Dose Escalation using a Bayesian Model: Rational decision rules

Helene Thygesen

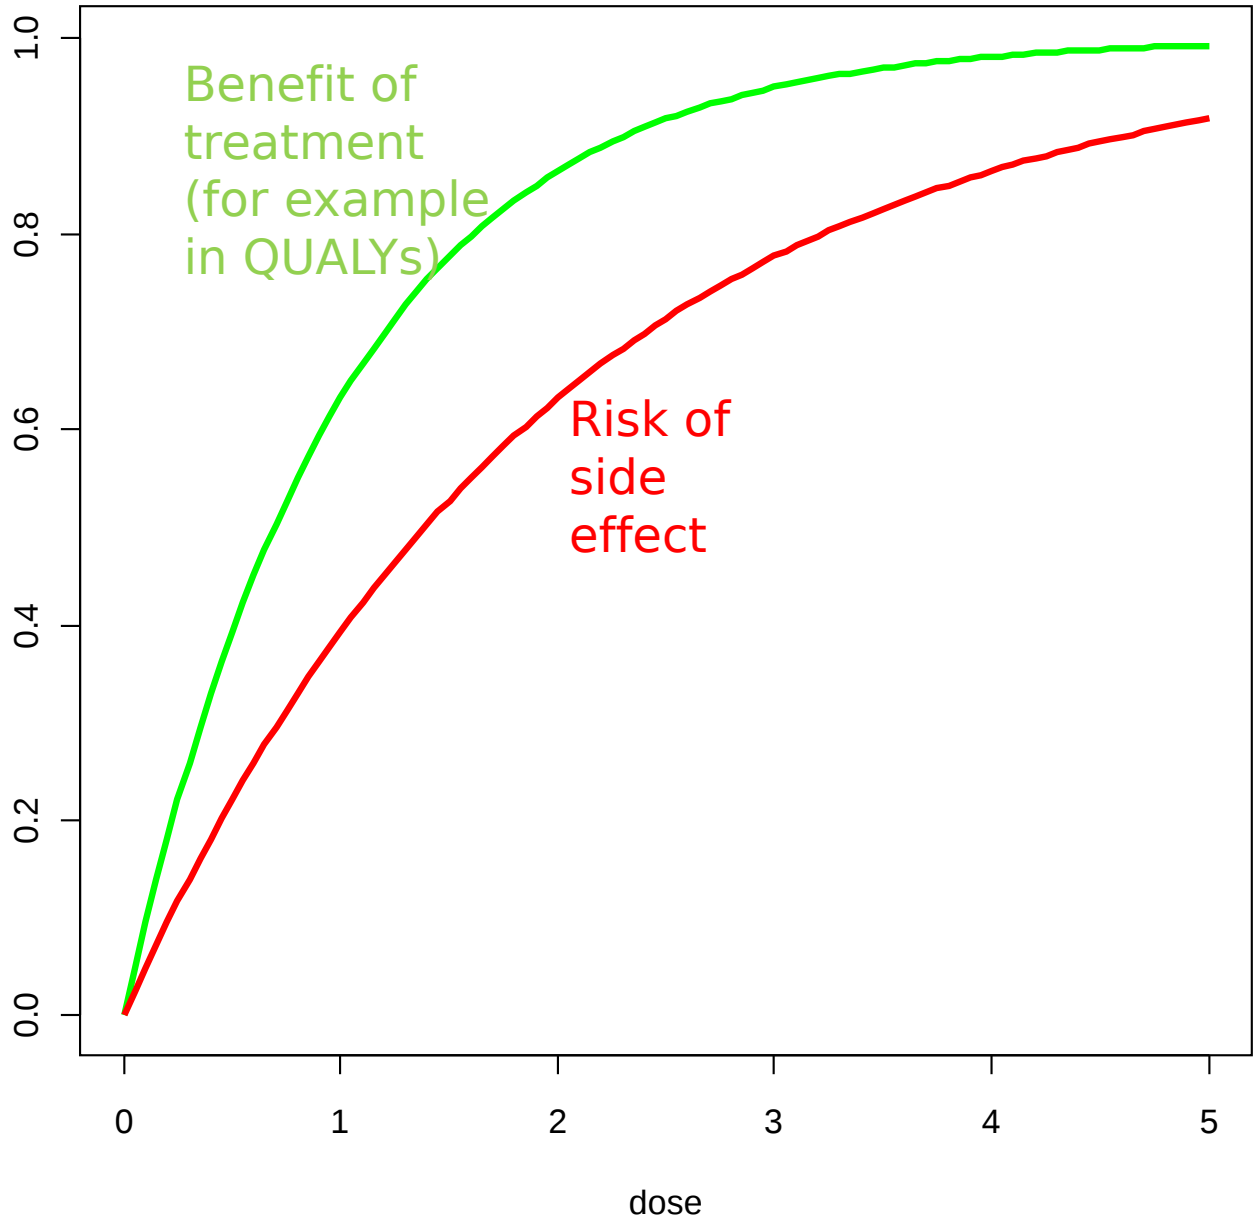
Benefit of
treatment
(for example
in QUALYs)
Risk of
side
effect
0.0
0.2
0.4
0.6
0.8
1.0
0
1
2
3
4
5
dose

If we decide on a constant c we can write utility= Benefit – c SideEffectRisk, and find the dose that optimizes the utility:

