicity. This is a case where the sponsor should be aware of the low toxicity via simulations and should therefore adjust the efficacy and/or toxicity thresholds. The simulations clearly show that, regardless of time and budget constraints, the sponsor has always interest in running large phase II studies to make accurate decisions regarding the termination of the development program or the selection of the dose. But, in practice, the sample size of the phase II study is necessarily limited by budget and time constraints: those simulations show that for some efficacy and safety profiles, for phase II study of reasonable size (i.e. 250 patients), the probability of making erroneous decision (like wrongly terminate the drug development in phase II) is not negligible (varies between 35% and 44%), especially if inadequate choices of efficacy/toxicity thresholds are made, as it is the case here.

Concerning the dose selection, the probability of selecting the right dose (conditional on sponsor's decision to go to phase III) also increases as the sample size increases. For those efficacy and safety profiles that show a clear peak of utility value for one given dose, accurate dose selection can be achieved with limited sample size. In case several adjacent doses show similar utility values, the identification of the optimum dose is more challenging and requires more patients.

Type I error assessment and risk of wrong decision

An important point is the assessment of the type I error, in order to verify that the clinical trial stops for lack of activity, and not because of excessive toxicity. It appears that this probability can be as high as 16% for the smallest phase II study in a scenario where the toxicity is very low. But again, this probability of false decision decreases as the sample size increases. Regarding the sequential designs, this probability does not exceed 10% when an interim analysis is conducted with $N'_2 = 250$ patients, and only reaches 6% at most, when

an interim analysis is conducted with $N'_2 = 500$ patients, which globally implies a stricter control of the type I error. The efficacy and safety thresholds we have used to specify the utility functions and decision rules can be determined and calibrated by the sponsor in order to maintain the type I error below a desired level. In order to improve this type I error, the sponsor should conduct some simulations to identify the most relevant efficacy and safety thresholds for the targeted, or expected, drug profile, as previously discussed. In fact, a bad choice of these thresholds can lead to an undesired increase in the type I error.

Interim analysis

We have seen that for some safety and efficacy profiles, it is necessary to run a large phase II study to make good decisions, whereas for others, a phase II study of moderate sample size is sufficient to make decisions with acceptable risk of mistakes (including type II error), between 25% and 35%, in other words, with acceptable phase II power, between 65% and 75% (success rate of phase II is usually between 40%-50%). An appealing strategy could be to plan upfront a large sample phase II study and perform an interim analysis, when half of the patients are enrolled, and try to make the selection at this stage. For some scenarios, in particular when the best dose shows a clear benefit in utility as compared to the others, this approach has good properties: with a quite large probability of study termination at interim analysis, it enables to reduce the sample size while maintaining the properties of the fixed large sample size design. For some other scenarios, it is less useful as the study is rarely terminated at the interim analysis, the sponsor being unable to clearly identify the best dose at interim analysis. This could be seen as a safe approach aiming to choose the optimal dose when half of the patients are enrolled, only if these analyzes are reliable and clearly identify this dose as the best one among the others. In all the chosen scenarios, the sponsor decides to stop the trial when at interim analysis, $\mathbb{P}[U(d^*) > U(d_j) \text{ for all the other doses } d_j | \text{data}] \geq l$. This threshold of l has to be chosen by the sponsor: we tested several values and the threshold of $l = 0.80$ seemed to show the best compromise between quality of dose selection (with a high threshold the choice of dose is more accurate) and frequency of early termination (with a too high threshold the studies are rarely terminated at interim analysis which reduces the interest of the method). Also, in our simulations, we concluded that those interim analyses only slightly increased the risk of wrongly taking the decision to go to phase III. For the Sigmoid scenario with a progressive toxicity profile for instance, the probability of taking the wrong decision with an interim analysis at $N'_2 = 250$ only increased by 0.2% compared to the fixed design with $N_2 = 500$ (see Table 4.4).

6.2 Perspectives

Added value of dose-response modeling

The decision rules we have proposed, are based on the estimation of two parametric dose-response models for efficacy and safety, using Emax and probit models respectively. Despite the wide use of these models in regression analyses, they might not necessarily be the "true" one or the best dose-response models in all applications. In addition, the use of parametric dose-response models is not necessary to apply the decision-making framework we have pro-

posed. In this latter case the mean efficacy response for each dose, the probability of adverse events of each dose would be specified as independent parameters. This non-parametric approach is, by definition, less sensitive to model misspecification but it complicates the use of the Bayesian approach we have described and leads to a power loss: in our simulation for instance, instead of the 5 parameters defining the two dose-response models, with non model based efficacy and safety responses, 10 parameters would need to be specified, and then 10 prior distribution would need to be defined. An interesting complementary work would be to compare our proposed dose-response modeling approach with this non-parametric approach, in the case when the model are well specified, but more interestingly, when the Emax and probit models we have considered are not the good ones.

Model Averaging approach

Still in relation to the study of the method robustness with respect to model misspecification, a further development of this work would be to assess the interest of Model Averaging. Model Averaging approaches (see Section 2.2.2) can be easily specified within a Bayesian framework: by specifying priors, $Pr(M_k)$, on the potential parametric models M_k themselves. In this framework, quantities such as $\mathbb{P}(d_j = \textit{optimal dose} | \textit{data})$ (defined in Section 4.1.2.2) can be written as follows:

$\mathbb{P}(d_j = \textit{optimal dose} | \textit{data}) = \sum_k \mathbb{P}(d_j = \textit{optimal dose} | \textit{data}; M_k) \times \mathbb{P}(M_k | \textit{data})$. The terms $\mathbb{P}(d_j = \textit{optimal dose} | \textit{data}; M_k)$ can be computed using standard MCMC algorithms for each of the candidate models. More difficult to calculate, on the computational point of view, are quantities such as $\mathbb{P}(M_k | \textit{data})$. Indeed they can be written as: $\mathbb{P}(M_k | \textit{data}) = \frac{f(\textit{data} | M_k) Pr(M_k)}{\sum_l f(\textit{data} | M_l) Pr(M_l)}$, where $f(\textit{data} | M_k)$ represents the full likelihood, accounting for the prior of the parameters of model M_k , of the data with model M_k ; this expression is difficult to compute in practice as each term $f(\textit{data} | M_l)$ involves integrals over the model parameters space. Fortunately, it is possible to use approximate formula [Davison2003], based on Laplace approximation, linking $f(\textit{data} | M_k)$ to the BIC (Bayes Information Criterion) of model M_k : $f(\textit{data} | M_k) \simeq e^{-\frac{1}{2}BIC(M_k)}$. To conclude, $\mathbb{P}(d_j = \textit{optimal dose} | \textit{data})$ could be approximated by $\sum_k \mathbb{P}(d_j = \textit{optimal dose} | \textit{data}; M_k) \times \frac{e^{-\frac{1}{2}BIC(M_k)} Pr(M_k)}{\sum_l e^{-\frac{1}{2}BIC(M_l)} Pr(M_l)}$,

and the whole method we have proposed could then be further developed in this multi-model framework. Improvements brought by such a Model Averaging approach on the robustness of the method could be assessed through simulations.

Impact of prior distributions

Another topic that could be developed is related to the way the Bayesian analyses are conducted. Risks of wrongly taking the decision to go to phase III are illustrative of the technical difficulty of simultaneously estimating two complex dose-response models with enough accuracy to properly rank doses using a utility function combining the two. In our simulation example, the sponsor's approach is Bayesian using informative and non-informative priors for efficacy (and informative priors for toxicity) as it is usually the case in such context. This choice was driven by the will to have a "conservative" approach leading to choose priors that

minimizes "subjectivity" as compared to the information included in the data. But in practice, as long as those analyses are made for internal decision-making, the sponsor could try to leverage the information available before the phase II was conducted to improve decisions. Maybe, for further development, it would be interesting to assess (through simulations) what level of information brought by the prior would be sufficient to improve the decisions; those considerations could guide the sponsor with respect to the nature of information to collect, in pre-clinical development or phase I studies, to inform those priors and then improve the utility-based decisions and dose selections. This could be done by using more informative priors related to the available information:

- (i) For efficacy, based on previous studies (like a proof of concept phase IIa trial) some information could be available related to the E_{max} parameter for instance: a prior $N(E_{max}, \sigma_{E_{max}}^2)$ not centered on 0, with a not too much inflated variability could be used
- (ii) For the safety, some precise knowledge could be available such as the probability of occurrence of toxicity in the control group information that can be translated in an informative prior on the intercept of the Probit model

Interim analysis

In our work, the interim and final analyses are conducted the same way. But in fact, according to sponsor's objectives related to the interim analysis, they could be conducted completely differently. We could also organize and reflect upon simulations around practical questions such as: in case the interim analysis does not identify "for sure" a better dose, but clearly identifies useless doses (because they are not effective or toxic), is it then possible to abandon one or two doses for the final analysis, without unduly diminishing the final utility?

On the other hand, if the only aim of the interim analysis is to assess if the drug shows some efficacy or not (with no further objective to identify the optimal dose), then a specific decision rule could be built in relation to the efficacy of the largest dose only (for instance the decision rule could be defined as a minimal PoS in phase III for the largest dose; studies would be stopped if efficacy of the largest dose is insufficient). In this example, studies would be stopped only for futility (we only stop for failure, never for success).

Transposition to oncology

An interesting perspective to work on is to transpose our proposed utility-based approach to oncology, for a phase I/phase II clinical development. However, applying a similar approach to oncology would require some significant modification of the methodology. In general, the efficacy criterion used in phase II is different from the efficacy criterion used in phase III. Very often, Best Overall Response is the phase I or phase II criterion whereas the phase III criterion is the Progression Free Survival and/or the Overall Survival. Therefore, unless basing calculations on strong assumptions, it would be difficult to assess the PoS of a dose in phase III only based on a phase I/phase II study. Phase II oncology studies with group parallel designs (including various doses or often various dose regimen) exist, but they are rare: very often the choice of dose is based on phase I dose escalation studies. Accordingly, an interesting application of our

approach would be to guide the dose escalation (choice of the next dose cohort) using a utility-based approach, see [Thall2012a]; [Thall2012b]; [Thall and Cook2004]; [Stallard et al.2004]. A possible approach for a phase I dose escalation study would be to define a utility function having the following form $U(d) = \mathbb{P}(\text{Response rate}(d) \geq \pi_1)^h \mathbb{P}(\text{Toxicity rate}(d) \leq \pi_2)^k$. Then, after each cohort is enrolled, an optimal dose would be chosen, and would be the dose of the next cohort (other complementary safety rules could be taken into account in addition). Such a definition of utility is only applicable if we can define probability distribution for the model parameters: the Bayesian framework is the most suitable for this purpose.

Frequentist approach for design optimality

Another approach to investigate is how the objective of identifying the optimal dose, with respect to a utility function, can also drive the design of the phase II study. Design optimality consideration are more easily defined in a frequentist framework. Even though our analysis framework is Bayesian, we can nevertheless base the definition of the design-optimality criterion on the variance of the maximum likelihood estimates of the utility function, at the doses of interest. This can be done using the delta-method [Davison2003] for instance. Calling $\hat{\Theta}$ the maximum likelihood estimates of the parameter of the two efficacy and safety models:

$$\begin{aligned} \text{var}(U(d; \hat{\Theta}, \text{design})) &\simeq \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right)^T \text{var}(\hat{\Theta}, \text{design}) \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right) \\ &= \text{trace} \left(\text{var}(\hat{\Theta}, \text{design}) \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right) \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right)^T \right) \end{aligned}$$

where $\text{var}(\hat{\Theta}, \text{design})$ can be approximated by the inverse of the Fisher information matrix. Since the utility has to be maximized within a certain range of dose, $[d_m, d_M]$, the approach of [Miller et al.2007] can be used in defining the design optimality criterion by:

$$\Phi(\text{design}) = \text{trace} \left(\text{var}(\hat{\Theta}; \text{design}) \int_{d_m}^{d_M} \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right) \left(\frac{\partial U(x; \hat{\Theta})}{\partial \Theta}\right)^T dx \right).$$

6.3 General conclusion

We have proposed a decision-making framework based on a utility function that, following a phase II study, can drive sponsor's decision with respect to the continuation, or not, of the drug development as well as the selection of the best dose for the phase III. With a utility function considering simultaneously the efficacy and safety drug profiles, we believe that our approach is flexible enough to be used for most of the therapeutic areas and indications and potentially allow to integrate other sources of information than the phase II study data. We think also that our utility-based dose selection participates to the new methodologies in hands of the sponsors to rationalize and improve the dose selection and then improve the efficiency of drug development to provide better therapeutic solutions to the patients.

Bibliography

- [Akaike1975] Akaike, H. (1975). A new look at the statistical model identification. IEEE Transactions on Automatic Control, 19(6):716–723.
- [Alexopoulos and Seila1996] Alexopoulos, C. and Seila, AF. (1996). Implementing the batch means method in simulation experiments. Proceedings of the 28th Conference on Winter Simulation, pages 214–221.
- [Alvarez2012] Alvarez, J. (2012). Introduction to Adaptive Experimental Design. Department of Biostatistics.
- [Ankolekar et al.2015] Ankolekar, S., Antonijevic, Z., Beckman, R., Bobulsky, S., Bolognese, J., Burman, CF., Chen, C., Chuang-Stein, C., David, F., Huml, R., Ireland, B., Jennison, C., Kim, K., Kimber, M., Kloeber, J., Marchenko, O., Ng-Cashin, J., Nixon, R., Patel, N., Persinger, C., Pinheiro, J., Sax, F., Schulz, K., Stojanovic, A., and Sun, L. (2015). Optimization of Pharmaceutical RD Programs and Portfolios: Design and Investment Strategy. Springer.
- [Antonijevic et al.2013] Antonijevic, Z., Manner, D., Pinheiro, J., Kimber, M., Burman, CF., and Bergenheim, K. (2013). Optimizing Drug Development Programs: Type 2 Diabetes Case Study. Therapeutic Innovation Regulatory Science, 47(3):363–374.
- [Antonijevic et al.2010] Antonijevic, Z., Pinheiro, J., Fardipour, P., and Lewis, RJ. (2010). Impact of Dose Selection Strategies Used in Phase II on the Probability of Success in Phase III. Statistics in Biopharmaceutical Research, 2(4):469–486.
- [Azevedo-Filho and Shachter2013] Azevedo-Filho, A. and Shachter, R. (2013). Laplace’s Method Approximations for Probabilistic Inference in Belief Networks with Continuous Variables. Artificial Intelligence.
- [Azriel et al.2011] Azriel, D., Mandel, M., and Rinott, Y. (2011). The treatment versus experimentation dilemma in dose finding studies. Elsevier, 141:2759–2768.
- [Bartoń2018] Bartoń, K. (2018). R Package ‘MuMIn’: Multi-Model Inference. pages 1–74.
- [Bauer et al.2014] Bauer, P., Glimm, E., Koenig, F., and Graf, A. (2014). Maximum type 1 error rate inflation in multiarmed clinical trials with adaptive interim sample size modifications. Biometrical Journal, 56(4).
- [Bauer and Hackl1985] Bauer, P. and Hackl, P. (1985). The application of Hunter’s inequality in simultaneous testing. Biometrical Journal, 27(1):25–38.

- [Bauer and Kieser1999] Bauer, P. and Kieser, M. (1999). Combining different phases in the development of medical treatments within a single trial. Statistics in Medicine, 18(14):1833–1848.
- [Bellman1957] Bellman, E. (1957). Dynamic Programming. Princeton University Press.
- [Berger1985] Berger, r. (1985). Statistical Decision Theory and Bayesian Analysis. Springer Science and Business Media New York.
- [Bornkamp2012] Bornkamp, B. (2012). Functional uniform priors for nonlinear modeling. Biometrics, 68:893–901.
- [Bornkamp et al.2007] Bornkamp, B., Bretz, F., Dmitrienko, A., Enas, G., Gaydos, B., Hsu, CH., König, F., Krams, M., Liu, Q., Neuenschwander, B., Parke, T., Pinheiro, J., Roy, A., Sax, R., and Shen, F. (2007). Innovative approaches for designing and analyzing adaptive dose-finding trials. Journal of Biopharmaceutical Statistics, 17:965–995.
- [Bornkamp et al.2009] Bornkamp, B., Pinheiro, J., and Bretz, F. (2009). MCPMod: An R Package for the Design and Analysis of Dose-Finding Studies. Journal of Statistical Software, 29(7):1–23.
- [Brannath et al.2009] Brannath, W., Zuber, E., Branson, M., Bretz, F., Gallo, P., Posch, M., and Racine-Poon, A. (2009). Confirmatory adaptive designs with Bayesian decision tools for a targeted therapy in oncology. Statistics in Medicine, 28(10):1445–1463.
- [Bretz et al.2005] Bretz, F., Branson, M., and Pinheiro, J. (2005). Combining Multiple Comparisons and Modeling Techniques in Dose-Response Studies. Biometrics, 61(3):738–748.
- [Bretz et al.2010] Bretz, F., Dette, H., and Pinheiro, J. (2010). Practical considerations for optimal designs in clinical dose finding studies. Statistics in Medicine, 29(7-8):731–742.
- [Bretz et al.2008] Bretz, F., Hsu, J., Pinheiro, J., and Liu, Y. (2008). Dose Finding - A Challenge in Statistics. Biometrical Journal, 50(4):480–504.
- [Bretz et al.2009] Bretz, F., Koenig, F., Brannath, W., Glimm, E., and Posch, M. (2009). Adaptive designs for confirmatory clinical trials. Statistics in Medicine, 28(8):1181–1217.
- [Bretz et al.2006] Bretz, F., Schmidli, H., König, F., Racine, A., and Maurer, W. (2006). Confirmatory Seamless Phase II/III Clinical Trials with Hypotheses Selection at Interim: General Concepts. Biometrical Journal, 48(4):623–634.
- [Chakrabarti and Ghosh2011] Chakrabarti, A. and Ghosh, JK. (2011). AIC, BIC and Recent Advances in Model Selection. Philosophy of science, 7:583–605.
- [Chaloner and Verdinelli1995] Chaloner, K. and Verdinelli, I. (1995). Bayesian Experimental Design: A Review. Statistical Science, 10(3).
- [Chen and Beckman2009] Chen, C. and Beckman, R. (2009). Optimal cost-effective designs of Phase II proof of concept trials and associated go-no go decisions. Journal of Biopharmaceutical Statistics, 19:424–436.

