Semi-Parametric Dose Finding Methods.

Matthieu Clertant [1] and John O'Quigley

Novametrics, Early phase dose finding methods, November 15th 2017

[1] pr.clertant@gmail.com, Lancaster University

Plan

Semi-Parametric Models

Prior Model and CRM Model

Theoretical Properties

Simulations II

Plan

Semi-Parametric Models

Prior Model and CRM Model

Theoretical Properties

Simulations II

Modeling Context

- $X \in D = \{1, \ldots, m\} \qquad Y \in \{0, 1\}.$
- $Y|d$ is Bernoulli with parameter β_d, and:

$$\beta_1 \ < \ \ldots \ < \ \beta_m \ .$$

- The target or threshold: α.

$$d^* = \text{MTD} = \arg\min_{d \in D} |\alpha - \beta_d|.$$

Sample of scenarios (uniform spacings)

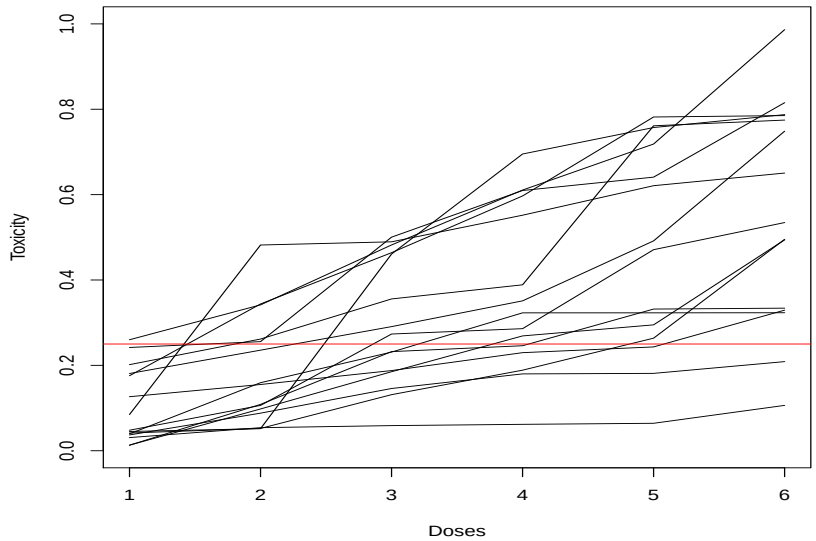

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

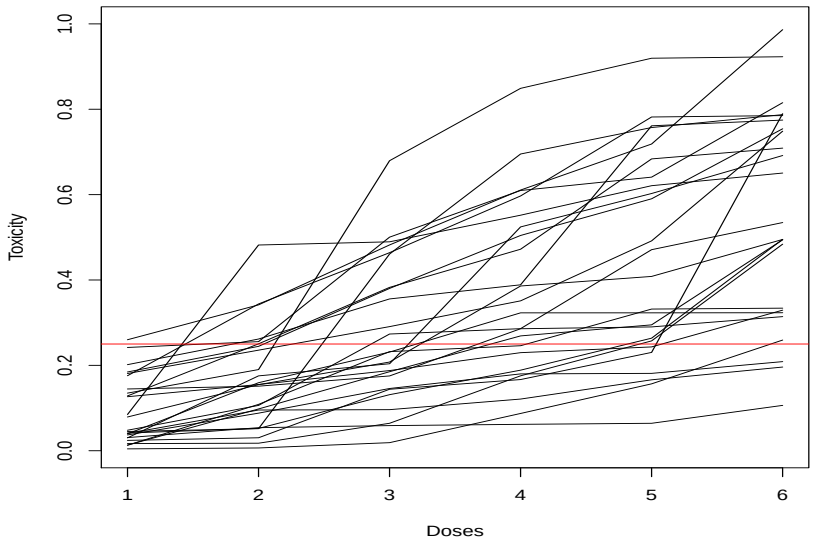

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

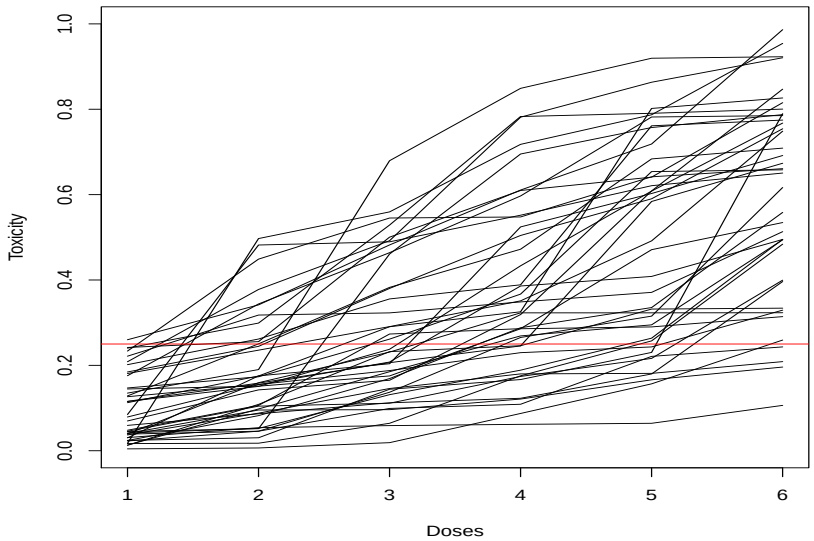

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

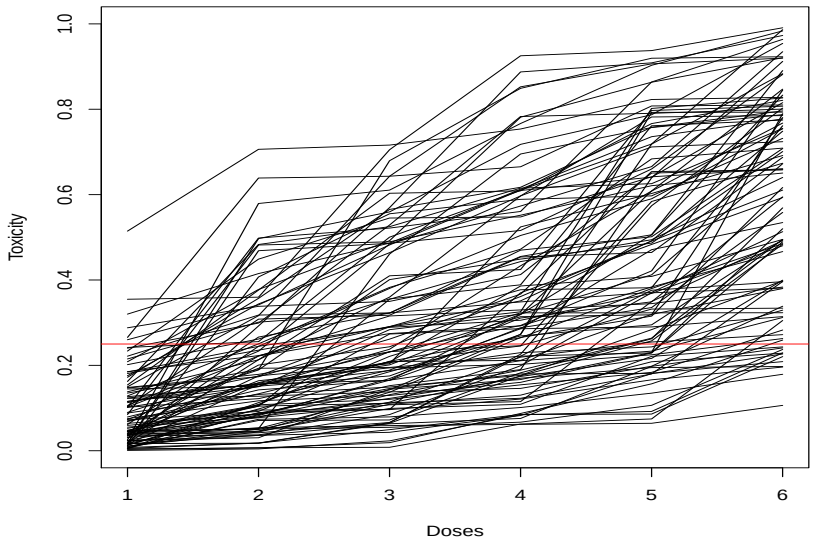

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

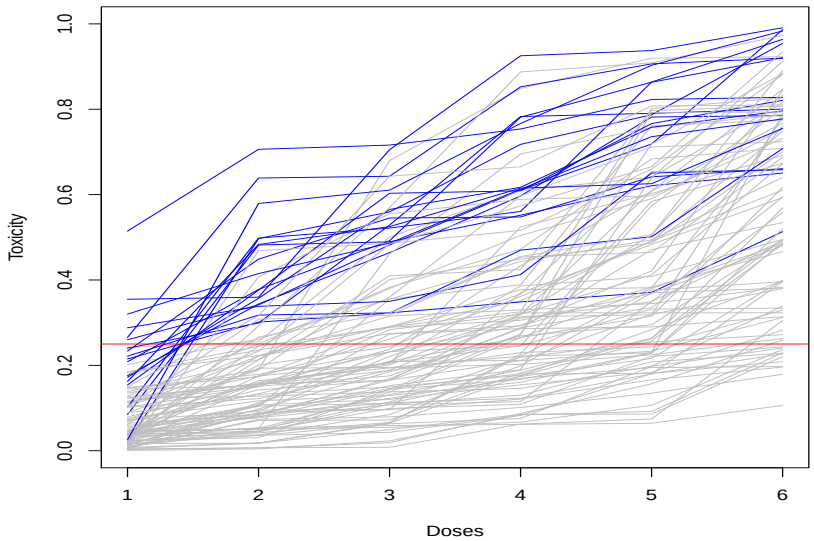

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

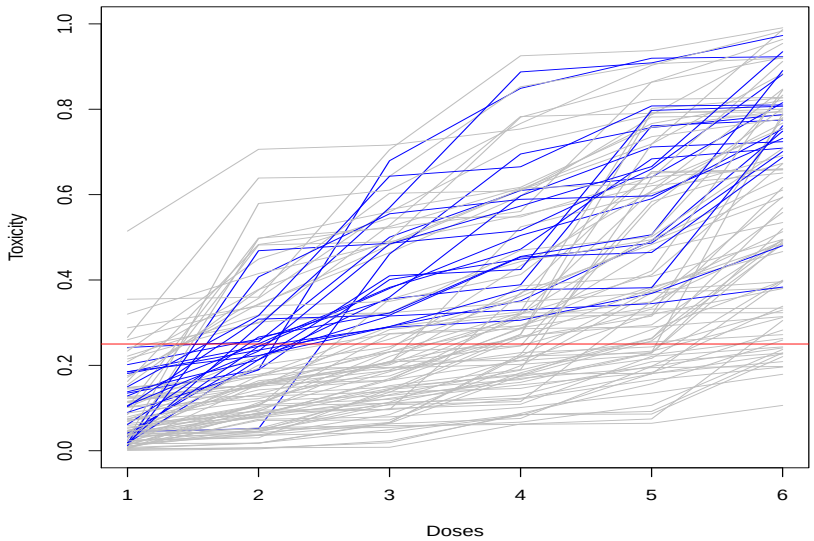

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

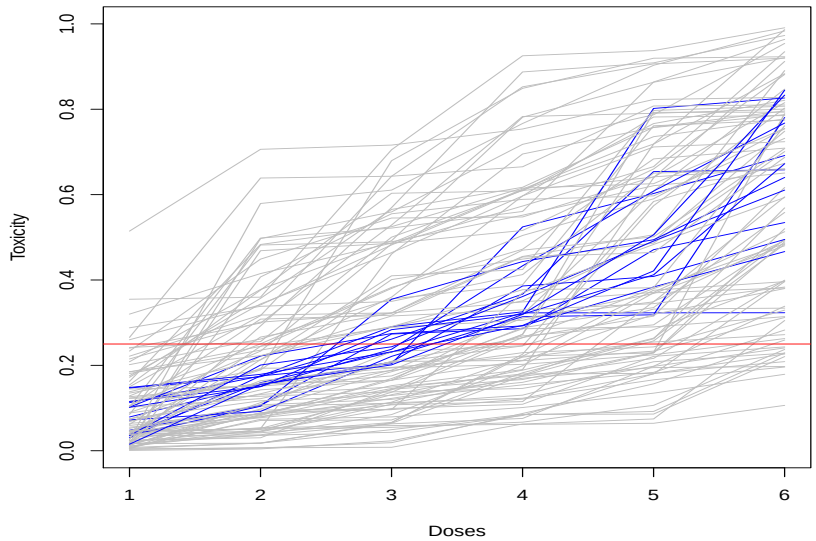

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights on the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

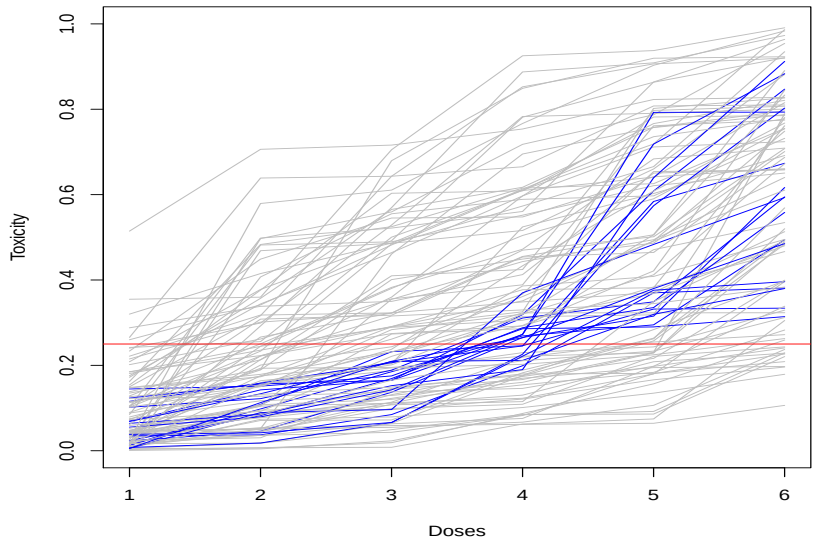

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights on the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

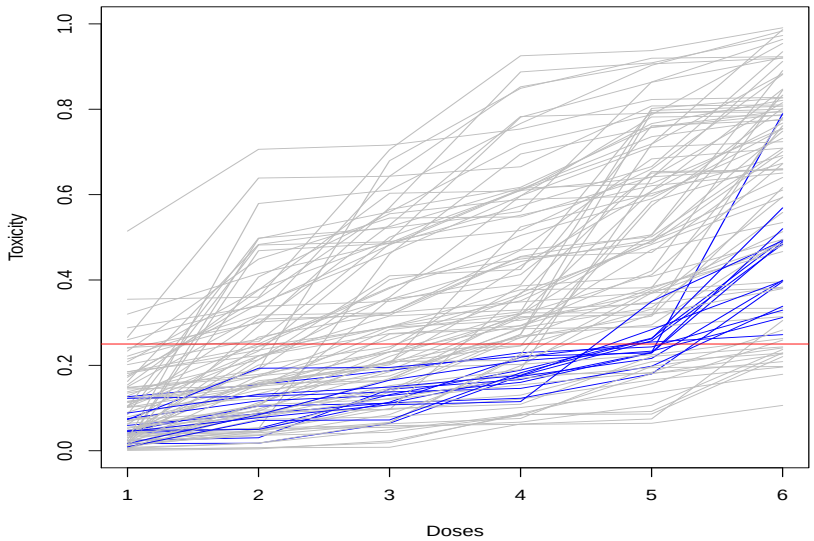

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights on the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

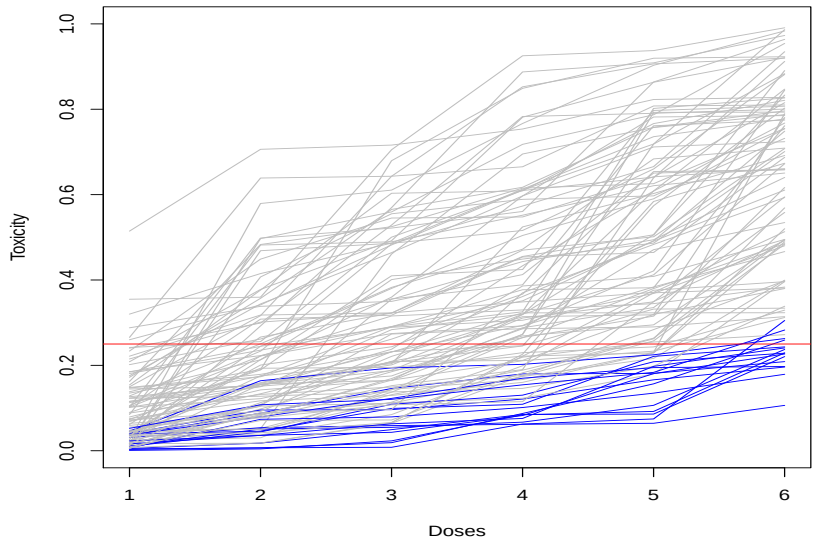

$p \in [0,1]^m$ is a vector of toxicity probabilities on the space of doses, with p_d the toxicity probability at dose d.

$\theta \in D$ is the MTD parameter and Π its prior.

$\Lambda = (\Lambda_\theta)_{\theta \in D}$ is the prior model. Λ_θ weights on the vector p which satisfies the following constraints:

$$\forall d \in D\,,\ |p_\theta - \alpha| \leq |p_d - \alpha|.$$

Sample of scenarios (uniform spacings)

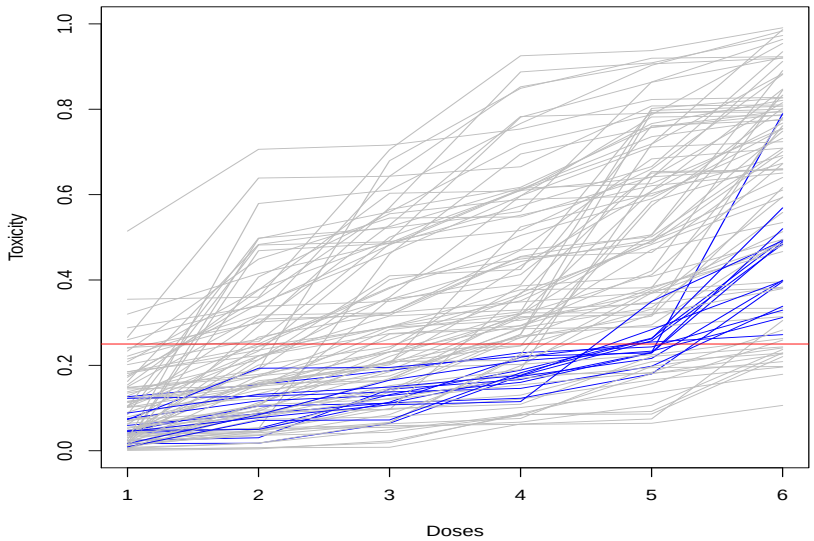

$\mathcal{L}_{(X_1^n,Y_1^n)}(p)$ is the likelihood according to the sample (X_1^n, Y_1^n).

Posterior on the parameter of interest:

$$\Lambda_\theta(dp|(X_1^n, Y_1^n)) \times \Pi(\theta|(X_1^n, Y_1^n)) \propto \mathcal{L}_{(X_1^n,Y_1^n)}(p) \times \Lambda_\theta(dp)\,\Pi(\theta).$$

$$\Pi(\theta|(X_1^n, Y_1^n)) \propto \int \mathcal{L}_{(X_1^n,Y_1^n)}(p)\,\Lambda_\theta(dp) \times \Pi(\theta).$$

Semi-Parametric Methods

1. $\Pi_n(\theta) = \Pi(\theta | X_1^n, Y_1^n) \propto \left[\int \mathcal{L}_{(X_1^n, Y_1^n)}(p) \Lambda_\theta(dp) \right] \Pi(\theta)$,
2. $X_{n+1} = \widehat{\theta}_n = \arg\max \Pi_n(\theta)$.

Other estimators: Estimated probability toxicities at each dose

$$\tilde{\beta}_{j,n} = \mathbb{E}_{(\Lambda \otimes \Pi)_n} [p_j] = \sum_{\theta=1}^{m} \left[\int p_j \Lambda_{\theta,n}(dp_j) \right] \Pi_n(\theta) .$$

Semi-Parametric Methods

1. $\Pi_n(\theta) = \Pi(\theta | X_1^n, Y_1^n) \propto \left[\int \mathcal{L}_{(X_1^n, Y_1^n)}(p) \Lambda_\theta(dp) \right] \Pi(\theta)$,
2. $X_{n+1} = \widehat{\theta}_n = \arg\max \Pi_n(\theta)$.

The posterior distribution on the class θ given (X_1^n, Y_1^n) is then:

$$\Lambda_{\theta,n}(p) \propto \mathcal{L}_{(X_1^n, Y_1^n)}(p) \Lambda_\theta(p) .$$

We then have:

$$\Pi_{n+1}(\theta) \propto \left[\int \mathcal{L}_{(X_{n+1}, Y_{n+1})}(p) \Lambda_{\theta,n}(p) \right] \Pi_n(\theta) .$$

Interval and ϵ-Sentivity

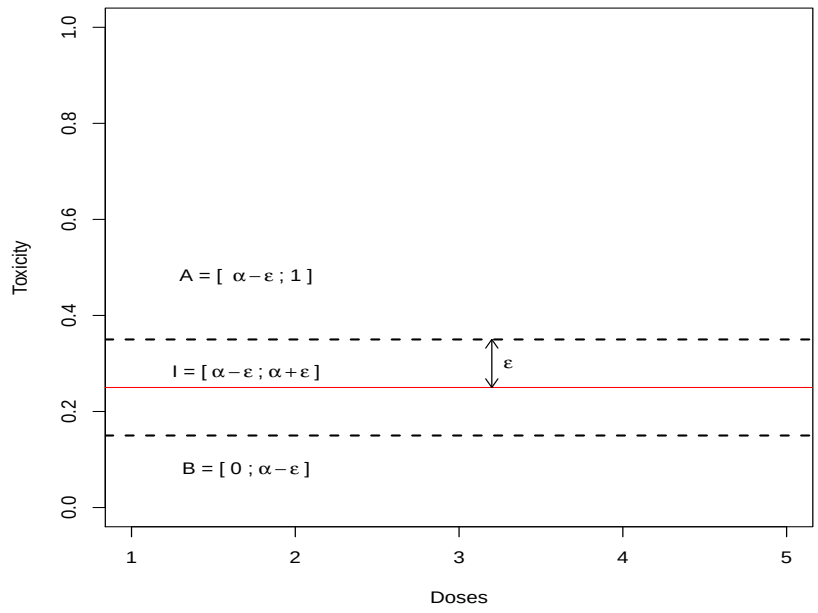

Assumptions on Λ

1. **Structure:** The support of Λ_θ is: $S_\theta = B^{\theta-1} \times I \times A^{m-\theta}$.
2. **Independence:** Λ_θ is a product of unidimensional distributions at each dose: $\Lambda_\theta = \Lambda_\theta^1 \times \ldots \times \Lambda_\theta^m$.

Support of Λ_3

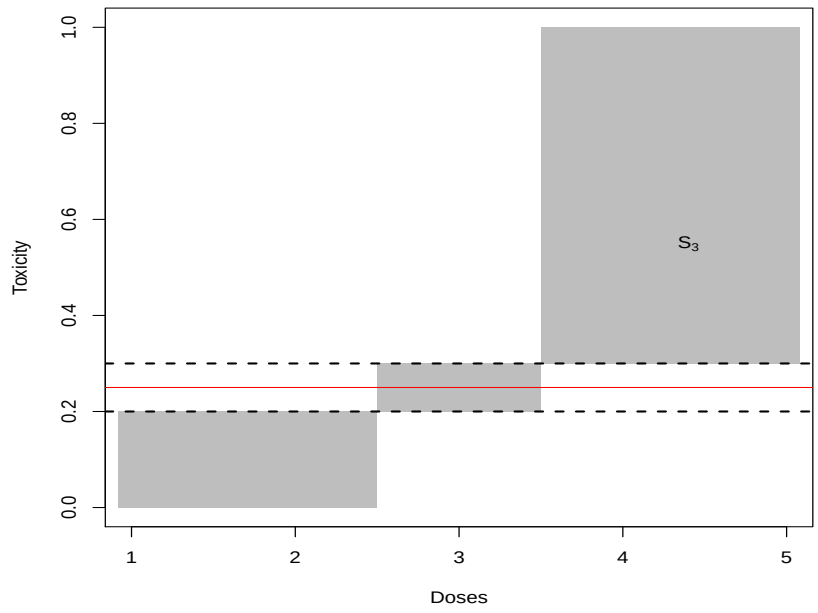

Assumptions on Λ

1. **Structure:** The support of Λ_θ is: $S_\theta = B^{\theta-1} \times I \times A^{m-\theta}$.
2. **Independence:** Λ_θ is a product of unidimensional distributions at each dose: $\Lambda_\theta = \Lambda_\theta^1 \times \ldots \times \Lambda_\theta^m$.

Plan

Semi-Parametric Models

Prior Model and CRM Model

Theoretical Properties

Simulations II

Continual Reassessment Method

A one-dimensional model of toxicity probability:

$$\mathcal{P}(Y = 1|X = d) = \Psi(d|a)$$

Algorithm (Bayes)

Step 1: Update the prior G of the parameter a; G_n is the posterior according to (X_1^n, Y_1^n)

Step 2: Calculate the estimators of toxicity probability at each doses:

$$\tilde{\beta}_{n,d} = \mathbb{E}_{G_n}[\Psi(d|a)].$$

The next dose is $X_{n+1} = \arg\min_{d \in D} |\tilde{\beta}_{d,n} - \alpha|$.