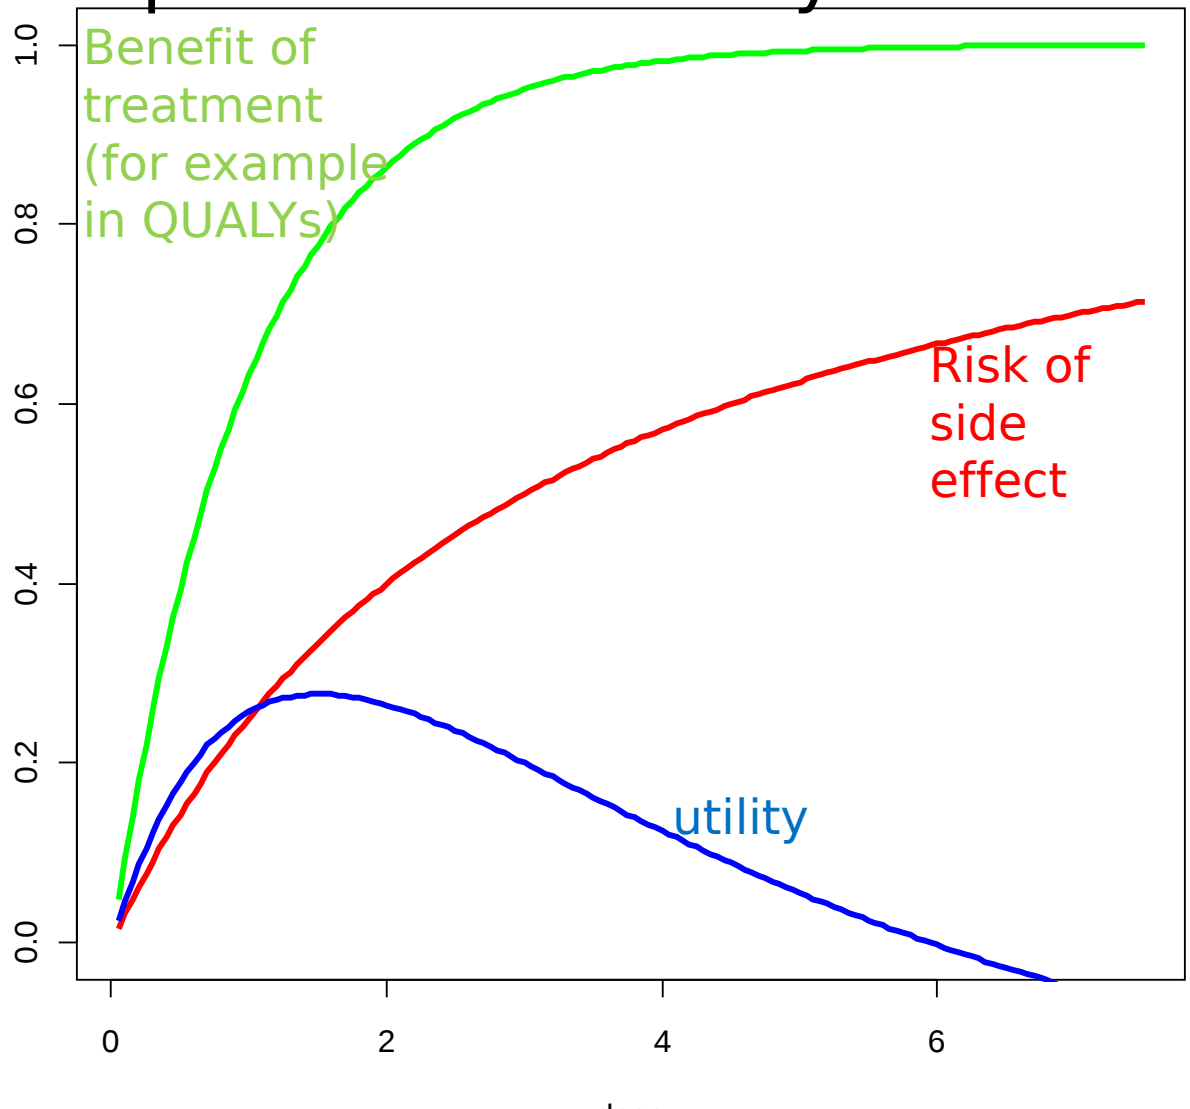

Assume:
Benefit = $1-e^{-x}$
SideEffectRisk = $1-e^{-\lambda x}$
x= dose level

We will conduct an experiment aimed at identifying λ and thereby the optimal dose level.

Start

Postulate prior belief w.r.t. λ

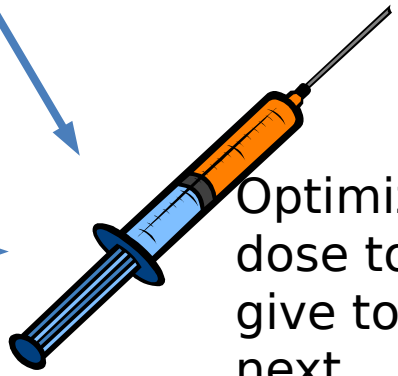

Optimize dose to give to next patient

If the optimal decision is to halt the trial, do so

Update belief:
$P_{post}=P_{pri}L/\Sigma P_{pri}L$

Side effect or not?

Which dose to give to the first patient?

Dose Escalation Wizard: Optimize dose for patient 1

Benefit of dose x:	1-exp(-x)
Side effect risk of dose x:	1-exp(-λx)
Prior belief w.r.t. λ:	Gamma(α,β)
α:	1.0
β:	3.0
Costs of side effect in benefit units:	1.5

OK

Dose Escalation Wizard: Optimize dose for patient 1

Benefit of dose x:	1-exp(-x)
Side effect risk:	1-exp(-λx)
Prior belief w.r.t. λ:	Gamma(α,β)
α:	1.0
β:	3.0
Costs of side effect in benefit units:	1.5

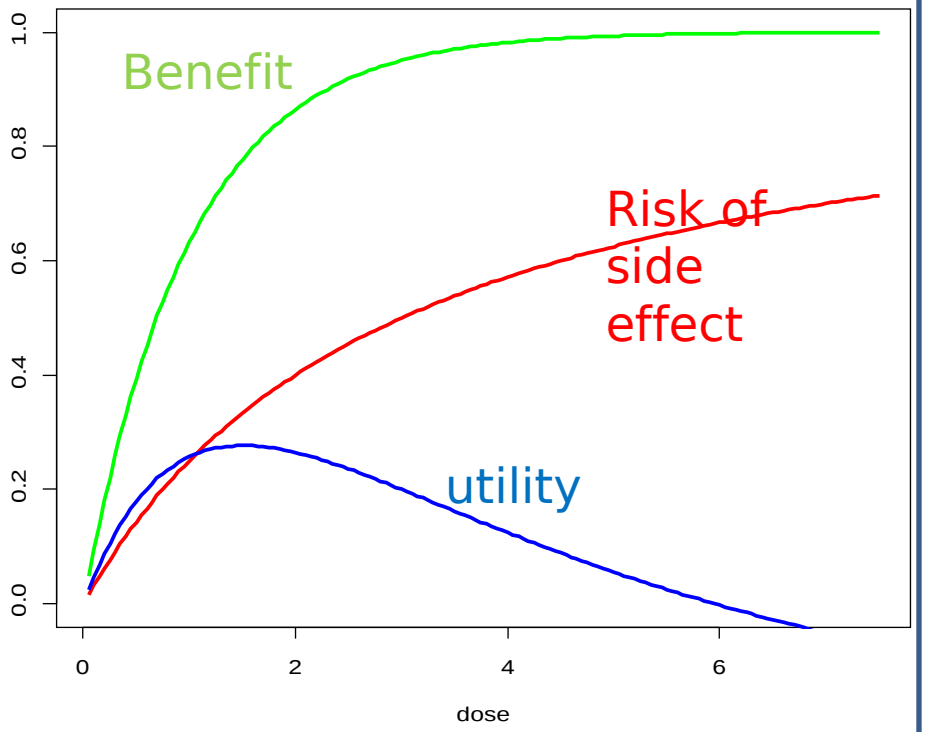

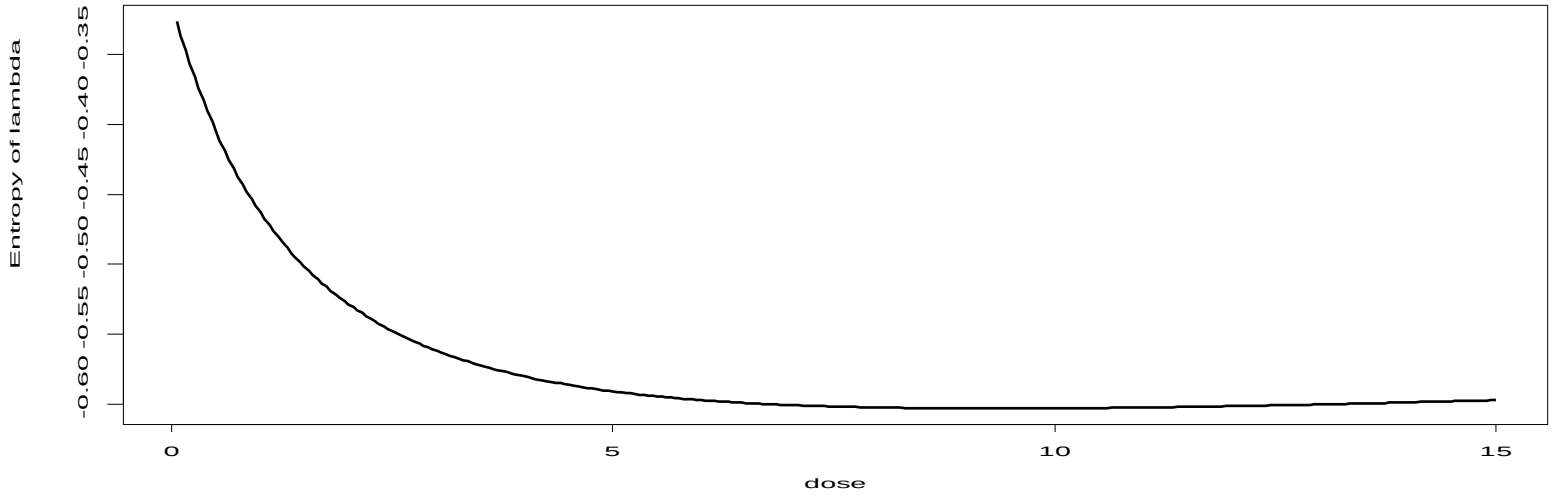

How much human suffering is equivalent to 1 bit of information?