- [Chen2017] Chen, Y. (2017). A Tutorial on Kernel Density Estimation and Recent Advances. Biostatistics and Epidemiology, 1(10.1080/24709360.2017.1396742).
- [Christen et al.2004] Christen, J.A., Müller, P., Wathen, K., and Wolf, J. (2004). A bayesian randomized clinical trial: A decision theoretic sequential design. Canadian Journal of Statistics, 32(4):387–402.
- [Chuang-Stein2006] Chuang-Stein, C. (2006). Sample size and the probability of a successful trial. Pharmaceutical Statistics, 5:305–309.
- [Cohen1988] Cohen, J. (1988). Statistical Power Analysis for the Behavioral Sciences. Hillsdale, N.J: L. Erlbaum Associates, (ISBN: 0-8058-0283-5).
- [Comets2010] Comets, E. (2010). Etude de la réponse aux médicaments par la modélisation des relations dose-concentration-effet. HAL, Médicaments. Université Paris-Diderot - Paris VII(tel-00482970):1–84.
- [Davison2003] Davison, A. (2003). Statistical Models (Cambridge Series in Statistical and Probabilistic Mathematics). Cambridge: Cambridge University Press.
- [Dette et al.2013] Dette, H., Bornkamp, B., and Bretz, F. (2013). On the efficiency of two-stage response-adaptive designs. Statistics in Medicine, 32:1646–1660.
- [Dette et al.2015] Dette, H., Titoff, S., Volgushev, S., and Bretz, F. (2015). Dose response signal detection under model uncertainty. Biometrics, 71(4):996–1008.
- [Dmitrienko and Tamhane2007] Dmitrienko, A. and Tamhane, AC. (2007). Gatekeeping procedures with clinical trial applications. Pharmaceutical Statistics, 6(3):171–180.
- [Dragalin et al.2010] Dragalin, V., Bornkamp, B., Bretz, F., Miller, F., Padmanabhan, SK., Patel, N., Perevozskaya, I., Pinheiro, J., and Smith, JR. (2010). A simulation study to compare new adaptive dose-ranging designs. Statistics in Biopharmaceutical Research, 2:487–512.
- [Dunnett1955] Dunnett, CW. (1955). A multiple comparison procedure for comparing several treatments with a control. Journal of the American Statistical Association, 50(272):1096–1121.
- [Dutta et al.1996] Dutta, S., Matsumoto, Y., and Ebling, WF. (1996). Is It Possible To Estimate the Parameters of the Sigmoid Emax Model with Truncated Data Typical of Clinical Studies? Journal of Pharmaceutical Sciences, 85(2):232–239.
- [Fishman and Yarberrry1997] Fishman, GS. and Yarberrry, LS. (1997). An implementation of the batch means method. INFORMS Journal on Computing, 9(3):231–318.
- [Foo and Duffull2017] Foo, LK. and Duffull, S. (2017). Designs to balance cost and success rate for an early phase clinical study. Journal of Biopharmaceutical Statistics, 27(1):148–158.
- [Forster2000] Forster, MR. (2000). Key Concepts in Model Selection: Performance and Generalizability. Journal of Mathematical Psychology, 44(1):205–231.

- [Gajewski et al.2015] Gajewski, BJ., Berry, SM., Quintana, M., Pasnoor, M., Dimachkie, M., Herbelin, L., and Barohn, R. (2015). Building efficient comparative effectiveness trials through adaptive designs, utility functions, and accrual rate optimization: finding the sweet spot. Statistics in medicine, 34(7):1134–1149.
- [Geiger et al.2012] Geiger, MJ., Skrivanek, Z., Gaydos, B., Chien, J., Berry, SM., Berry, D., and Anderson, JH. (2012). An Adaptive, Dose-Finding, Seamless Phase 2/3 Study of a Long-Acting Glucagon-Like Peptide-1 Analog (Dulaglutide): Trial Design and Baseline Characteristics. Journal of Diabetes Science and Technology, 6(6):1319–1327.
- [Geyer et al.2011] Geyer, CJ., Robert, C., Casella, G., Fan, Y., Sisson, SA., Rosenthal, JS., Neal, RM., Gelman, A., Shirley, K., Flegal, JM., Jones, GL., Craiu, RV., Meng, XL., Huber, M., Hobert, JP., Thompson, E., Caffo, B., Bowman, D., Eberly, L., Bassett, SS., Dyk, DV., Park, T., Higdon, D., Reese, CS., Moulton, JD., Vrugt, JA., Fox, C., King, R., Haran, M., Park, JH., Peng, R., Dominici, F., Louis, TA., Zeger, S., Fearnhead, P., Levy, R., Mislevy, RJ., Behrens, JT., Millar, RB., Garip, F., and Western, B. (2011). Handbook of markov chain monte carlo. Chapman Hall/CRC, (ISBN: 9781420079425):3–592.
- [Ghosh et al.2007] Ghosh, JK., Delampady, M., and Samanta, T. (2007). An Introduction to Bayesian Analysis: Theory and Methods. Springer, (ISBN-13: 978-0387-40084-6).
- [Giatting et al.2007] Giatting, G., Gletting, P., Reske, SN., Hohl, K., and Ring, C. (2007). Choosing the optimal fit function: Comparison of the Akaike information criterion and the F-test. Medical Physics, 34(11):4285–4292.
- [Glimm2013] Glimm, E. (2013). Adaptive designs - beyond p-value combination methods. Novartis Pharma.
- [Grieve2017] Grieve, AP. (2017). Response-adaptive clinical trials: case studies in the medical literature. Pharmaceutical Statistics, 16(1):64–86.
- [Grieve et al.2013] Grieve, AP., Chow, SC., Curram, J., Dawe, S., Harnisch, LO., Henig, NR., Hung, HMJ., Ivy, DD., Kawut, SM., Rahbar, MH., Xiao, S., and MR, W. (2013). Advancing clinical trial design in pulmonary hypertension. Pulmonary Circulation, 3(1):217–225.
- [Gutjahr and Bornkamp2017] Gutjahr, G. and Bornkamp, B. (2017). Likelihood ratio tests for a dose-response effect using multiple nonlinear regression models. Biometrics, 73(1):197–205.
- [Hocking1976] Hocking, RR. (1976). The Analysis and Selection of Variables in Linear Regression. Biometrics, 32(1):1–49.
- [Hoeting et al.1999] Hoeting, JA., Madigan, D., Raftery, AE., and Volinsky, CT. (1999). Bayesian Model Averaging: A Tutorial. Statistical Science, 14(4):382–417.
- [Houede et al.2010] Houede, N., Thall, PF., Nguyen, H., Paoletti, X., and Kramar, A. (2010). Utility-Based Optimization of Combination Therapy Using Ordinal Toxicity and Efficacy in phase I/II trials. Biometrics, 66(2):532–540.
- [Hsu1996] Hsu, J. (1996). Multiple Comparisons, Theory and methods. Department of Statistics, The Ohio State University, USA, pages 1–296.

- [Jones et al.2011] Jones, B., Layton, G., Richardson, H., and Thomas, N. (2011). Model-based Bayesian Adaptive Dose-Finding Designs for a Phase II Trial. Statistics in Biopharmaceutical Research, 3(2):276–287.
- [Julious and Swank2005] Julious, S. and Swank, D. (2005). Moving statistics beyond the individual clinical trial: applying decision science to optimize a clinical development plan. Pharmaceutical Statistics, 4:37–46.
- [Kirchner et al.2016] Kirchner, M., Kieser, M., Götte, H., and Schüler, A. (2016). Utility-based optimization of phase II/III programs. Statistics in Medicine, 35(2):305–316.
- [Koenig et al.2007] Koenig, F., Brannath, W., Bretz, F., and Posch, M. (2007). Adaptive Dunnett tests for treatment selection. Statistics in Medicine, 27:1612–1625.
- [Koenig et al.2008] Koenig, F., Brannath, W., Bretz, F., and Posch, M. (2008). Adaptive Dunnett tests for treatment selection. Statistics in Medicine, 27(10):1612–1625.
- [Li et al.2017] Li, DH., Whitmore, JB., Guo, W., and Ji, Y. (2017). Toxicity and Efficacy Probability Interval Design for Phase I Adoptive Cell Therapy Dose-Finding Clinical Trials. Clinical Cancer Research, 23(1):13–20.
- [Li et al.2002] Li, G., Shih, WJ., Xie, T., and Lu, J. (2002). A sample size adjustment procedure for clinical trials based on conditional power. Biostatistics, 3(2):277–287.
- [Lin et al.2016] Lin, J., Lin, LA., and Sankoh, S. (2016). A General Overview of Adaptive Randomization Design for Clinical Trials. Biometrics and Biostatistics, 7(2):294.
- [Liukko2007] Liukko, K. (2007). Model averaging with Akaike weights. Turku PET Centre Modelling report, pages 1–2.
- [Maca et al.2006] Maca, J., Bhattacharya, S., Dragalin, V., Gallo, P., and Krams, M. (2006). Adaptive Seamless Phase II/III Designs-Background, Operational Aspects, and Examples. Drug information journal, 40:4.
- [MacCallum and Bornkamp2015] MacCallum, E. and Bornkamp, B. (2015). Accounting for parameter uncertainty in two-stage designs for Phase II dose-response studies. Modern Adaptive Randomized Clinical Trials: Statistical and Practical Aspects, pages 427–450.
- [Mahajan and Gupta2010] Mahajan, R. and Gupta, K. (2010). Adaptive design clinical trials: Methodology, challenges and prospect. Indian Journal of Pharmacology, 42(4):201–207.
- [Meibohm and Derendorf1997] Meibohm, B. and Derendorf, H. (1997). Basic concepts of pharmacokinetic/pharmacodynamic (PK/PD) modeling. International journal of clinical pharmacology and therapeutics, 35(10):401–413.
- [Miller et al.2014] Miller, F., Björnsson, M., Svensson, O., and Karlsten, R. (2014). Experiences with an adaptive design for a dose-finding study in patients with osteoarthritis. Contemporary clinical trials, 37:189–199.
- [Miller et al.2007] Miller, F., Guilbaud, O., and Dette, H. (2007). Optimal Designs for Estimating the Interesting Part of a Dose-Effect Curve. Journal of Biopharmaceutical Statistics, 17(6):1097–1115.

- [Morris et al.2019] Morris, TP., White, IR., and Crowther, MJ. (2019). Using simulation studies to evaluate statistical methods. Statistics in Medicine, 38(11):2074–2102.
- [Mullen2014] Mullen, K. (2014). Continuous Global Optimization in R. Journal of Statistical Software, 60(6).
- [Müller and Schäfer2001] Müller, HH. and Schäfer, H. (2001). Adaptive group sequential designs for clinical trials: combining the advantages of adaptive and of classical group sequential approaches. Biometrics, 57(3):886–891.
- [Musuamba et al.2017] Musuamba, FT., Manolis, E., Holford, N., Cheung, S., Friberg, LE., Ogungbenro, K., Posch, M., Yates, J., Berry, S., Thomas, N., Corriol-Rohou, S., Bornkamp, B., Bretz, F., Hooker, AC., Van der Graaf, PH., Standing, JF., Hay, J., Cole, S., Gigante, V., Karlsson, K., Dumortier, T., Benda, N., Serone, F., Das, S., Brochot, A., Ehmann, F., Hemmings, R., and Rusten, IS. (2017). Advanced Methods for Dose and Regimen Finding During Drug Development: Summary of the EMA/EFPIA Workshop on Dose Finding (London 4-5 December 2014). CPT: Pharmacometrics and Systems Pharmacology, 6(7):418–429.
- [Nash2014] Nash, J. (2014). On Best Practice Optimization Methods in R. Journal of Statistical Software, 60(2).
- [Nguefack-Tsague2014] Nguefack-Tsague, G. (2014). On Optimal Weighting Scheme in Model Averaging. American Journal of Applied Mathematics and Statistics, 2(3):150–156.
- [Pallmann et al.2018] Pallmann, P., Bedding, AW., Choodari-Oskoei, B., Dimairo, M., Flight, L., Hampson, LV., Holmes, J., Mander, AP., Odondi, L., Sydes, MR., Villar, SS., Wason, J., Weir, CJ., Wheeler, GM., Yap, C., , and Jaki, T. (2018). Adaptive designs in clinical trials: why use them, and how to run and report them. BMC Medicine, 16(1):29.
- [Parmigiani et al.2009] Parmigiani, G., Inoue, L., and Lopes, H. (2009). Decision Theory, Principles and Approaches. Wiley series in probability and statistics.
- [Patel et al.2012] Patel, N., Bolognese, J., Chuang-Stein, C., Hewitt, D., Gammaitoni, A., and Pinheiro, J. (2012). Designing phase II trials based on program-level considerations: a case study for neuropathic pain. Drug Information Journal, 46(4):439–454.
- [Patel et al.2013] Patel, NR., Ankolekar, S., Antonijevic, Z., and Rajicic, N. (2013). A mathematical model for maximizing the value of phase III drug development portfolios incorporating budget constraints and risk. Statistics in Medicine, 32(10):1763–1777.
- [Peddada and Haseman2005] Peddada, SD. and Haseman, JK. (2005). Analysis of Nonlinear Regression Models: A Cautionary Note. Dose-response, 3(3):342–352.
- [Pinheiro et al.2014] Pinheiro, J., Bornkamp, B., Glimm, E., and Bretz, F. (2014). Model-based dose finding under model uncertainty using general parametric models. Statistics in Medicine, 33(10):1646–1661.
- [Proschan and Hunsberger1995] Proschan, MA. and Hunsberger, SA. (1995). Designed extension of studies based on conditional power. Biometrics, 51(4):1315–1324.

- [Pulkstenis et al.2017] Pulkstenis, E., Patra, K., and Zhang, J. (2017). A bayesian paradigm for decision-making in proof-of-concept trials. Journal of Biopharmaceutical Statistics, 27(3):442–456.
- [Rady et al.2009] Rady, EA., Abd El-Monsef, MME., and Seyam, MM. (2009). Relationships among several optimality criteria. Interstat, 15(6):1–11.
- [Rafter et al.2002] Rafter, JA., Abell, ML., and Braselton, JP. (2002). Multiple Comparison Methods for Means. Society for Industrial and Applied Mathematics, 44(2):259–278.
- [Ravenzwaaij et al.2018] Ravenzwaaij, DV., Cassey, P., and Brown, SD. (2018). A simple introduction to markov chain monte-carlo sampling. Psychonomic Bulletin and Review, 25(1):143–154.
- [Robert and Casella2004] Robert, C. and Casella, G. (2004). Monte Carlo Statistical Methods. Springer, (ISBN: 978-1-4757-4145-2).
- [Roberts and Tweedie1996] Roberts, GO. and Tweedie, RL. (1996). Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. Biometrika, 83(1):95–110.
- [Robertson et al.1991] Robertson, T., Wright, FT., and Dykstra, RL. (1991). Order Restricted Statistical Inference. Journal of Applied Econometrics, 6(1):105–107.
- [Sacks et al.2014] Sacks, LV., Shamsuddin, HH., Yasinskaya, YI., Bouri, K., Lanthier, ML., and Sherman, RE. (2014). Scientific and regulatory reasons for delay and denial of FDA approval of initial applications for new drugs, 2000-2012. JAMA, 311(4):378–384.
- [Savage1954] Savage, LJ. (1954). The Foundations of Statistics. Dover publications, INC., (ISBN-13: 978-0486623498).
- [Schmeiser and Song1996] Schmeiser, BW. and Song, WT. (1996). Batching methods in simulation output analysis: What we know and what we don't. Proceedings of the 28th Conference on Winter Simulation, pages 122–127.
- [Schwarz1978] Schwarz, G. (1978). Estimating the dimension of a model. The Annals of Statistics, 6(2):461–464.
- [Snipes and Taylor2014] Snipes, M. and Taylor, DC. (2014). Model selection and Akaike Information Criteria: An example from wine ratings and prices. Wine Economics and Policy, 3(1):3–9.
- [Stallard et al.2004] Stallard, N., Thall, PF., and Whitehead, J. (2004). Decision Theoretic Designs for Phase II Clinical Trials with Multiple Outcomes. Biometrics, 55(3):971–977.
- [Stallard et al.2005] Stallard, N., Whitehead, J., and Cleall, S. (2005). Decision-making in a phase II clinical trial: a new approach combining bayesian and frequentist concepts. Pharmaceutical Statistics, 4:119–128.
- [Temple2012] Temple, J. (2012). Adaptive designs for Dose-finding trials. University of Bath, Department of Mathematical Sciences, pages 1–206.