CRM: Support of the Model

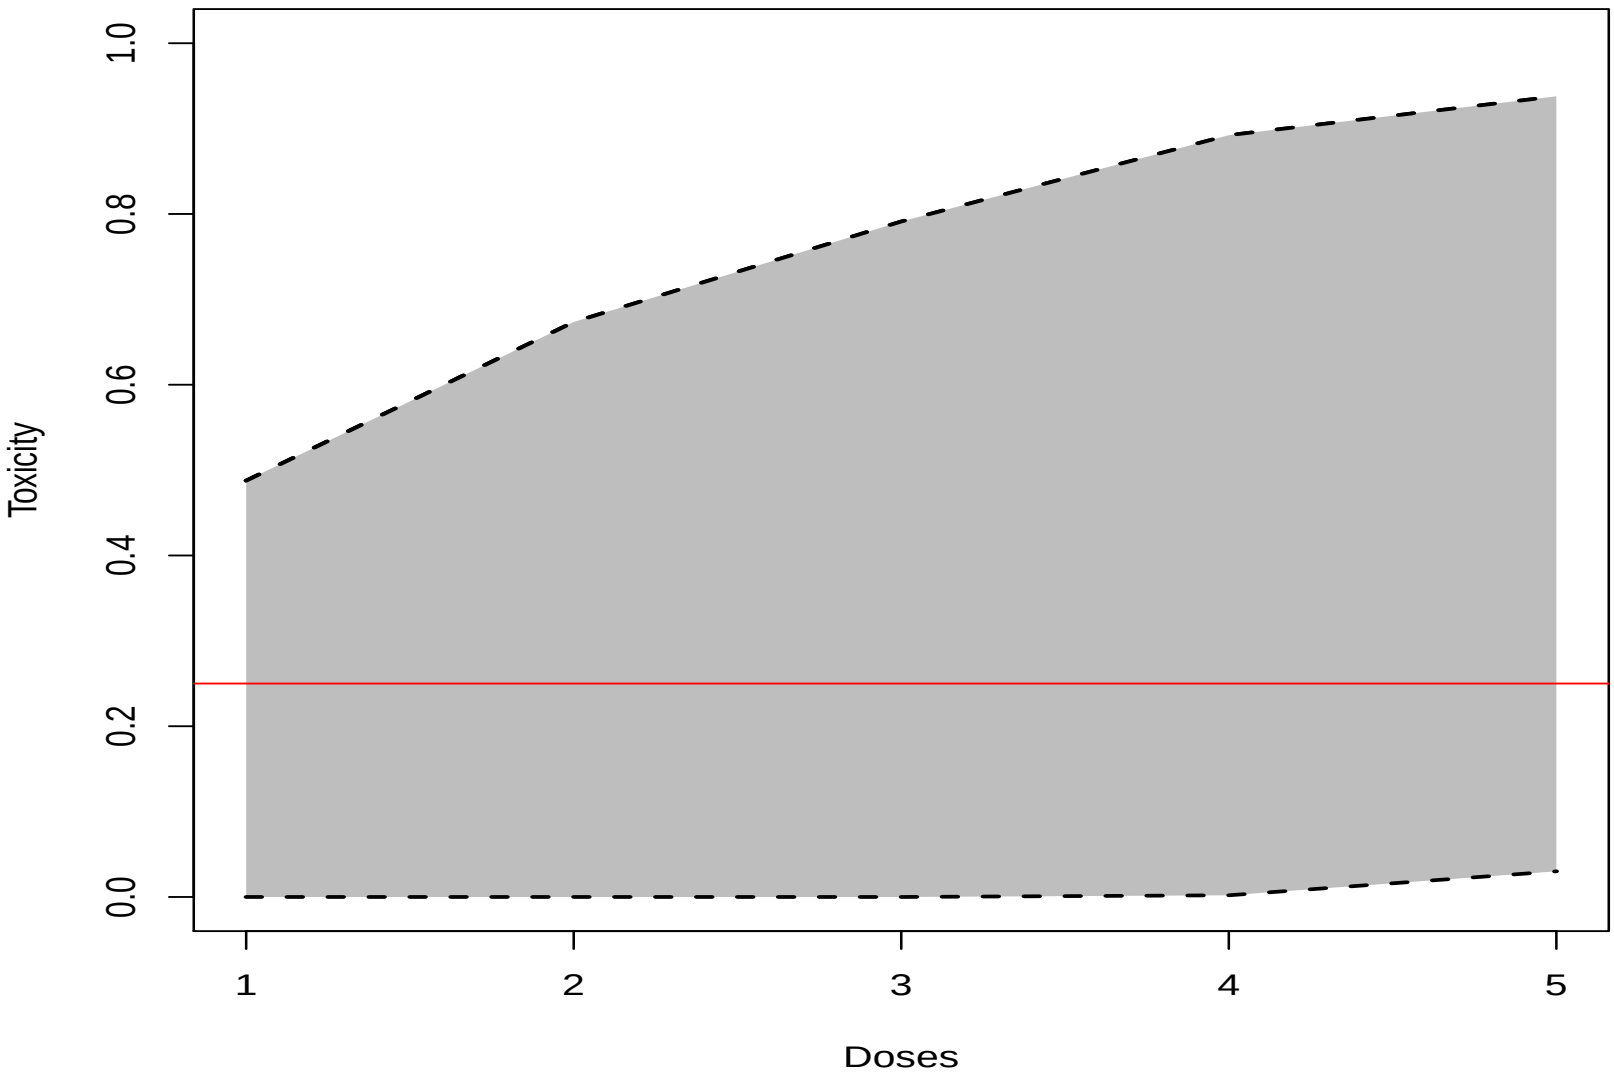

CRM: Support of the Model

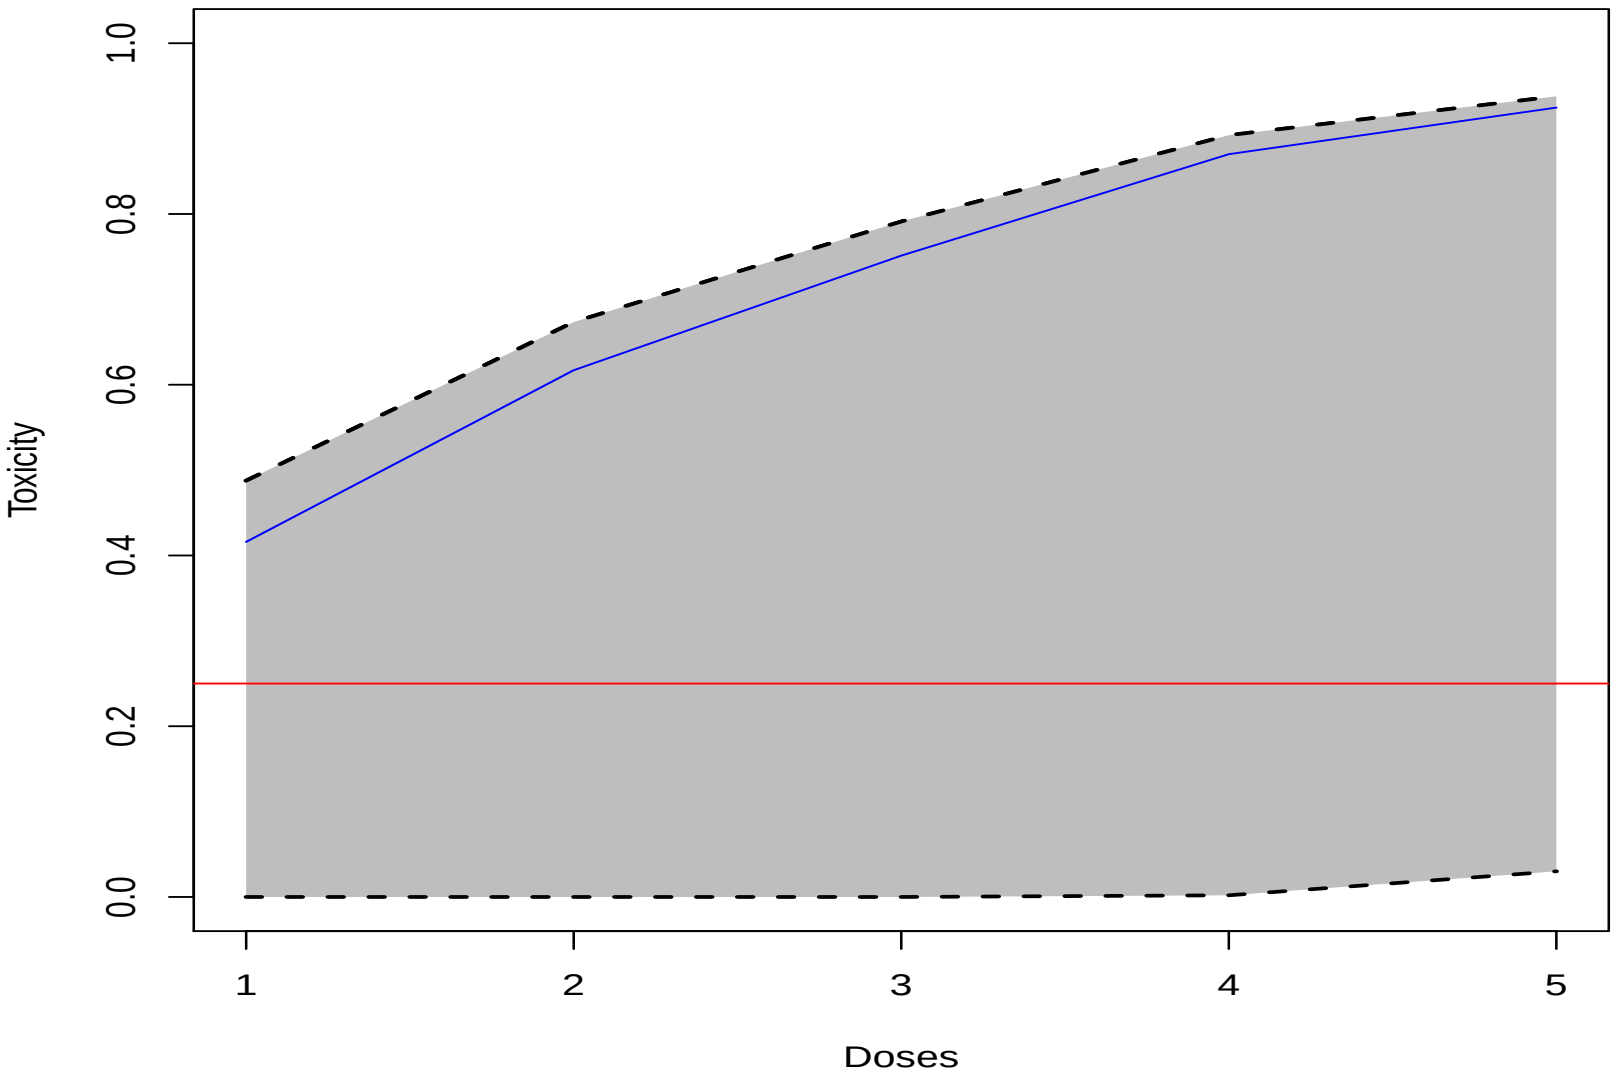

CRM: Support of the Model

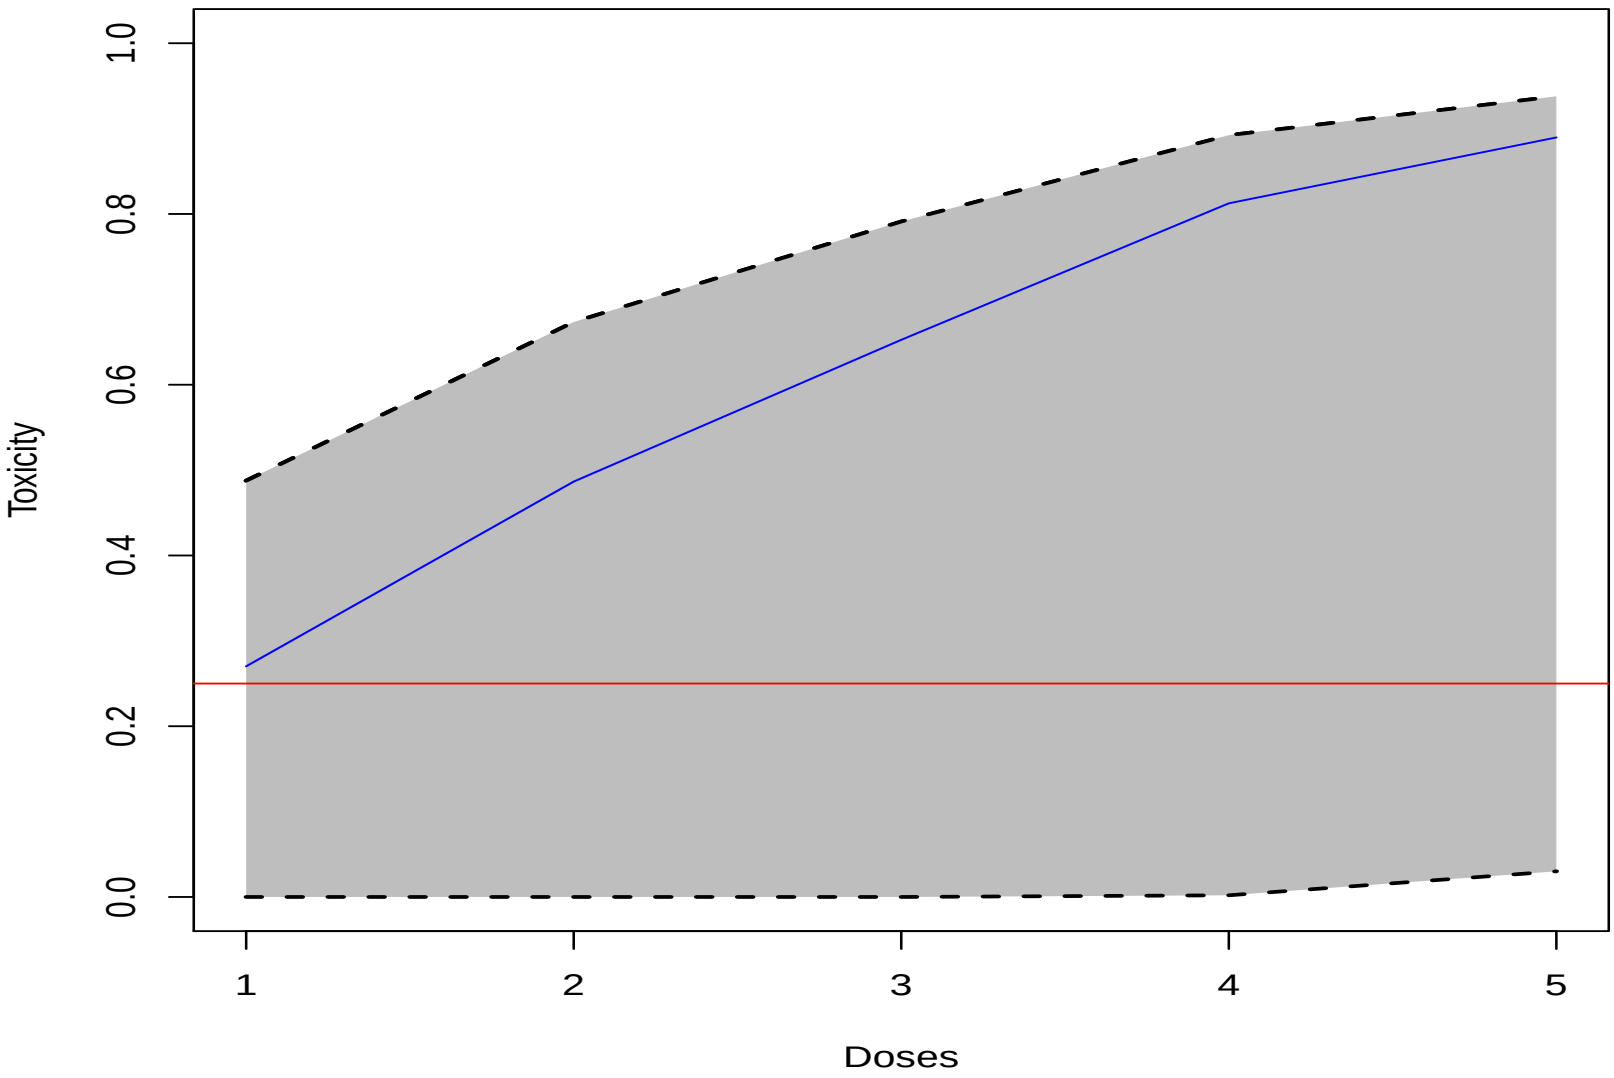

CRM: Support of the Model

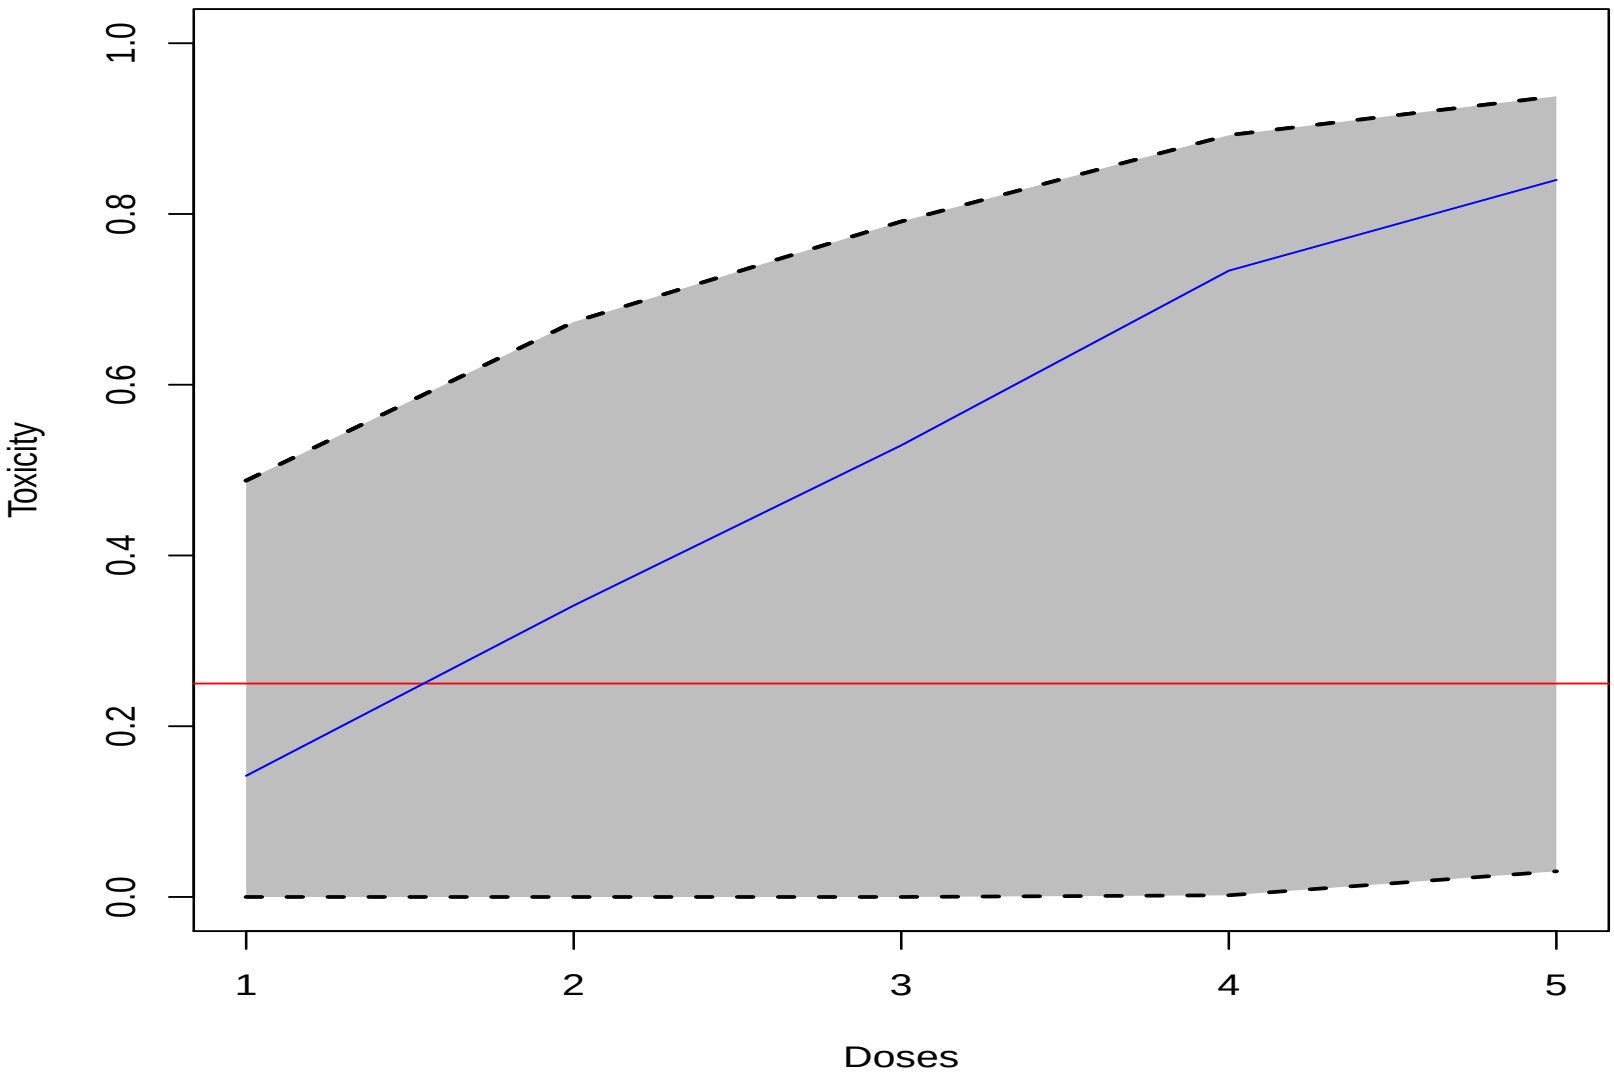

CRM: Support of the Model

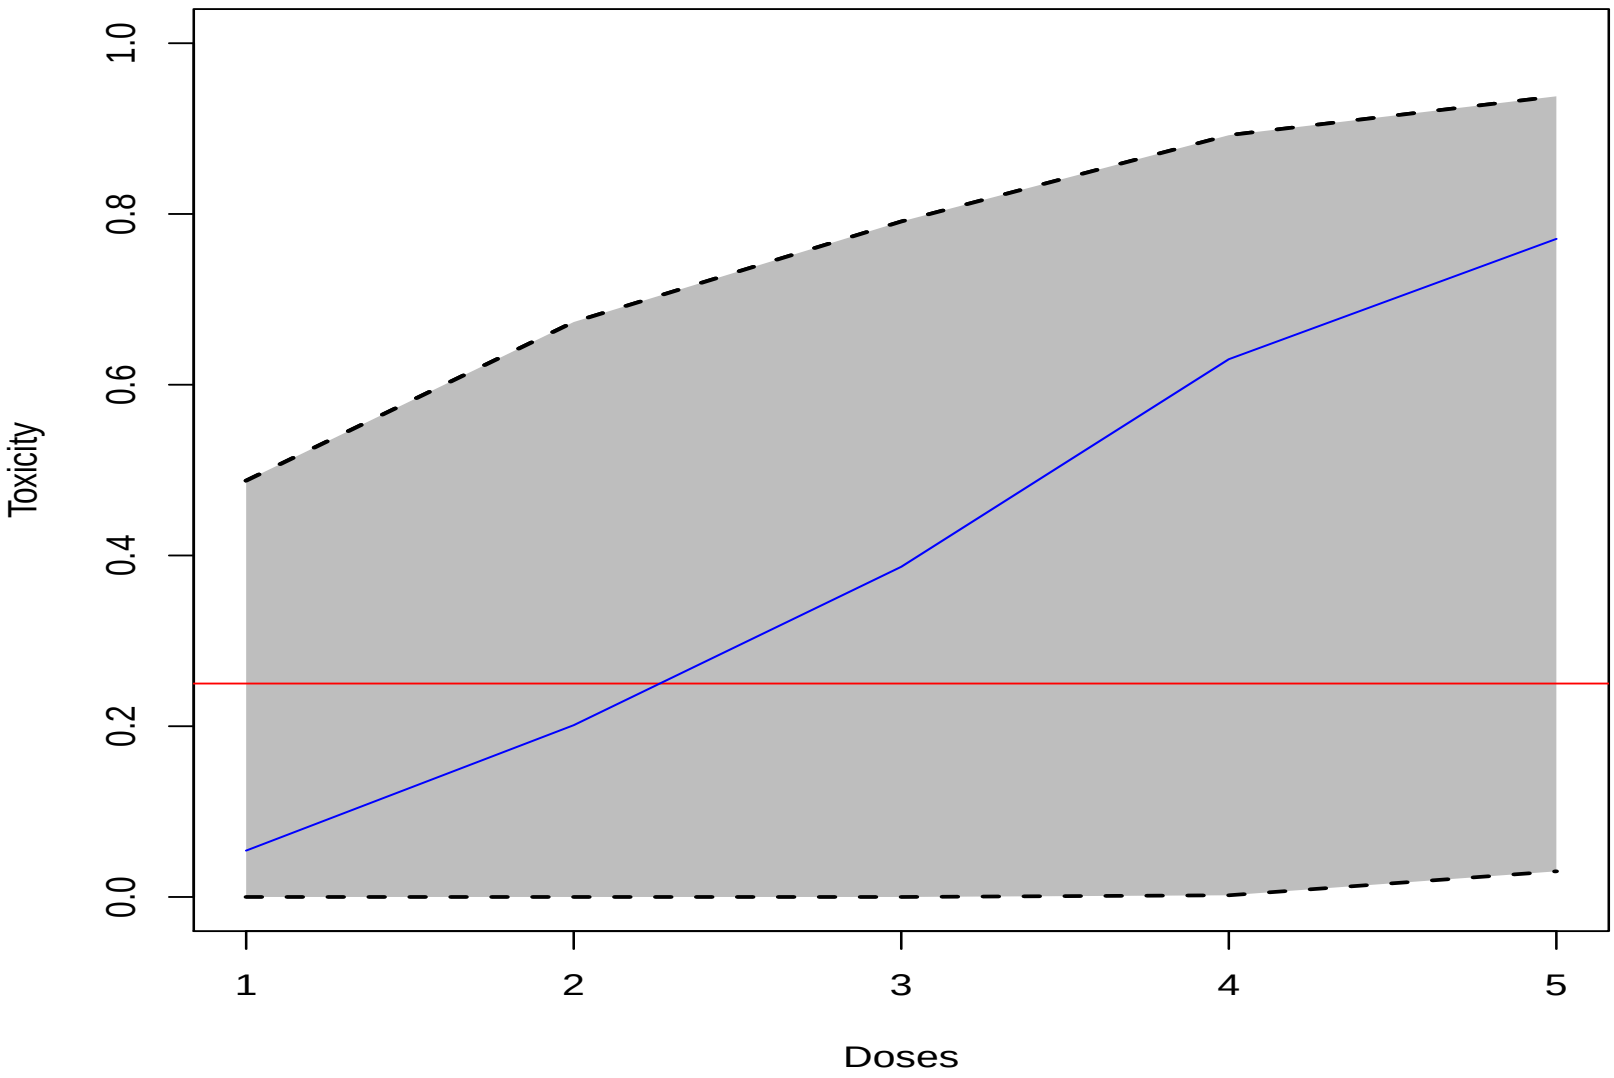

CRM: Support of the Model

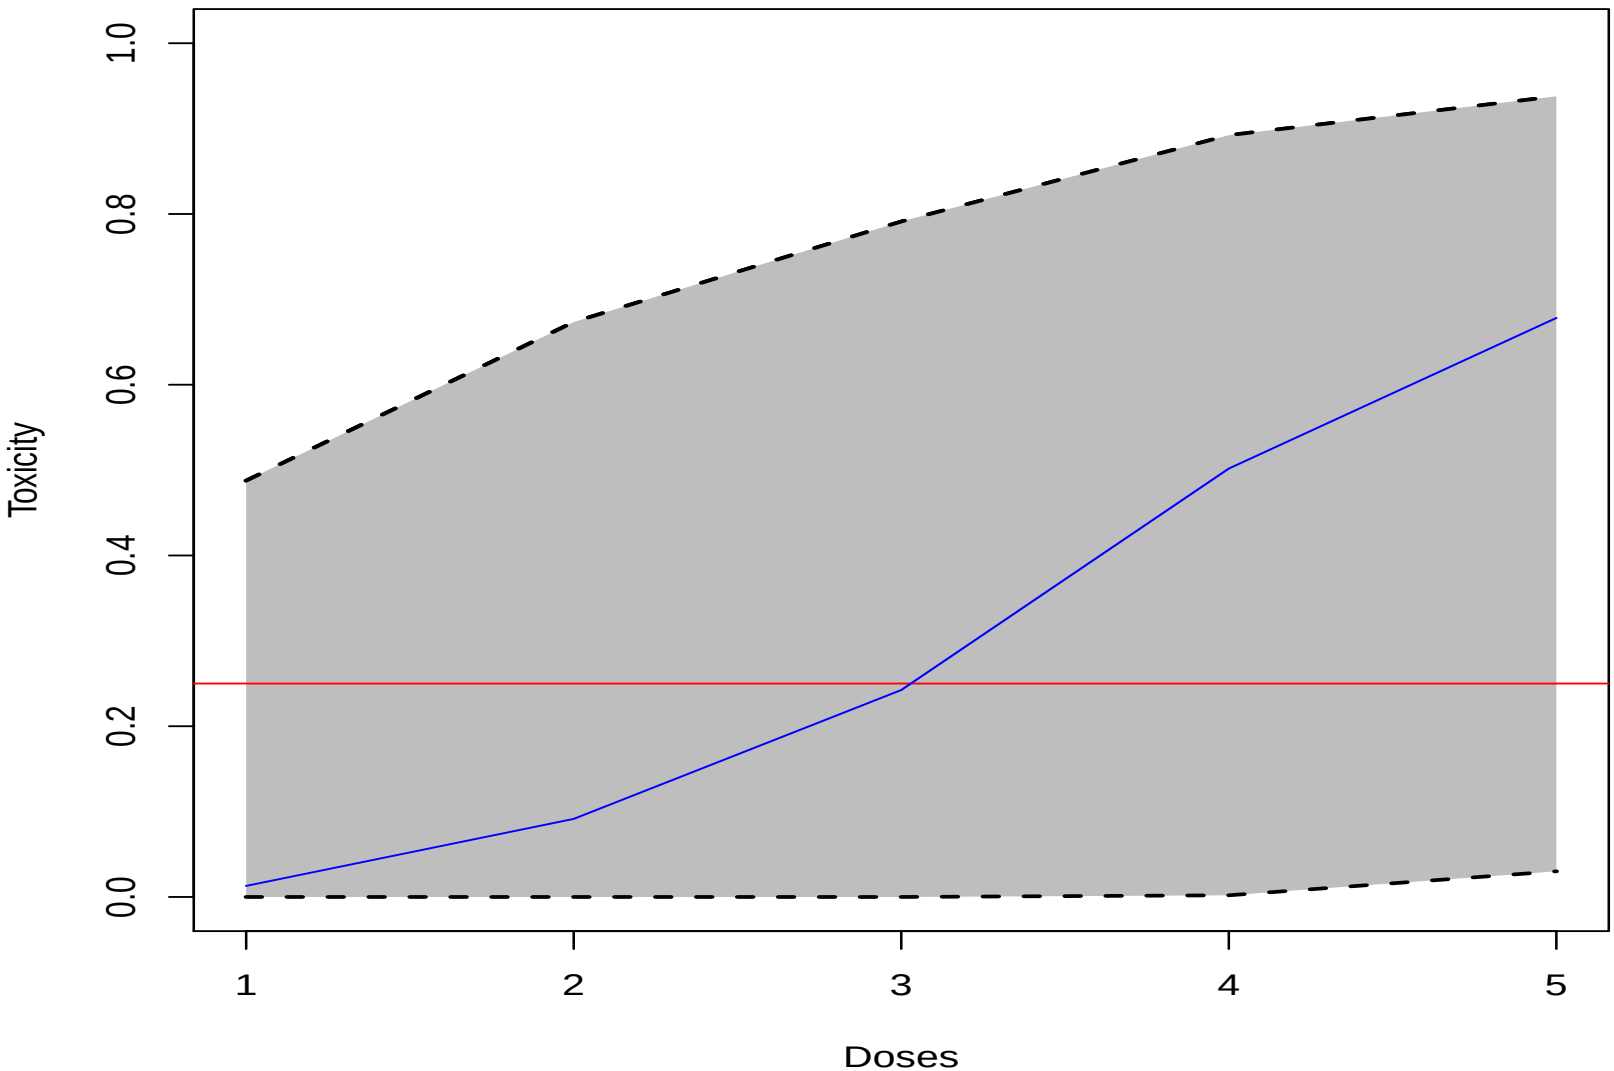

CRM: Support of the Model

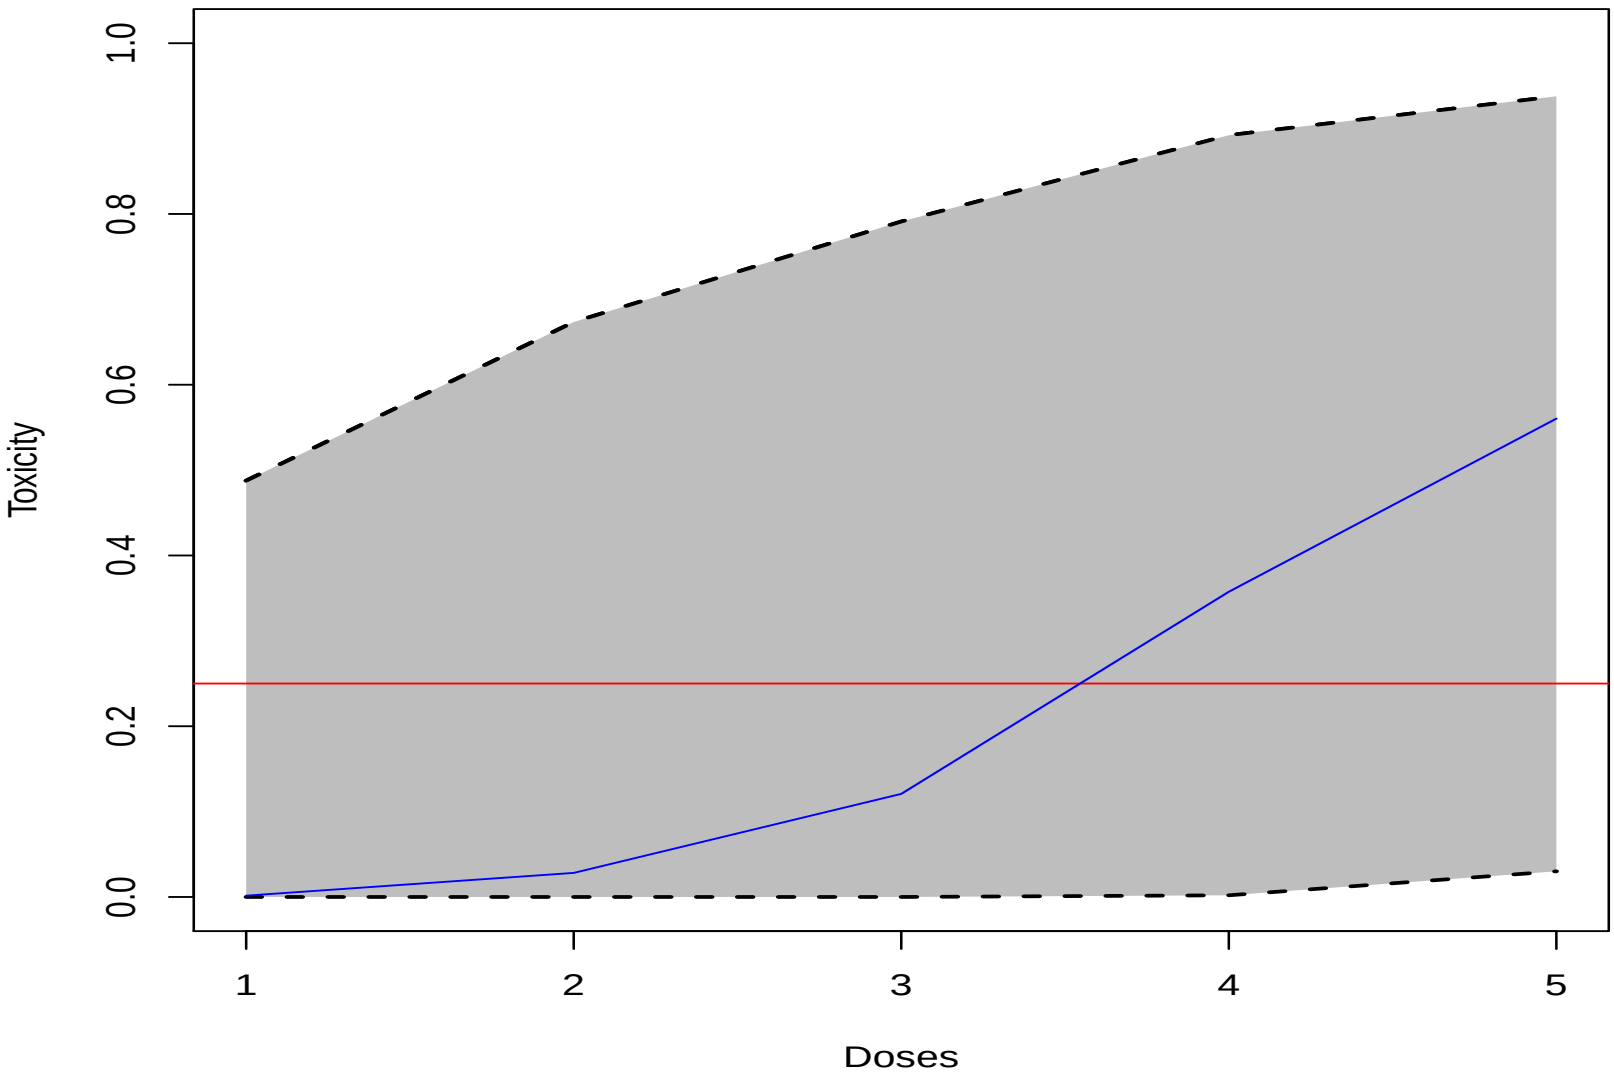