Ethical invoice for experimental use of dose x
(side effect risk = ∫1-exp(-λx)db(λ)) in one patient:

Text	#	Unit utility		Total item utility
Benefit to trial patient:	1	1-exp(x)	+	1-exp(x)
Side effect in trial patient:	1	c∫1-exp(-λx)db(λ)	-	c∫1-exp(-λx)db(λ)
Improvement of utility for future patients in clinical practice, who get a better dose as a result of the experiment:	100	y		+ 100y
				---
Total:				Total utility

(I should have put utility for subsequent trial patients on the bill, too)

Calculating the expected utility to future patients (clinical practice)

Suppose, for example, we give the trial patient a dose x=2. Assume (for the sake of the argument) that the drug has to go into clinical practice after that single experiment!

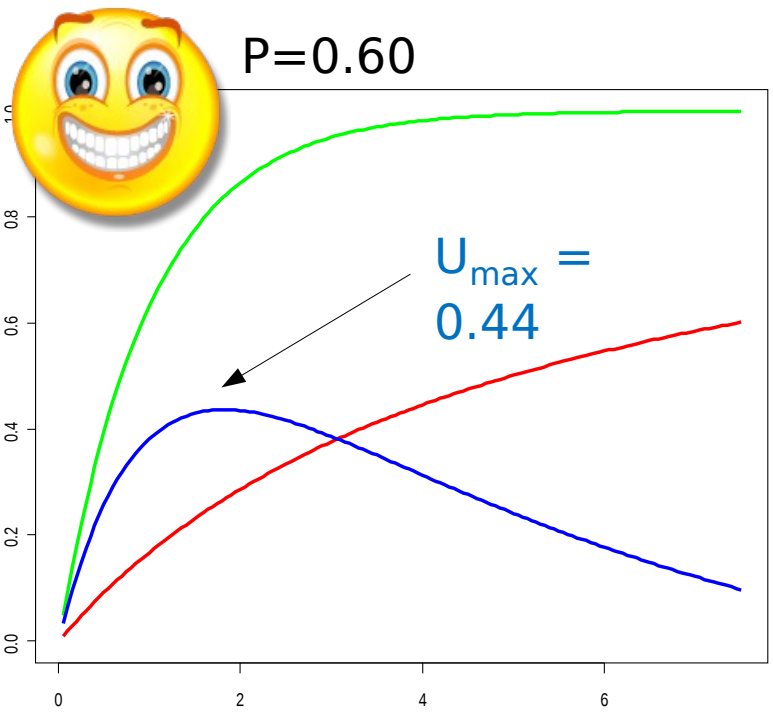

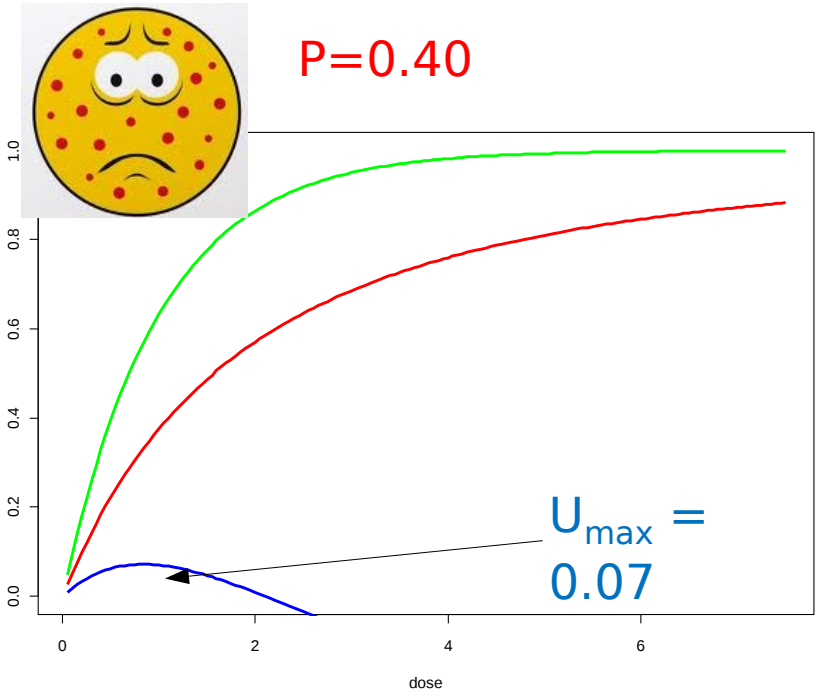

So the expected utility of the drug to patients in clinical practice is 0.60*0.44 + 0.40*0.07 = 0.29 if we give the trial patient the dose x=2.

Taking future clinical practice into

Dose Escalation Wizard: Optimize dose for patient 1

	Trial ethics:	Clinical practice utility:
Benefit of dose x:	1-exp(-x)	1-exp(-x)
Side effect risk of dose x:	1-exp(-λx)	
Prior belief w.r.t. λ:	Gamma(α,β)	
α:	1.0	
β:	3.0	
Costs of side effect in benefit units:	1.5	1.5
#patients affected:	1	100