- [Thall2012a] Thall, PF. (2012a). Adaptive Randomization to Improve Utility-Based Dose-Finding with Bivariate Ordinal Outcomes. Journal of Biopharmaceutical Statistics, 22(4):785–801.
- [Thall2012b] Thall, PF. (2012b). BAYESIAN ADAPTIVE DOSE-FINDING BASED ON EFFICACY AND TOXICITY. Journal of Statistical Research, 46(2):187–202.
- [Thall and Cook2004] Thall, PF. and Cook, JD. (2004). Dose-Finding Based on Efficacy-Toxicity Trade-Offs. Biometrics, 60(3):684–693.
- [Thomas2017] Thomas, N. (2017). Understanding MCP-MOD dose finding as a method based on linear regression. Statistics in Medicine, 36(27):4401–4413.
- [Thomas and Dooti2017] Thomas, N. and Dooti, R. (2017). Analysis of Clinical Dose-Response in Small-Molecule Drug Development: 2009-2014. Statistics in Biopharmaceutical Research, 9(2):137–146.
- [Thomas et al.2014] Thomas, N., Sweeney, K., and Somayaji, V. (2014). Meta-Analysis of Clinical Dose-Response in a Large Drug Development Portfolio. Statistics in Biopharmaceutical Research, 6(4):302–317.
- [Ting2006] Ting, N. (2006). Dose Finding in drug development. Statistics for Biology and Health, pages 149–153.
- [Tsai2006] Tsai, CL. (2006). Two-Stage Designs for Phase II Clinical Trials. Institute of Statistics, National University of Kaohsiung, pages 1–47.
- [Turkheimer et al.2003] Turkheimer, FE., Hinz, R., and Cunningham, VJ. (2003). On the Undecidability Among Kinetic Models: From Model Selection to Model Averaging. Journal of Cerebral Blood Flow and Metabolism, 23(4):490–498.
- [Vandemeulebroecke2006] Vandemeulebroecke, M. (2006). An investigation of two-stage tests. Statistica Sinica, 16(3):933–951.
- [Vandenberg et al.2012] Vandenberg, LN., Colborn, T., Hayes, TB., Heindel, JJ., Jacobs, DR., Lee, DH., Shioda, T., Soto, AM., Saal, FS., Welshons, WV., Zoeller, RT., and Myers, JP. (2012). Hormones and Endocrine-Disrupting Chemicals: Low-Dose Effects and Nonmonotonic Dose Responses. Endocrine Reviews, 33(3):378–455.
- [Verhagen et al.1998] Verhagen, AP., de Vet, HC., de Bie, RA., Kessels, AG., Boers, M., Bouter, LM., and Knipschild, PG. (1998). The Delphi List: A Criteria List for Quality Assessment of Randomized Clinical Trials for Conducting Systematic Reviews Developed by Delphi Consensus. Journal of clinical epidemiology, 51(12):1235–1241.
- [Watson and Holmes2016] Watson, J. and Holmes, C. (2016). Approximate Models and Robust Decisions. Statistical Science, 31(4).
- [Willan and Pinto2005] Willan, A. and Pinto, E. (2005). The value of information and optimal clinical trial design. Statistics in Medicine, 24:1791–1806.

- [Wong1994] Wong, WK. (1994). Comparing robust properties of A, D, E and G-optimal designs. Computational Statistics and Data Analysis, 18(4):441–448.
- [Yu et al.2018] Yu, J., Kong, X., Ai, M., and Tsui, KL. (2018). Optimal designs for dose-response models with linear effects of covariates. Computational Statistics and Data Analysis, 127:217–228.

Appendices

A.1. Complementary developments on utility-based seamless design optimization

The aim of this appendix is the following:

1. illustrate influence of relative sample size of phase II (with respect to total sample size) on the expected utility $U0$, as well as influence of safety (as the dose grows) on the PoS of $U0$
2. illustrate how utility $U5$ and $U9$ can be used to optimize the seamless phase II/phase III design study

1. Expanded discussion on utility $U0$

The aim of this section is to detail the calculations and the statistical methodology to compute the expected utility for $U0$. And discuss influence of the phase II sample size on $\mathbb{E}(U0)$.

General context

- Seamless phase II/phase III study context: fixed total sample size N_{tot} , phase II $\rightarrow N_2 = f \times N_{tot}$, phase III $\rightarrow N_3 = (1 - f) \times N_{tot}$
- Frequentist approach: one 'true' parameter θ_0
- Emax mean dose-response profile
- As an example, we consider here the $U0$ utility function to illustrate this approach, based on PoS and safety modeling, efficacy is mitigated by safety: for a given dose, d , $\mathbb{P}(\text{safety issue}) = sa(d/d_J)^2$ where d_J is the maximal dose in the design and sa is the probability of safety problems
- Utility = Reward(if success) - Cost
 - success \rightarrow significant test + lack of safety issue
- Decision rule of the sponsor only based on point estimate of model parameter $\hat{\theta}$

– after phase II sponsors decides to 'Go' into phase III if the 'best' selected dose d^* has $PoS_{adj}(d^*) \geq 0.3$ & expected utility (=gain) > 0 , where:

* $\bar{\Delta}(d)$ is the difference between chosen dose d arm & placebo in phase III and $SE^2 = 2\sigma^2/(N_3/2)$

* $\mathbb{P}_{H_1}(\bar{\Delta}_d \geq 1.96 \times \sqrt{2SE^2}) = \Phi\left(\frac{m(d; \theta) - m(0; \theta) - 1.96 \times \sqrt{2SE^2}}{\sqrt{2SE^2}}\right)$;

$m(d; \theta) = E_0 + \frac{E_{max} \cdot d}{ED_{50} + d}$ (three-parameter Emax model for the efficacy)

* expected utility, for the sponsor, after phase II: equal $-\gamma N_2$ if 'NoGo', $-\gamma N_{tot} + R\mathbb{E}_{\hat{\theta}}(Success)$ if go

• $U = U_0(d, f) = -\gamma N_2 \mathbf{1}(NoGo) + \mathbf{1}(Go)(-\gamma N_{tot} + R \times PoS_{adj}(d))$

• Before study starts: sponsor's general strategy is to maximize (in phase II design, w (w is the design, i.e. patients allocation per dose), and N_2/N_{tot} ratio, f) the expected utility.

$$\mathbb{E}_{\theta_0} U = \mathbb{E}_0(U_1 \times (1 - Go(\hat{\theta})) + U_2 \times Go(\hat{\theta})) = \mathbb{E}_0^{(\hat{\theta})}(\mathbb{E}_0(U_1 \times (1 - Go(\hat{\theta})) + U_2 \times Go(\hat{\theta}) | \hat{\theta})) = \mathbb{E}_0^{(\hat{\theta})}(\mathbb{E}_0(U_1 | \hat{\theta}) \times (1 - Go(\hat{\theta})) + \mathbb{E}_0(U_2 | \hat{\theta}) \times Go(\hat{\theta}))$$

– $\mathbb{E}_0(U_1 | \hat{\theta}) = -\gamma \times f \times N_{tot}$: loss due to cost of phase II if 'NoGo'

– $\mathbb{E}_0(U_2 | \hat{\theta}) = -\gamma \times N_{tot} + R \times PoS_{adj(\theta_0)}(d(\hat{\theta}))$: expected gain minus total cost if 'Go'

Therefore, sponsor's strategy is to maximize in w, f the function:

$$\mathbb{E}_0 U = \mathbb{E}_0^{(\hat{\theta})} \left(((-\gamma \times f \times N_{tot}) \times (1 - Go(\hat{\theta})) + (-\gamma \times N_{tot} + R \times PoS_{adj(\theta_0)}(d(\hat{\theta}))) \times Go(\hat{\theta})) \right),$$

where $\hat{\theta} \sim N(\theta_0, (f \times N_{tot} \sum_d w_d I_d)^{-1})$, $Go(\hat{\theta})$ is the decision based on parameter estimate and $PoS_{adj(\theta_0)}(d(\hat{\theta}))$ is the true PoS, depends on true parameter θ_0 , of the dose chosen (based on the estimated parameter $\hat{\theta}$); $I = -\mathbb{E}_{\theta_0} \partial^2 l / \partial \theta^2$ is the Fisher information matrix.

Sponsor's strategy

• After phase II:

– compute $\mathbb{E}(U(d, f) | phase II)$ for each dose d

– compute $d^* = \arg \max_d \mathbb{E}(U(d, f) | phase II)$

– decide if worth going into phase III: if $PoS_{adj}(d^*) \geq 0.30$ & expected utility (=gain) > 0

• Before phase II:

– choose $N_2 (= f \times N_{tot})$ sample size of phase II

– choose the design w

$$(w_*) = \arg \max \mathbb{E}_{w,f}^{(phaseII)} \mathbb{E}(U(d^*)|phaseII)$$

or

$$(f_*) = \arg \max \mathbb{E}_{w,f}^{(phaseII)} \mathbb{E}(U(d^*)|phaseII)$$

Computation of expectation

The expectation is computed via numerical integration: $E(U) = \int (((-\gamma \times fNtot) \times (1 - Go(\theta)) + (-\gamma \times Ntot + R \times Pos_{adj}(\theta_0)(d(\theta))) \times Go(\theta)))p(\theta)d\theta$, where $p(\theta)$ is the density of a Gaussian distribution centered at the true value of the parameters and with covariance matrix equal to the inverse of the Fisher matrix.

In practice, the computation was done, in R, as following: the hessian was calculated for each dose using the symbolic differentiation operator 'D', and the three dimensional integral was computed using the 'cubature' package in R.

Computation of Fisher information matrix

- The fisher information matrix (FIM), denoted by \mathcal{I} , is equal to $-\mathbb{E}^{(y)} \partial^2(\log-likelihood)/\partial\theta^2$
- For a patient receiving dose d , Y follows a normal distribution centered in $m(d, \theta)$ with variance equal to σ^2
- For one given patient with data y , the likelihood is $\frac{1}{\sqrt{2\pi\sigma^2}} \exp(-\frac{1}{2\sigma^2} \times (y - m(d, \theta))^2)$, and by forgetting the π , the log-likelihood is $-\frac{1}{2\sigma^2} \times (y - m(d, \theta))^2$
- So for some patients N_d at some dose d :

$$\mathcal{I}_\theta(d) = \frac{N_d}{\sigma^2} \times \frac{1}{2} \mathbb{E}^{(Y)} \left[\partial^2 (y - m(d; \theta))^2 / \partial\theta^2 \right]$$
- For one design, w , and a total of N_2 patients in phase II:

$$\mathcal{I}_\theta = \frac{N_2}{2} \sum_d w_d \times \frac{1}{2} \mathbb{E}^{(Y)} \left[\partial^2 (y - m(d; \theta))^2 / \partial\theta^2 \right]$$
- The second derivative $\partial^2 (y - m(d; \theta))^2 / \partial\theta^2$ is linear in y
 $\Rightarrow \mathbb{E}^{(y)} \partial^2 (y - m(d; \theta))^2 / \partial\theta^2 = \partial^2 (y - m(d; \theta))^2 / \partial\theta^2$ evaluated at $y = \mathbb{E}Y = m(d; \theta)$

Concretely, this is how we compute the Fisher information matrix (FIM) with a three-parameter Emax model:

Log-likelihood, FIM for one patient, $ll : \theta = (E_0, E_{max}, ED_{50})$

$$\rightarrow -\frac{1}{2\sigma^2} \times \left(y - E_0 - E_{max} \frac{d}{ED_{50} + d} \right)^2.$$

Fisher information matrix:

$$\begin{aligned} \text{FIM} &= \mathcal{I}_\theta(d) = -\mathbb{E}(\partial^2 l / \partial \theta^2) \\ &= \begin{pmatrix} 1 & d/(\sigma^2(ED_{50}+d)) & d \times E_{max}/(\sigma^2(ED_{50}+d)^2) \\ d/(\sigma^2(ED_{50}+d)) & d^2/(\sigma^2 \times (ED_{50}+d)^2) & -(d^2 \times E_{max})/(\sigma^2 \times (ED_{50}+d)^3) \\ d \times E_{max}/(\sigma^2(ED_{50}+d)^2) & -(d^2 \times E_{max})/(\sigma^2 \times (ED_{50}+d)^3) & (d^2 \times E_{max}^2)/(\sigma^2 \times (ED_{50}+d)^4) \end{pmatrix} \end{aligned}$$

For a design w , and for N_2 patients in phase II, the global FIM is the sum $\frac{N_2}{\sigma^2} \sum_d w_d I_d$.

In the following examples, computation done using symbolic differentiation function 'D' in R.

Global computation

- The large sample approximation of the estimate, $\hat{\theta}$, distribution (based on Fisher information matrix) is used: this enables to avoid simulating the whole set of individual phase II data.
- The θ parameter is three-dimensional: $\theta_1 = E_0, \theta_2 = E_{max}, \theta_3 = ED_{50}$. The E_0 parameter is supposed to have a value of zero (no placebo effect) in our chosen scenarios. The ED_{50} parameter is necessarily > 0 (otherwise the relationship is meaningless and not defined for the dose equal to opposite of ED_{50}); therefore to avoid risk of negative \widehat{ED}_{50} , the model is reparametrized in: $\theta_2 = E_{max}, \theta_3 = \log(ED_{50}) \Leftrightarrow ED_{50} = e^{\theta_3}$

Programming

- Computation of Fisher matrix: for each dose d , the Hessian matrix $\partial^2 l(d) / \partial \theta^2$ is computed using R symbolic differentiation operator 'D' (l is the log-likelihood)
- For computing the expectation, with respect to $\hat{\theta}$, of the conditional utility, we applied a numerical integration
 - for this the 'cubature' package was used

Results for U0

Mean dose-response

- The maximum effect size simulated in the example is 0.4
- Costs/Reward parameter values set to the values in J.Temple thesis [Temple2012]:
 - R =reward=12000
 - γ =cost per patient=1
- The function for safety assumes that the maximum probability of phase III failing due to safety is sa
- $N_{tot} = 2000$
- σ =residual variability=1

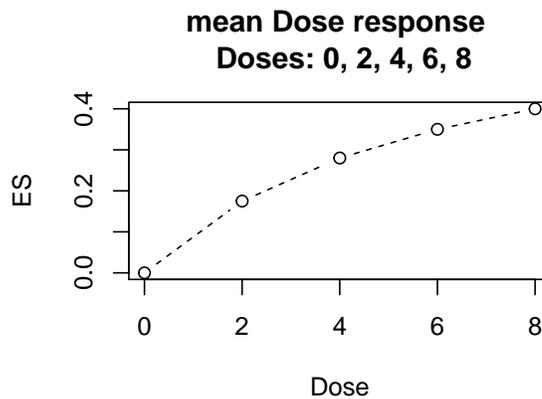


Figure 6.2: Effect size as a function of the dose.

Impact of safety

The following graph shows the PoS by dose for various values of the safety parameter and for 600 patients in phase III. For good safety profiles (low value of 'sa') the PoS increases then reaches a plateau, whereas for less good safety profiles the PoS increased to an optimum value and then decreases.

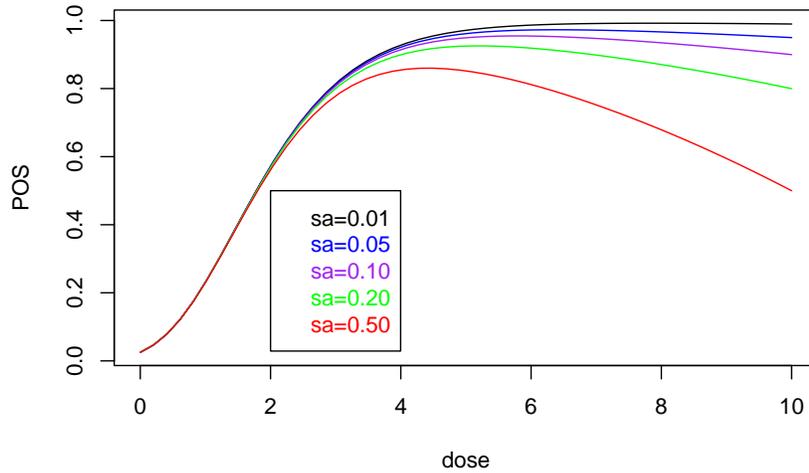


Figure 6.3: PoS by dose for various values of the safety parameter.