CRM: Support of the Model

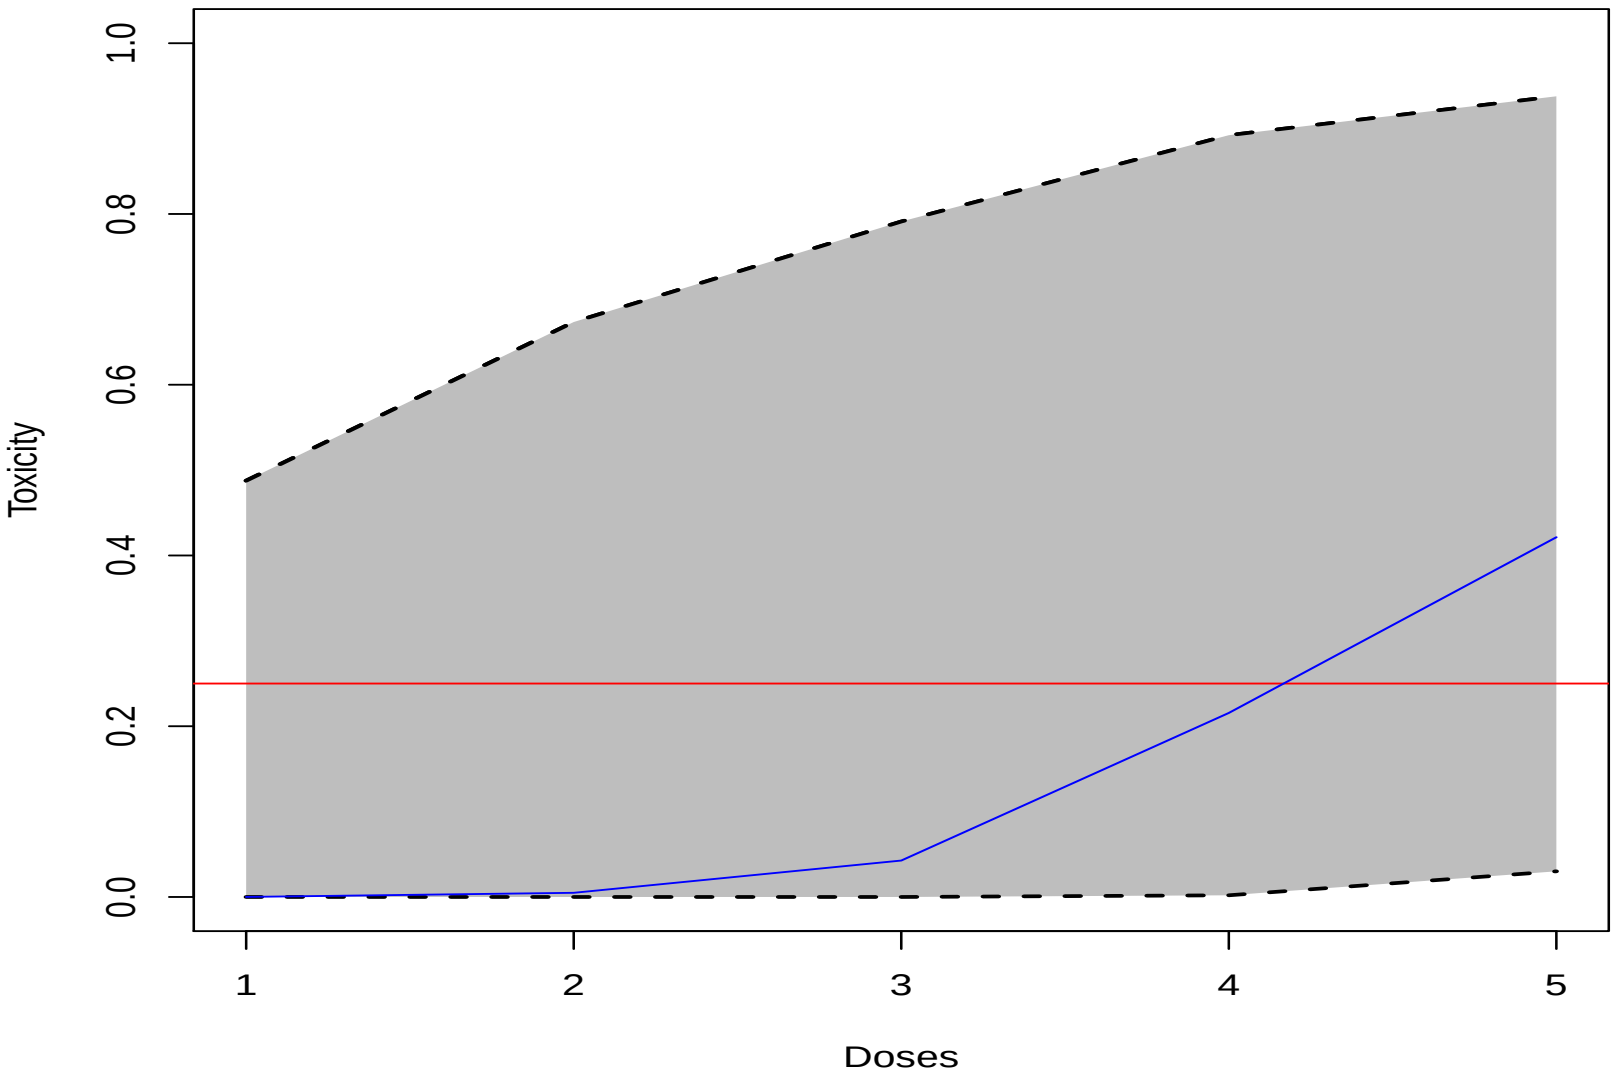

CRM: Support of the Model

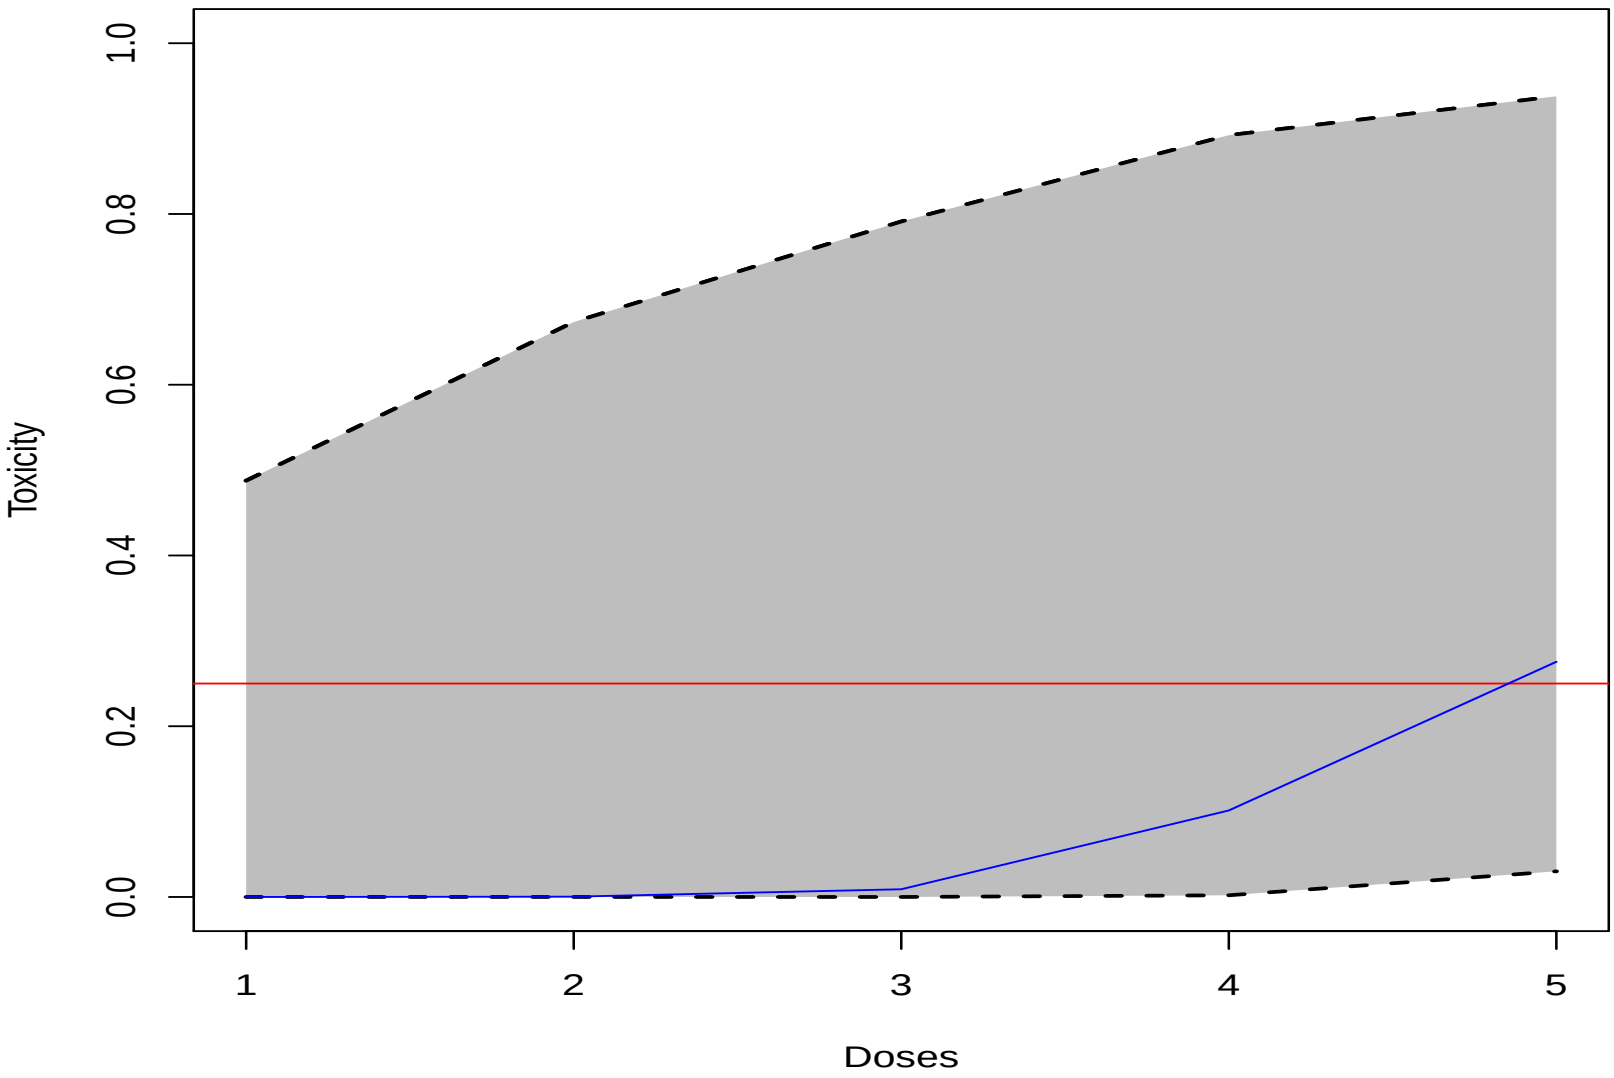

CRM: Support of the Model

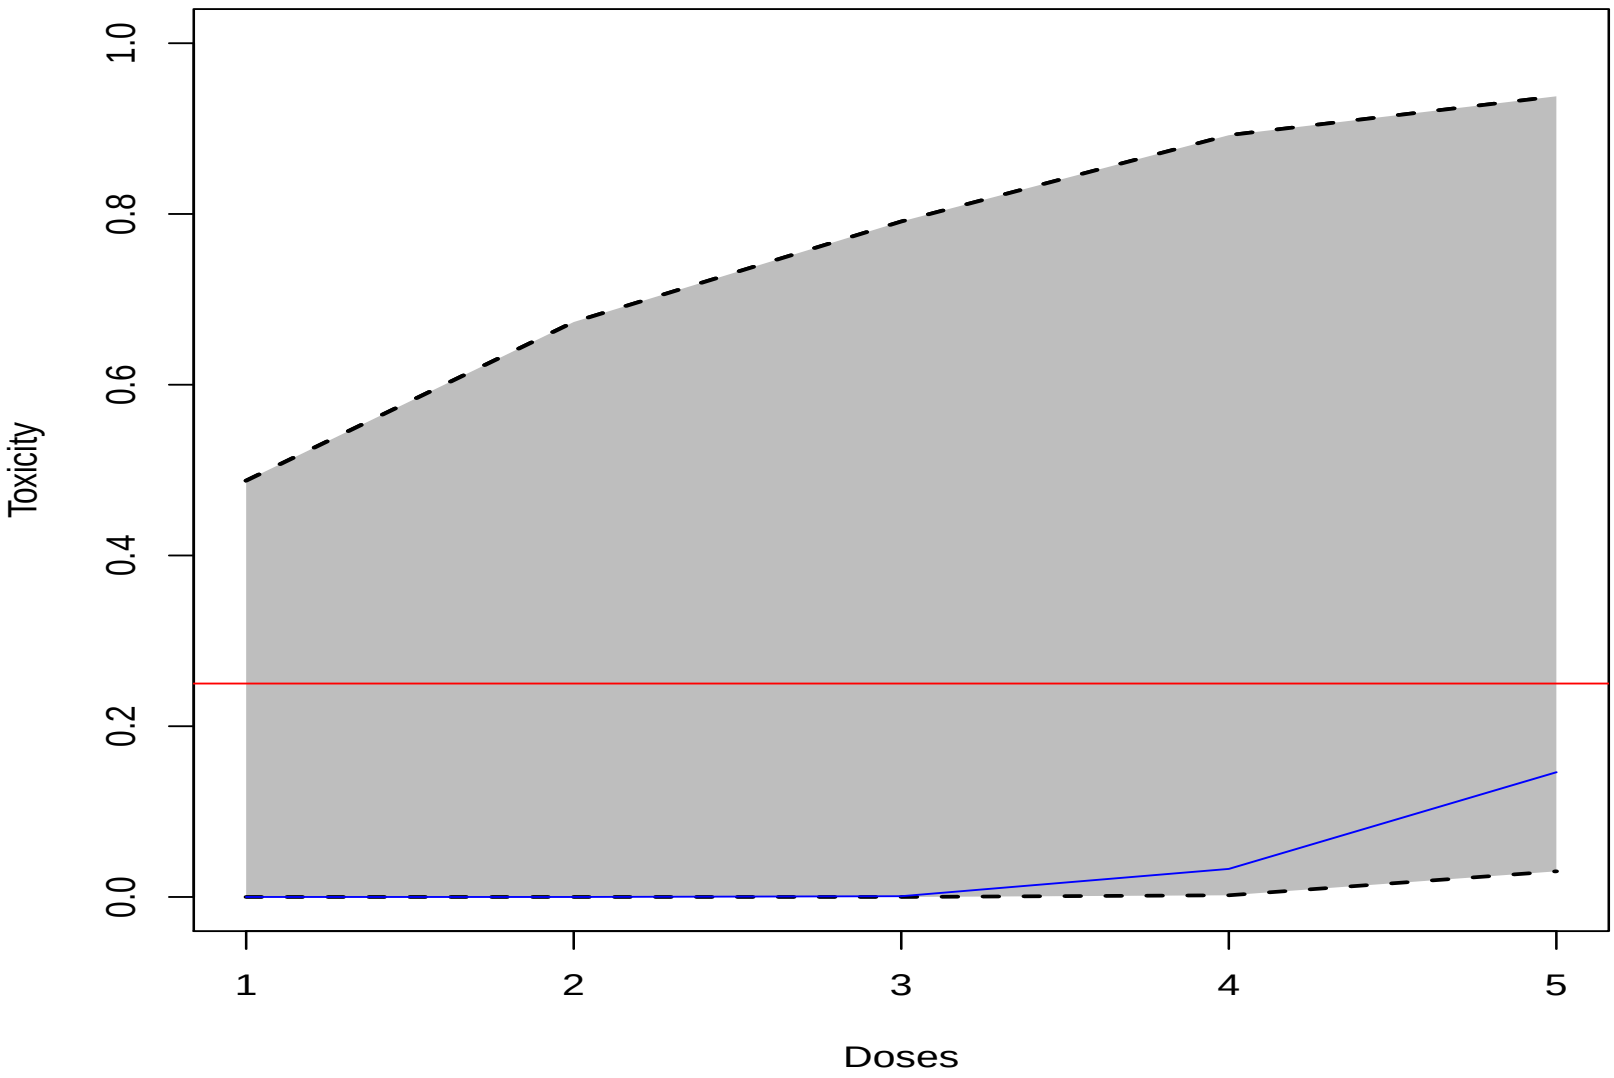

CRM: Support of the Model

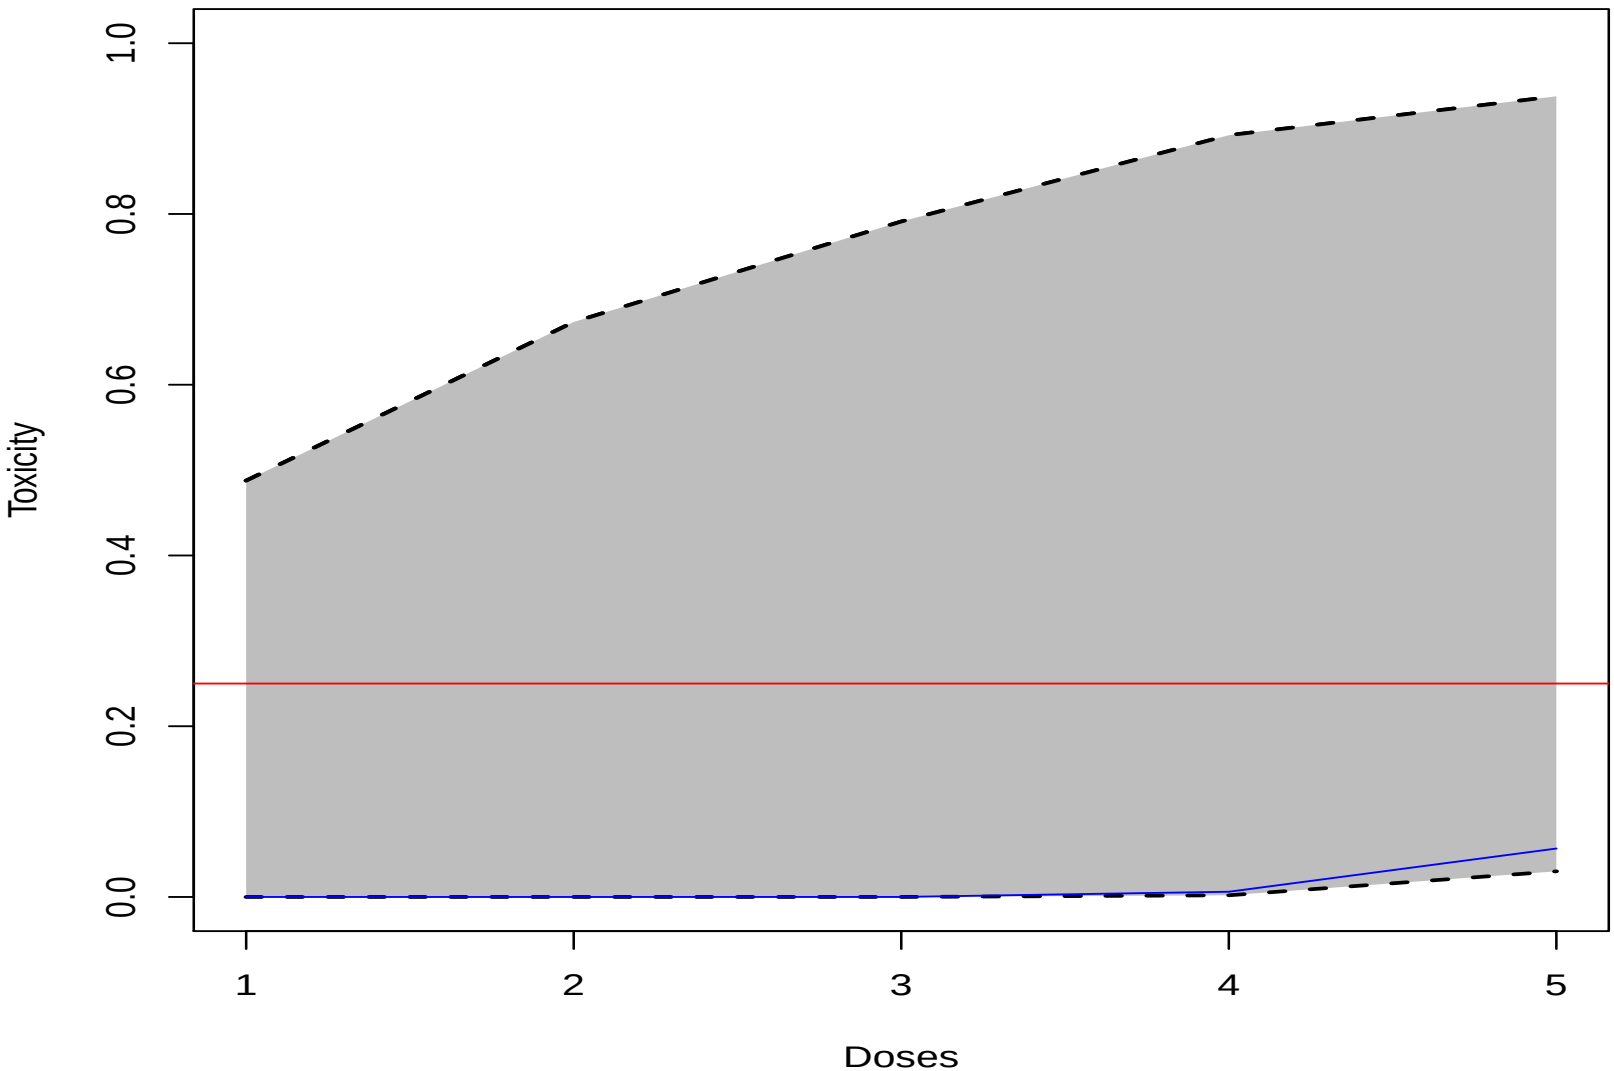

CRM: Support of the Model

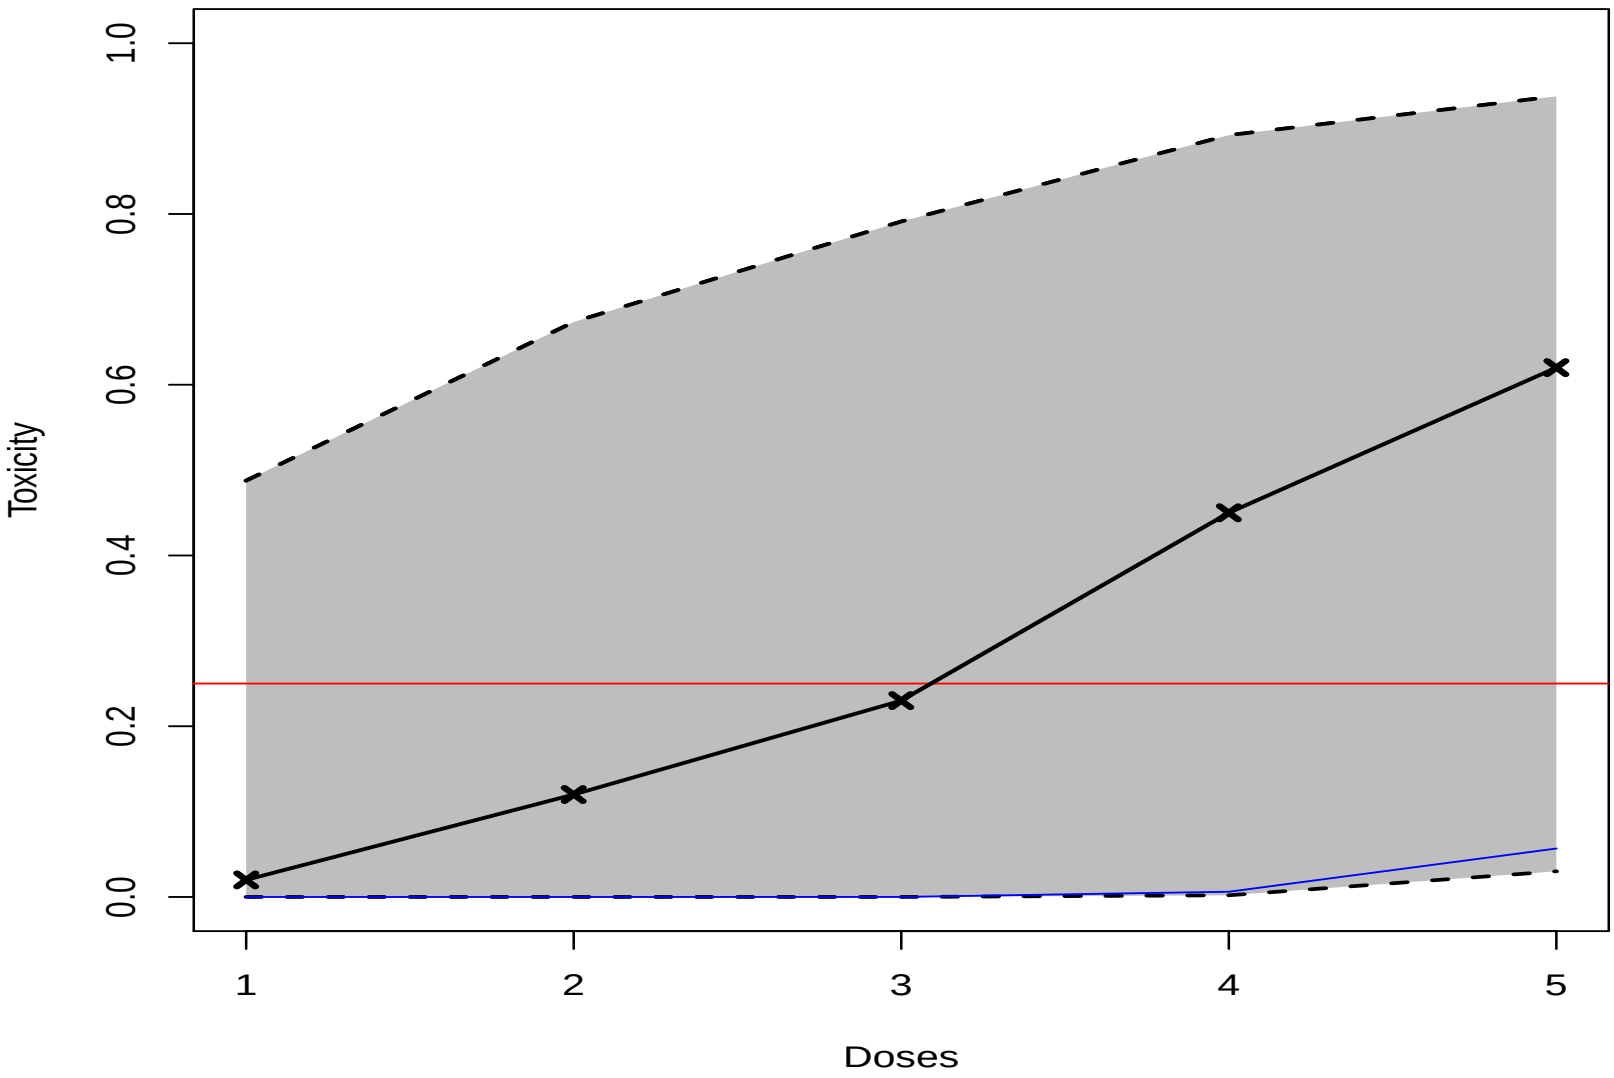

CRM: Support of the Model

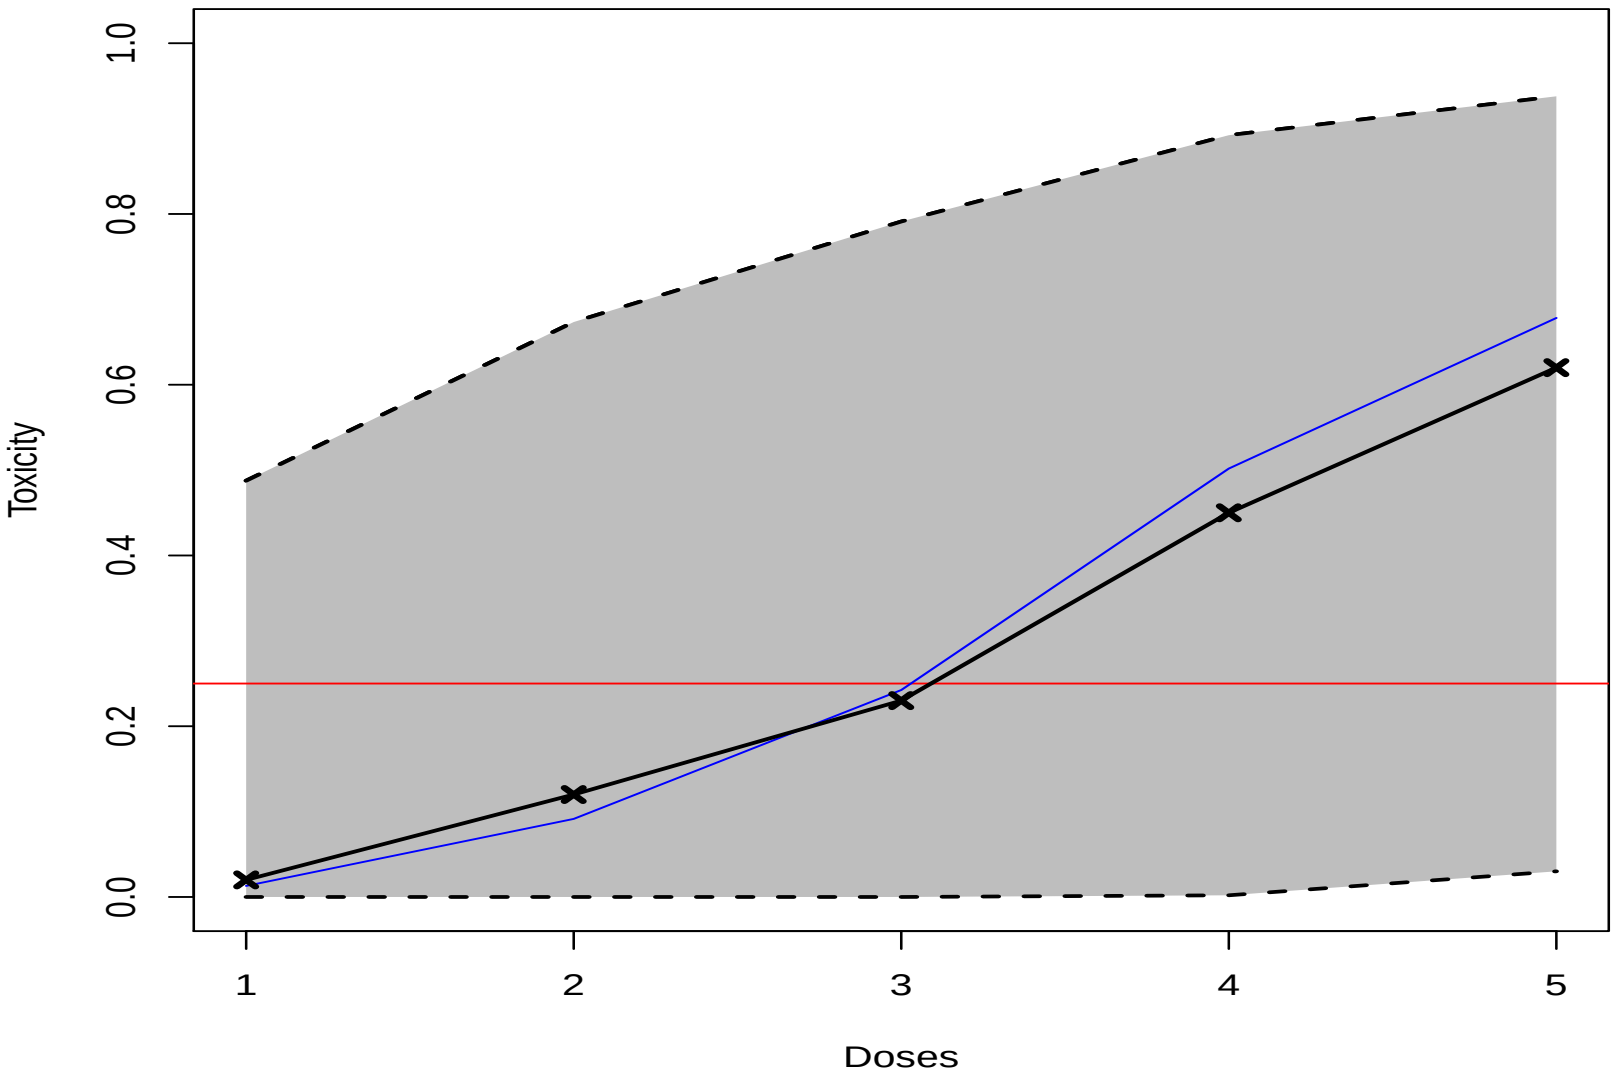

CRM: Support of the Model