OK

Assuming the first trial patient is the last one

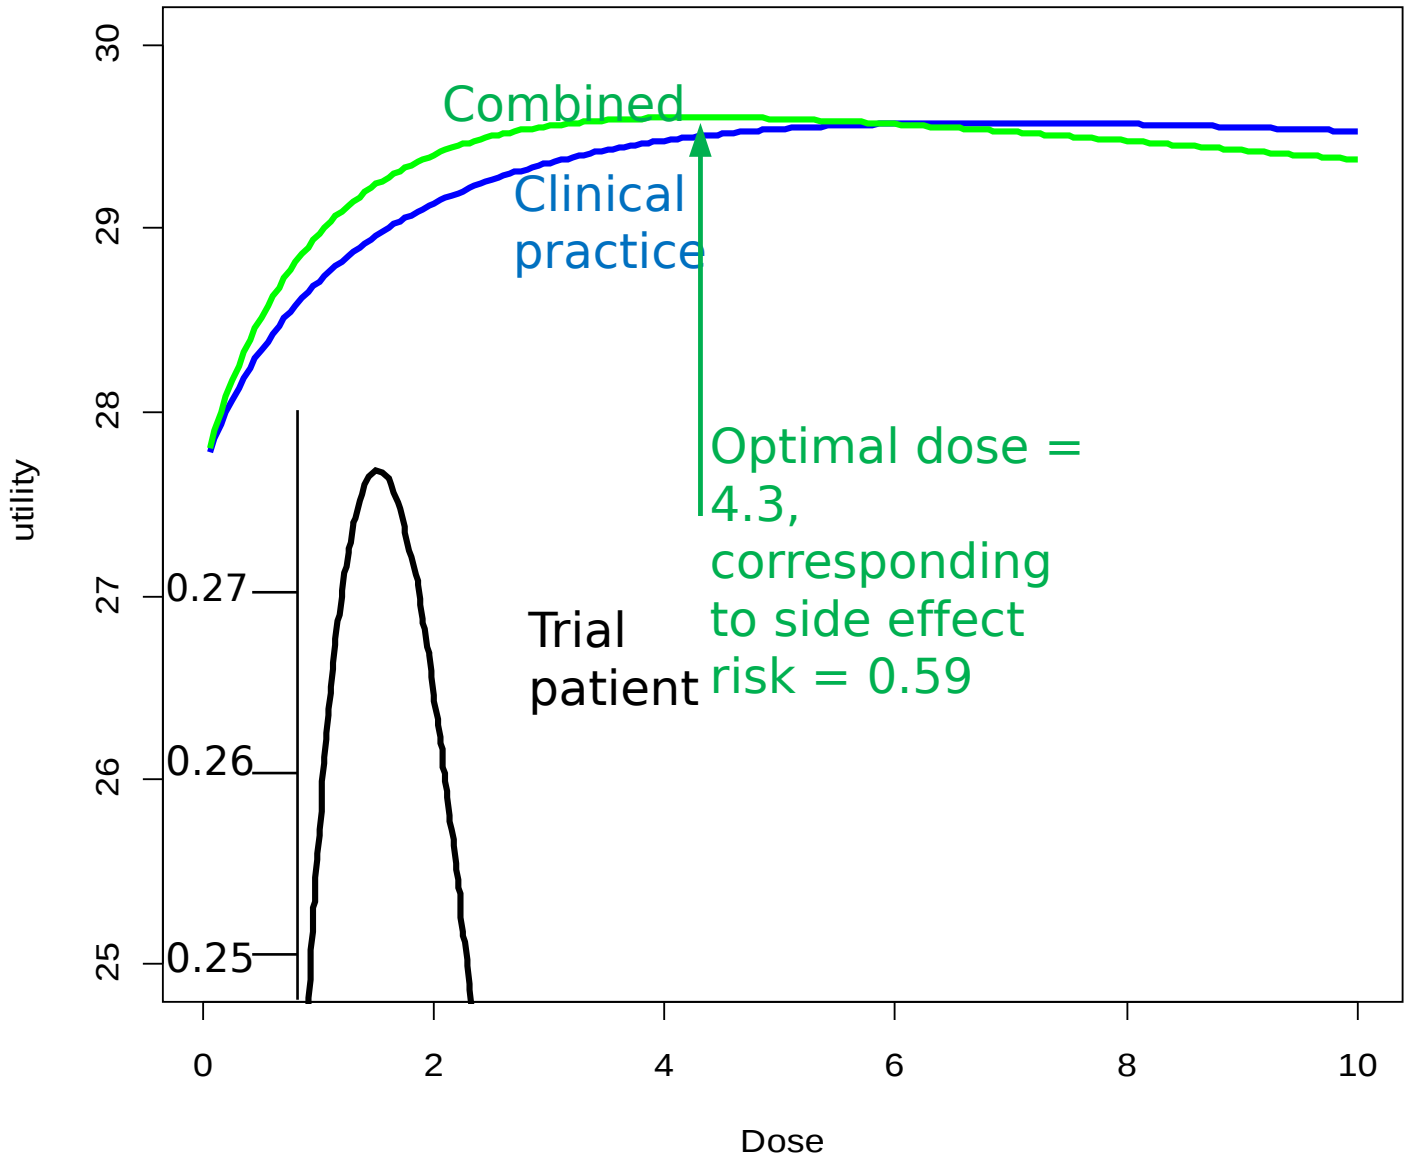

Assuming the first trial patient is the penultimate