Computation of mean utility

Two methods of integration were used:

- One based on a quadrature method for multidimensional integrals (cubature package)
- The other one based on successive calls of the R "integrate" function
- The first method seems to be the fastest
- They seem more appropriate than use of Monte Carlo simulation to compute the expectation of U (a lot of simulations are required for having 'smooth' results: otherwise when plotting $\mathbb{E}U$ as a function of phase II sample size, there are oscillations). Monte Carlo could be more efficient than numerical integration when there is a larger number of parameters (> 3 parameters for instance, this method was applied in the next section) [possible theme or research: use Laplace approximation method to compute the integrals when optimizing]

For balanced design, utility as a function of phase II sample size

In the following, we plotted $E(U)$ graphs as a function of N_2 , for several scenarios (where scenarios refer here to Sigmoid efficacy profile and several values of 'sa'). These graphs are represented for a balanced design.

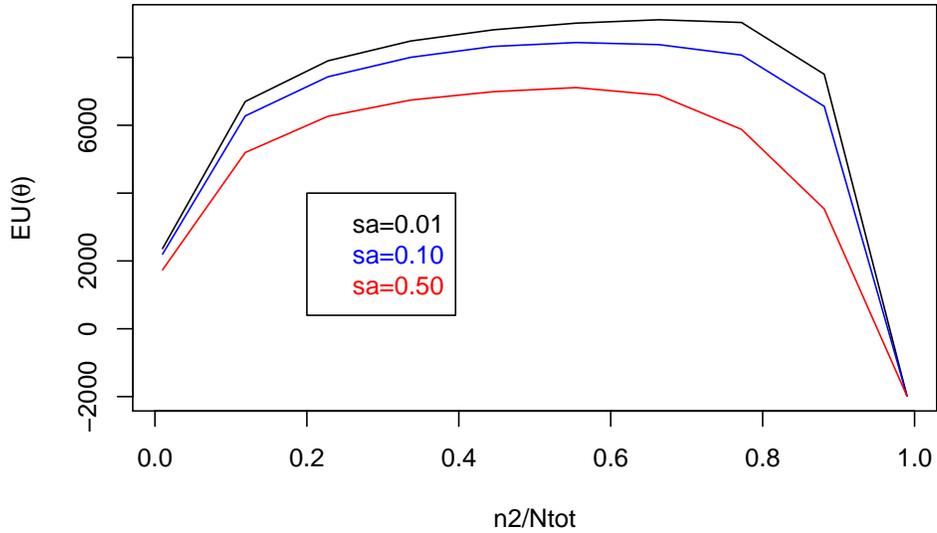


Figure 6.4: Utility as a function of phase II sample size and safety sa .

Based on Figure 6.4, we noticed that whatever the safety profile, the expected utility increases and then reaches a plateau and a maximum value before rapidly decreasing for very low phase III sample size (when N_2/N_{tot} is large). Also we can notice that the optimal phase II sample size decreases as the safety gets worse: this is due to the fact that optimal dose decreases as the safety gets worse and then requires more patients in phase III to be successful .

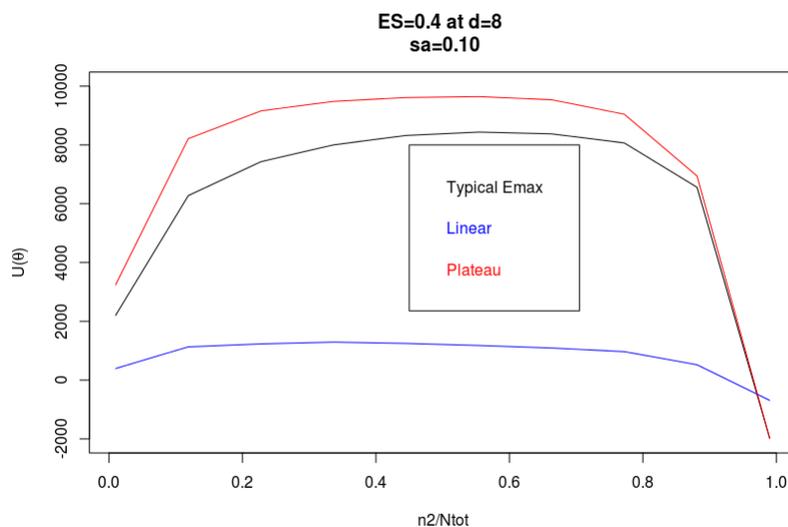


Figure 6.5: Utility as a function of phase II sample size for three dose-response profiles.

In Figure 6.5, we see the expected utility as a function of N_2 sample size, for various profile of the dose-response (but the safety and mean response of the highest dose are the same: Effect size=0.4 for $d = 8$)

- The linear profile is not favourable: the highest dose has some safety issue whereas the efficacy of the previous dose is far behind
- For the two other scenarios, the situation is better: lowest and safer doses have efficacy profiles closer to the highest dose and can be selected for phase III \Rightarrow increased utility

Note that running the R code was extremely long because of this last graph above, especially when the model is close to the linear one; the reason is that when the model is almost linear, the Emax model is badly estimated with enormous variability \Rightarrow the integration domain is very large \Rightarrow huge computational time.

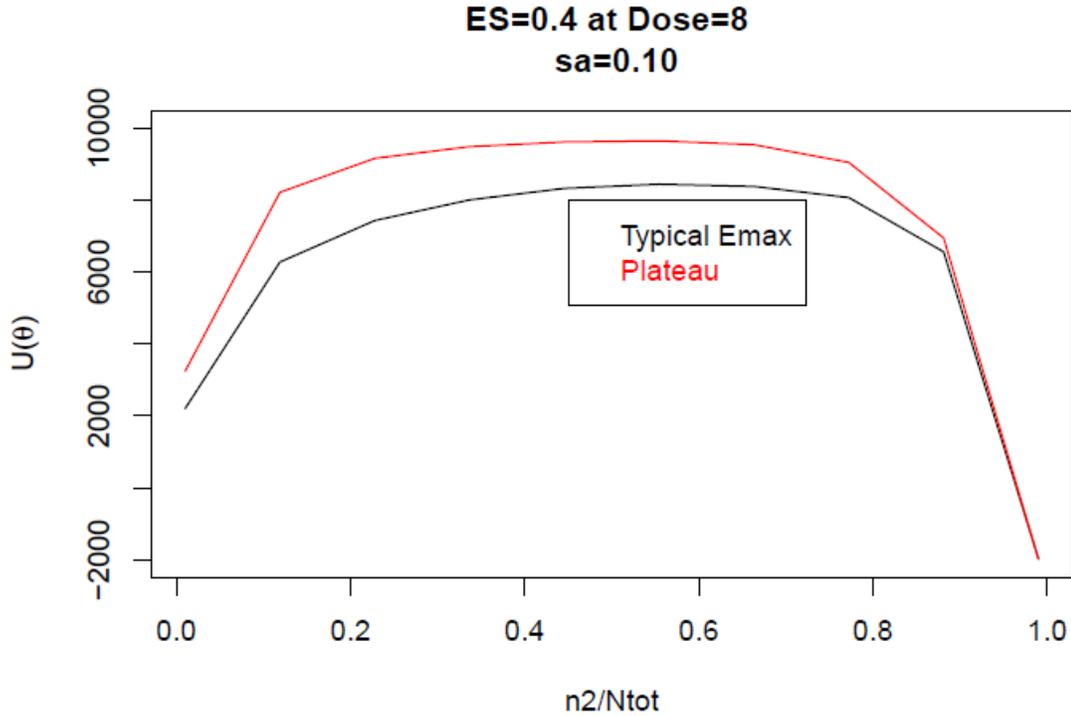


Figure 6.6: Utility as a function of phase II sample size for two dose-response profiles.

In Figure 6.6, I represented the graph comparing the efficiency profiles but this time, I abandoned the almost linear profile ($E_{max} = 1000 \times 0.4/8 = 50$, $ED_{50} = 1000$); I chose these parameters to have a very large $ED_{50} = 1000$ compared to $d_J = 8$, with an effect approximately equal to $E_{max} \times 8/ED_{50} = 0.4$ at $d = 8$; the conclusion is that the Emax model is badly estimated and has no points towards the maximum effect; if the profile is linear or almost linear, then it is better to have a linear model and calculate the average utility in this framework.

The plateau profile ($E_{max} = 0.4$, $ED_{50} = 0.50$) is more favourable than a more typical Emax model ($E_{max} = 0.7$, $ED_{50} = 6$): lower doses with similar efficacy than the maximum dose, but with a better safety, can be selected \Rightarrow increased utility.

Conclusion

From these analyses we noticed that, at least for balanced design, the optimal sample size of the phase II part of the seamless design can be quite large.

In the following, we will tackle the problem of the optimisation of the design (optimal repartition of the patients between treatment arms) of the phase II part, in addition to the problem of optimisation of its total sample size. For this purpose, as we noticed that the computation time of the expectations EU with numerical integration routines were very long, we will, for these next parts, rather use simulations.

2. Optimisation of the seamless design based on $U5$ and $U9$ utility functions

It is recalled that $U5$ and $U9$ are defined as follows:

- Utility 5

$$U5(d, f) = PoS(d) \times (1 - c \times \delta)$$

- Utility 9

$$U9(d, f) = PoS(d) \times (1 - c \times (\frac{d}{d_J})^2)$$

Methodology

For these utility examples, we have decided to:

- Work in a frequentist context
- Analyse phase II with a parametric model (E_{max} model with a parameter θ)
- Work with a four-parameter sigmoid E_{max} model, reflecting our mean dose-response function: $m(d, \theta) = E_0 + E_{max} \times \frac{d^g}{ED_{50}^g + d^g} = E_0 + (E_m - E_0) \times \frac{d^g}{ED_{50}^g + d^g}$
 - E_0 is the placebo effect ($E_0 = 0$ for simulations)
 - $E_{max} = E_m - E_0$ is the maximum difference against placebo, where E_m is the maximal effect ($E_m - E_0 = 0.22$ for simulations)
 - ED_{50} is the dose that gives half of the maximum effect ($ED_{50} = 6$ for simulations)
 - g is the Hill exponent that describes and determines the steepness of the curve ($g = 3$ for simulations)
- 'c' coefficient is set to 0.8: it was calibrated so that the highest tested dose is located after the peak of the utility curve, this exemplifies the model's behavior, and shows that it does not necessarily select the highest dose all the time
- Express conditional PoS only with the estimated parameter in phase II, two possible approaches:

- Conditional PoS calculated only with the estimate point of phase II: $\bar{Y}_d - \bar{Y}_0 \sim N(m(d; \hat{\theta}_{ML}) - m(0; \hat{\theta}_{ML}), 2\sigma^2/N_3)$, where ML refers to the maximum likelihood and I^{-1} refers to the reverse of the Fisher information matrix; $E_0 = 0$ in simulations so $m(0; \hat{\theta}_{ML}) = 0$
- Conditional PoS computed by taking into account the uncertainty (around the estimate point), via Delta-method:

$$\Rightarrow \text{increased variability } \bar{Y}_d - \bar{Y}_0 \sim N(m(d; \hat{\theta}_{ML}) - m(0; \hat{\theta}_{ML}), \frac{\partial m}{\partial \theta} I^{-1} \frac{\partial m}{\partial \theta} + 2\sigma^2/N_3).$$

PoS given variability:

$$PoS = \mathbb{P}_{H_1}(\bar{\Delta}_d \geq 1.96\sqrt{2SE^2}), \text{ where } \bar{\Delta}_d \sim N(m(d; \theta) - m(0; \theta), 2SE^2 + V^2),$$

$$SE^2 = \sigma^2/n_{arm}, n_{arm} = N_3/2 \text{ (we divide } N_3 \text{ by 2 because the PoS is computed for phase III, and in phase III we have the chosen dose and the placebo), } \bar{\Delta}_d \text{ is the difference in effects between dose } d \text{ and the placebo after phase III, and } V^2 \text{ is its the variability, } V^2 = \frac{\partial m}{\partial \theta} I^{-1} \frac{\partial m}{\partial \theta}$$

- In the following simulation results, we only retained the first approach (without taking the uncertainty into account), because with the second approach, variability increases and consequently, the PoS decreases, compared to the first approach

Computation of expectation

The difficulty here is that the calculation of the expectation $E(U)$ is done with an integral of four dimensions instead of three. This is why we have decided to compute the expectation via simulations, as follows:

$\mathbb{E}(U(d^*)|phaseII)$ is a function, \mathcal{U} , of $\hat{\theta} \Rightarrow \mathbb{E}_{w,f}^{(phaseII)} \mathbb{E}(U(d^*)|phaseII) = \mathbb{E}^{(\hat{\theta})} \mathcal{U}(\hat{\theta})$ with $\hat{\theta} \sim N(\theta_0, \mathcal{I}_{\theta_0}^{-1})$.

Then $\mathbb{E}_{w,f}^{(phaseII)} \mathbb{E}(U(d^*)|phaseII)$ can be estimated by:

$$\frac{1}{N_{sim}} \sum_{r=1}^{N_{sim}} \mathcal{U}(\hat{\theta}_r)$$

where the $\hat{\theta}_r$ are sampled from $N(\theta_0, \mathcal{I}_{\theta_0}^{-1})$.

In the following results, a graph highlighting the theoretical curves related to the Sigmoid scenario is drawn for each utility function, where the blue curve is the PoS, the dotted curve is the penalty and the black curve is the utility, i.e. the product of PoS \times penalty, and a table summarizing all the simulation results is given. This table contains the following: w is the design (patients allocation per dose), f is the parameter representing the distribution between phase II and phase III, 'Go' is the probability of going to phase III with the chosen dose, 'doses' represents probabilities of choosing $d = 2$, $d = 4$, $d = 6$ and $d = 8$ respectively among the 'Go', POS(go) is the POSs mean among the 'Go' with the chosen dose and $E(U)$ is the expected utility of the chosen dose for 10000 simulated phase II studies among 'Go' and 'NoGo' decisions (utility is set to 0 when it is a 'NoGo' decision). Optimisations of patient

allocation to doses and global allocation ratio between phase II and phase III are conducted separately (for $U5$, w is optimized while f is fixed, and for $U9$, f is optimized while w is fixed). For optimization purpose, as we anticipated very long computation time of the expectations $E(U)$ with numerical integration routines, we used simulations instead (see Section 3.4 and discussion above).

Results for U5

For $U5$, below is a plot containing the three theoretical curves of the utility function:

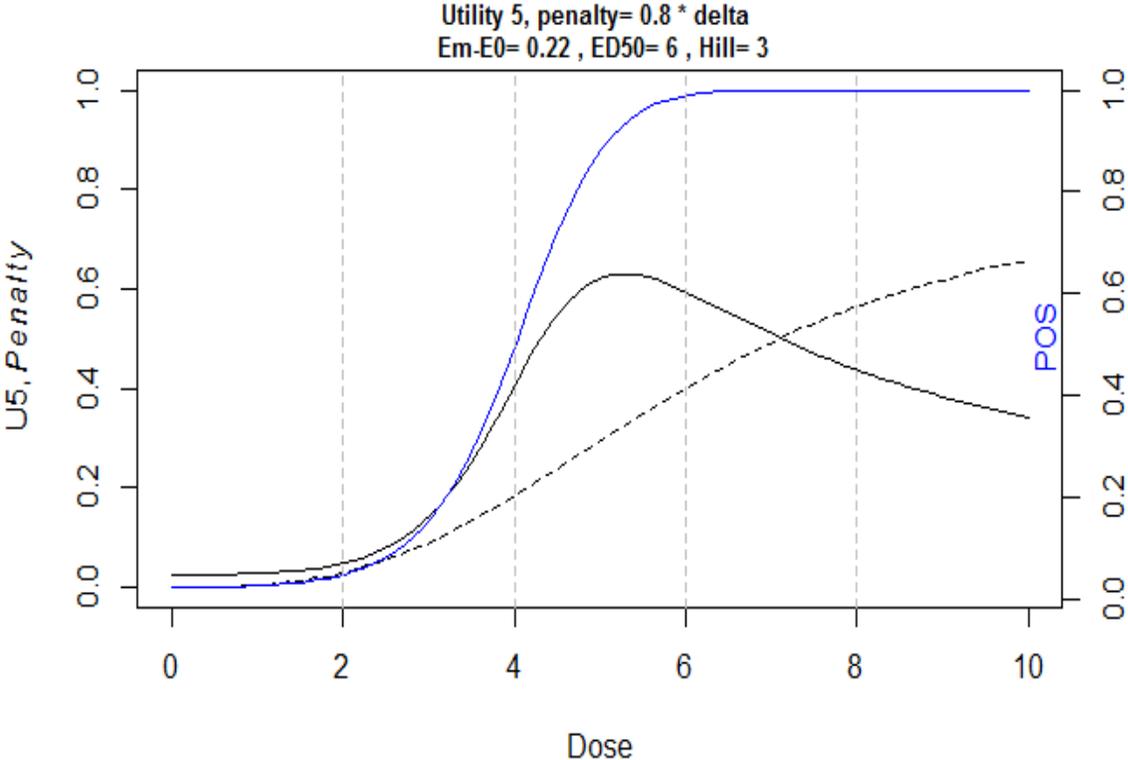


Figure 6.7: Theoretical curves for U5, Sigmoid scenario.