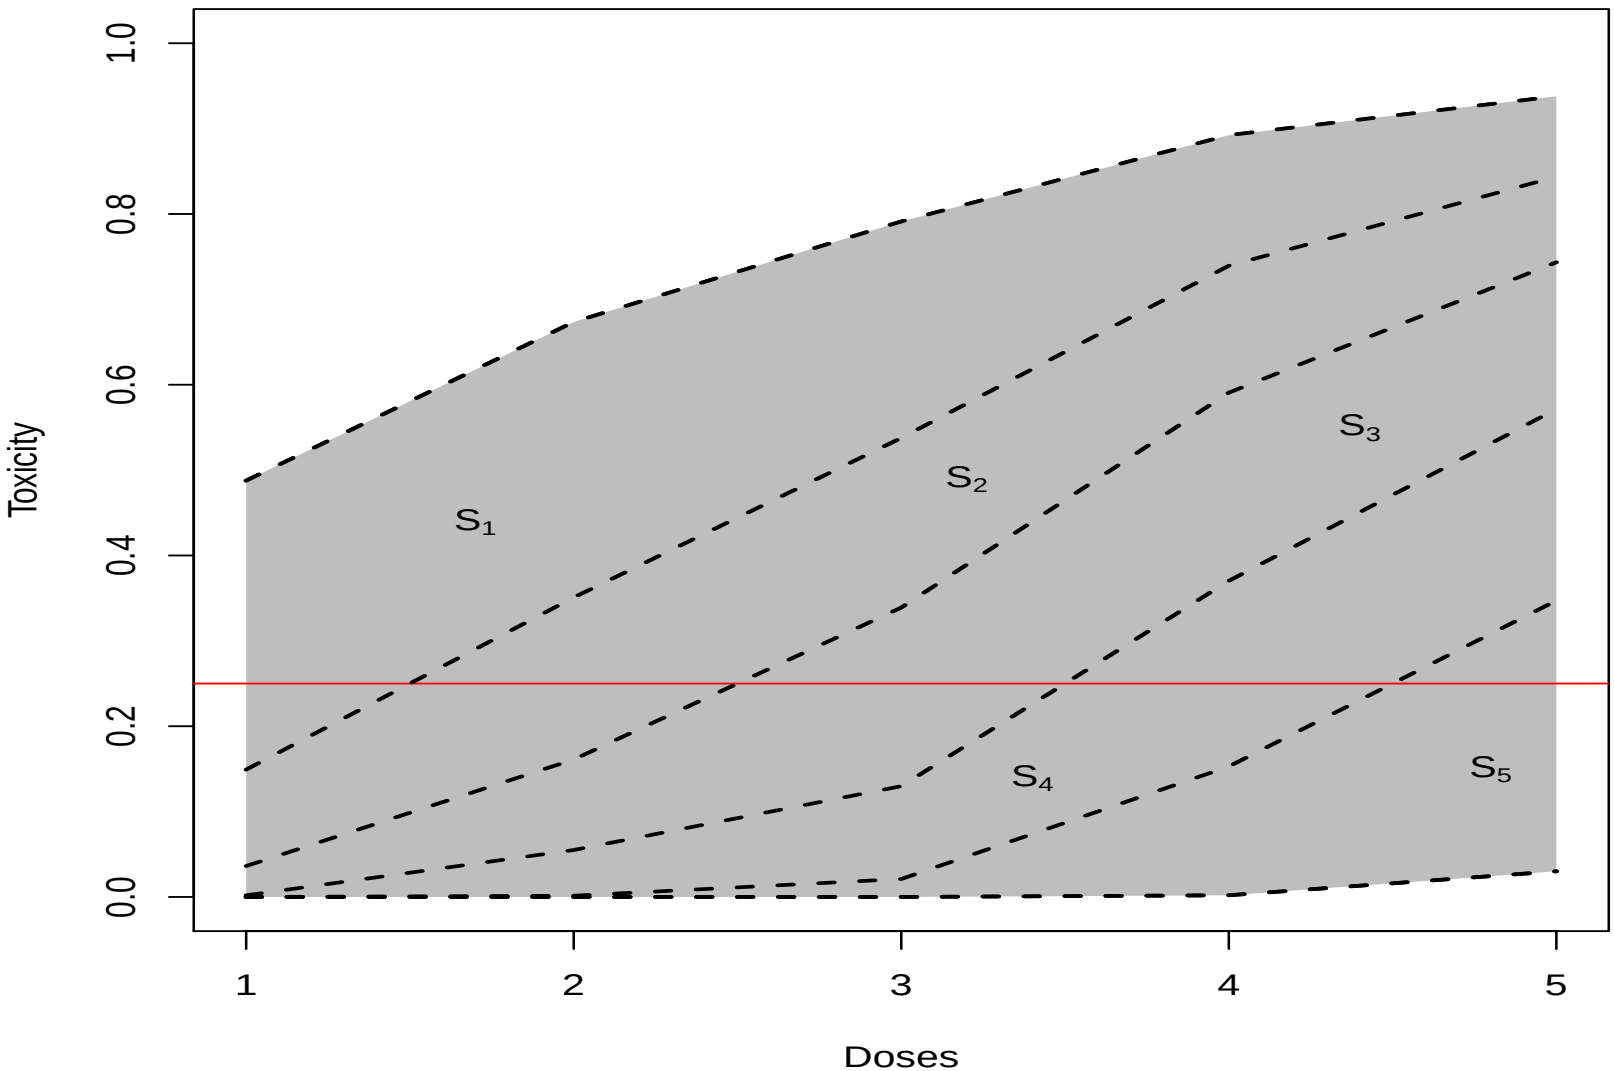

Good specification (Shen and O'quigley, 1996)

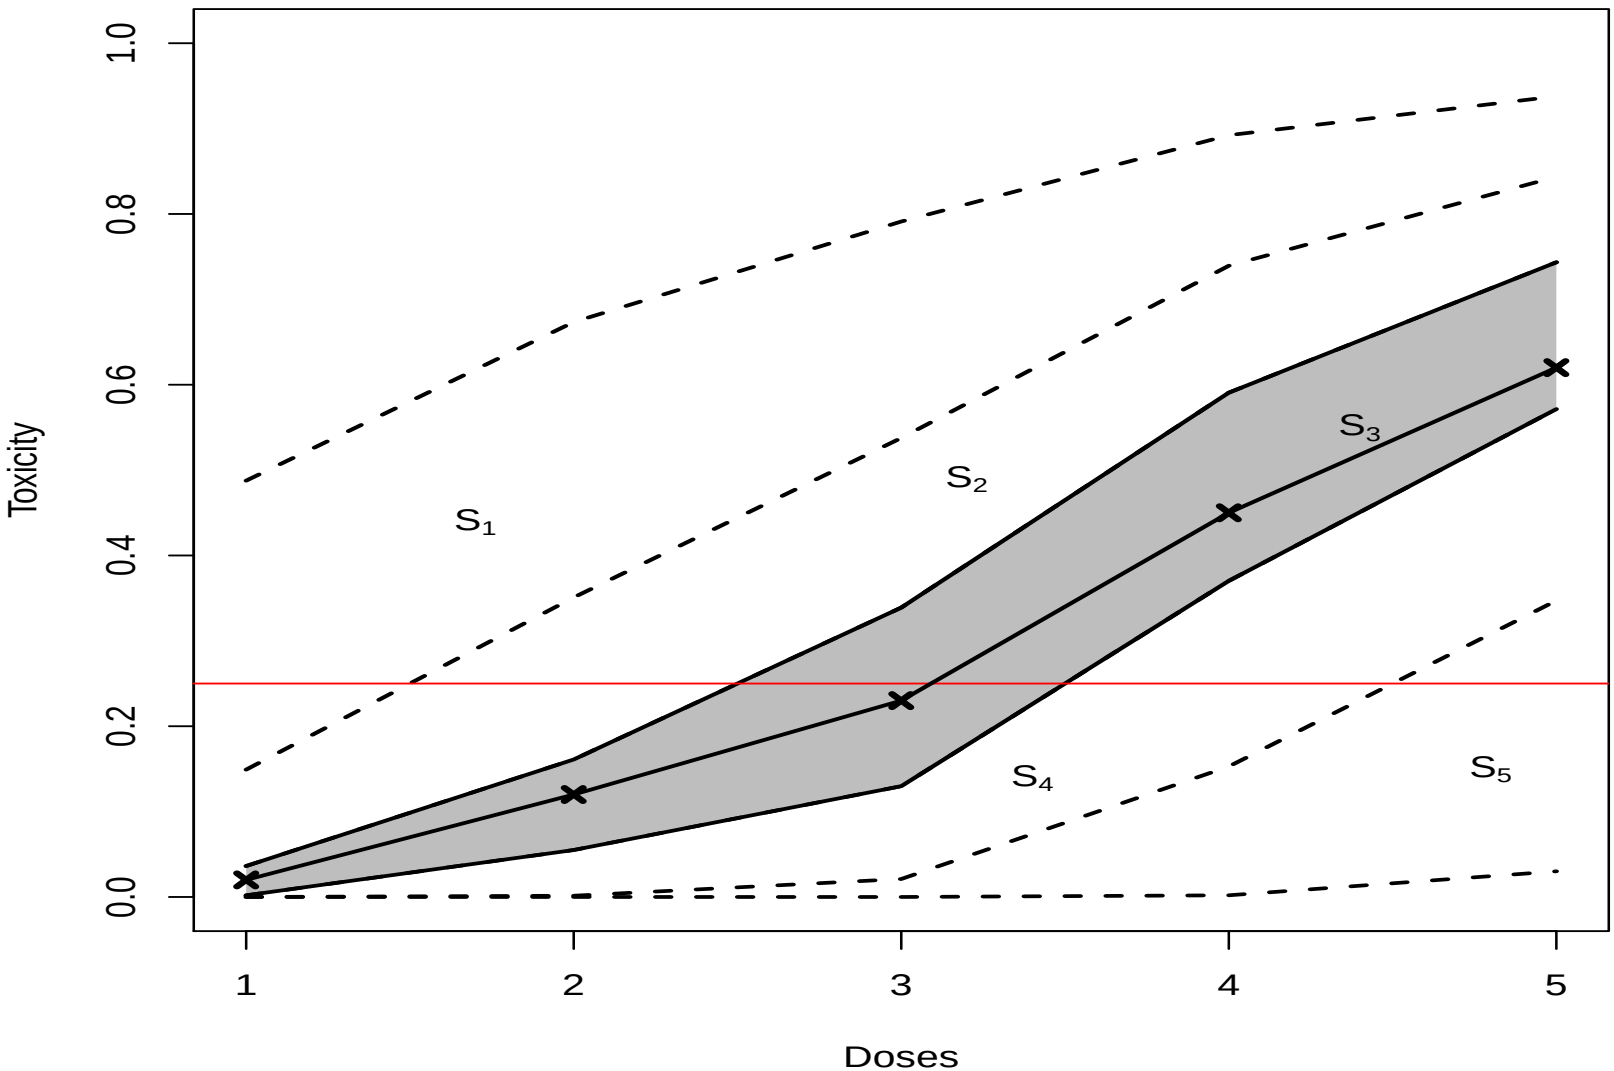

Poor specification (Cheung and Chappell, 2002, Azriel, 2012)

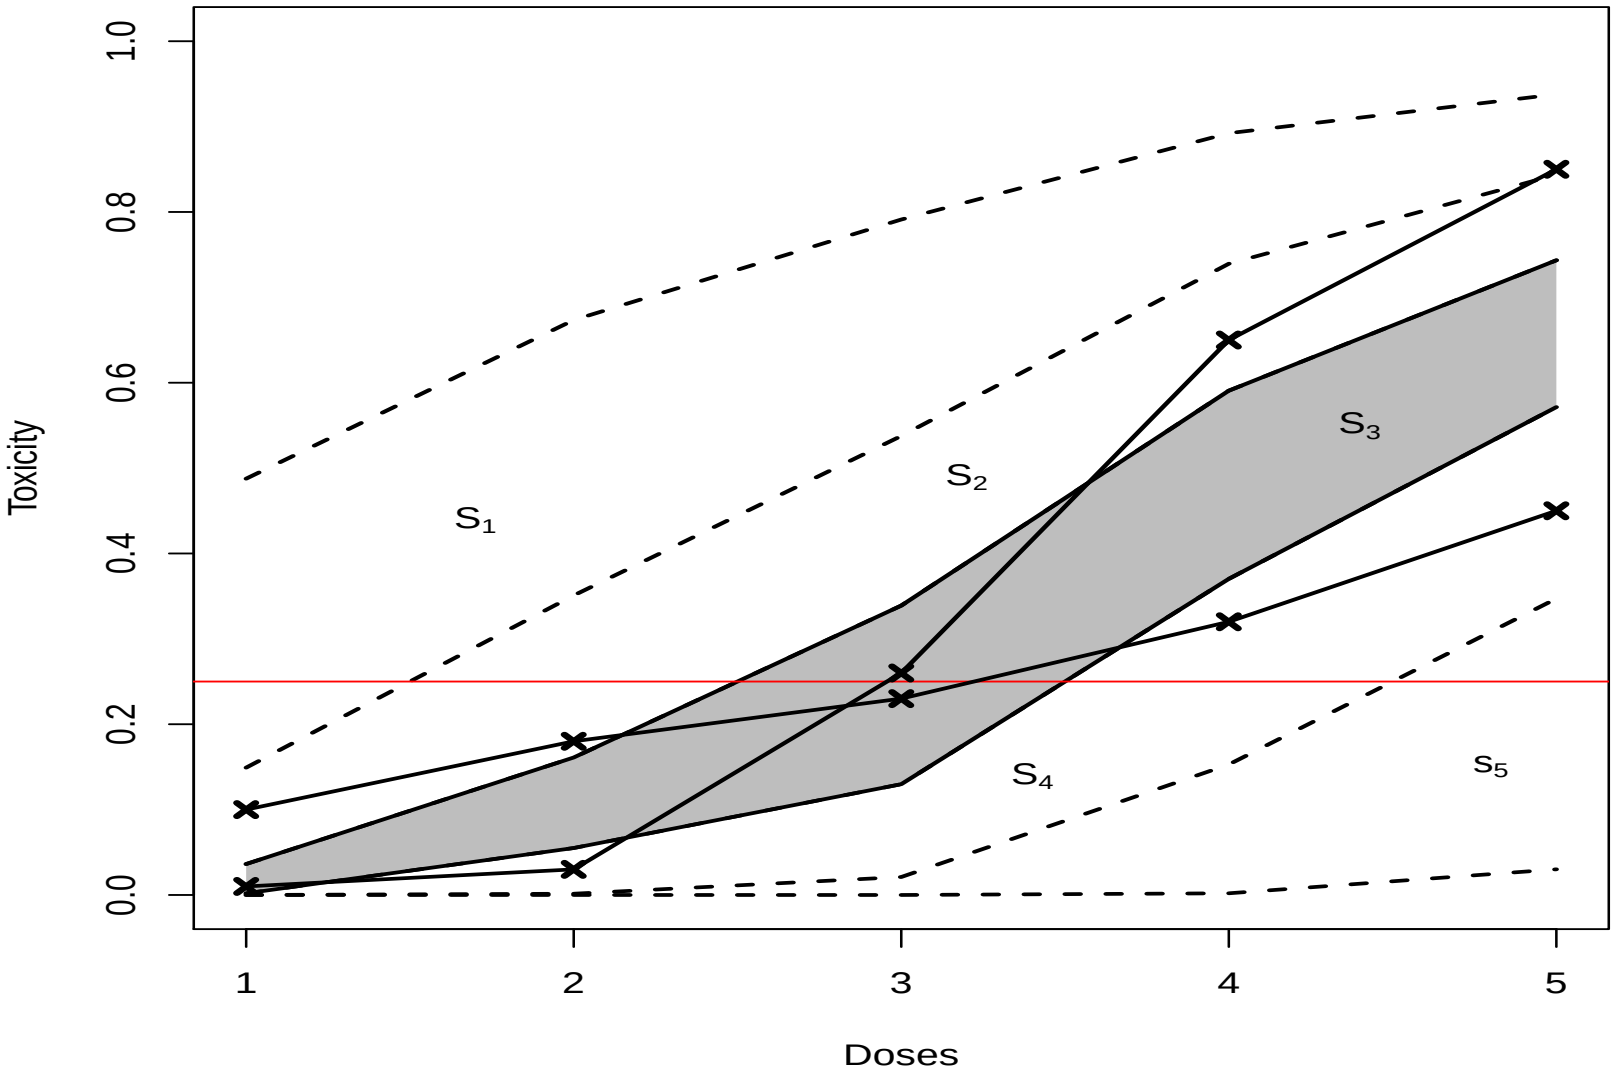

From CRM model to prior model

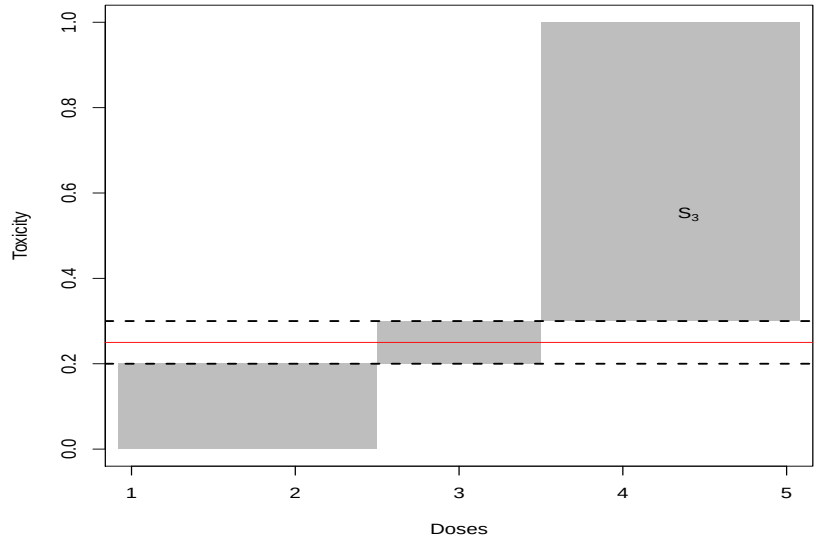

Truncated beta distribution are used. We set: $v^\theta = (\Psi(d|\alpha_\theta))_{d\in D}$, with α_θ such that: $\Psi(\theta|\alpha_\theta) = \alpha$. The dispersion parameter is c.