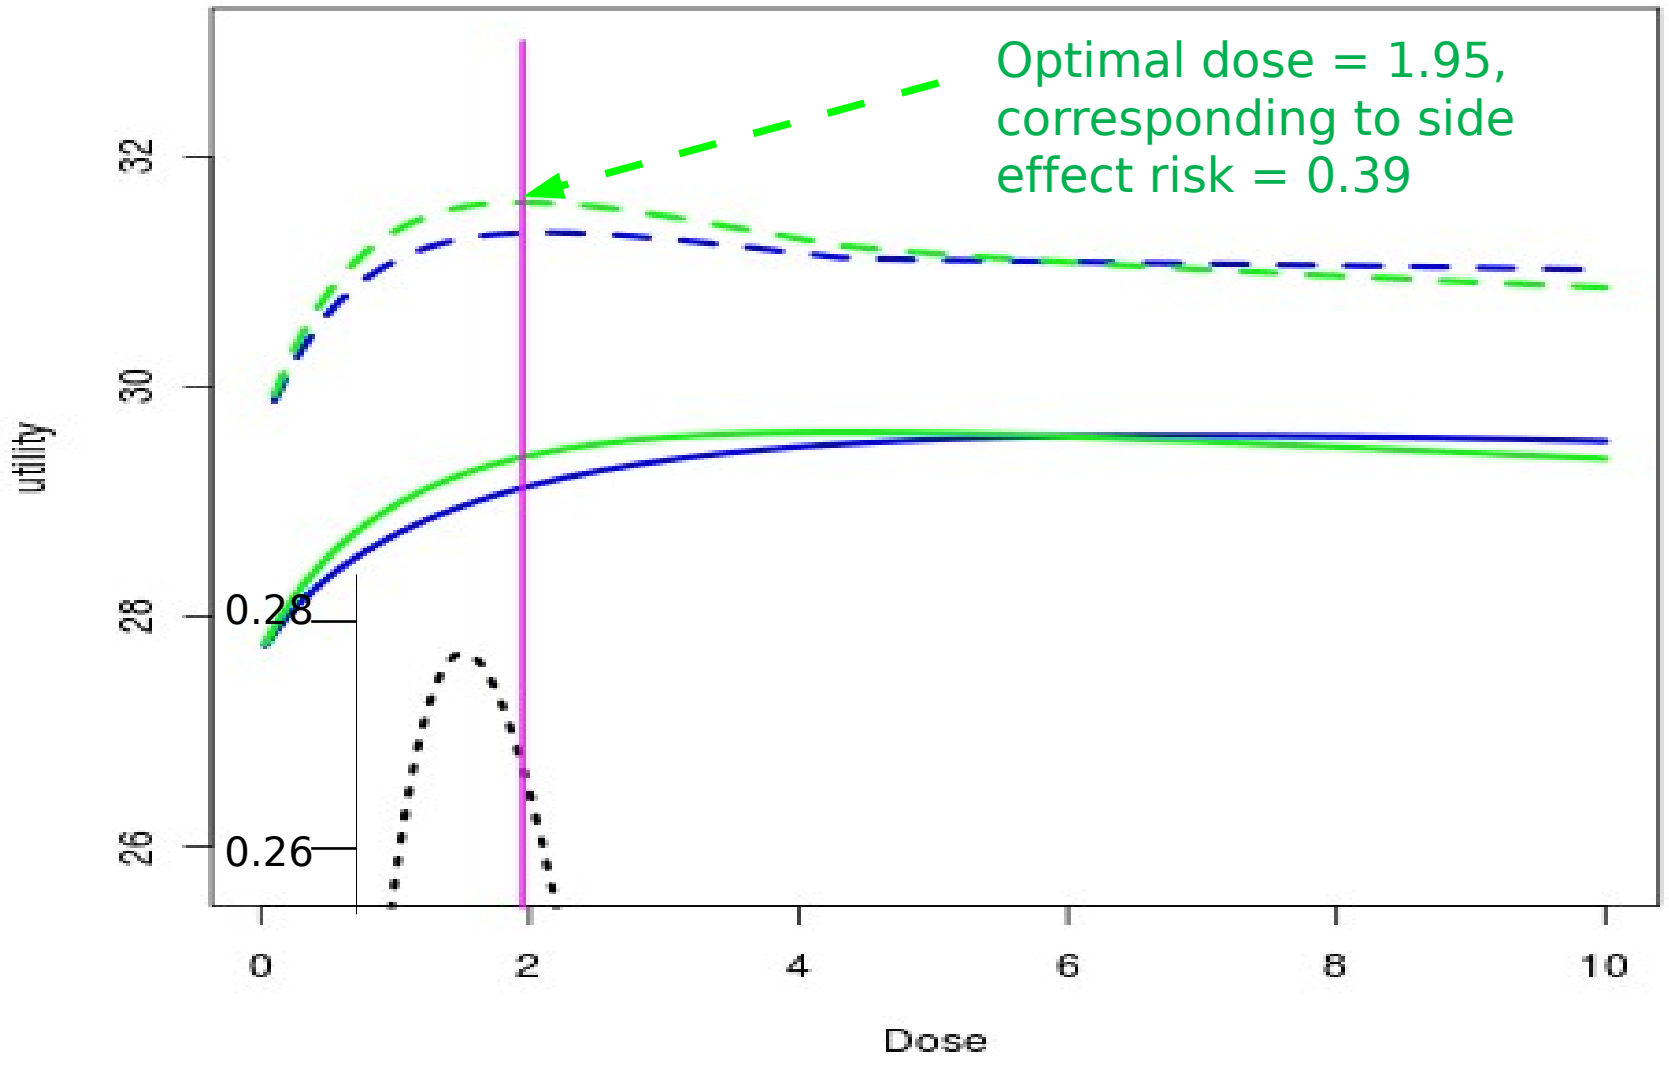

Assuming the first trial patient is the penultimate