Concerning the balanced design (i.e. patients are equally allocated to doses), we have:

Non optimal

sigmoid	plateau
$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.25$	$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.25$
Go=:63%	Go=:94%
doses= 0.47, 0.21, 0.19, 0.14	doses= 0.66, 0.21, 0.04, 0.09
POS(go)=45%	POS(go)=90%
E(U)= 0.16	E(U)= 0.40

Table 6.1: Balanced design -patients are equally allocated to doses- for $U5$, Sigmoid and Plateau scenarios.

By comparing results to the theoretical utility graph, we can see that $d = 2$ is recommended with a very high probability: $d = 2$ is chosen in 47% of cases if 'Go', in Sigmoid scenario.

Below are the optimization results in w , with $f = 0.20$:

Optimal design (f=0.20)

sigmoid	Plateau
$w = (0.20, 0.19, 0.19, 0.21, 0.20)$, $f = 0.20$	$w = (0.22, 0.21, 0.18, 0.20, 0.19)$, $f = 0.20$
Go=:62%	Go=:93%
doses= 0.51, 0.20, 0.15, 0.14	doses= 0.66, 0.18, 0.05, 0.11
POS(go)=42%	POS(go)=92%
E(U)= 0.14	E(U)= 0.40

Table 6.2: Optimal design -optimizing the dose allocation ratio- for $U5$, Sigmoid and Plateau scenarios, with $f = 0.20$.

The balanced design is almost the optimal design. In addition, there is no real gain brought by the optimization compared to the balanced design, the probabilities of 'Go', the choice of doses, the PoSs mean and the expectations of utility are almost the same (for the Sigmoid scenario for example, design optimization decreases the average utility by 0.02% compared to the balanced design).

Below are the optimal design results for both Sigmoid and Plateau scenarios, with $f = 0.25$:

Optimal design

sigmoid	Plateau
$w = (0.21, 0.21, 0.21, 0.19, 0.18)$, $f = 0.25$	$w = (0.22, 0.18, 0.19, 0.21, 0.20)$, $f = 0.25$
Go=:63%	Go=:94%
doses= 0.46, 0.22, 0.18, 0.14	doses= 0.65, 0.21, 0.05, 0.09
POS(go)=45%	POS(go)=90%
E(U)= 0.16	E(U)= 0.39

Table 6.3: Optimal design -optimizing the dose allocation ratio- for U_5 , Sigmoid and Plateau scenarios, with $f = 0.25$.

We can see that recommended doses do not seem to be consistent with Sigmoid and Plateau scenarios: $d = 2$ is too much recommended, whereas the best two doses $d = 4$ and $d = 6$ according to the theoretical utility. Note that for the Plateau scenario, the optimal design performs slightly better in selecting less often the one of the first two doses, in addition, design optimization does not increase the average utility compared to the balanced design; in such a favorable scenario -Plateau- there is little/less room for improvement anyway compared to the Sigmoid scenario. There was no noticed gain either for the Sigmoid scenario: design optimization does not increase the average utility compared to the balanced design, it remains almost the same for both designs (optimal and balanced designs).

Verification / validation of these results

According to the 15000 simulations, we find Prob(Go) and the probability of choosing each of the doses:

- Probability to go to phase III = 63.20%
- Probability of choosing $d = 2$, if go, = 46.61%
- Probability of choosing $d = 4$, if go, = 20.92%
- Probability of choosing $d = 6$, if go, = 18.97%
- Probability of choosing $d = 8$, if go, = 13.50%

We have similar results for Sigmoid scenario, $d = 2$ is chosen in 46.61% of cases.

Interpretation / Explanation of these results

I think that these bad dose choices are due to the fact that Phase II sample size is too small (500 patients) for this sigmoid model \Rightarrow often bad estimation \Rightarrow bad decisions.

One can try to show that if Phase II were larger, the decisions would be better and the U5 means would be closer to the theoretical curve previously shown.

First approach, increase the f

Indeed, it is possible to increase the phase II by increasing the f , but by increasing f , we decrease phase III: so if we increase the f , we cannot compare ourselves to this curve above because it is based on theoretical PoS for a phase III of 1500 patients. Nevertheless, we can verify that if f increases, the probability of choosing $d = 2$ will decrease, but the utility will eventually decrease also by lack of patients in phase III (see tables below):

sigmoid
$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.25$ Go=:63% doses= 0.47, 0.21, 0.19, 0.14 POS(go)=45% E(U)= 0.16
sigmoid
$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.50$ Go=:68% doses=0.27, 0.23, 0.31, 0.19 POS(go)=57% E(U)= 0.22
sigmoid
$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.75$ Go=:70% doses= 0.12, 0.11, 0.37, 0.41 POS(go)=66% E(U)= 0.22
sigmoid
$w = (0.2, 0.2, 0.2, 0.2, 0.2)$, $f = 0.95$ Go=:44% doses= 0.01, 0.00, 0.03, 0.96 POS(go)=33% E(U)= 0.06

Table 6.4: Simulation results for $U5$ with the balanced design (patients are equally allocated to doses), Sigmoid scenario, by increasing f .

Based on these tables, we can clearly see that by increasing the sample size of phase II, we reduce the probability of choosing $d = 2$ (but we also make bad choices because the more f increases, the more we choose the highest dose: we compensate the loss of the number of Patients in phase III by the selection of the most effective dose).

Second approach

I examined obtained results when we increase the phase II, by considering $N_2 = 2000$ patients, and by fixing phase III sample size, $N_3 = 1500$ patients.

With $N_2 = 2000$ patients, we obtain:

- Prob(choosing $d = 2$) = 17%
- Prob(choosing $d = 4$) = 36%
- Prob(choosing $d = 6$) = 41%
- Prob(choosing $d = 8$) = 9%

→ **This time better decisions are made: $d = 2$ is rarely chosen (17% of cases), but $d = 4$ or $d = 6$ are very often chose (one or the other is chosen in 77% of the cases, optimal dose being $d = 6$ according to theory).**

It is very important to note that the "theoretical" utility depends on the size of the phase III, and therefore, the optimal dose depends on the size of the phase III: optimal dose increases when f increases: since N_{tot} is constant, when N_2 increases, N_3 decreases which induces that higher doses are necessary to have a sufficiently high PoS, see graph below.

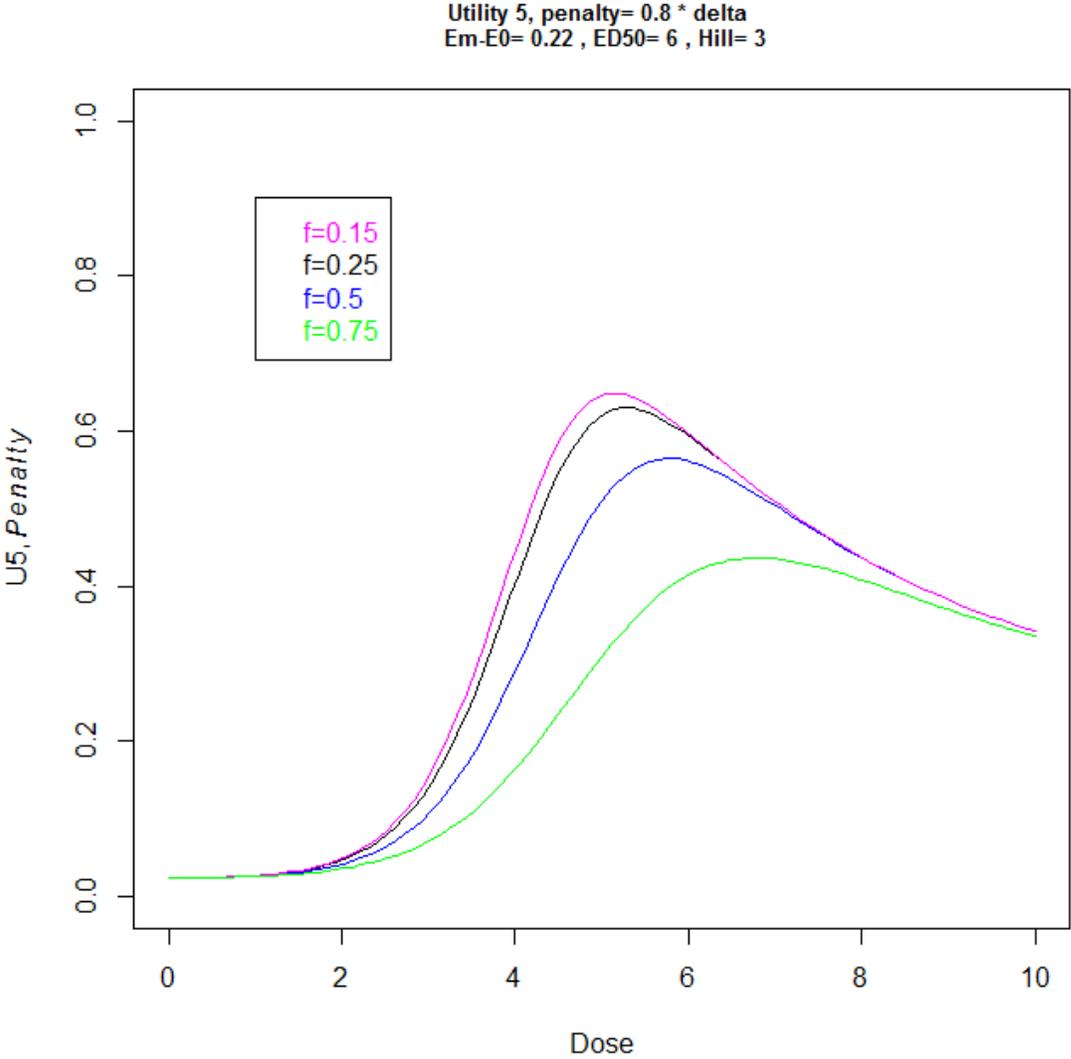


Figure 6.8: Theoretical utilities as a function of the dose and f .

Modification of optimization strategy

In the following, we have decided to work on a global patient allocation between phase II and phase III optimization, rather than dose allocation optimization, with the $U9(d, f) = PoS(d) \times (1 - c \times (\frac{d}{d_J})^2)$ utility. In fact, according to all previous results, there was no difference between the optimal and balanced designs when it comes to patients allocation to doses, and no real gain was noticed regarding the PoS, the global utility and the 'Go' proportion in phase III. So now, we will only work with a balanced w design, but this time we will seek to optimize the patients between phase II and phase III, that is to say, we will proceed an overall optimization of the patients allocation between phase II and phase III (while maintaining a fixed total number as before, $N2 + N3 = 2000$).

The simulations above clearly suggest that the estimation of a four-parameter Emax sigmoid model is a very challenging task, requiring a large number of patients to obtain sufficiently accurate estimates. Therefore we decided for the major part of the thesis to drop the Hill exponent from the dose-response models considered.

It is worth noting that in the dose-finding framework, despite the fact that the Emax model is the author's favorite and the most used one, it is known that this model (especially in its sigmoid version) can be very difficult to estimate, with much imprecision, if the design is not rich enough (the model is too rich for a design with very few doses and patients, in our case, we only have four different doses, a placebo, and 500 patients in phase II, i.e. 100 patients per arm, which is not enough to correctly estimate such a complex model). Speaking of the difficulty of estimating the Emax model, it is specified in [Dutta et al.1996] that parameters estimation of a sigmoid Emax model definitely leads to an eventual imprecision (i.e. the amount or degree of random error in a calculation, usually represented by the standard deviation, coefficient of variation, or range), and if the concentration-effect relationship extremes are not empirically investigated, this imprecision could be enormous.

In the following, we also included a second constraint in the decision rule: the PoS must be > 0.30 and the effect difference between placebo and the recommended dose must be > 0.04 (to eliminate low doses).

In fact, the threshold here (0.04) was chosen on the basis of the theoretical effect of $d = 2$ (i.e. the lowest dose) obtained with the three-parameter Emax model (E_0 , E_{max} and ED_{50} only), which was 0.055. In general, these thresholds are pre-clinically defined, but here, for our simulations, we considered a threshold equal to 0.04.

Results for U9

Below are the theoretical utility, PoS and penalty curves for $U9$.

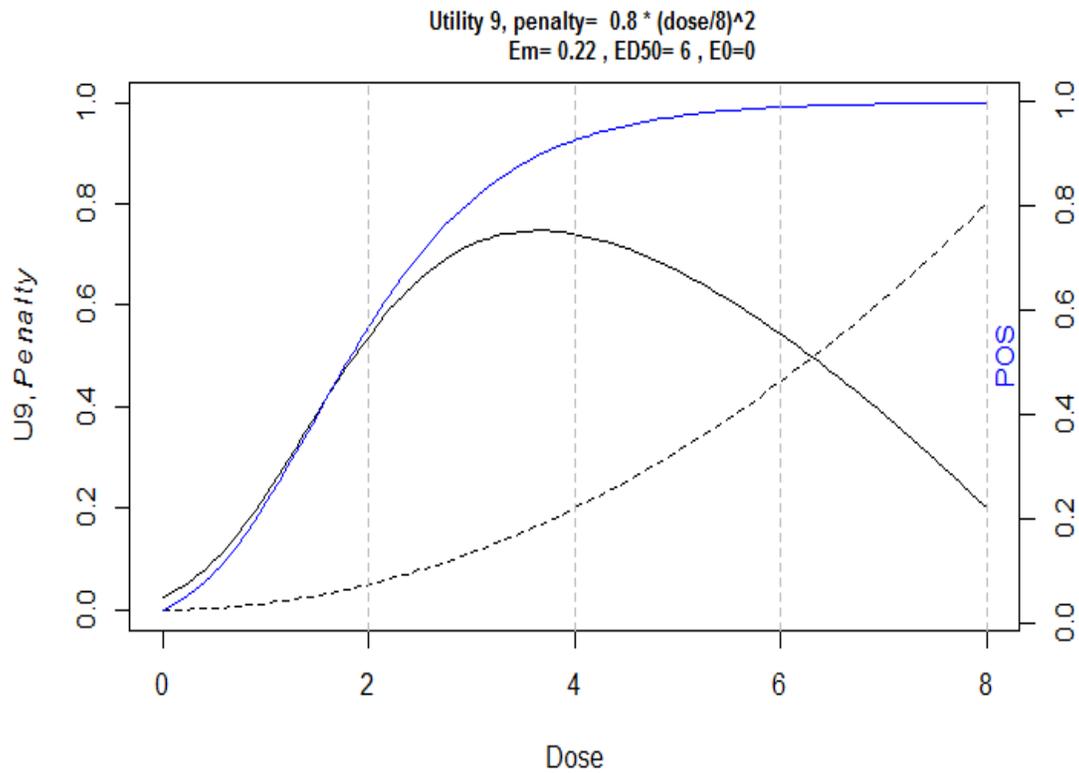


Figure 6.9: Theoretical curves for U9, Sigmoid scenario.

Below are the probability results of selecting each dose, as well as the probability of 'Go', based on 15000 simulations:

- Prob(Go)=0.7352667
- Prob(choosing $d = 2$)=32.13
- Prob(choosing $d = 4$)=40.39
- Prob(choosing $d = 6$)=20.94
- Prob(choosing $d = 8$)=6.54

We can see that the dose choices are good compared to the theoretical utility curve above. Here, $d = 2$ dose is not far from the optimal one, according to the theory.

Below are the comparisons of the Sigmoid and Plateau scenarios, between the non-optimal design (i.e. fixed f between phase II and phase III, I compared different fixed values of f), and the optimal design (optimizing the patients allocation between phase II and phase III).
Reminder: here we are working with a balanced design, so we fix the w and we no longer optimize it, we only optimize f :

Non optimal

	sigmoid	plateau
f=0.10	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=51% doses= 0.44, 0.32, 0.17, 0.07 POS(go)=83% E(U)= 0.31	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=66% doses=0.74, 0.16, 0.07, 0.03 POS(go)=99% E(U)= 0.56
f=0.25	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=74% doses= 0.32, 0.41, 0.21, 0.06 POS(go)=83% E(U)= 0.44	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=86% doses= 0.69, 0.21, 0.08, 0.03 POS(go)=97% E(U)= 0.72
f=0.50	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=86% doses= 0.17, 0.52, 0.25, 0.05 POS(go)=77% E(U)= 0.46	$w = (0.2,0.2,0.2,0.2,0.2)$ Go=94% doses= 0.61, 0.29, 0.08, 0.02 POS(go)=90% E(U)= 0.71

Optimal

sigmoid	plateau
$w = (0.2,0.2,0.2,0.2,0.2)$, f=0.40 Go=83% doses= 0.23, 0.49, 0.23, 0.05 POS(go)=80% E(U)= 0.47	$w = (0.2,0.2,0.2,0.2,0.2)$, f=0.37 Go=92% doses=0.65, 0.25, 0.08, 0.02 POS(go)=95% E(U)=0.74

Table 6.5: Optimal versus non-optimal design for U_9 , where optimal design consists here in optimizing the global patient allocation between phase II and phase III, Sigmoid and Plateau scenarios.

We can conclude that with a three-parameter (E_0 , E_{max} and ED_{50}) without the 'Hill' exponent, parameters are well estimated and correct decisions and dose choices are made.

On the other hand, according to the optimal design above, it is recommended to have more patients in phase II to make a better choice, which amounts to the idea that I tried to prove

previously with $U5$, by increasing the sample size of phase II.