$$\Lambda_\theta = \prod_{j=1}^{\theta-1} \mathcal{B}_B(c\, v_j^\theta + 1,\, c(1 - v_j^\theta) + 1) \times \mathcal{B}_{I_\epsilon}(c\, v_\theta^\theta + 1, c(1 - v_\theta^\theta) + 1)$$
$$\times \prod_{j=\theta+1}^{m} \mathcal{B}_A(c\, v_j^\theta + 1, c(1 - v_j^\theta) + 1)\,.$$

From CRM model to prior model

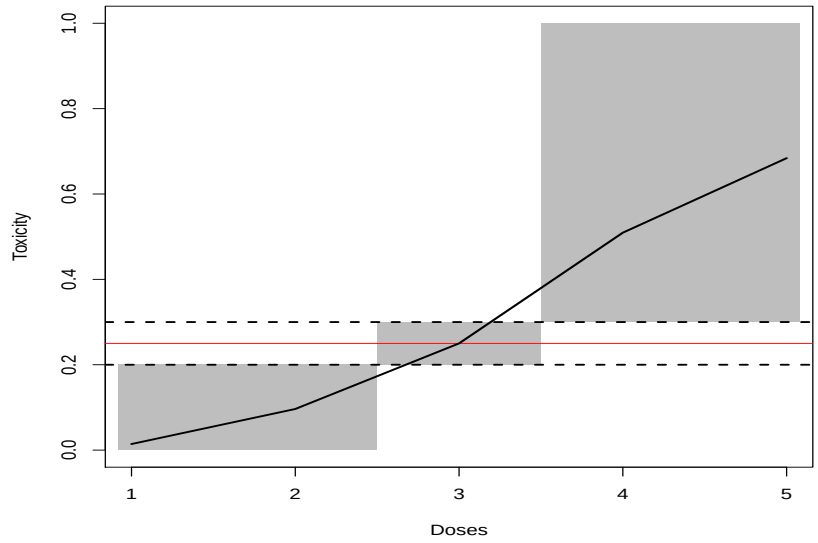

Truncated beta distribution are used. We set: $v^\theta = (\Psi(d|\alpha_\theta))_{d\in D}$, with α_θ such that: $\Psi(\theta|\alpha_\theta) = \alpha$. The dispersion parameter is c.

$$\Lambda_\theta = \prod_{j=1}^{\theta-1} \mathcal{B}_B(c\, v_j^\theta + 1,\, c(1 - v_j^\theta) + 1) \times \mathcal{B}_{I_\epsilon}(c\, v_\theta^\theta + 1, c(1 - v_\theta^\theta) + 1)$$
$$\times \prod_{j=\theta+1}^{m} \mathcal{B}_A(c\, v_j^\theta + 1, c(1 - v_j^\theta) + 1)\,.$$

From CRM model to prior model

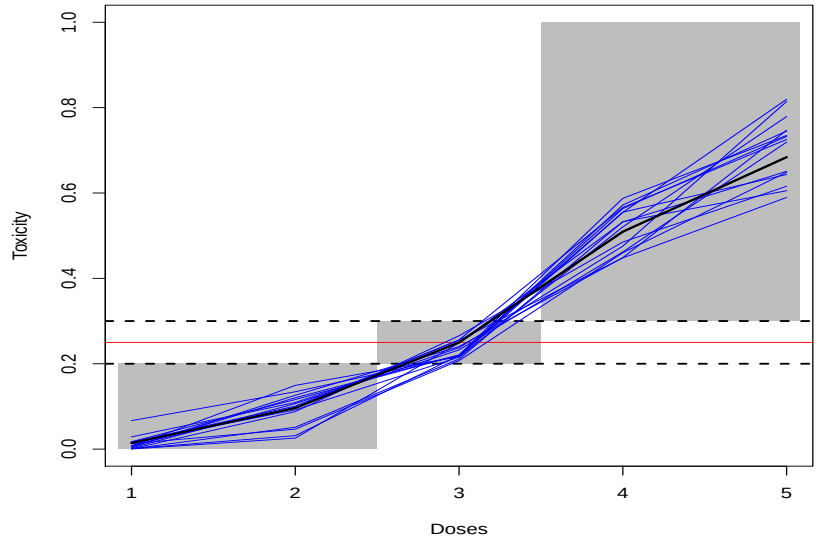

Truncated beta distribution are used. We set: $v^\theta = (\Psi(d|\alpha_\theta))_{d\in D}$, with α_θ such that: $\Psi(\theta|\alpha_\theta) = \alpha$. The dispersion parameter is c.

$$\Lambda_\theta = \prod_{j=1}^{\theta-1} \mathcal{B}_B(c\, v_j^\theta + 1,\, c(1-v_j^\theta)+1) \times \mathcal{B}_{I_\epsilon}(c\, v_\theta^\theta + 1, c(1-v_\theta^\theta)+1)$$

$$\times \prod_{j=\theta+1}^{m} \mathcal{B}_A(c\, v_j^\theta + 1, c(1-v_j^\theta)+1)\,.$$

Decomposition of the CRM parameter space

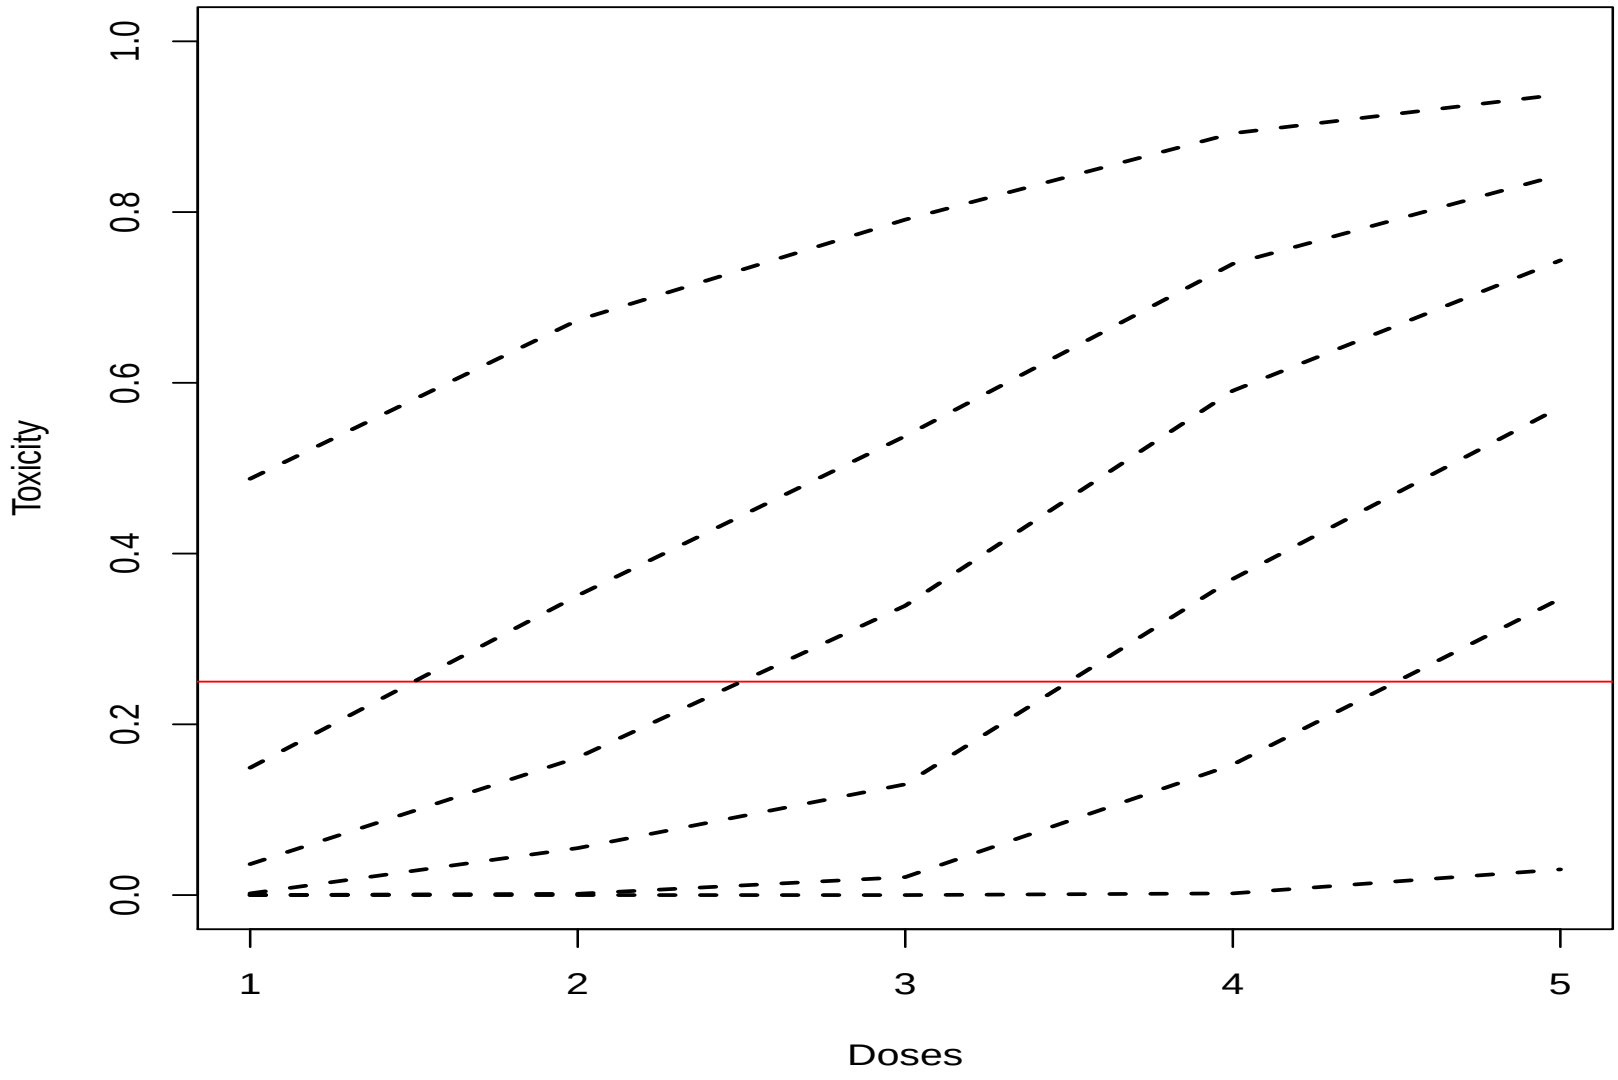

Sample under the prior model of the Semi-Parametric CRM

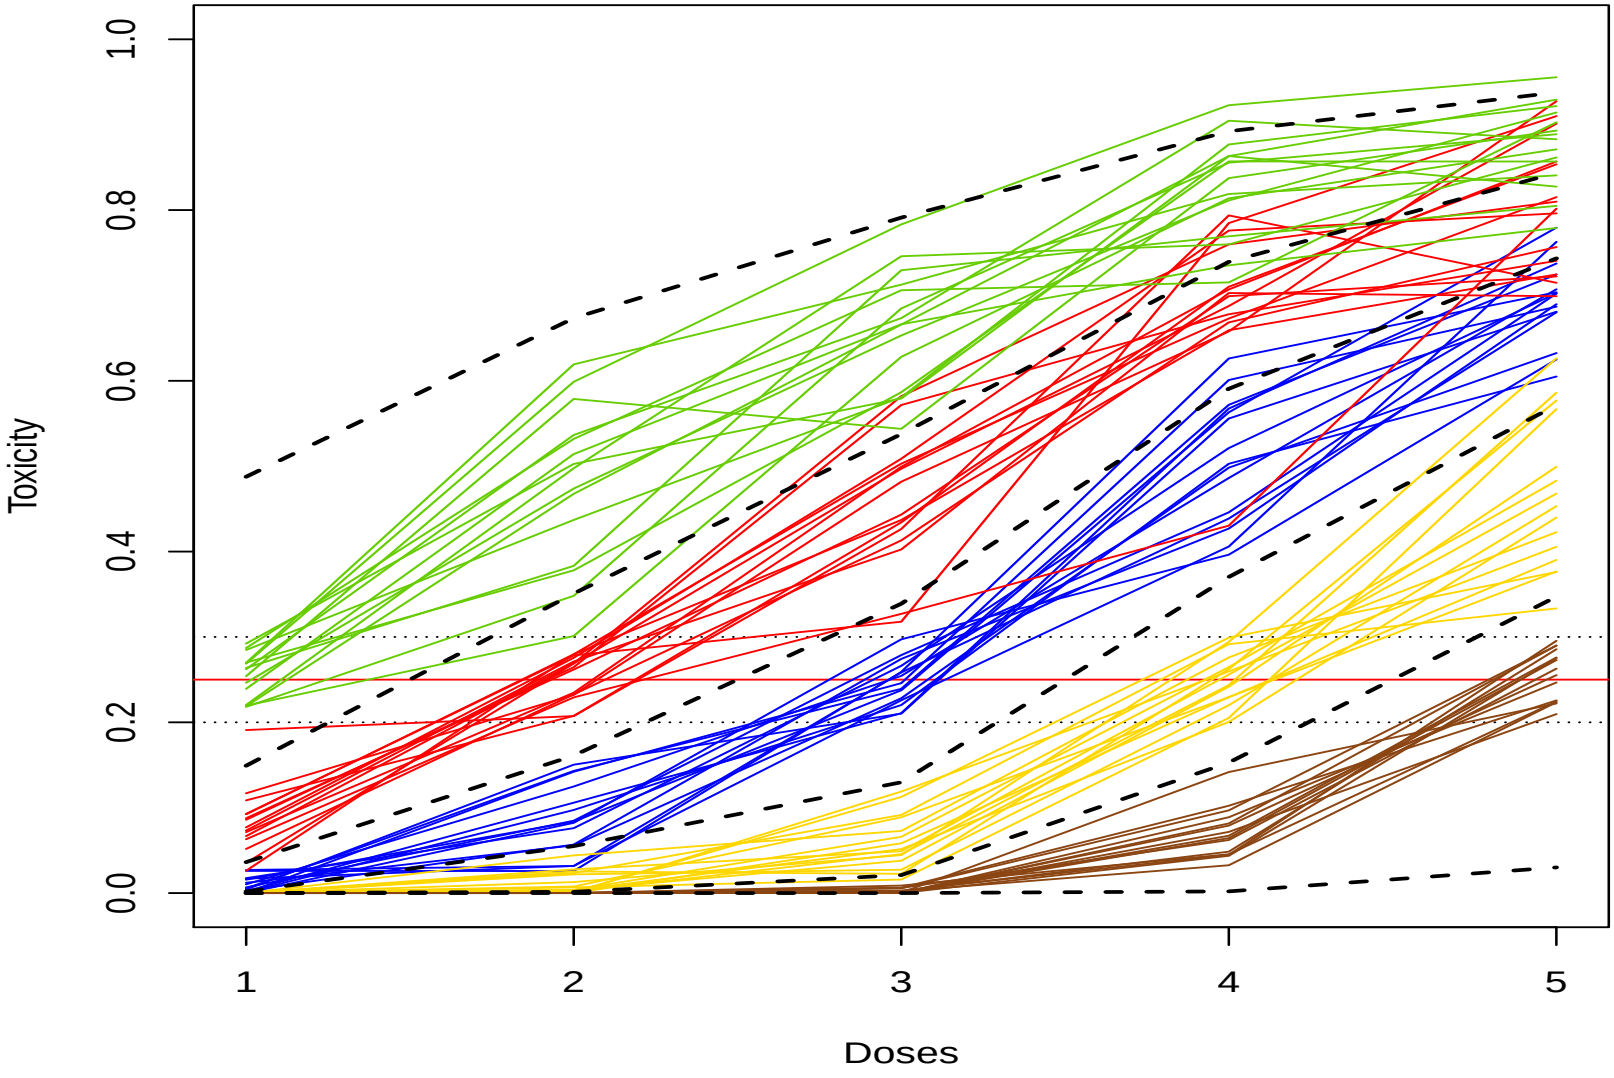

Plan

Semi-Parametric Models

Prior Model and CRM Model

Theoretical Properties

Simulations II

Asymptotic behavior

Impossibility Theorem (Azriel and al.)

Let $\mathcal{M}$ be a method and $\mathcal{M}(X_1^n, Y_1^n)$ the current estimator of the next dose. A scenario β satisfying the monotonicity exists such that:

$$\mathbb{P}_\beta(\exists N : \forall n > N,\, \mathcal{M}(X_1^n, Y_1^n) = MTD) < 1\,.$$

Asymptotic behavior

The two consecutive doses associated with a probability of toxicity on each part of the threshold α are a (above) and b (below).

Definition

Consider the collection of doses associated with a toxicity belonging to $I = [\alpha \pm \epsilon] : \mathcal{E}(I, \beta) = \{j \in D : \beta_j \in I\}$.
(a) A method is ϵ-sensitive, if for all β such that $\mathcal{E}(I, \beta) \neq \emptyset$, we have: $\mathbb{P}_\beta \left[\exists N\,, \forall n > N : X_n \in \mathcal{E}(I, \beta)\right] = 1$.
(b) A method is balanced, if for all β such that $\mathcal{E}(I, \beta) = \emptyset$, we have:

$$X_n \xrightarrow{S} \{a, b\}\,,\ a.s., \text{ i.e., } \sup_{x \in \{a,b\}} \left(\liminf_{n \to +\infty} d(X_n, x) \right) = 0\,,\ a.s.$$

Theorem

*Under assumptions **'Structure'**, **'Independence'** and some regularity properties on the prior model, the SPM is ϵ-sensitive and balanced.*

Balanced behavior

When a method is balanced, there exist estimators on the basis of observations which are almost surely convergent to the MTD.

Theorem

*Assume that the central interval I is chosen equal to $\{\alpha\}$. Under assumptions **'Structure'**, **'Independence'** and some regularity properties on the prior model, the SPM is always balanced and we have:*

$$\frac{n_b}{n_a} \underset{n\to+\infty}{\longrightarrow} \frac{D_{KL}(\beta_a||\alpha)}{D_{KL}(\beta_b||\alpha)}\ , \quad \textit{with}\ \ n_d = \sum_{k=1}^{n} \mathbb{1}_{\{X_k=d\}}.$$

Suppose that $\alpha = 0.25$:
if $\beta_a = 0.35$ and $\beta_b = 0.23$, then $n_a \approx 4.3\%$ of n_b, asymptotically;
if $\beta_a = 0.27$ and $\beta_b = 0.15$, then $n_b \approx 3.5\%$ of n_a, asymptotically.

Plan

Semi-Parametric Models

Prior Model and CRM Model

Theoretical Properties

Simulations II

Simulation Set-Up

Trial: 6 doses, $\alpha = 0.20$, patients enrolled in cohorts of size 1.

A two stage CRM (Bayesian and Likelihood):

- The prior is $\mathcal{N}(0, 1.34^2)$
- $t = (0.05, 0.10, 0.20, 0.35, 0.50, 0.70)$
- power model: $\Psi_t(x, a) = t(x)^{\exp(a)}$

A Semi-parametric CRM:

- $\epsilon = 0.015$
- $v^\theta = (\Psi_t(x, \alpha_\theta))_{x \in D}$
- $c = 48$
- $\Pi = (0.1929, 0.1928, 0.1755, 0.1704, 0.1518, 0.1165)$, the closest vector to a uniform distribution that still follows the same sequence of allocation as CRM when no DLTs are observed.

Table: Some particular scenarios.

Doses		1	2	3	4	5	6
Scenario 1		**0.05**	**0.10**	**0.20**	**0.35**	**0.50**	**0.70**
PS	SP-CRM	2.3	22.7	54.0	19.7	01.2	0.0
	CRM	02.4	22.2	53.9	20.2	01.3	0.0
PA	SP-CRM	10.8	24.3	39.0	19.0	05.9	00.7
	CRM	12.3	22.1	37.7	20.4	06.4	00.8
Scenario 2		**0.20**	**0.26**	**0.28**	**0.3**	**0.35**	**0.50**
PS	SP-CRM	49.4	21.5	13.2	9.6	5.4	0.6
	CRM	48.1	19.5	14.3	11.2	6.0	0.6
PA	SP-CRM	47.4	20.6	13.5	9.1	7.0	2.1
	CRM	47.6	17.6	14.1	10.7	7.6	2.2
Scenario 3		**0.01**	**0.02**	**0.05**	**0.09**	**0.18**	**0.40**
PS	SP-CRM	0.0	0.2	2.8	20.3	59.2	17.3
	CRM	0.0	0.1	3.4	21.8	58.4	16.1
PA	SP-CRM	4.6	6.0	10.5	19.9	40.7	17.9
	CRM	4.9	5.3	9.7	20.7	40.1	19.0

PS: Percentage of final selection at each dose among 10 000 trials.
PA: Percentage of patients treated at each dose among 10 000 of 25 patients.

Table: Scenarios that differ from the CRM model.

Doses		1	2	3	4	5	6
Scenario 4		**0.0**	**0.0**	**0.0**	**0.23**	**0.3**	**0.35**
PS	SP-CRM	0.0	0.0	10.2	56.8	23.6	9.2
	CRM	0.0	0.0	10.5	52.3	26.9	10.2
PA	SP-CRM	4.0	4.0	19	38.8	22.8	11.1
	CRM	4.0	4.0	16.9	37.8	24.4	12.7
Scenario 5		**0.0**	**0.0**	**0.16**	**0.3**	**0.35**	**0.4**
PS	SP-CRM	0.0	2.3	51.7	31.5	11.1	3.2
	CRM	0.0	3.5	46.7	33.6	12.6	3.6
PA	SP-CRM	4.0	11.8	40.3	24.3	13.7	5.8
	CRM	4.7	11.0	36.3	26.7	14.5	6.5
Scenario 6		**0.01**	**0.02**	**0.05**	**0.11**	**0.14**	**0.21**
PS	SP-CRM	0.0	0.1	3.2	15.7	31.0	49.8
	CRM	0.0	0.1	3.4	15.5	31.2	49.6
PA	SP-CRM	4.6	5.8	10.8	16.7	26.7	35.1
	CRM	4.9	5.3	10.2	16.7	25.9	36.0

PS: Percentage of final selection at each dose among 10 000 trials.
PA: Percentage of patients treated at each dose among 10 000 of 25 patients.

Percentage of correct selection for scenario 6

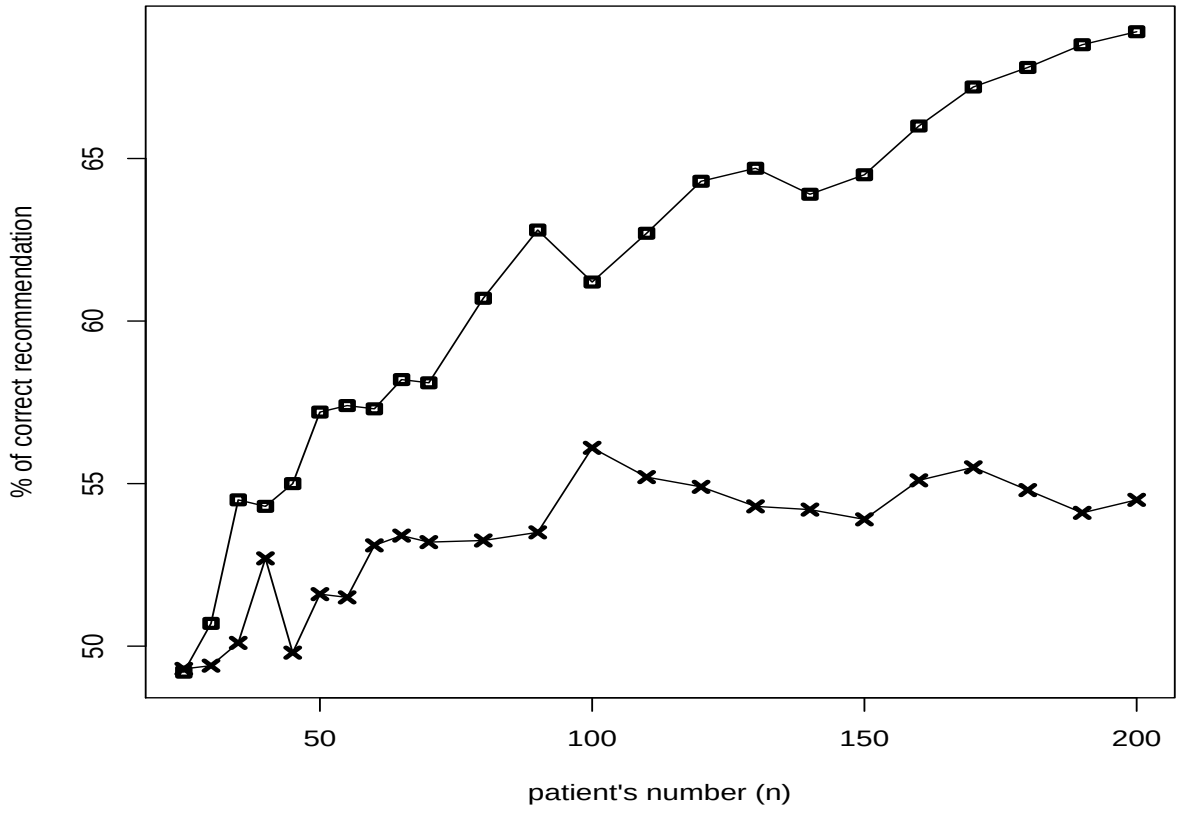

Figure: For scenario 6, (PCS) as a function of the number of included patients in the study. ■: SP-CRM ; ×: CRM.