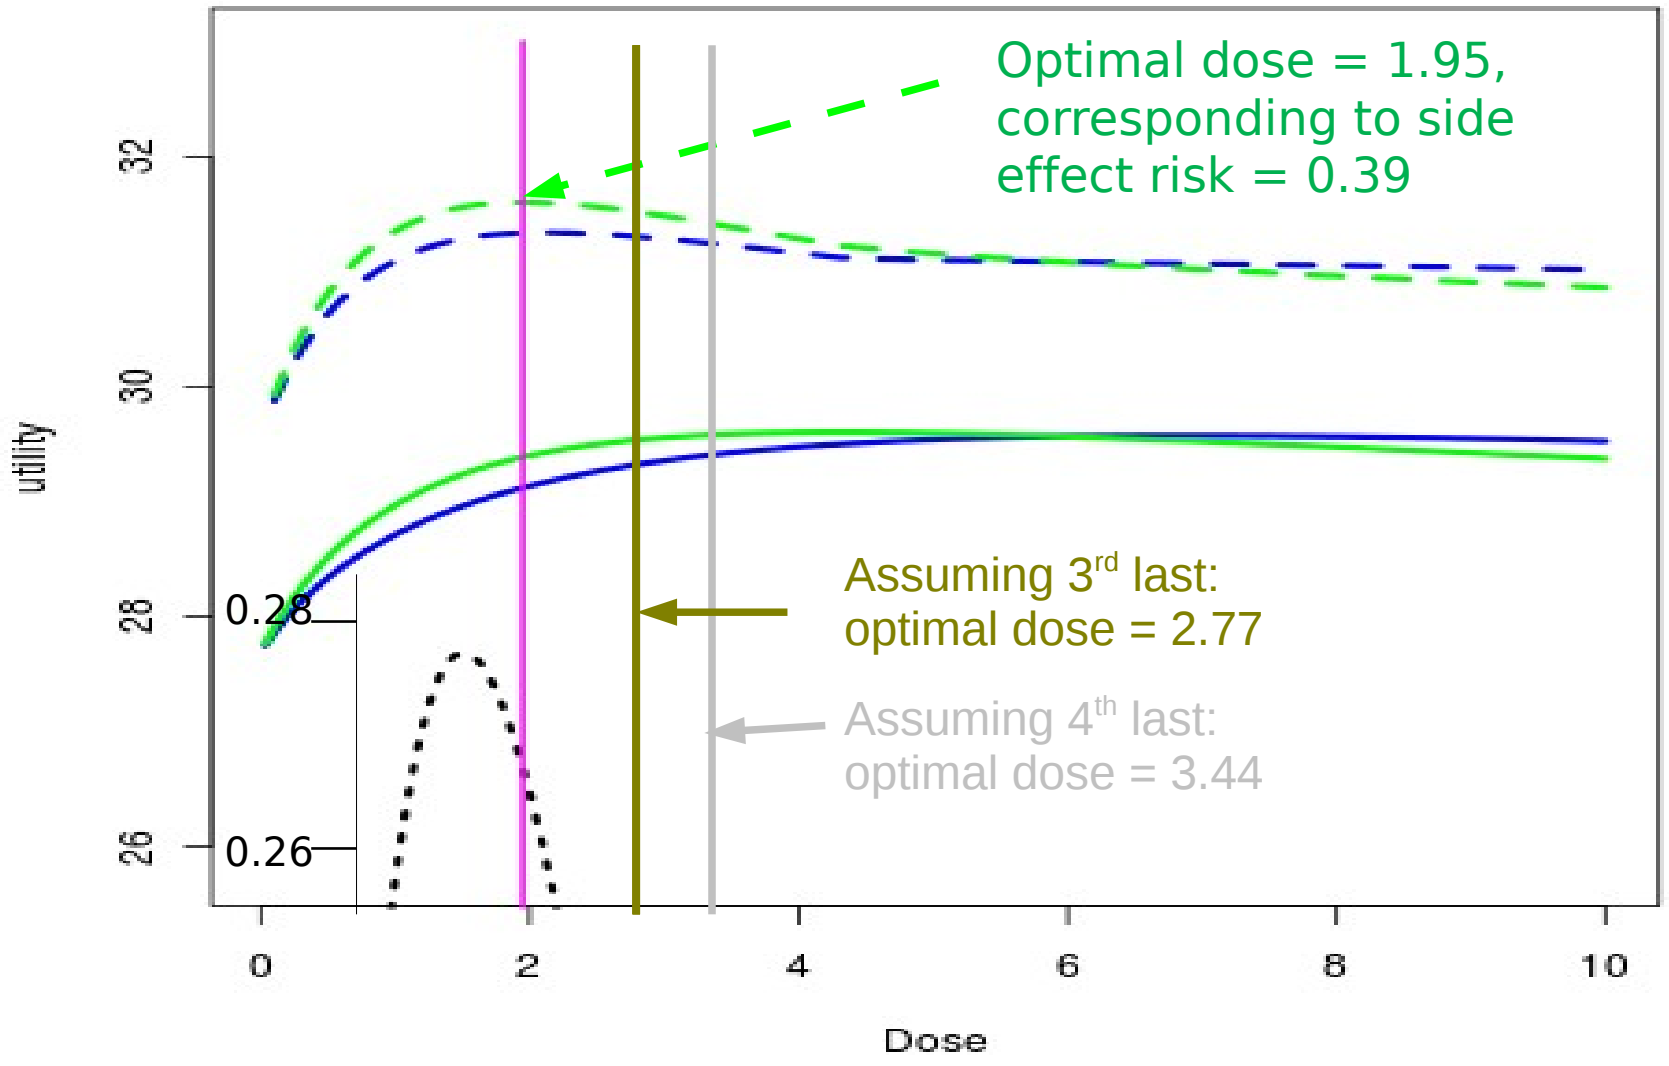

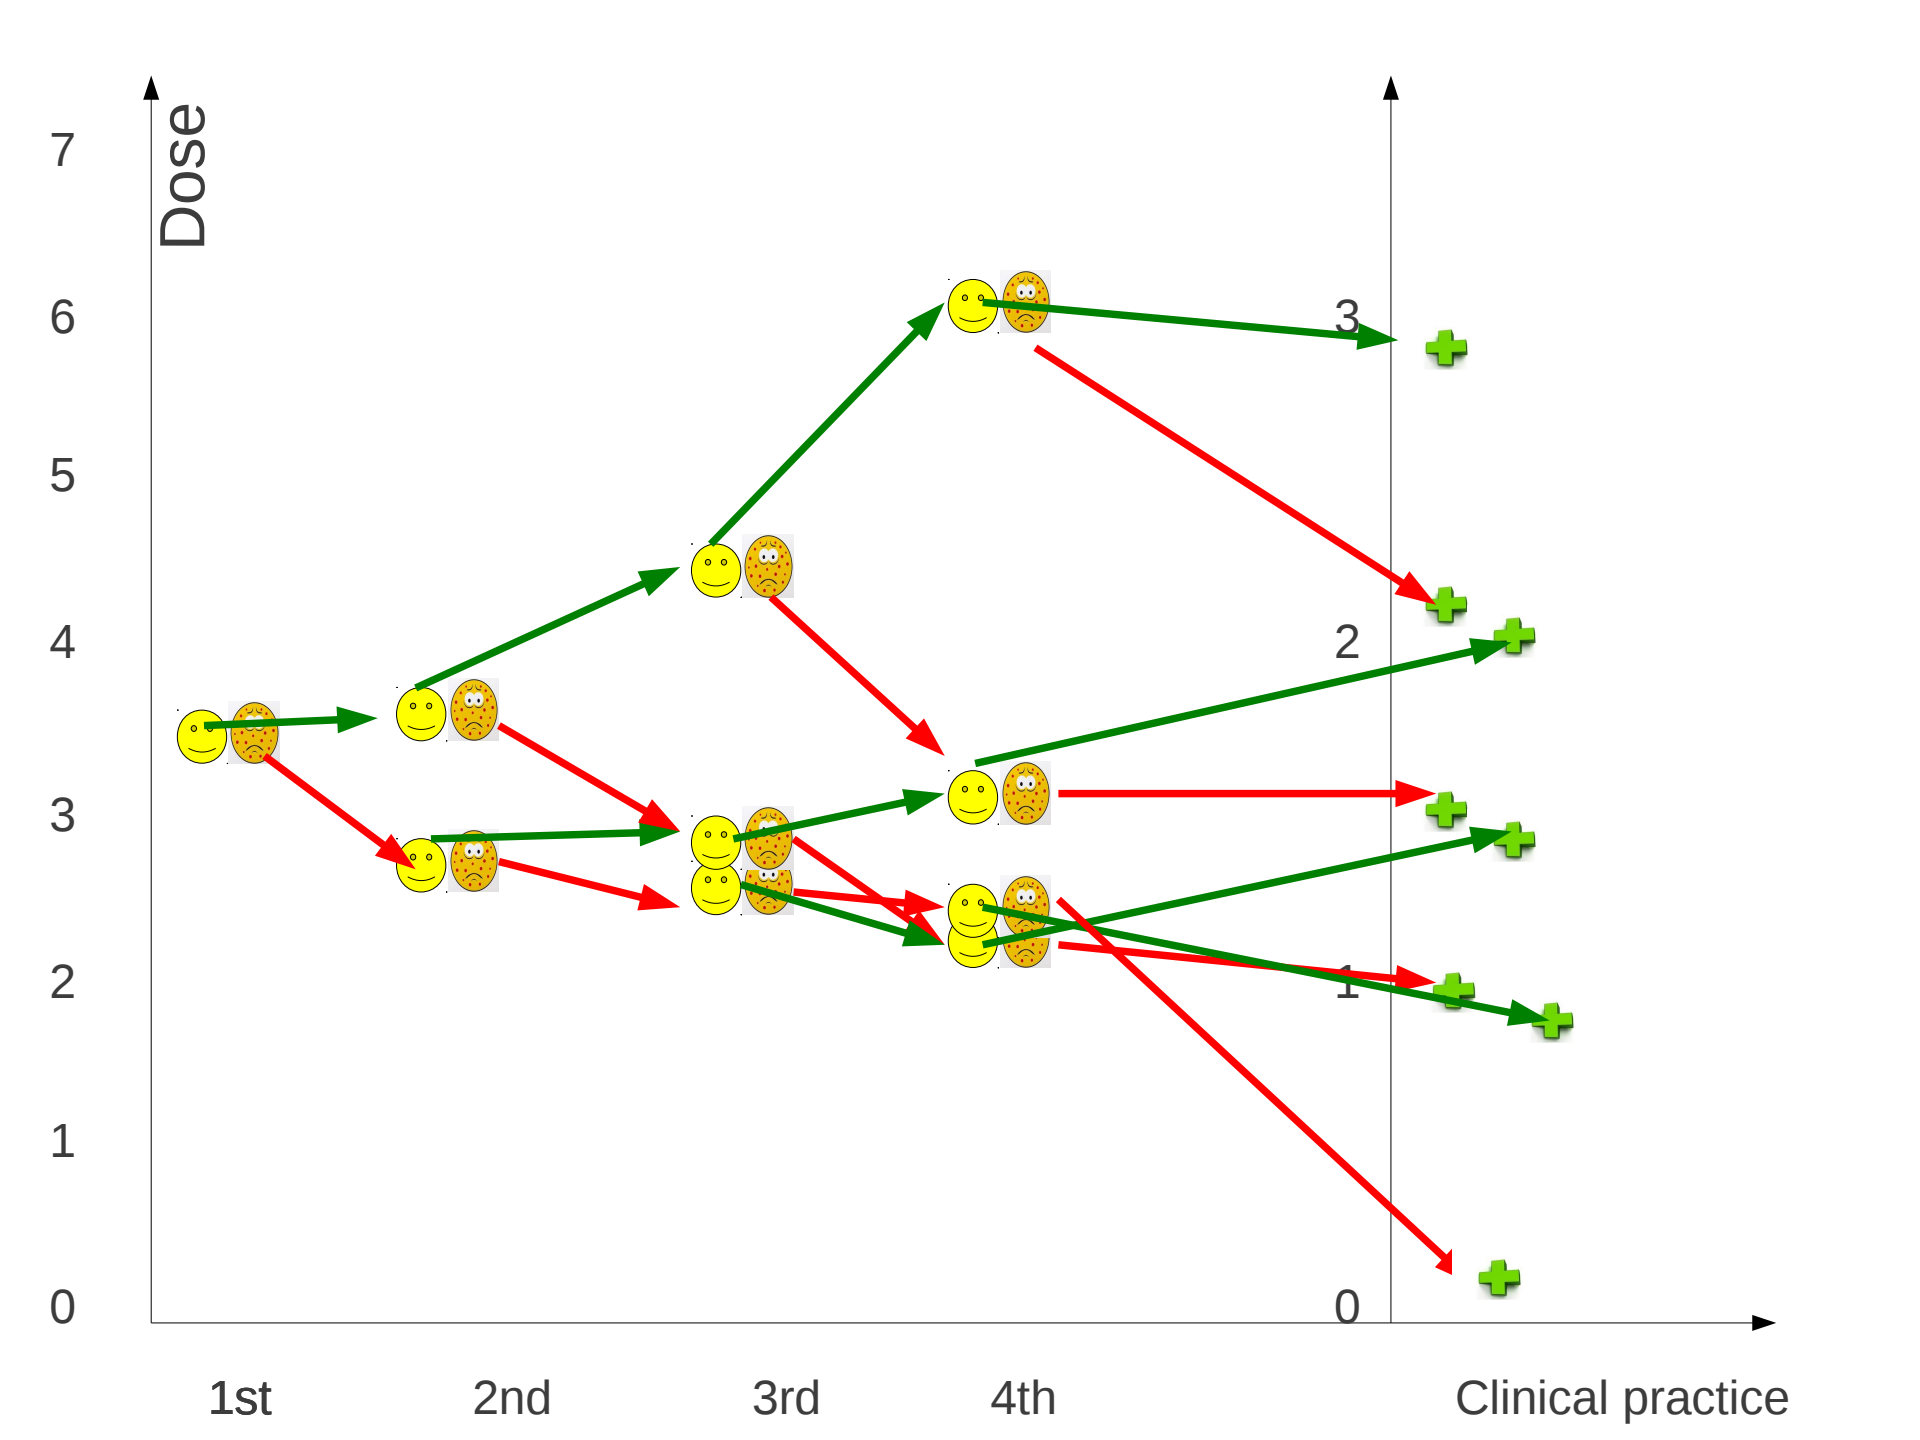
Dose
7
6
5
4
3
2
1
0
3
2
1
0
1st
2nd
3rd
4th
Clinical practice