Below is the graph of utility expectations (after the sponsor's choice: Go and dose choice) as a function of f , for the Sigmoid scenario:

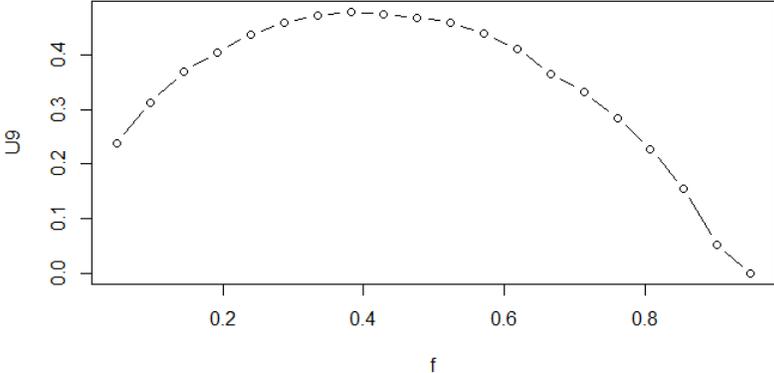


Figure 6.10: Utility expectations as a function of f , Sigmoid scenario.

Below is the graph of utility expectations (after the sponsor's choice: Go and dose choice) as a function of f , for the Plateau scenario:

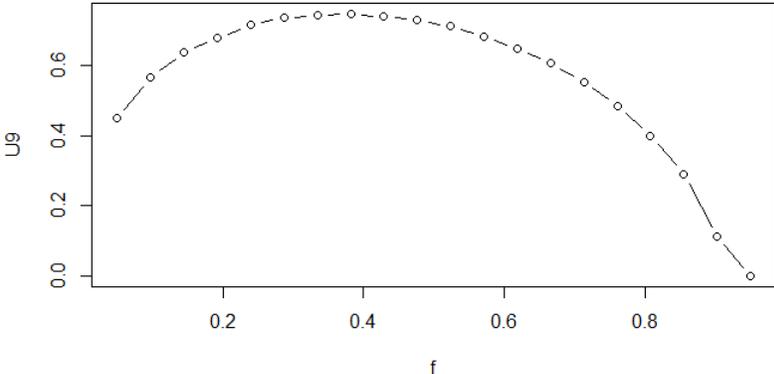


Figure 6.11: Utility expectations as a function of f , Plateau scenario.

A.2. Convergence and autocorrelations

For each MCMC iteration, we have a candidate parameter vector from which the PoS and the toxicity component are computed for each dose. In most applications, we might find that keeping a few thousand iterations is sufficient for reasonably accurate posterior inference. In all our procedures, the relationship between the requested number of iterations, the number of iterations kept, and the amount of batching is as follows:

$$kept = \left\lceil \frac{(requested - burnin)}{batching} \right\rceil, \text{ that is to say, } n = \left\lceil \frac{(n_{iter} - burnin)}{B} \right\rceil$$

where '*burnin*' is the number of iterations to be discarded, and $\lceil \cdot \rceil$ is the rounding operator.

To estimate safety and efficacy data, we generated $n_{iter} = 150000$ MCMC iterations for each study. Among these iterations, we discarded an initial portion of the Markov chain sample so that the effect of initial values on the posterior inference is minimized: $burnin = 150000/2 = 75000$ first iterations. When plotting MCMC iterations, we visually detected autocorrelations until lag 25 at least (with a slow and gradual decrease in autocorrelations). This does not mean that autocorrelations are negligible after lag 25. This diagnostic inference is reliable if the sampler is actually working (has nearly reached equilibrium) and worthless otherwise. Thus a batch length of 100 for instance should be sufficient, but let's use a length of 150 to be safe (i.e. $B = 150$).

Convergence were rechecked afterwards via visual plots.

A.3. Results for low efficacy threshold, $PoS > 0.30$

Sigmoid Scenario Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.30 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $P(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $P(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.607	0.693	0.677	0.736	0.737
Prob(choose(Go))	0.839	0.915	0.895	0.950	0.953
Distribution selected doses (Conditional to 'Go')	0.090 0.840 0.070 0.000	0.030 0.920 0.050 0.000	0.020 0.920 0.060 0.000	0.000 0.960 0.030 0.000	0.000 0.960 0.040 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.050 0.950 0.000 0.000	-	0 1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.020 0.890 0.090 0.000	-	0.010 0.900 0.100 0.000	-
POS (conditional to 'Go')	0.770	0.790	0.796	0.798	0.799
Prob(Stop at interim)	-	0.404	-	0.655	-
% Stop for futility	-	0.008	-	0.020	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	399	-	673	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.646	0.723	0.713	0.758	0.762

Table 6.6: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

Sigmoid Scenario Acceptable safety of all doses Except for d=8 Threshold.eff=0.30 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $P(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $P(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.672	0.734	0.728	0.779	0.784
Prob(choose(Go))	0.853	0.920	0.904	0.963	0.967
Distribution selected doses (Conditional to 'Go')	0.020 0.790 0.180 0.010	0.000 0.780 0.210 0.010	0.000 0.740 0.260 0.000	0.000 0.710 0.290 0.000	0.000 0.680 0.320 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.010 0.990 0.000 0.000	-	0 1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.640 0.350 0.010	-	0.000 0.470 0.530 0.000	-
POS (conditional to 'Go')	0.814	0.825	0.832	0.836	0.840
Prob(Stop at interim)	-	0.373	-	0.438	-
% Stop for futility	-	0.006	-	0.012	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	407	-	781	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.695	0.759	0.752	0.805	0.812

Table 6.7: Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).

A.4. Results for high efficacy threshold, $PoS > 0.90$

Sigmoid scenario Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.90 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.096	0.114	0.127	0.123	0.132
Prob(choose(Go))	0.134	0.150	0.165	0.158	0.168
Distribution selected doses (Conditional to 'Go')	0.170 0.810 0.020 0.000	0.070 0.920 0.010 0.000	0.030 0.950 0.020 0.000	0.010 0.970 0.020 0.000	0.000 0.980 0.020 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.140 0.860 0.000 0.000	-	0.010 0.990 0.000 0.000	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.030 0.960 0.010 0.000	-	0.020 0.920 0.060 0.000	-
POS (conditional to 'Go')	0.732	0.768	0.787	0.793	0.797
Prob(Stop at interim)	-	0.404	-	0.655	-
% Stop for futility	-	0.347	-	0.545	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	399	-	673	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.098	0.115	0.130	0.125	0.134

Table 6.8: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

Sigmoid scenario Acceptable safety of all doses Except for d=8 Threshold.eff=0.90 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.125	0.132	0.155	0.176	0.164
Prob(choose(Go))	0.162	0.167	0.193	0.219	0.203
Distribution selected doses (Conditional to 'Go')	0.060 0.850 0.080 0.010	0.010 0.920 0.070 0.010	0.000 0.840 0.160 0.000	0.000 0.820 0.180 0.000	0.000 0.730 0.270 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.010 0.990 0.000 0.000	-	0 1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.860 0.120 0.010	-	0.000 0.580 0.420 0.000	-
POS (conditional to 'Go')	0.784	0.804	0.817	0.820	0.832
Prob(Stop at interim)	-	0.373	-	0.438	-
% Stop for futility	-	0.302	-	0.315	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	407	-	781	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.127	0.134	0.158	0.180	0.169

Table 6.9: Simulation results, Sigmoid scenario with acceptable safety scenario (except for highest dose).

A.5. Boxplots related to the posterior distribution

Boxplots of posterior means of utilities: Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15)

In the following, we represent the boxplots of posterior means of utilities, for the 1000 simulated phase II studies, for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), and for several values of N_2 . The black line represents the median, and the golden triangle represents the true (theoretical) utility value.

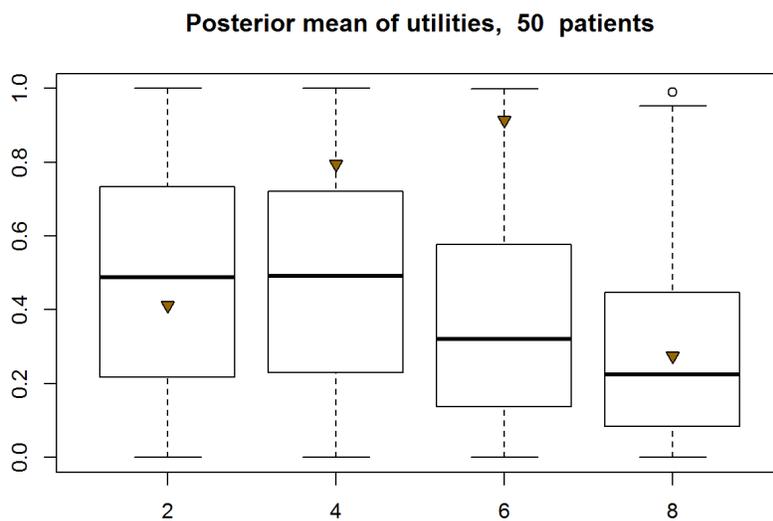


Figure 6.12: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.

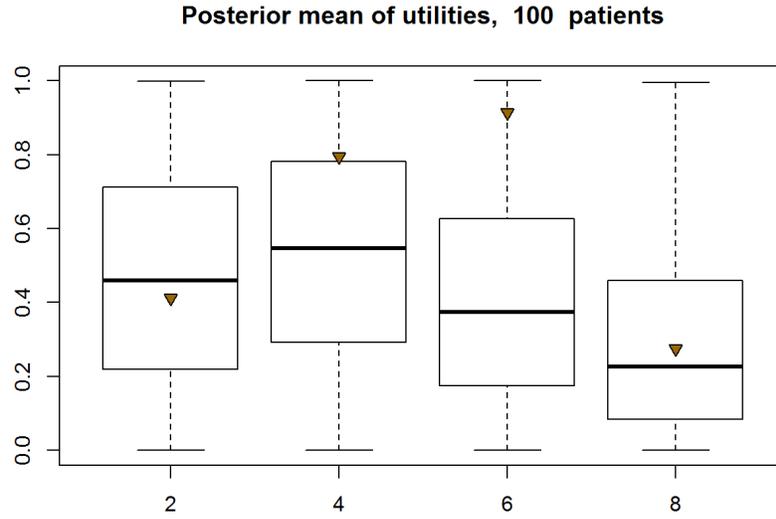


Figure 6.13: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.

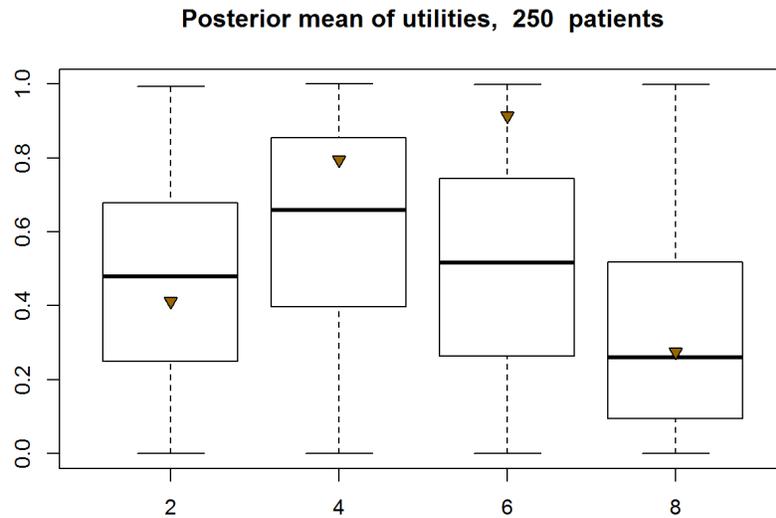


Figure 6.14: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.

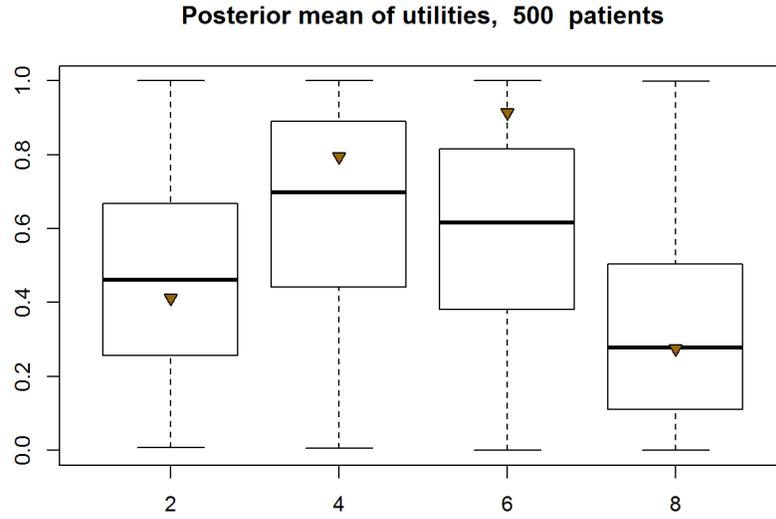


Figure 6.15: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.

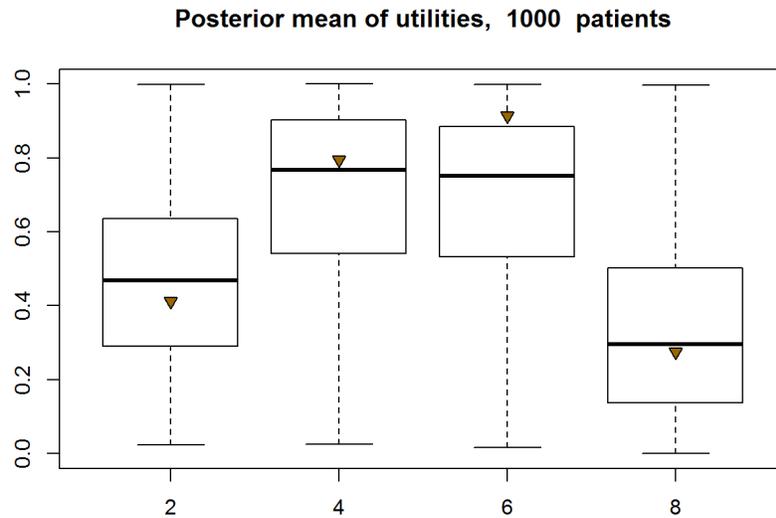


Figure 6.16: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.

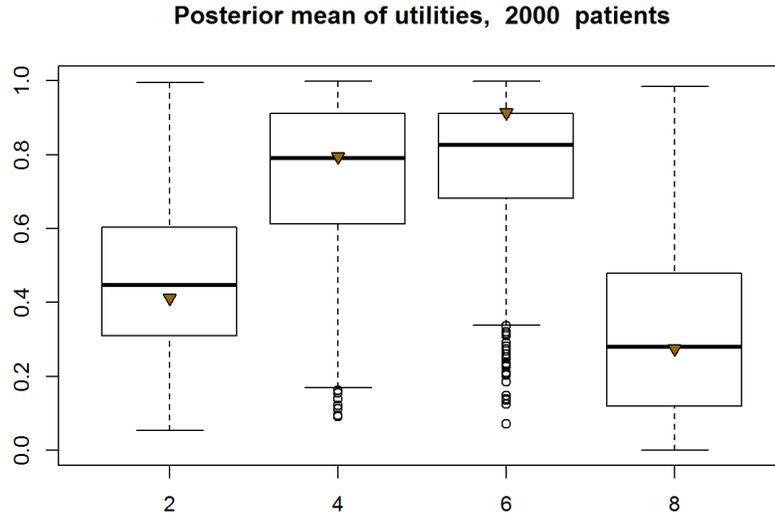


Figure 6.17: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$.

Concerning the posterior mean of the utilities, and for the smallest sample size, $N_2 = 50$ patients, the utility values of the two middle doses appear quite underestimated. As expected, the quality of the estimators improves as the sample size grows. But even for a phase II study with 1000 patients, the distribution of the estimates of the middle doses are still quite close, whereas the true value of the utility of $d = 6$ is slightly larger than for $d = 4$. This underestimation of the utility value for $d = 6$ can be due to a less stable estimation of the safety profile, which can be explained by a low probability of toxicity for the doses ≤ 6 , leading to an inaccurate estimation of the slope.

Boxplots of posterior medians of utilities: Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15)

In the following, we represent the boxplots of posterior medians of utilities, for the 1000 simulated phase II studies, and for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15).

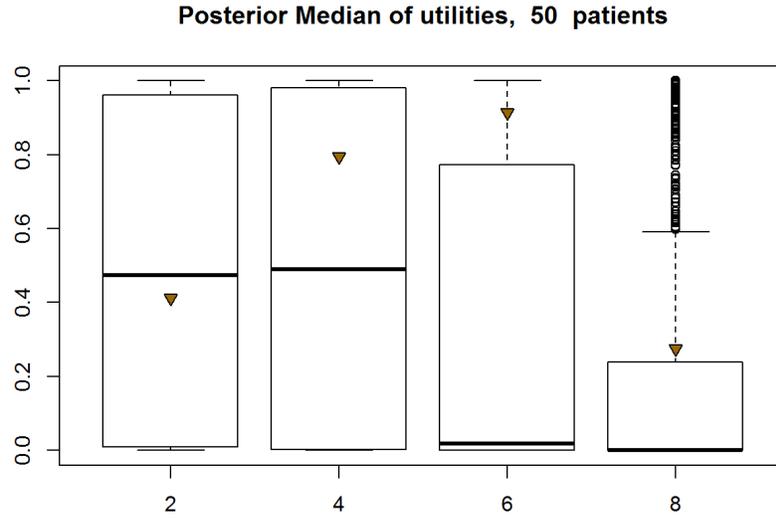


Figure 6.18: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 50$.