Percentage of correct selection for 100 000 scenarios[a]

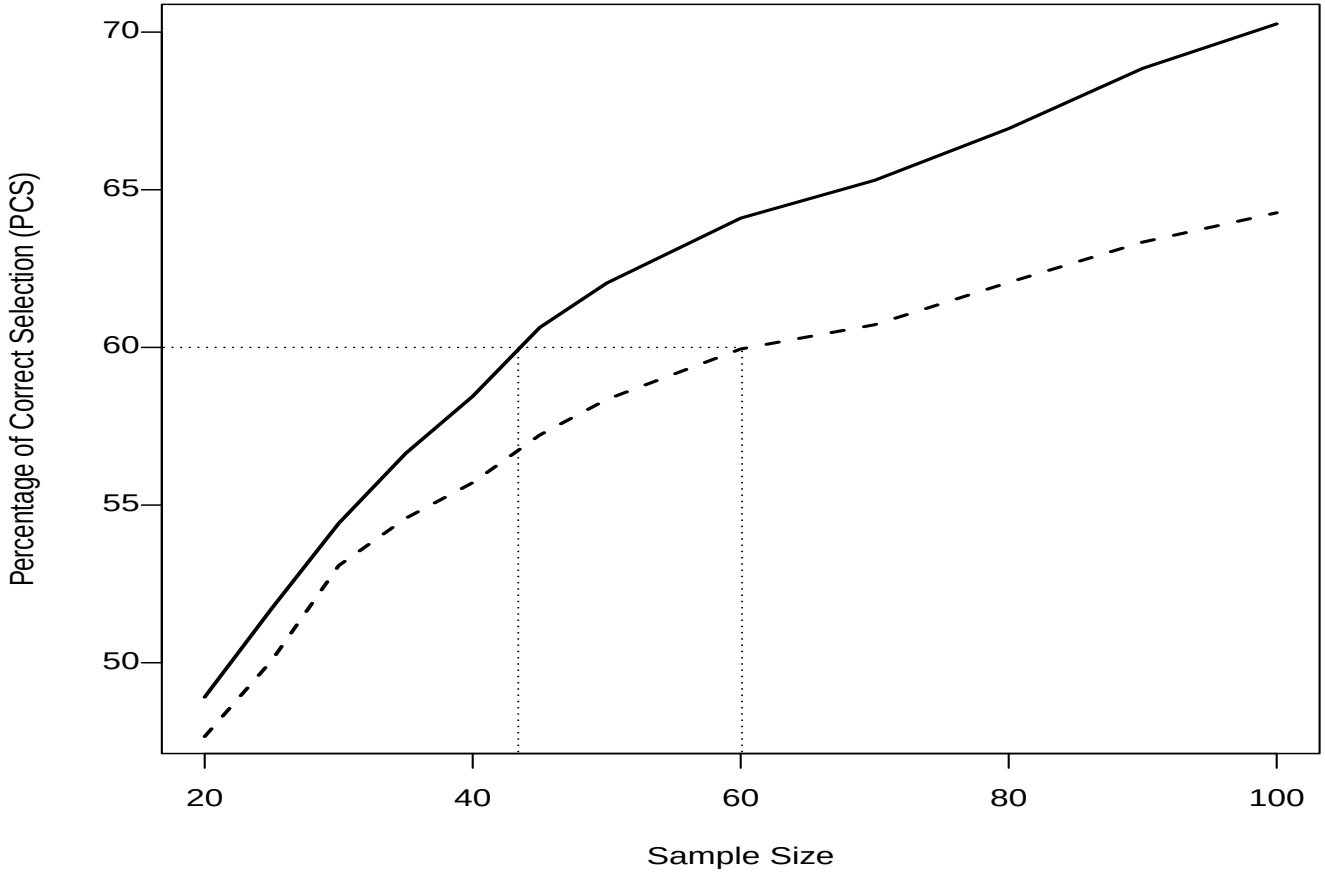

Figure: SPM: black curves; CRM: dotted curves.

[a] generated by a uniform spacing algorithm.

Percentage of correct allocation for 100 000 scenarios [a]

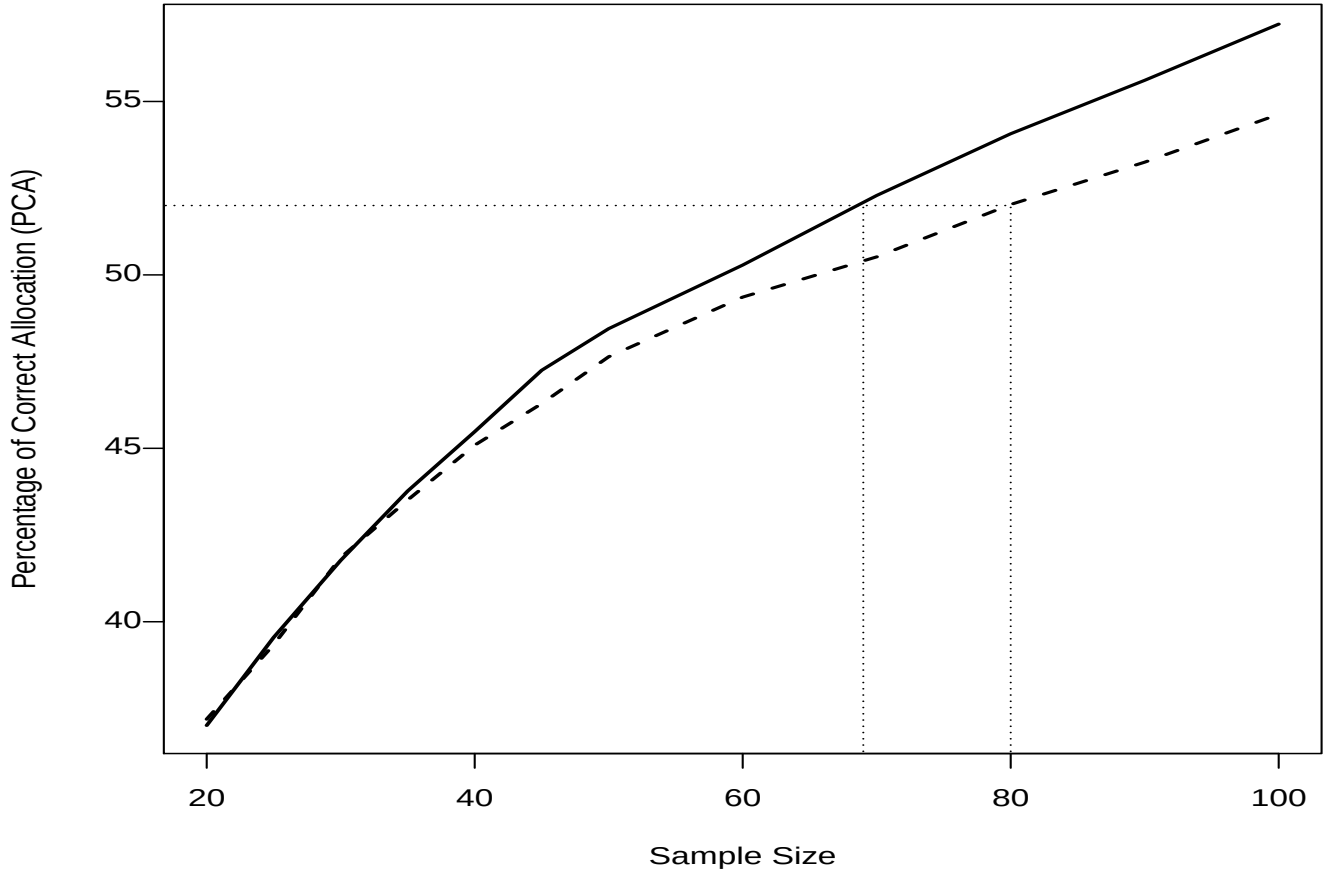

Figure: SPM: black curves; CRM: dotted curves.

[a] generated by a uniform spacing algorithm.

Simulation Set-Up

Trial: 8 doses, $\alpha = 0.25$, patients enrolled in cohorts of size 1.
A Semi-parametric model sliding on five doses:

- $\epsilon = 0$
- For all θ : $v^{\theta}_{\theta-1} = 0.6 \times \alpha$, $v^{\theta}_{\theta+1} = 1.4 \times \alpha$, $v^{\theta}_{j} = 0.4 \times \alpha$ when $j < \theta - 1$ and $v^{\theta}_{j} = 1.6 \times \alpha$ when $j > \theta + 1$.
- $c = 40$
- $\Pi \propto (1.000, 0.9428, 0.8494, 0.7653, 0.6895, 0.6213, 0.5598, 0.5043)$, a vector defined to respect some desirable properties at the beginning of the trial.

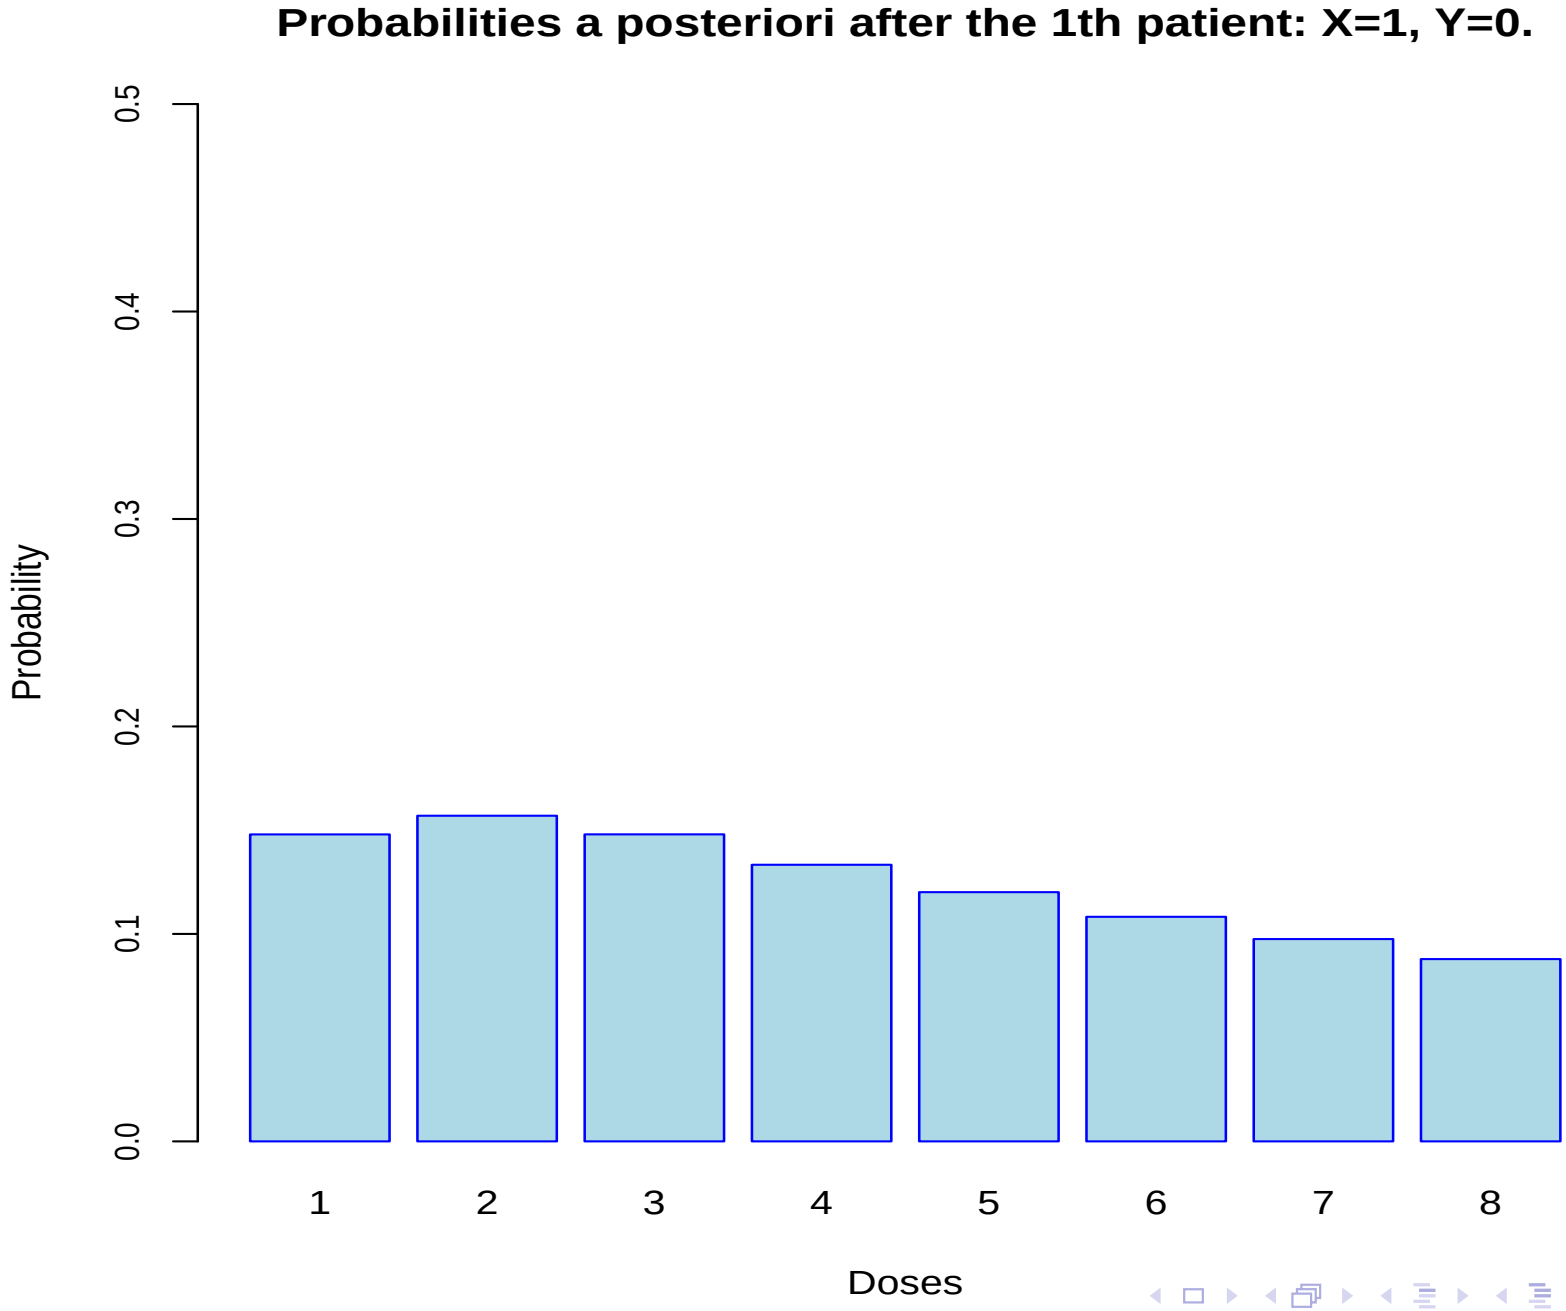
Probabilities a posteriori after the 1th patient: X=1, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

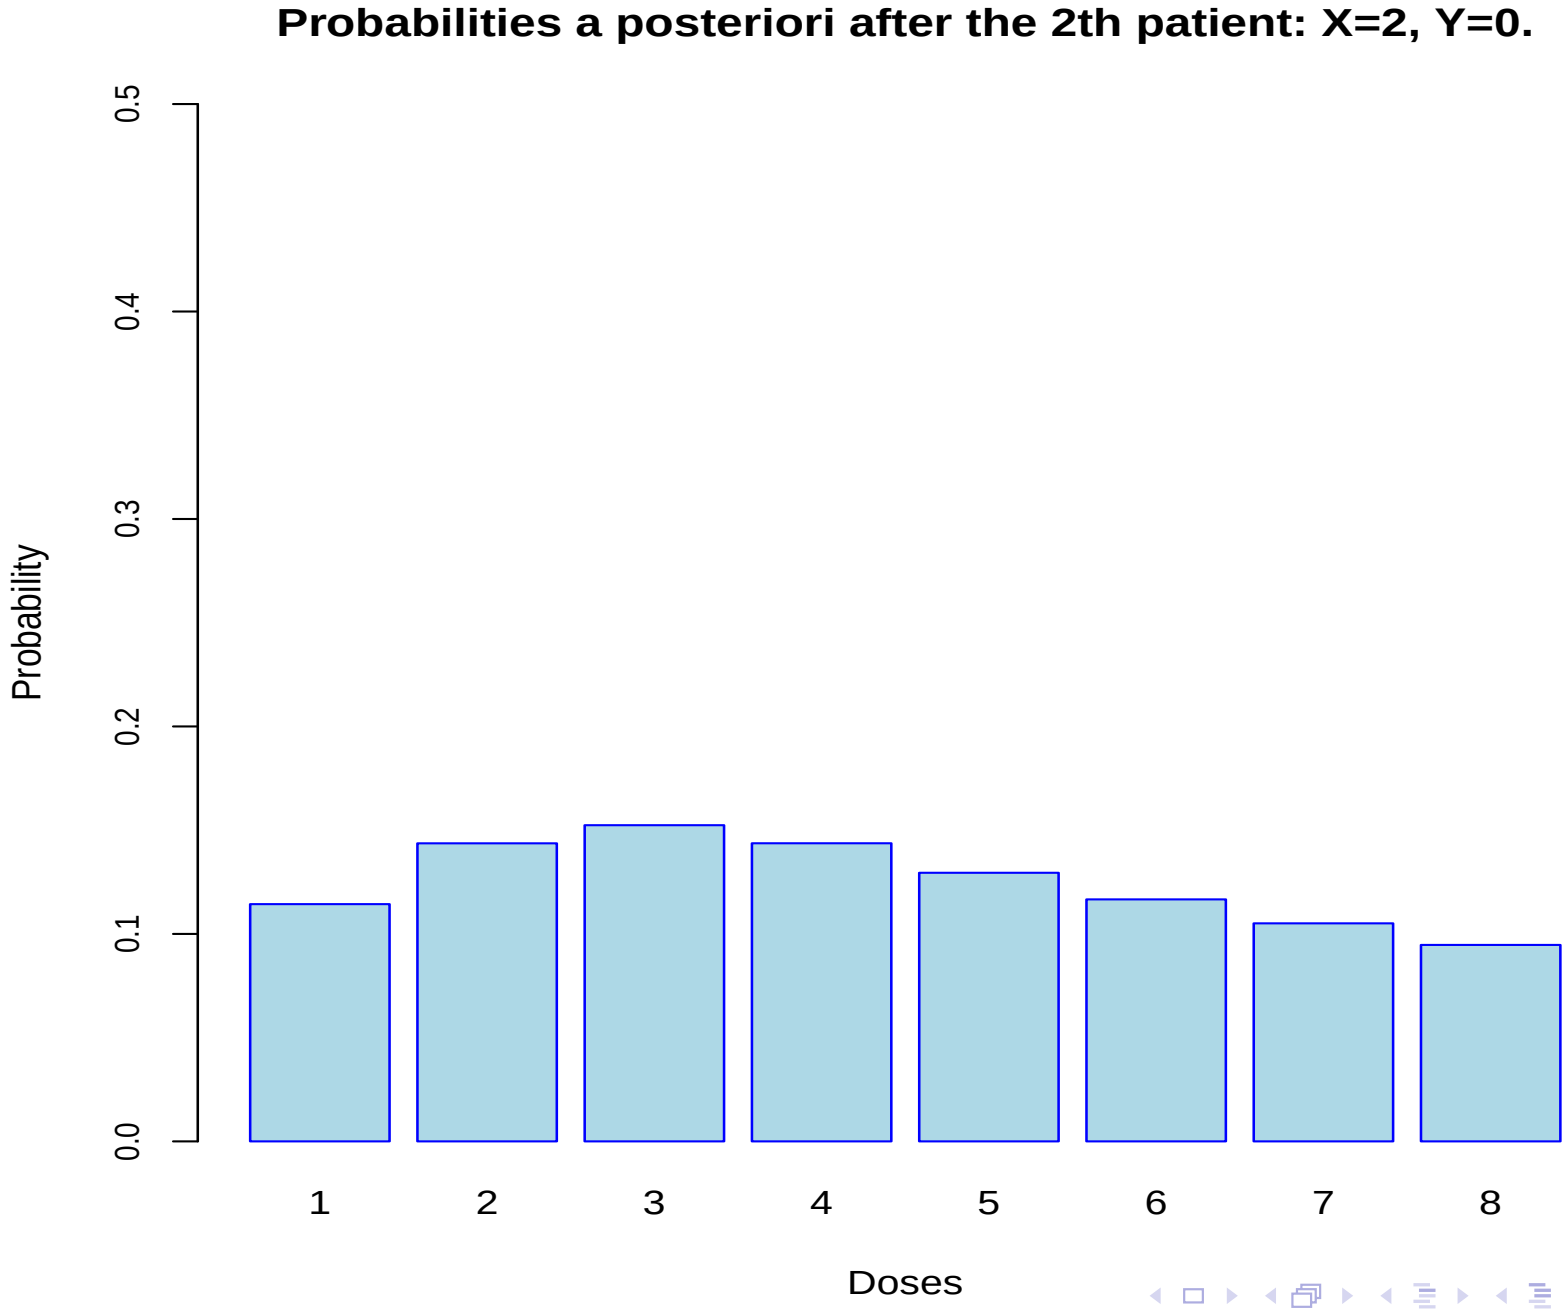
Probabilities a posteriori after the 2th patient: X=2, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

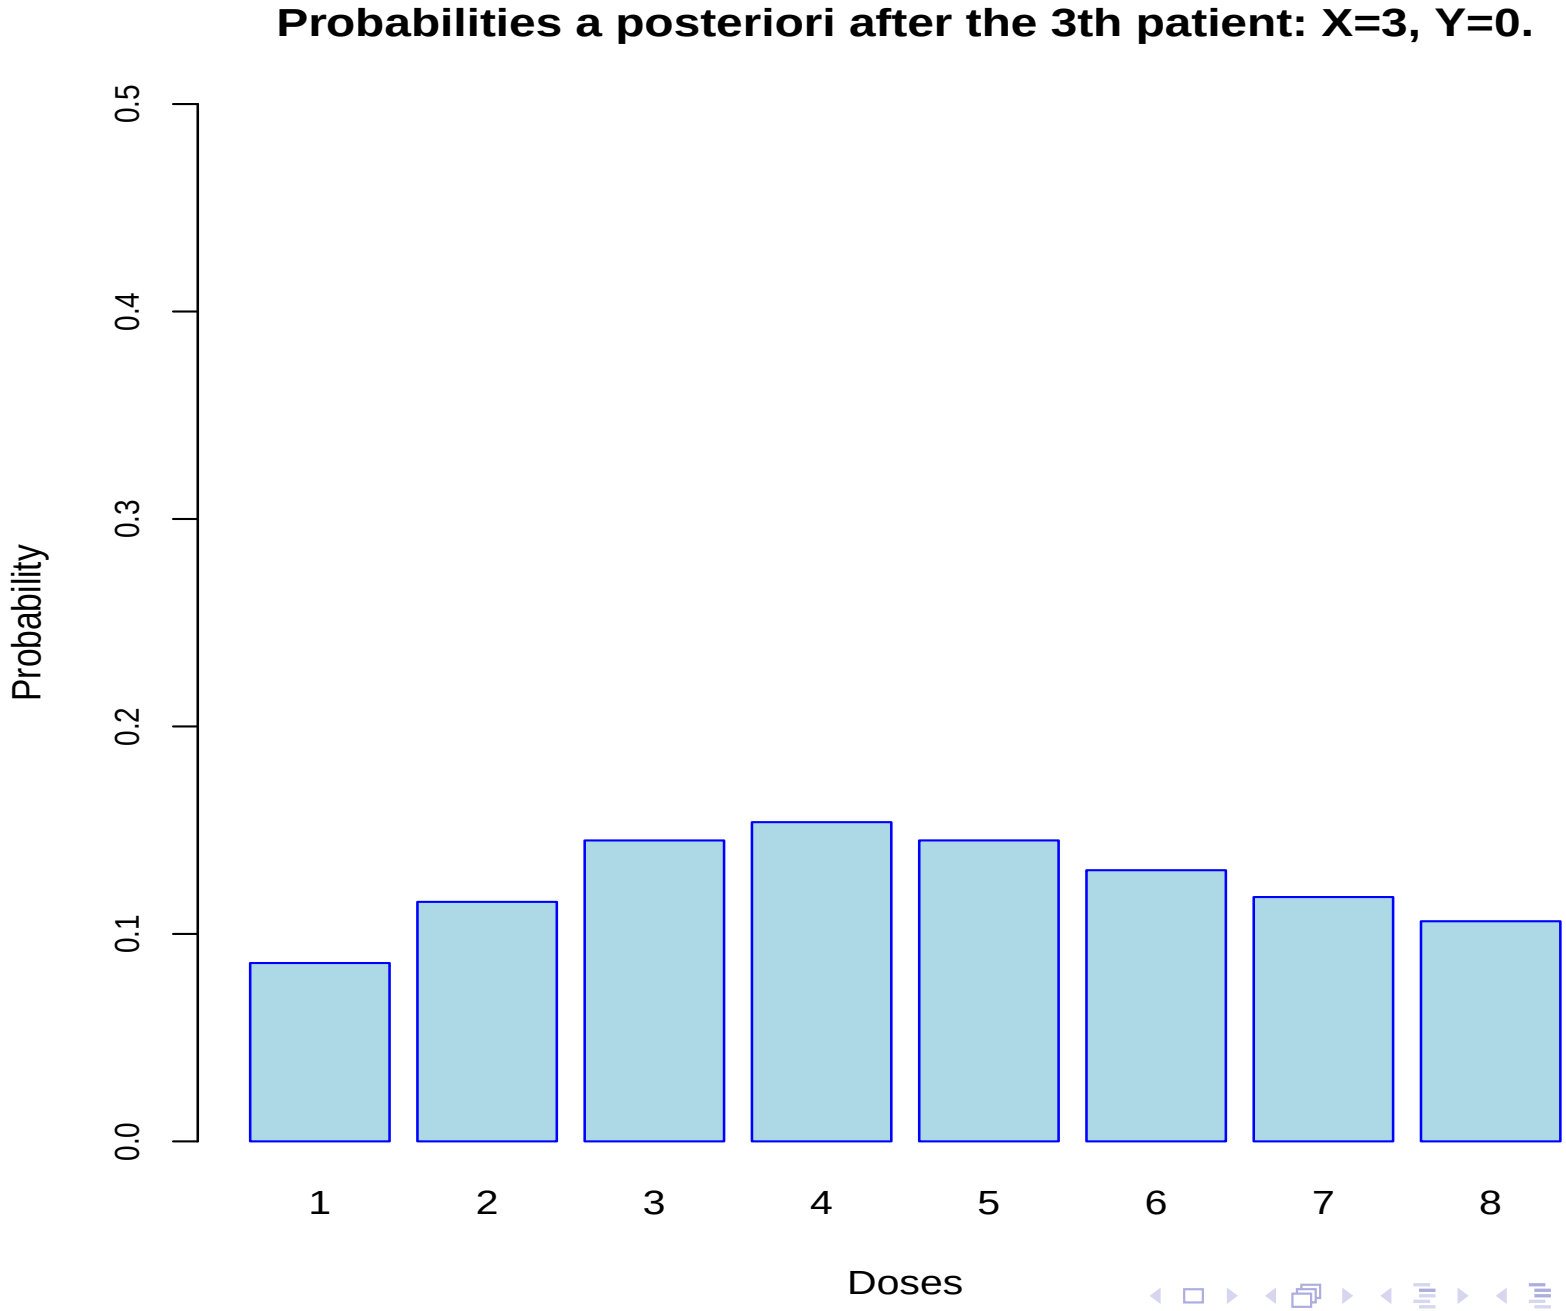
Probabilities a posteriori after the 3th patient: X=3, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