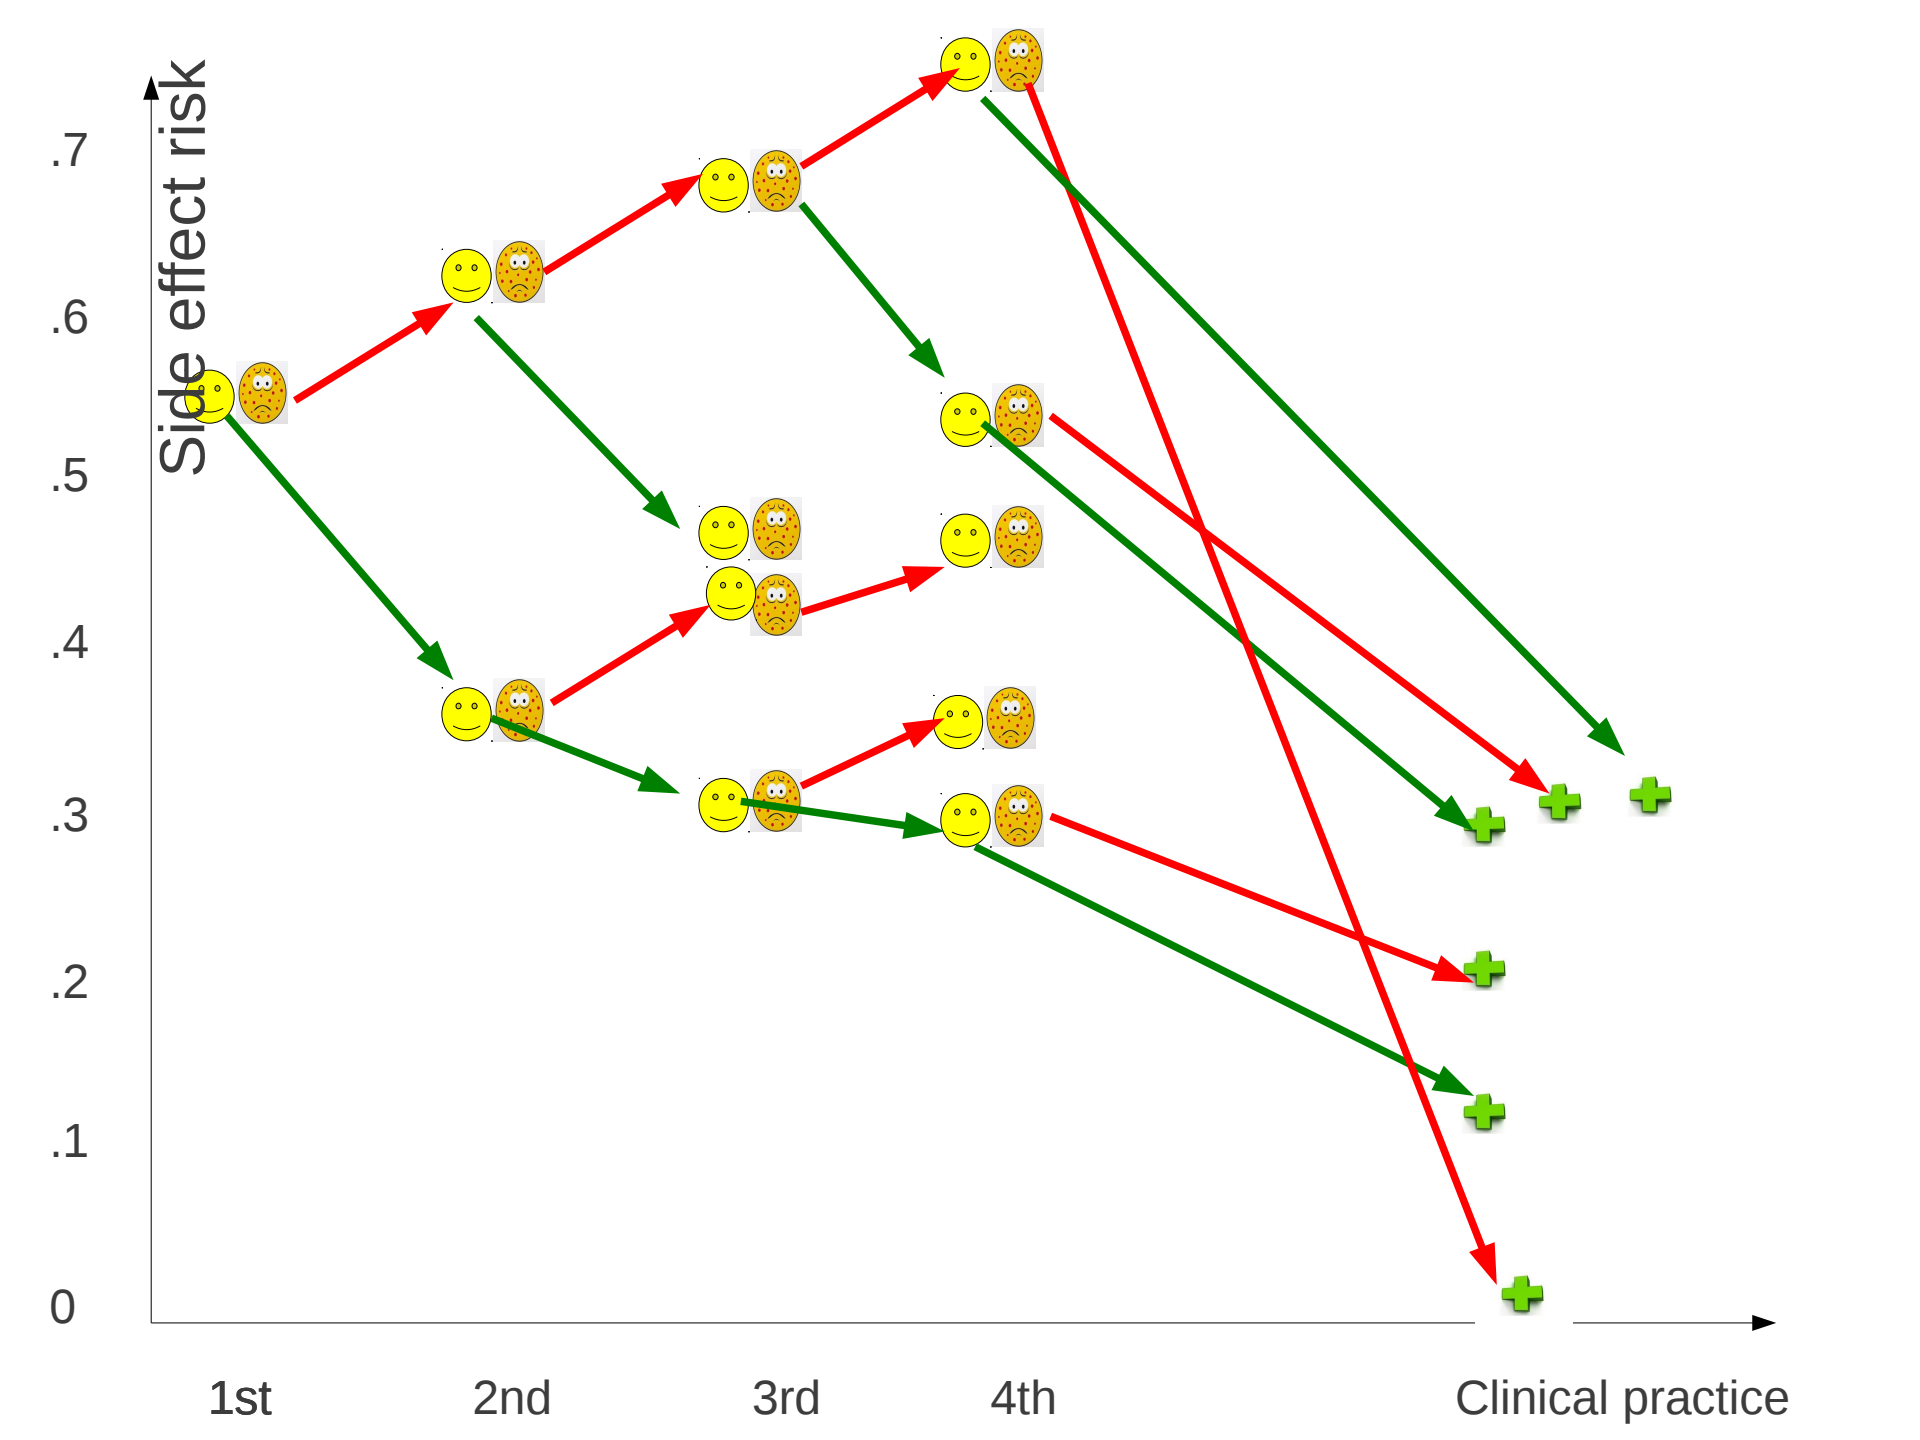
Side effect risk
.7
.6
.5
.4
.3
.2
.1
0
1st
2nd
3rd
4th
Clinical practice