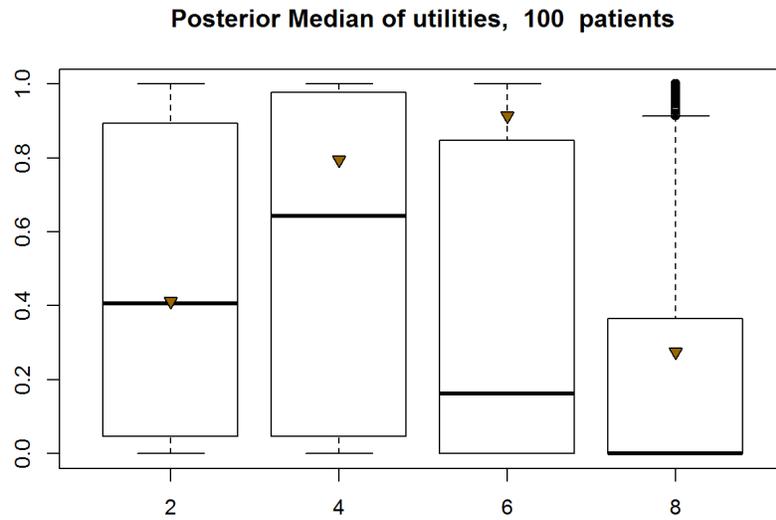


Figure 6.19: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 100$.

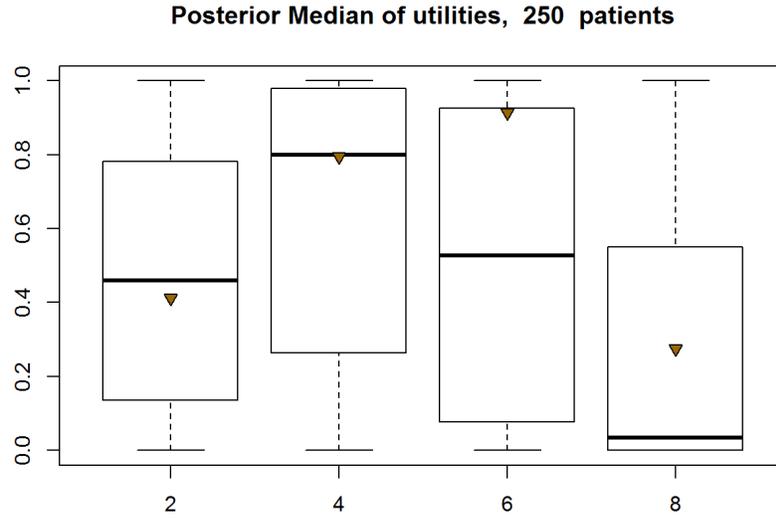


Figure 6.20: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 250$.

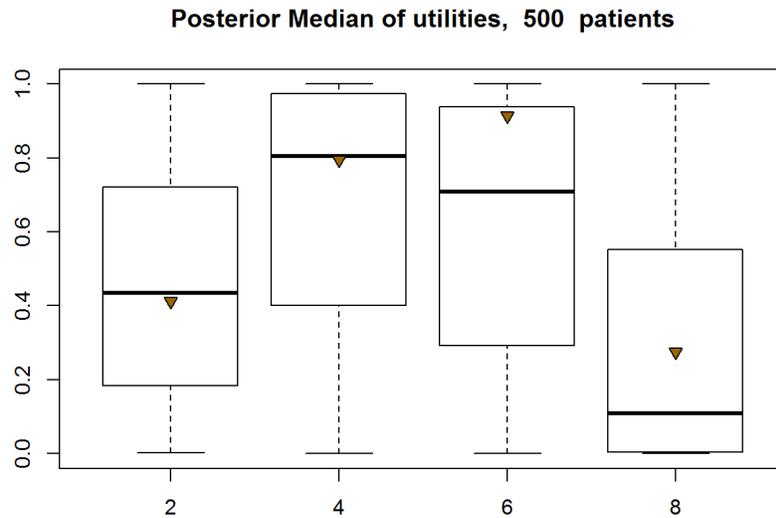


Figure 6.21: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 500$.

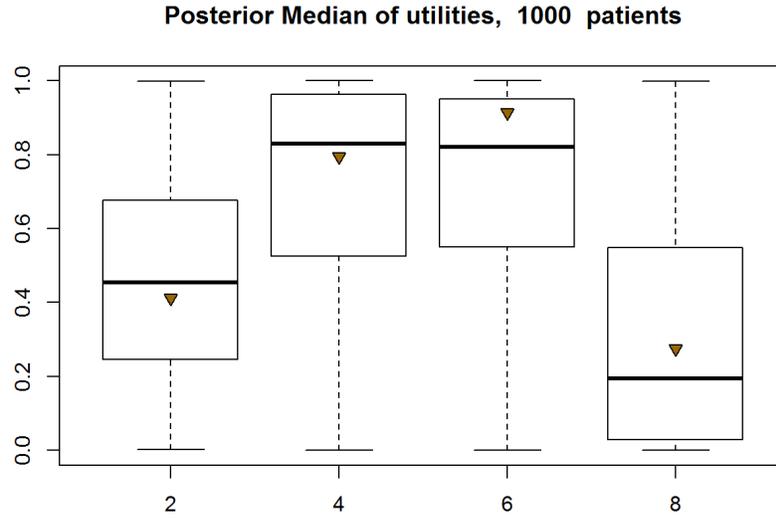


Figure 6.22: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 1000$.

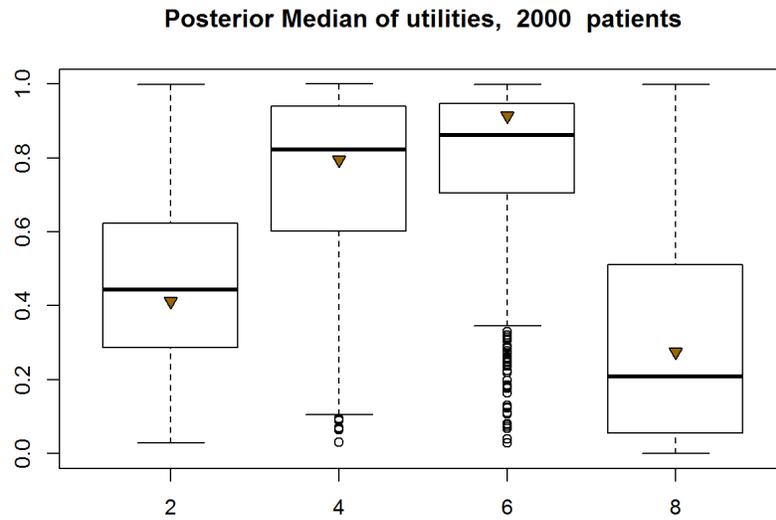


Figure 6.23: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.15), $N_2 = 2000$.

If we consider now the distribution of the posterior median, it appears that it has a larger range as compared to the posterior mean distribution: the posterior median showing clearly more variability than the posterior mean. But in terms of median of the distribution, the posterior mean and the posterior distribution median look similar. Both estimators have difficulty to capture the utility value of the $d = 6$: the posterior median appears wrongly more accurate due to its increased variability that makes the true value appear within the interquartile interval of the distribution, but in terms of median values, both estimators (posterior mean and posterior median) show a similar bias, even for the phase II study of 1000 patients. According to the boxplots, it takes 2000 patients to properly estimate the dose $d = 6$ with both the posterior mean and median, which is not very realistic according to real life conditions.

Conclusion

We conclude from these simulations, based on this very challenging scenario combination, that an accurate estimation of the utility values and a correct ranking of the doses can be very demanding in terms of sample size. Estimating the utility values is not the primary goal of the sponsor's analysis (the primary objective is to rank the doses with respect to the utility function), but the posterior mean appears to be a better estimate, the posterior median showing an inflated variability.

Boxplots of posterior means of utilities: Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

In the following, we represent the boxplots of posterior means of utilities, for the 1000 simulated phase II studies, for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), and for several values of N_2 .

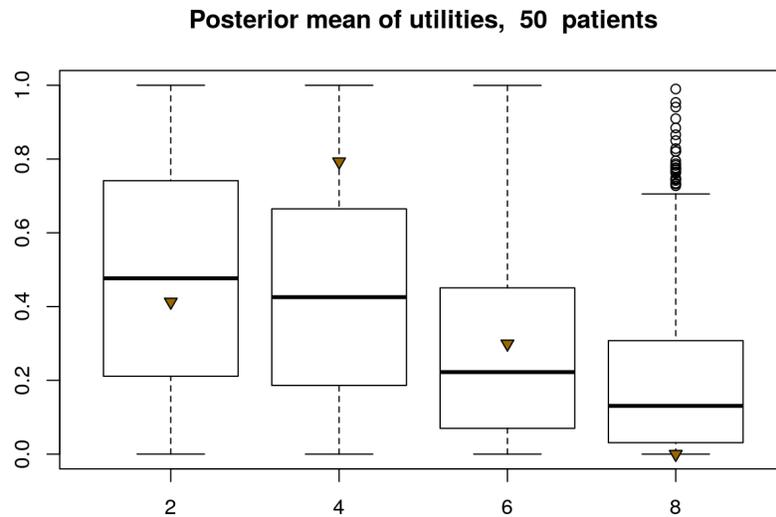


Figure 6.24: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$.

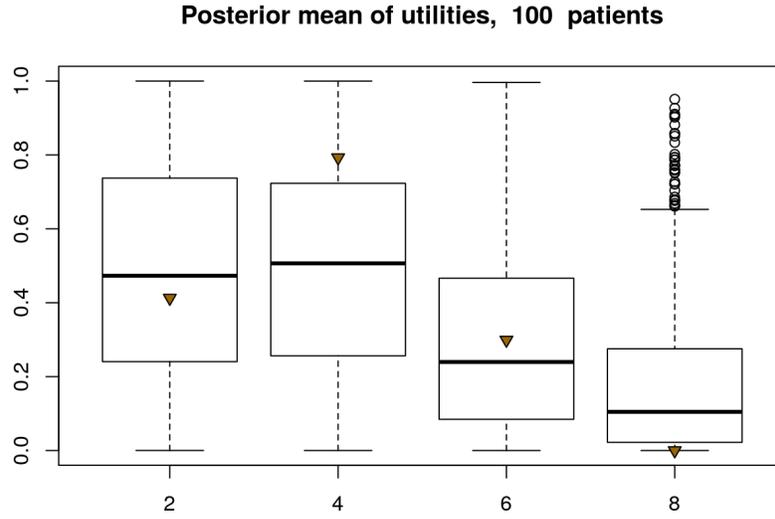


Figure 6.25: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$.

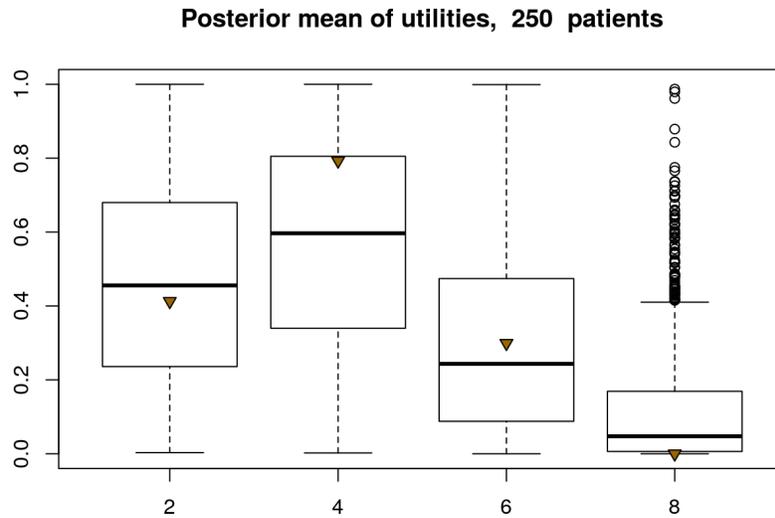


Figure 6.26: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$.

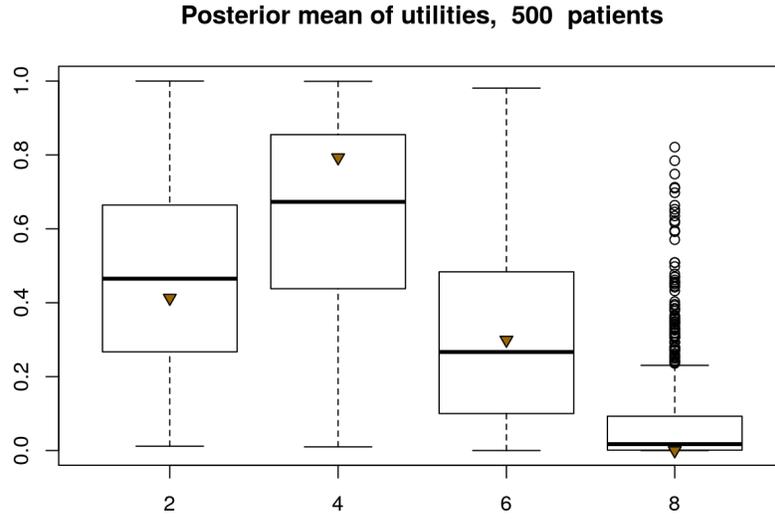


Figure 6.27: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$.

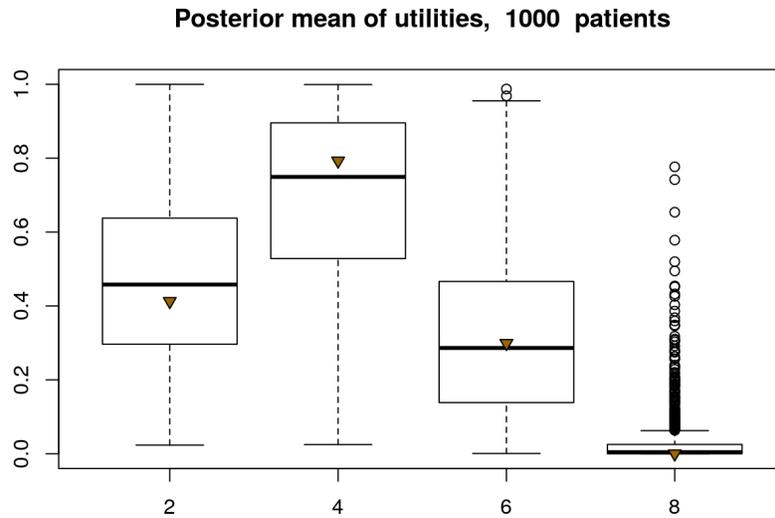


Figure 6.28: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$.

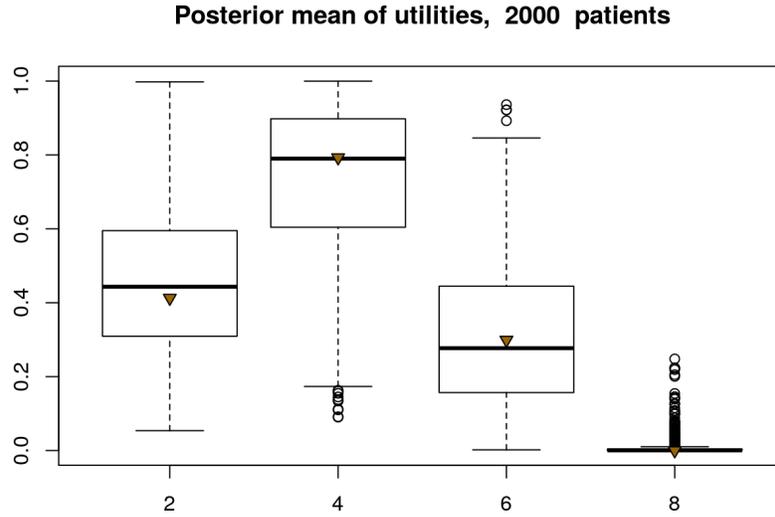


Figure 6.29: Boxplots of posterior means of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$.