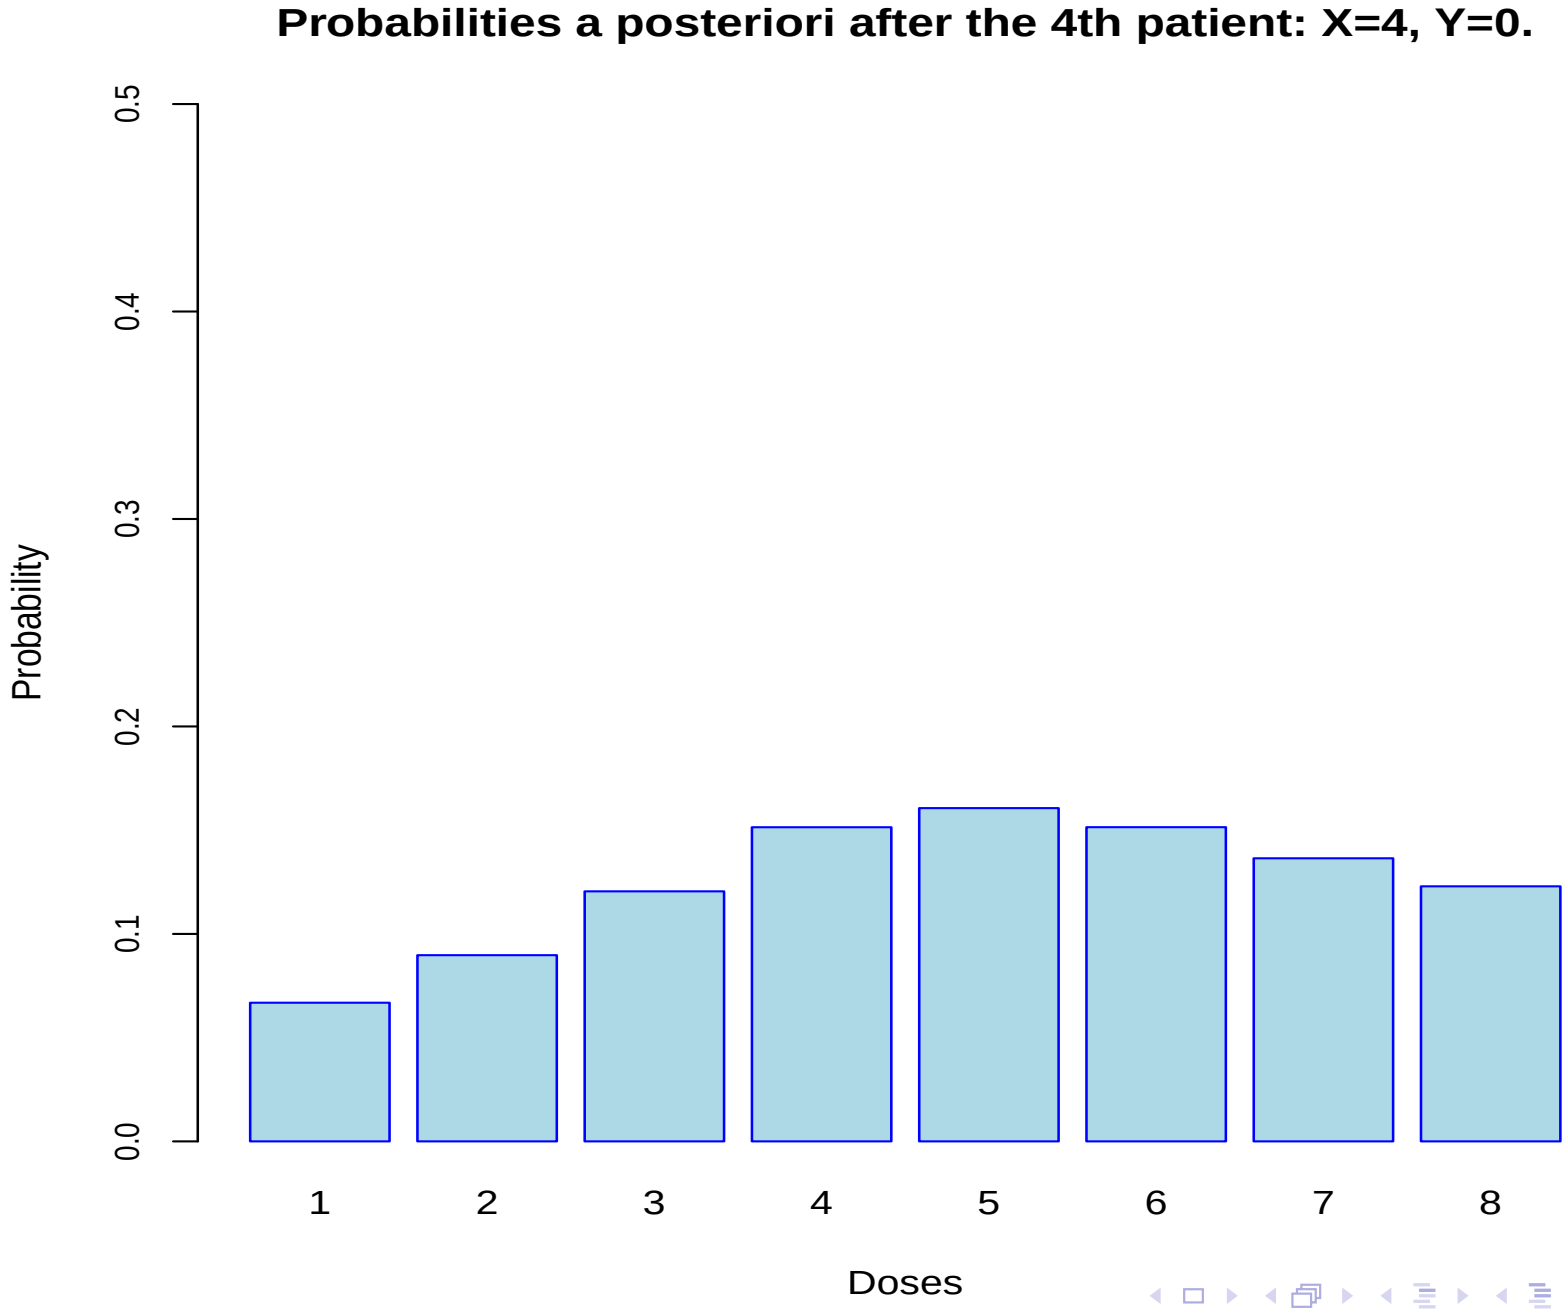
Probabilities a posteriori after the 4th patient: X=4, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

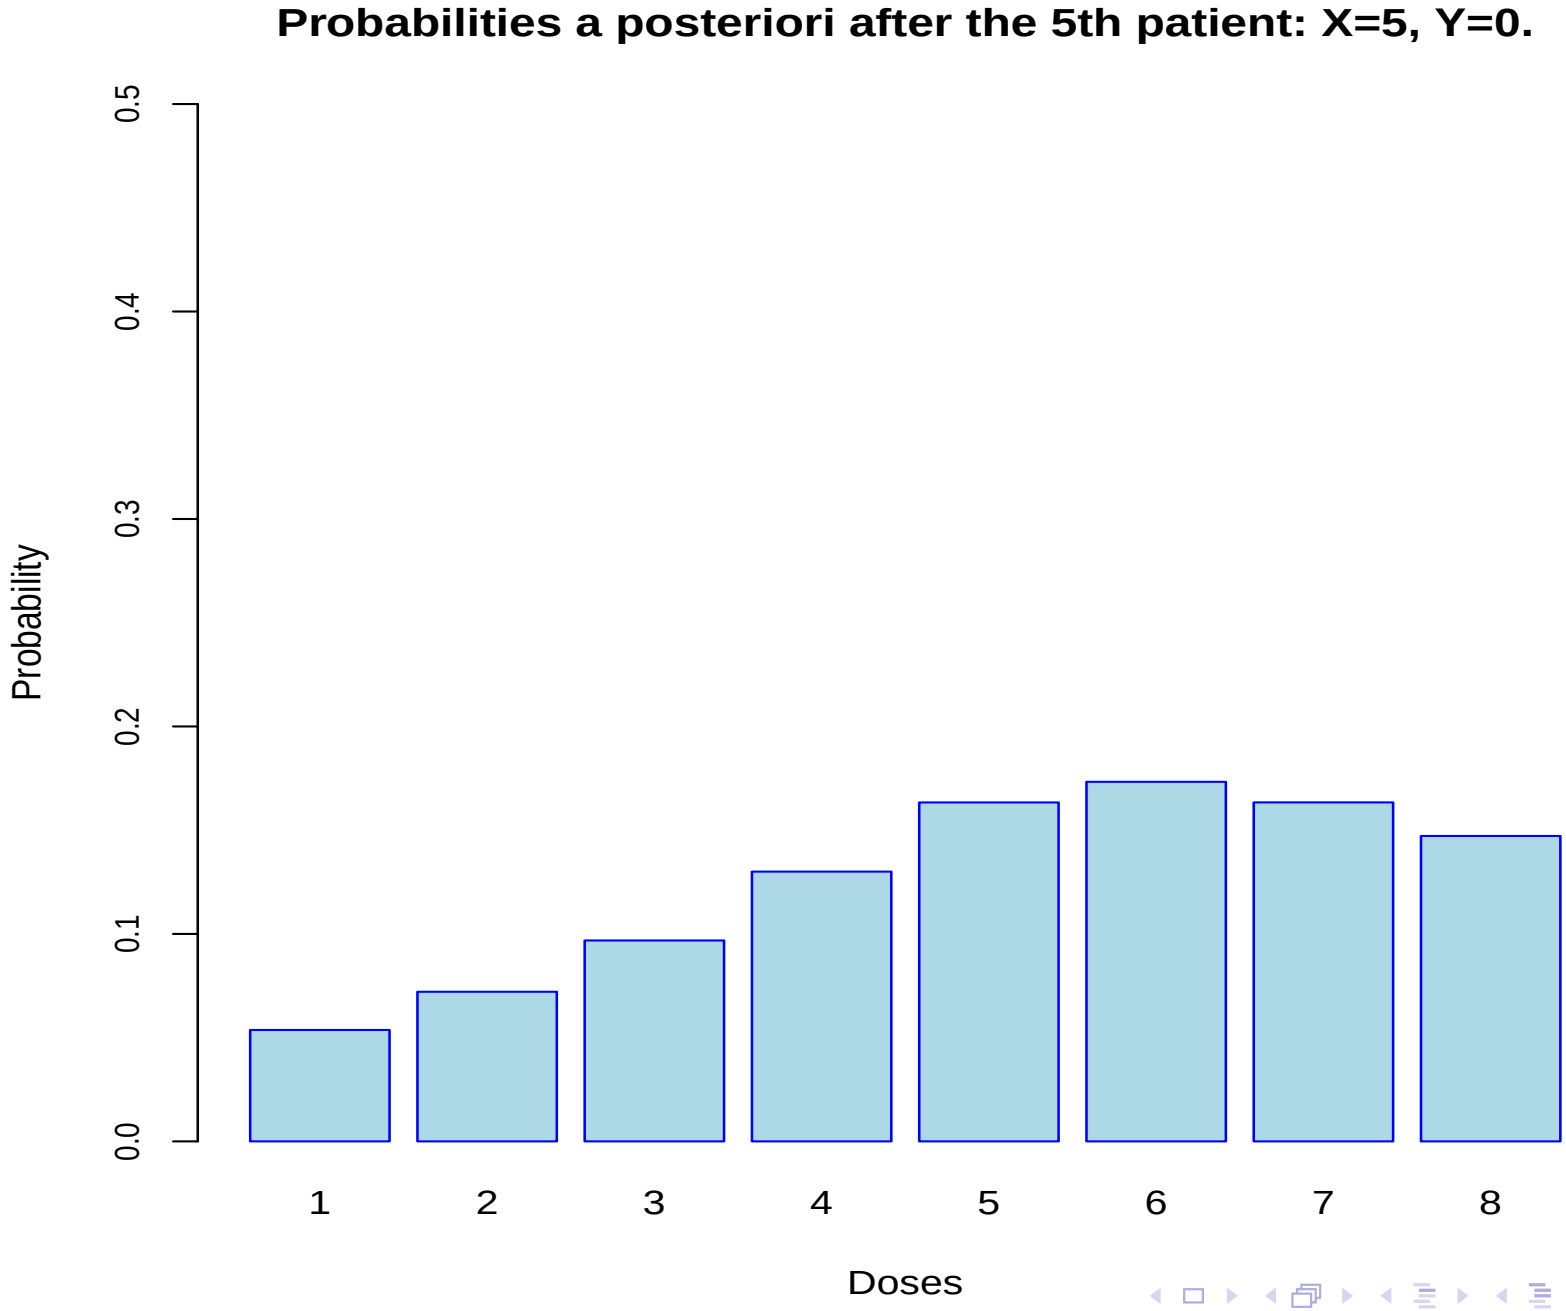
Probabilities a posteriori after the 5th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

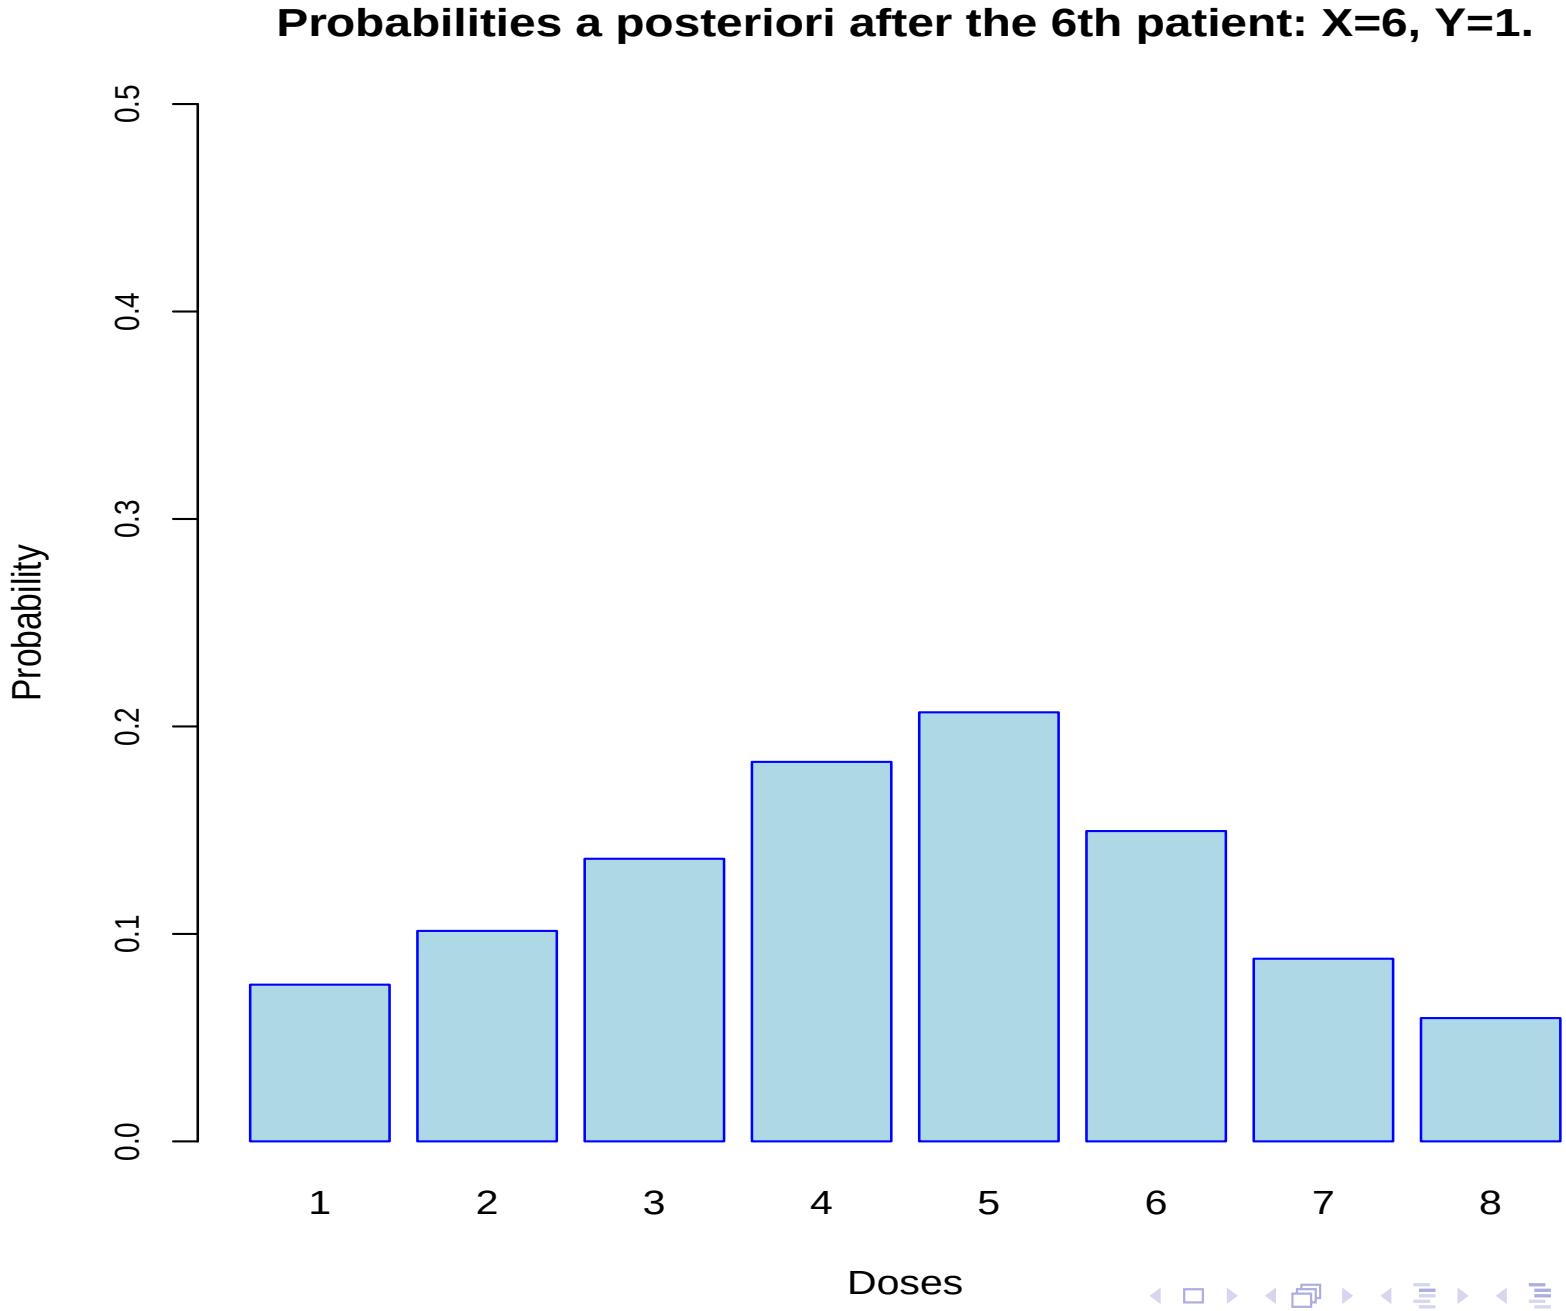
Probabilities a posteriori after the 6th patient: X=6, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

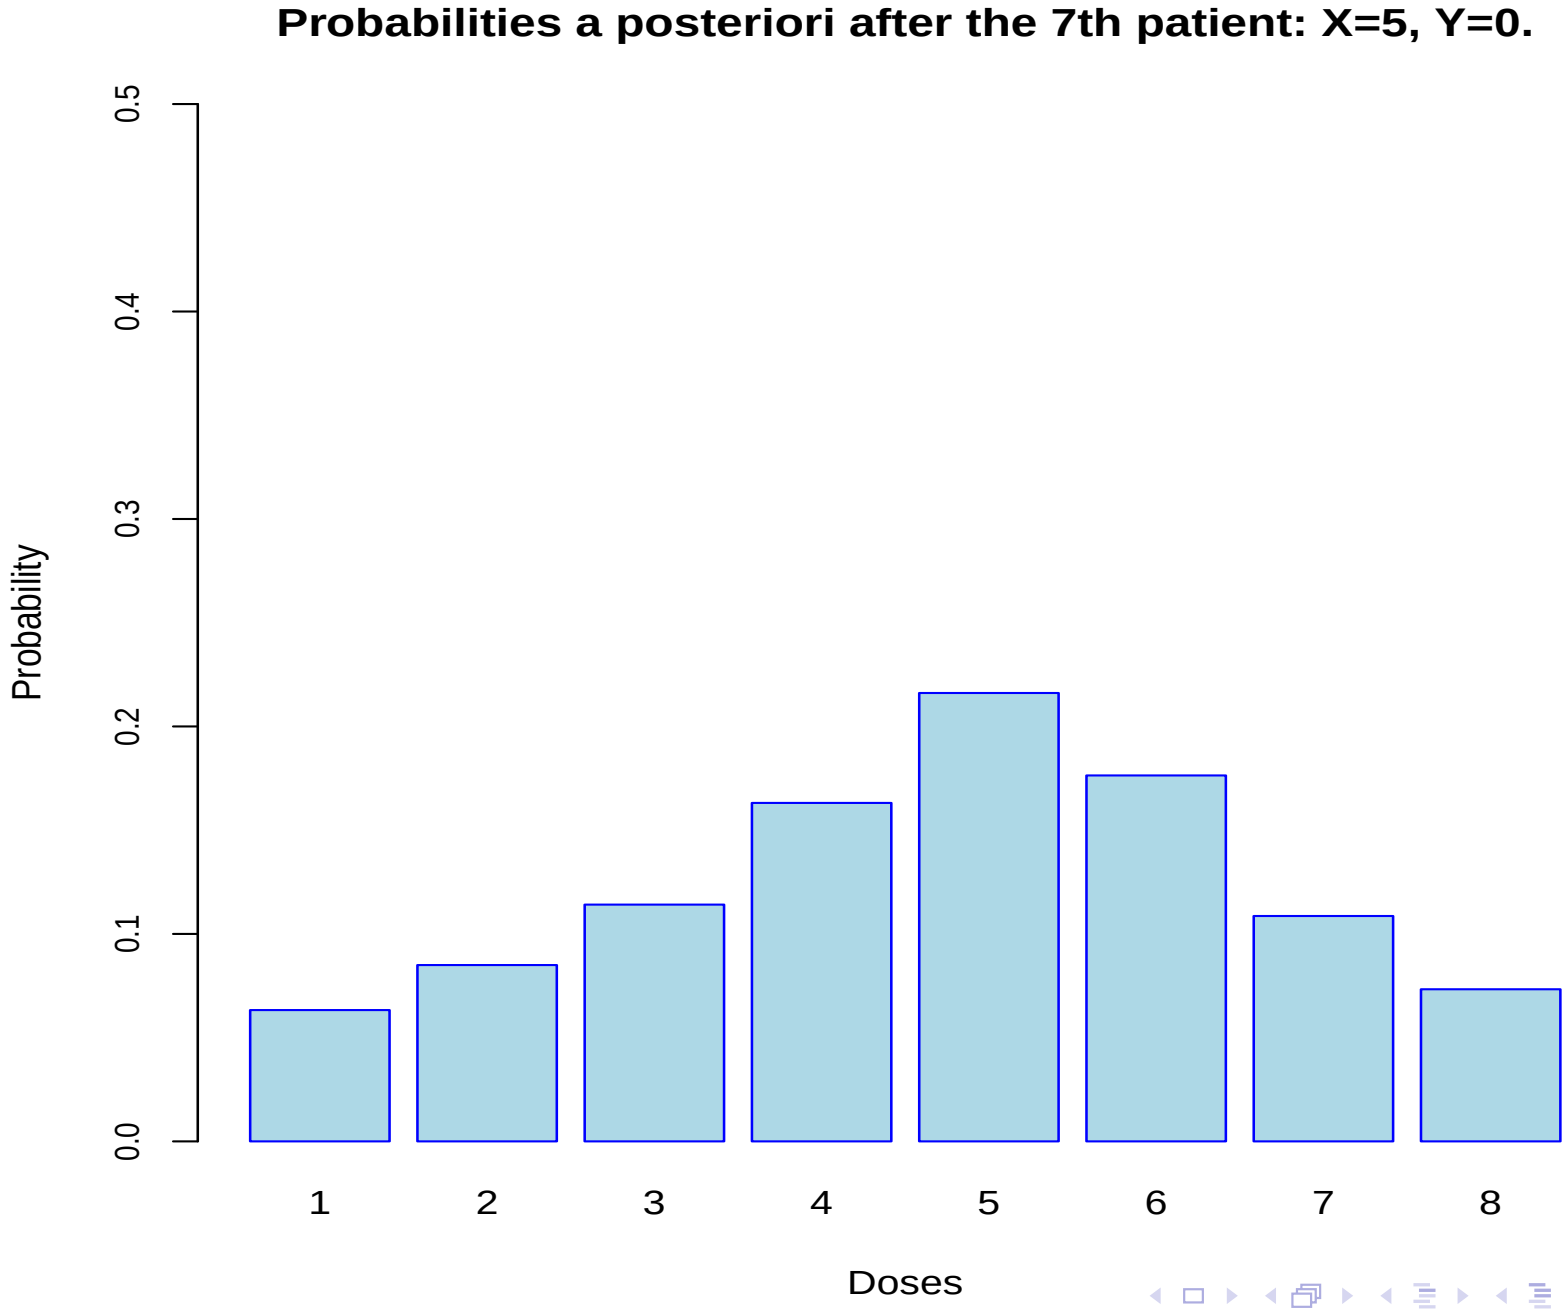
Probabilities a posteriori after the 7th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

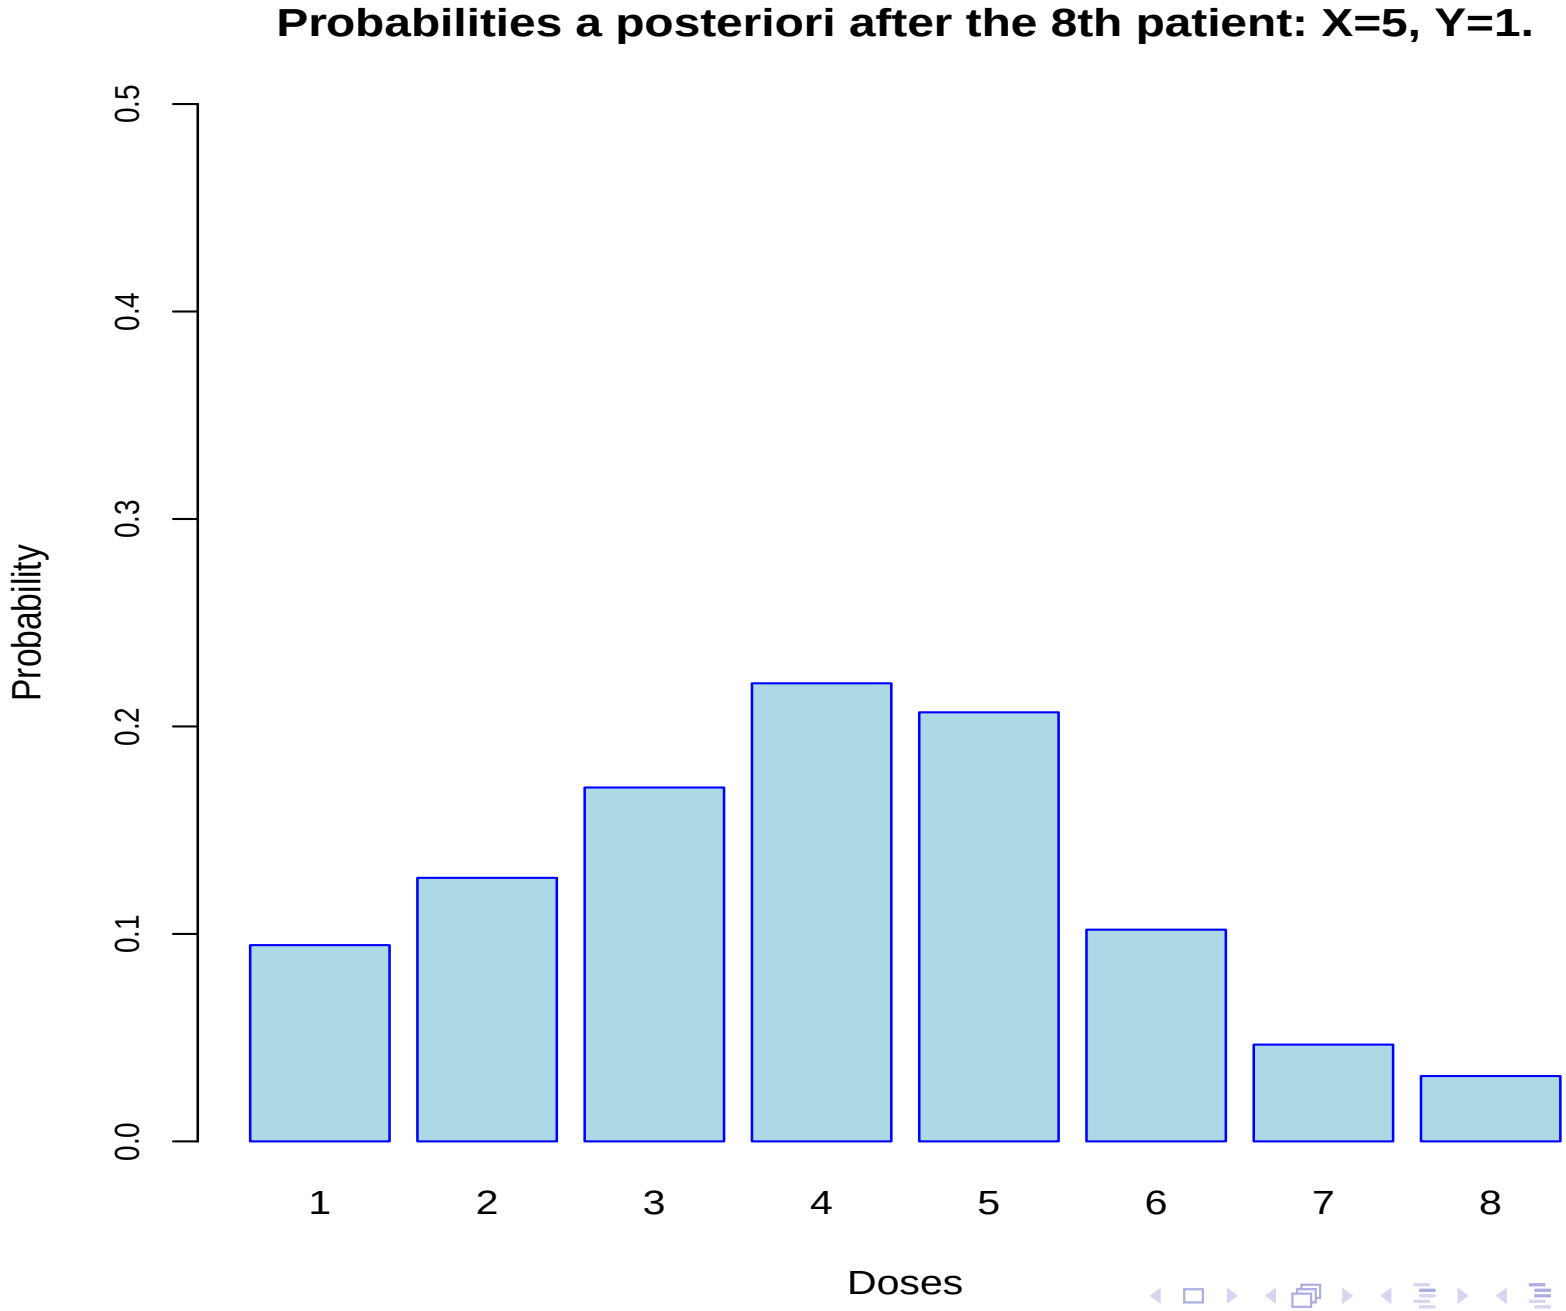
Probabilities a posteriori after the 8th patient: X=5, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