Comparing to a traditional chase-the-target-risk approach

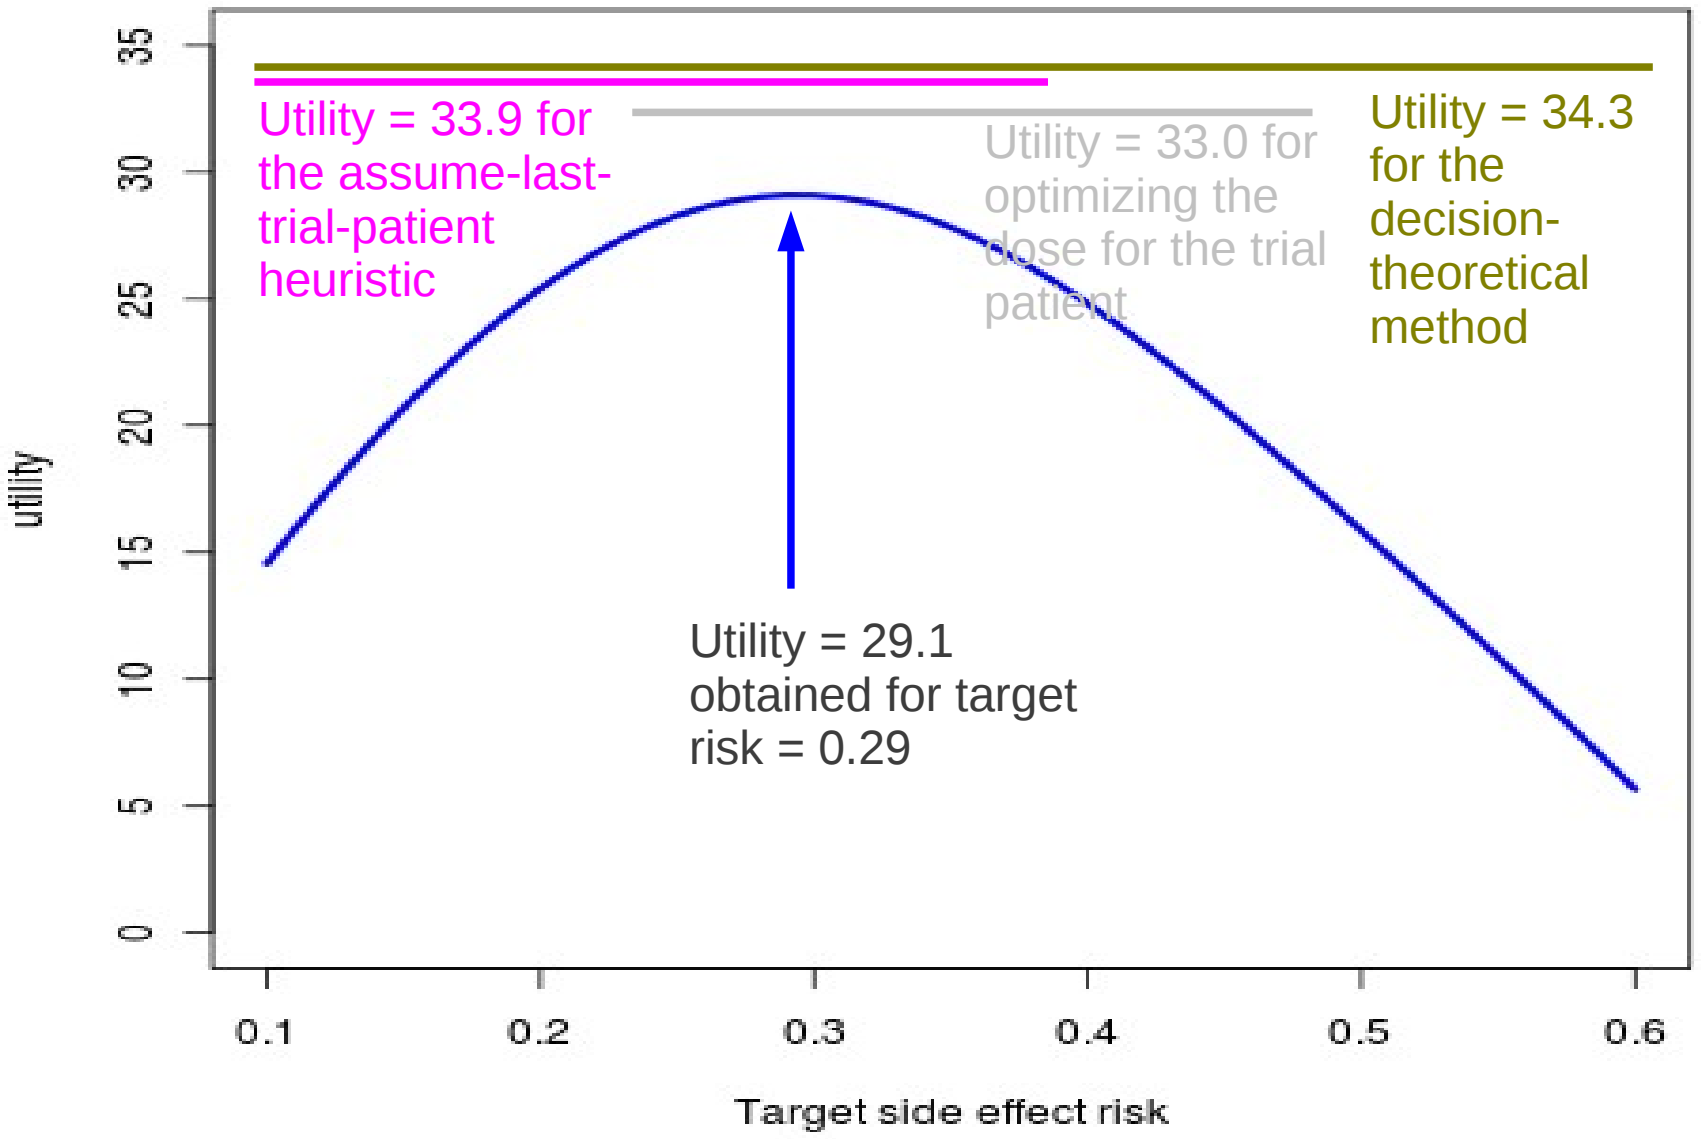

Beyond gamma-Bernoulli

Objective	Model	Computational feasibility
Avoid risks of extra-severe side effects	Gamma-Poisson: Side effects observed as integers	As long as selected doses have negligible change of eliciting large side effects ...
Avoid risks of extra-severe side effects	Gaussian: Side effects observed as continuous	Probably not feasible
Avoid risks of extra-severe side effects	Gamma-Bernoulli (or Poisson) with concave risk-utility	Straight forward
Inter-patient variability	Gamma-Bernoulli	May be necessary to fix the overdispersion parameter. Probably only makes sense in very large trials
Therapeutic benefit observed also	Bivariate Gamma-Bernoulli	Straight forward
Drug combinations	Interaction effect	As long as the pharmacologists don't force some hairy interaction model on us ..
Discrete dose levels	(unchanged)	Need to pre-specify levels

Stopping rule

- The procedure stops automatically if the optimal dose is zero. In the scenario used here this will happen if all first five patients have side effects.
- What is the ethical cost of continuing the trial? Say in each trial cycle 10 patients would have got the treatment in clinical practice but we only recruit one. Calculate the costs of denying treatment to the other 9.

Summary

• There may be no such thing as a target side effect risk – in general, the posterior mean of the side effect risk associated with the optimal dose changes as the study proceeds

• Maximizing the direct utility of the information from the next patient may not be so bad. For optimal design you will need to look at least 4 cohorts (patients) ahead, though. Probably more.

• We need to think about making the approach robust against model misspecification

• A constraint on the amount of sacrifice trial patients have to suffer for the common good may be appropriate – for example, you can require the utility to be positive for each individual

• In practice, the utility of “clinical practice” should be replaced with the utility of the subsequent later phase trials. So we would need to model that also.

• The utility parameters may be different for trial patients than for clinical practice. For example, the therapeutic benefit is zero for healthy volunteers.