Boxplots of posterior medians of utilities: Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20)

In the following, we represent the boxplots of posterior medians of utilities, for the 1000 simulated phase II studies, and for the Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

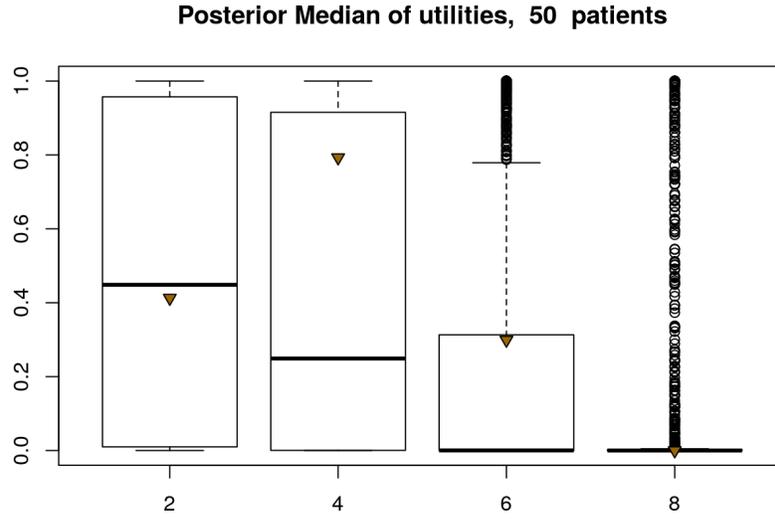


Figure 6.30: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 50$.

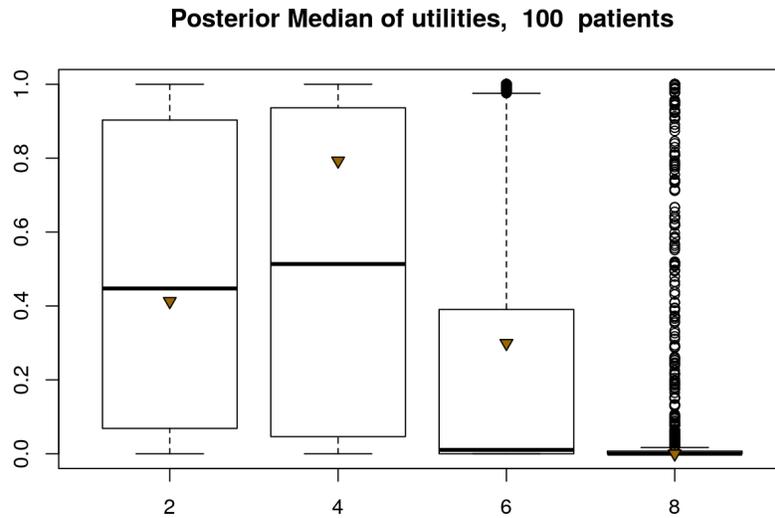


Figure 6.31: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 100$.

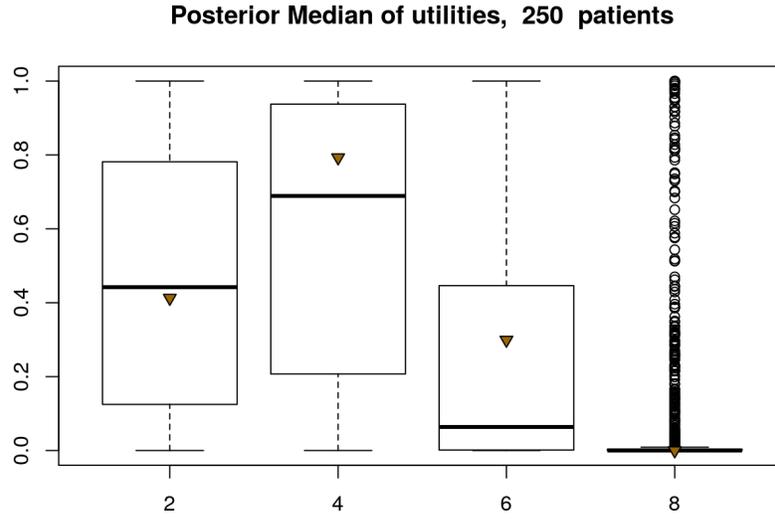


Figure 6.32: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 250$.

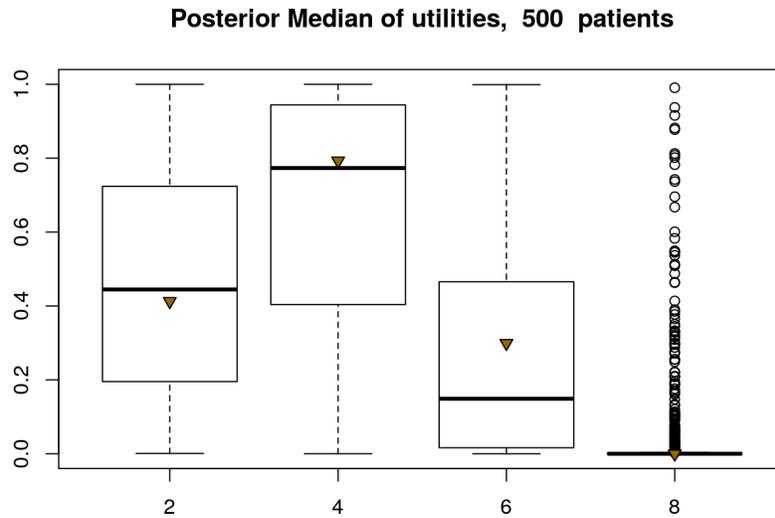


Figure 6.33: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 500$.

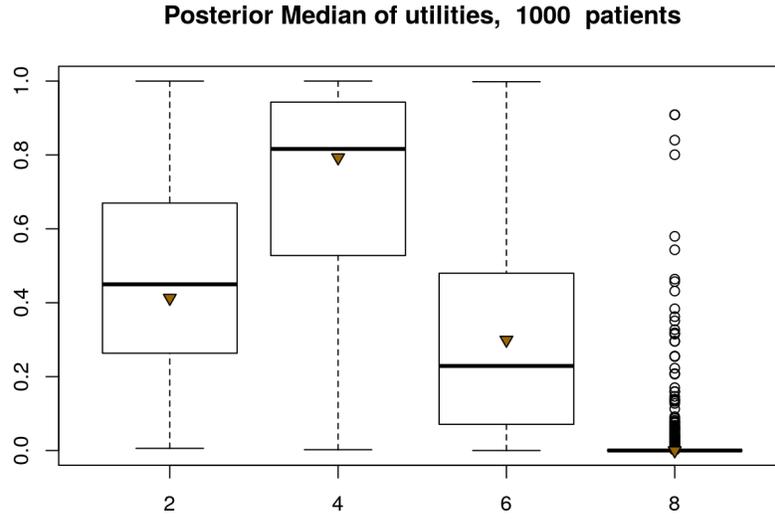


Figure 6.34: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 1000$.

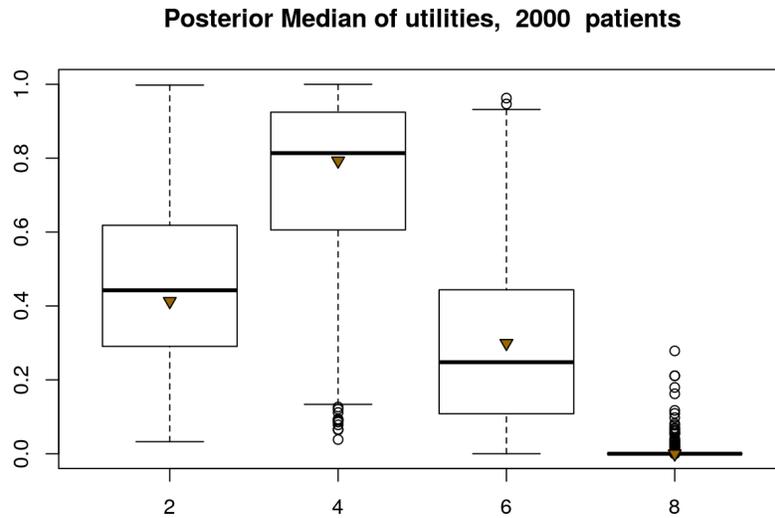


Figure 6.35: Boxplots of posterior medians of utilities, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $N_2 = 2000$.

A.6. Results for U2

In the following results, optimisations of patient allocation to doses and global allocation ratio between phase II and phase III are conducted simultaneously (w and f are optimized at the same time).

Non optimal

sigmoid	plateau
$w = (0.2, 0.2, 0.2, 0.2, 0.2), f = 0.25$	$w = (0.2, 0.2, 0.2, 0.2, 0.2), f = 0.25$
$Go = 88\%$	$Go = 95\%$
$doses = 0.29, 0.68, 0.02, 0.00$	$doses = 1.00, 0.00, 0.00, 0.00$
$POS(go) = 75\%$	$POS(go) = 100\%$
$E(U) = 4370.71$	$E(U) = 3761.01$

Optimal

sigmoid	plateau
$w = (0.15, 0.04, 0.30, 0.37, 0.14), f = 0.65$	$w = (0.08, 0.15, 0.33, 0.01, 0.43), f = 0.49$
$Go = 97\%$	$Go = 99\%$
$doses = 0.04, 0.94, 0.02, 0.00$	$doses = 1.00, 0.00, 0.00, 0.00$
$POS(go) = 95\%$	$POS(go) = 100\%$
$E(U) = 6502.97$	$E(U) = 3972.67$

Table 6.10: Optimal versus non-optimal design for $U2$, where optimal design consists here in optimizing the patient allocation to doses and the global patient allocation between phase II and phase III simultaneously, Sigmoid and Plateau scenarios.

Table 6.10 is a typical example highlighting the crucial importance of phase II sample size, where the optimal design is allocating more patients in phase II than in phase III, which is not quite realistic in real life.

A.7. Sensitivity analysis

Prior dose-response distribution plots

Plots of prior dose-response distributions are given in the following Figure 6.36.

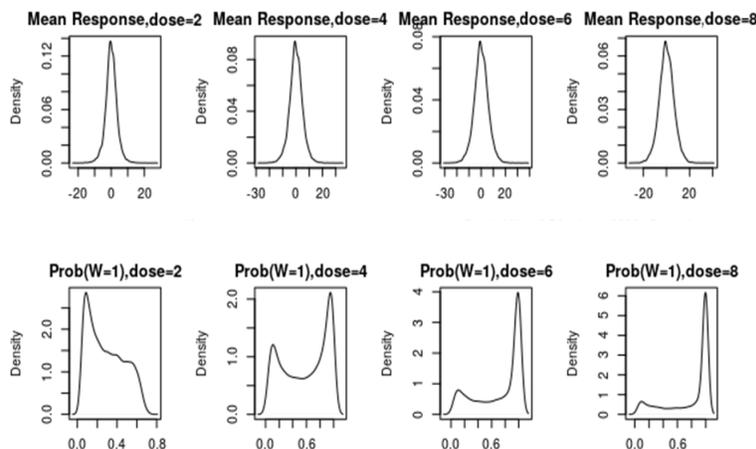


Figure 6.36: Prior dose-response distribution plots.

According to the chosen priors of Section 4.1.3.1, we can see that efficacy mean dose-response is concentrated around 0, and this is due to the non-informative prior chosen for E_{max} parameter. On the other hand, the toxicity probability increases with the dose, and this is due to the informative/conservative prior chosen for the slope b of the Probit model (the choice of this prior was motivated by a conservative approach, assuming that the incidence of toxicity was necessarily increasing with the dose).

Prior choices

In the following Table 6.11, we considered the following non-informative priors for the efficacy and toxicity model parameters (instead of informative priors as the ones considered in Section 4.1.3.1).

For efficacy: $E_0 \sim N(0, 100)$, $ED_{50} \sim U(0, 20)$, and $E_{max} \sim N(0, 100)$. We think that a log-normal distribution is not adapted for ED_{50} because it imposes a mode, i.e. a ED_{50} with a higher probability; therefore a Uniform distribution better represents uncertainty and we propose an upper bound of 20 because we think that it is not realistic to go far beyond this value.

For safety: $a \sim N(0, 1)$ (Uniform distribution for the placebo toxicity) and $b \sim N(0, 1)$.

In Section 4.1.3.1, we considered a positive Uniform distribution on b (conservative prior); we consider here a symmetrical distribution (allowing the product to have fewer adverse events than the placebo).

Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.361	0.425	0.428	0.466	0.489
Prob(choose(Go))	0.509	0.576	0.574	0.614	0.637
Distribution selected doses (Conditional to 'Go')	0.09 0.820 0.060 0.020	0.050 0.880 0.070 0.000	0.030 0.900 0.070 0.000	0.010 0.930 0.060 0.000	0.000 0.950 0.050 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.080 0.910 0.010 0.000	-	0.010 0.980 0.010 0.000	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.020 0.860 0.120 0.000	-	0.020 0.840 0.140 0.000	-
POS (conditional to 'Go')	0.771	0.787	0.793	0.799	0.800
Prob(Stop at interim)	-	0.410	-	0.613	-
% Stop for futility	-	0.162	-	0.227	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	398	-	694	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.392	0.453	0.455	0.491	0.510

Table 6.11: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

Sensitivity analyses conducted in Table 6.11 aimed at examining the performance of the designs with respect to different priors (by considering non informative priors for all model parameters). Results are promising as they are consistent with the ones obtained with the chosen informative priors (except for E_{max} parameter) in Section 4.1.3.1, but require more patients to reach similar properties and decision rule qualities.

Parameters h and k

In Table 6.12, we considered $k = 1$ instead of $k = 2$ (still with $h = 1$ and with the initial priors of Section 4.1.3.1).

Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.421	0.482	0.487	0.537	0.552
Prob(choose(Go))	0.562	0.631	0.632	0.694	0.712
Distribution selected doses (Conditional to 'Go')	0.070 0.860 0.070 0.000	0.020 0.900 0.090 0.000	0.010 0.920 0.070 0.000	0.000 0.930 0.070 0.000	0.000 0.930 0.070 0.000
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.030 0.970 0.000 0.000	-	0 1 0 0	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.010 0.840 0.160 0.000	-	0.000 0.810 0.190 0.000	-
POS (conditional to 'Go')	0.777	0.800	0.801	0.805	0.804
Prob(Stop at interim)	-	0.424	-	0.582	-
% Stop for futility	-	0.133	-	0.157	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	394	-	709	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.437	0.505	0.506	0.559	0.572

Table 6.12: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20).

As expected, with a lower k value (i.e. less constraint on safety), expected utilities, powers, optimal dose selection and interim analysis stopping frequencies, 'Go' decisions, and PoSs are globally higher compared to the ones obtained with $k = 2$, as we are less demanding in terms of safety control.

The choice of threshold t

In Figure 6.37 and Table 6.13, we considered $t = 0.30$ instead of $t = 0.15$ (with the initial h and k values, $h = 1$ and $k = 2$). The highest dose $d = 8$ is now the optimal one with this new threshold (see Figure 8).

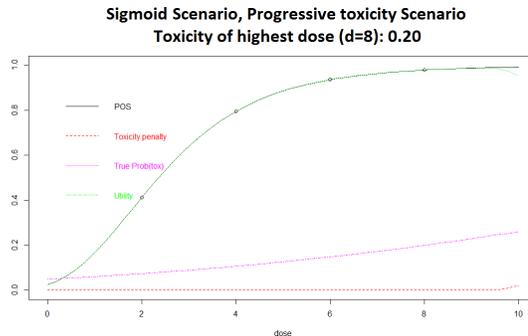


Figure 6.37: Theoretical curves, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $t = 0.30$.

Sigmoid Progressive toxicity scenario Tox(d=8)=0.20 Threshold.eff=0.60 Threshold.safe=0.50	N2=250	N2=500 interim at N2'=250 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=500	N2=1000 interim at N2'=500 ; stop if : $\mathbb{P}(\text{best dose} \text{data}) \geq 0.8$	N2=1000
E(U)	0.637	0.708	0.738	0.782	0.840
Prob(choose(Go))	0.663	0.730	0.761	0.802	0.861
Distribution selected doses (Conditional to 'Go')	0.000 0.010 0.380 0.610	0.000 0.000 0.200 0.800	0.000 0.000 0.180 0.820	0.000 0.000 0.070 0.930	0.000 0.000 0.050 0.950
Distribution selected doses (Conditional to 'Go') at interim analysis	-	0.000 0.000 0.190 0.810	-	0.000 0.000 0.090 0.910	-
Distribution selected doses (Conditional to 'Go') at final analysis	-	0.000 0.000 0.200 0.800	-	0.000 0.000 0.060 0.940	-
POS (conditional to 'Go')	0.961	0.970	0.970	0.975	0.976
Prob(Stop at interim)	-	0.361	-	0.596	-
% Stop for futility	-	0.102	-	0.143	-
Mean(N2)=N2x(1-prob(interim))+N2'xprob(interim)	-	410	-	702	-
Power=Prob(choose(Go))xPOS(conditional to 'Go')	0.637	0.708	0.738	0.782	0.840

Table 6.13: Simulation results, Sigmoid scenario with progressive toxicity scenario (and toxicity of highest dose = 0.20), $t = 0.30$.

With a higher t threshold, we are much less restrictive on safety issues, so we choose the highest dose. Simulation results are consistent with this new utility: according to Table 6.13, $d = 8$ is chosen with a significant probability even with the smallest phase II sample size, this probability reaches 95% with the largest phase II study. The expected utility, the probability of making the good decision to go to phase III, the PoS and the power are considerably high as compared to the previous results with $t = 0.15$. The interim analysis conducted when half of the planned 500 and 1000 patients are enrolled is useful as the probability to stop the trial is relatively high. As a consequence, this mid-term data inspection can be considered as an economic strategy, saving a large number of patients, accordingly, saving budget and time.