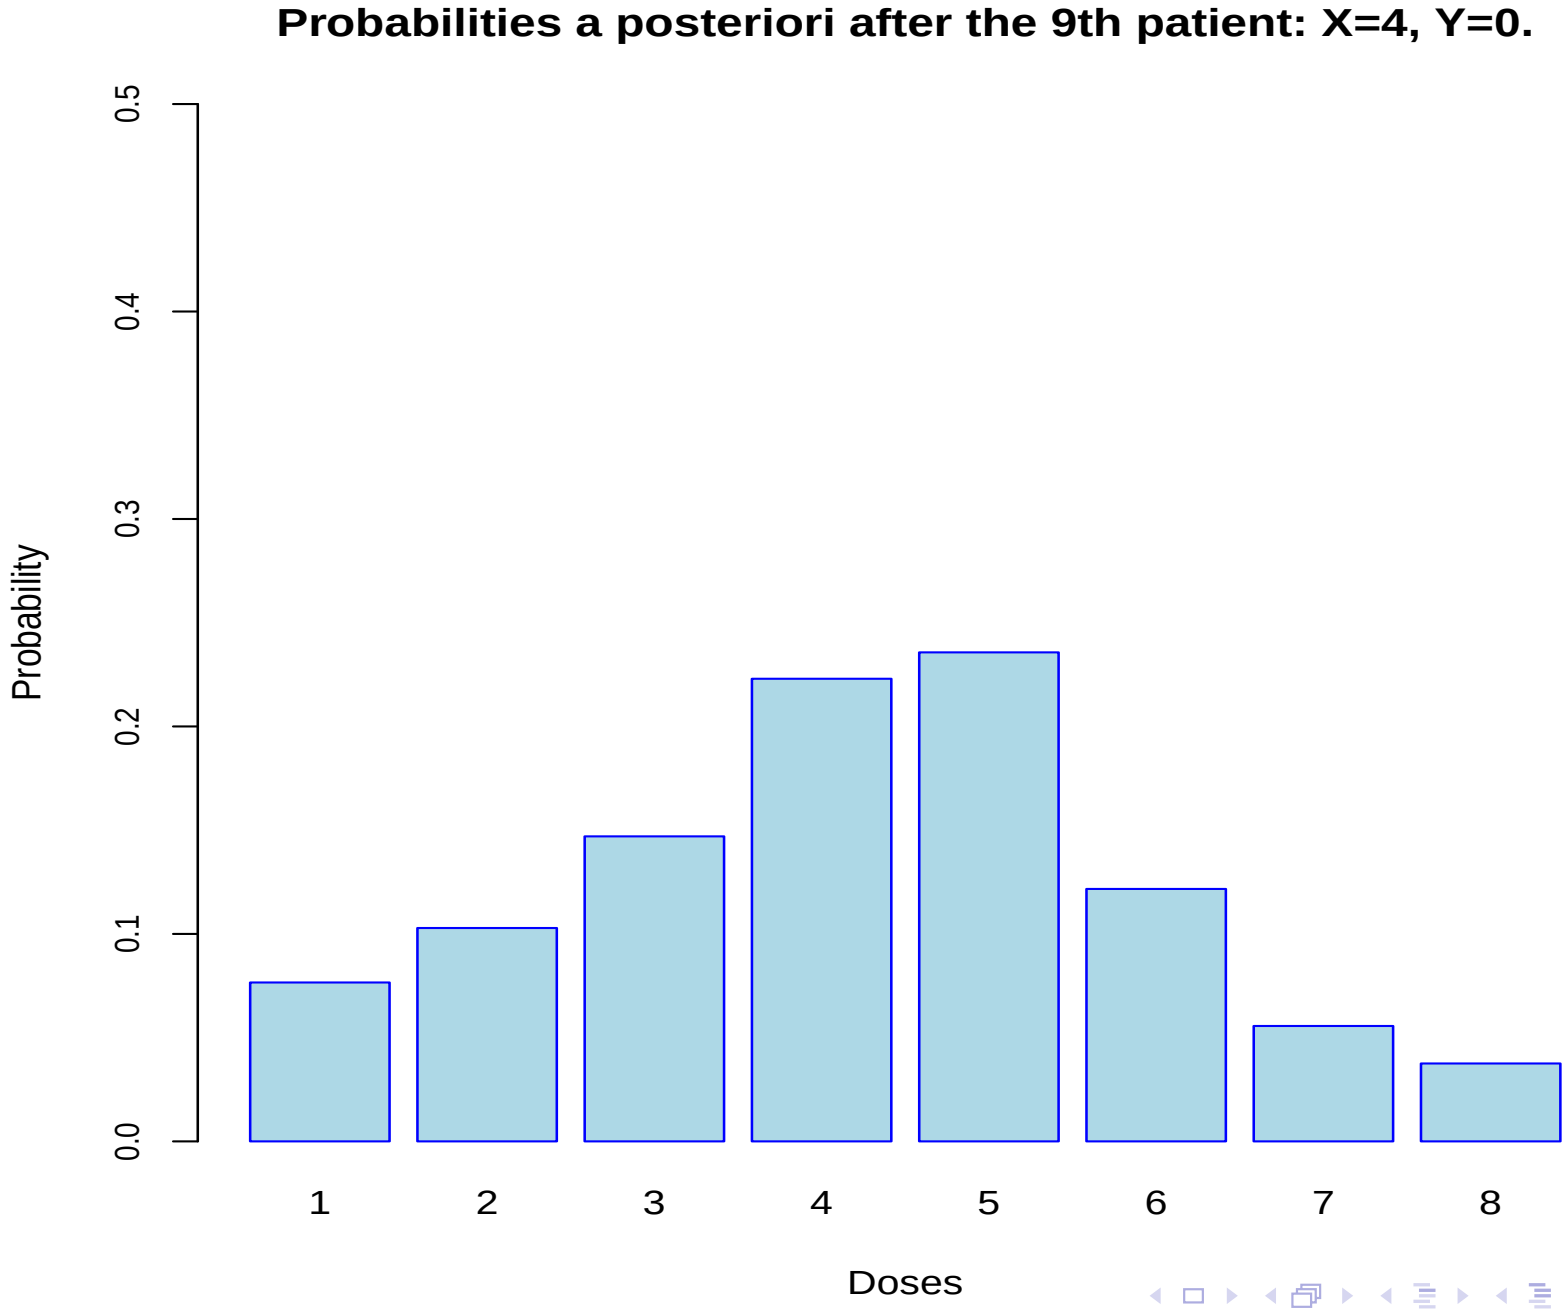
Probabilities a posteriori after the 9th patient: X=4, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

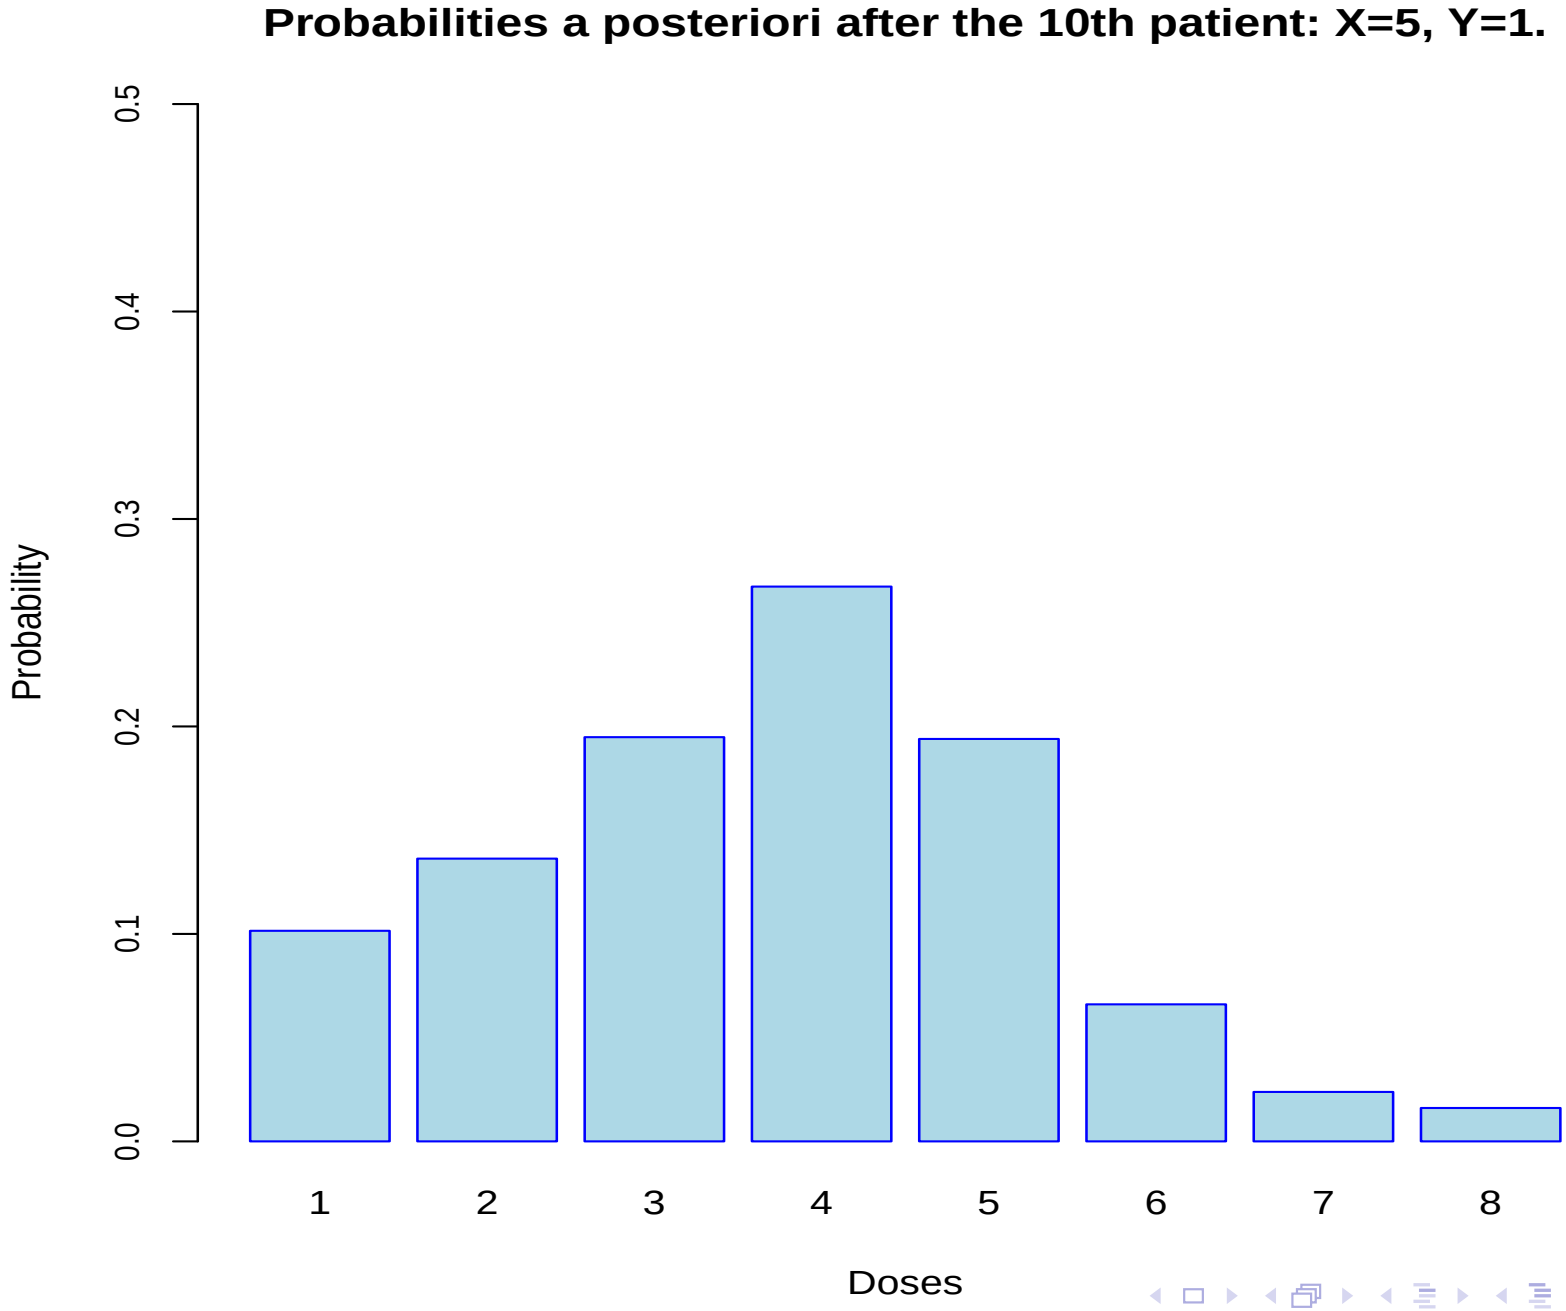
Probabilities a posteriori after the 10th patient: X=5, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

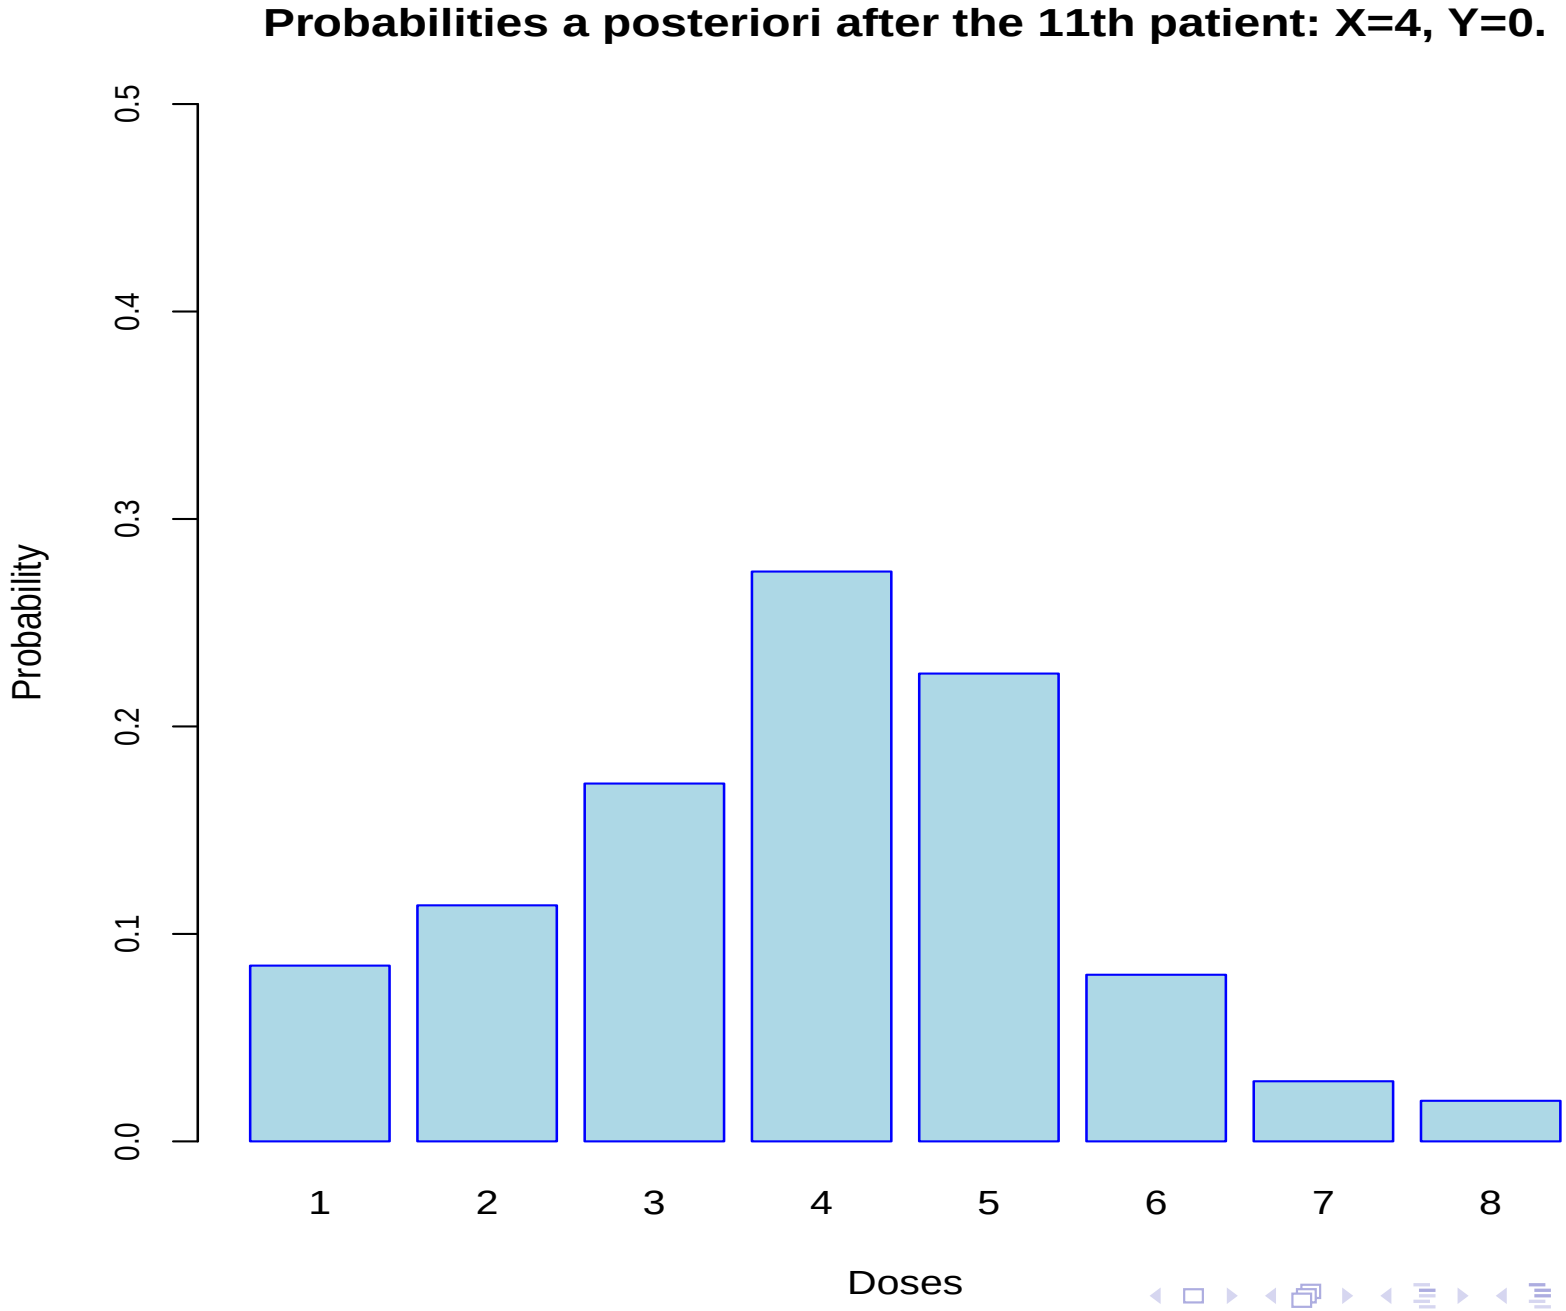
Probabilities a posteriori after the 11th patient: X=4, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

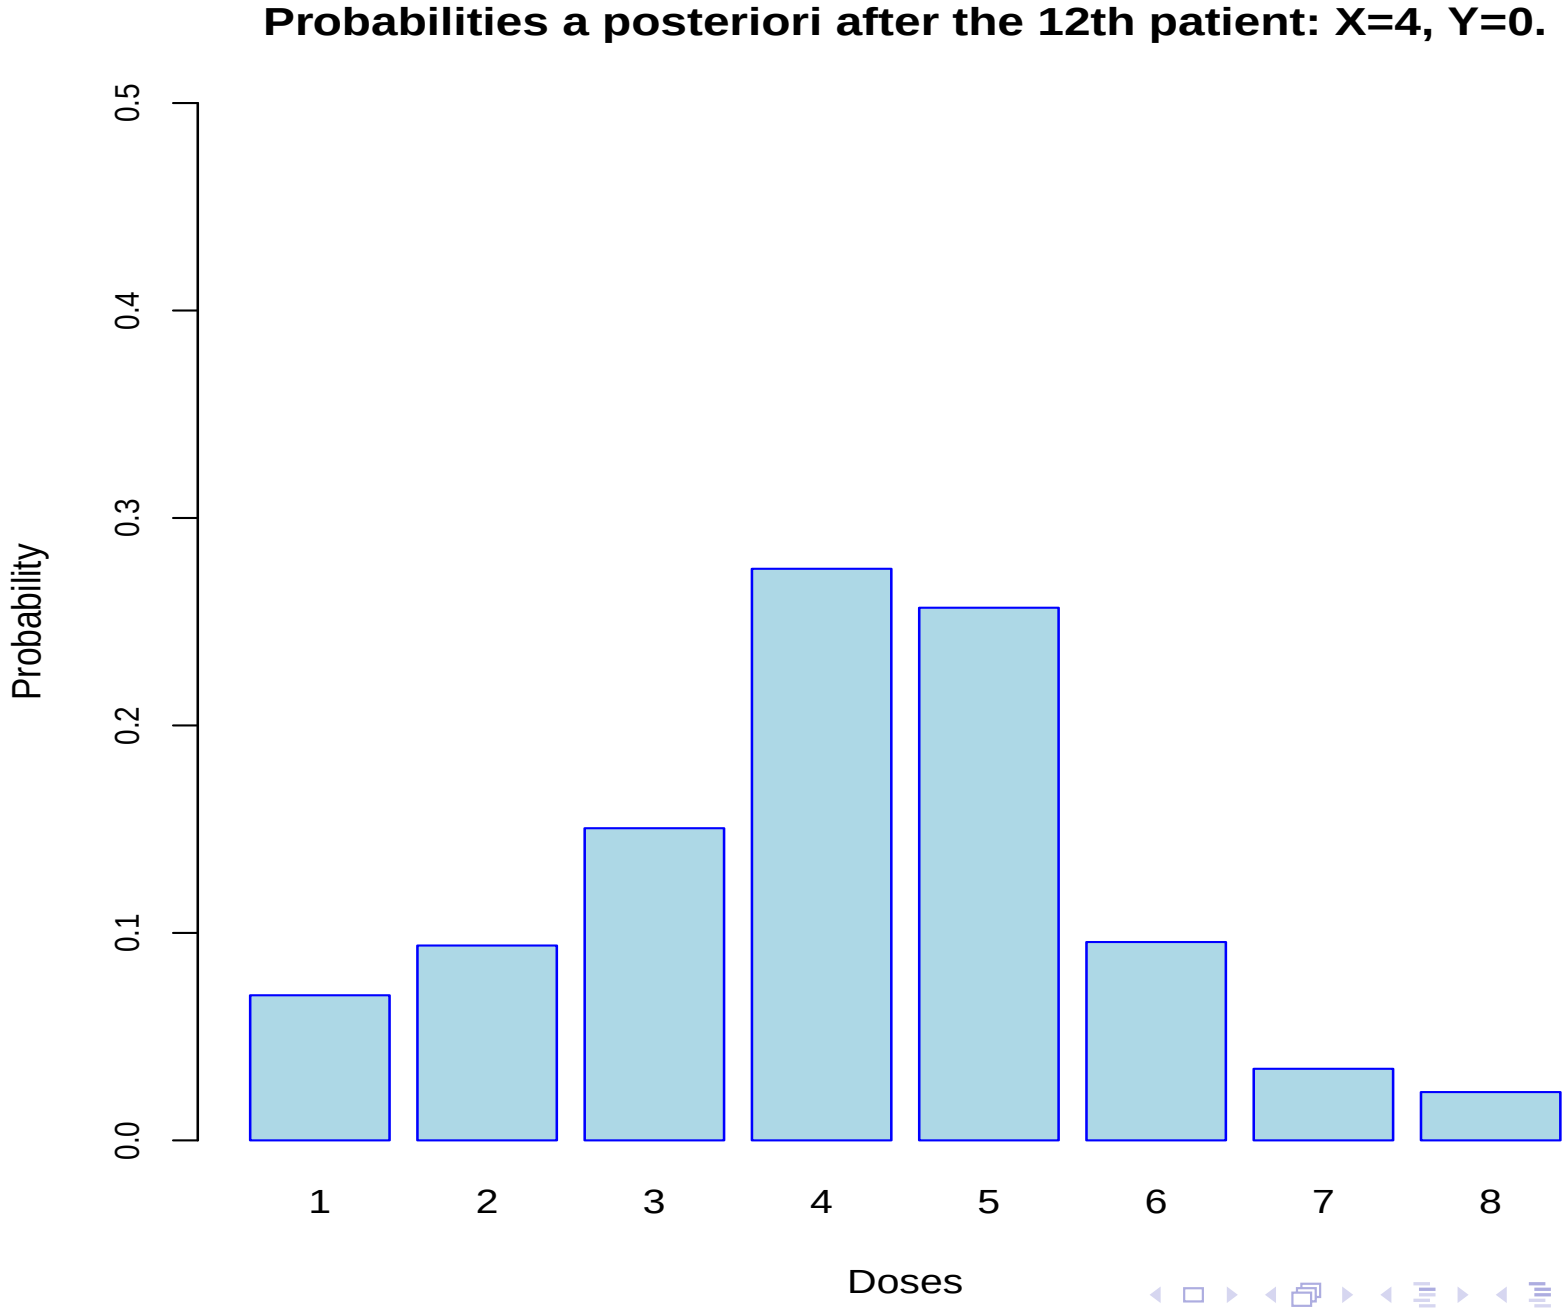
Probabilities a posteriori after the 12th patient: X=4, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

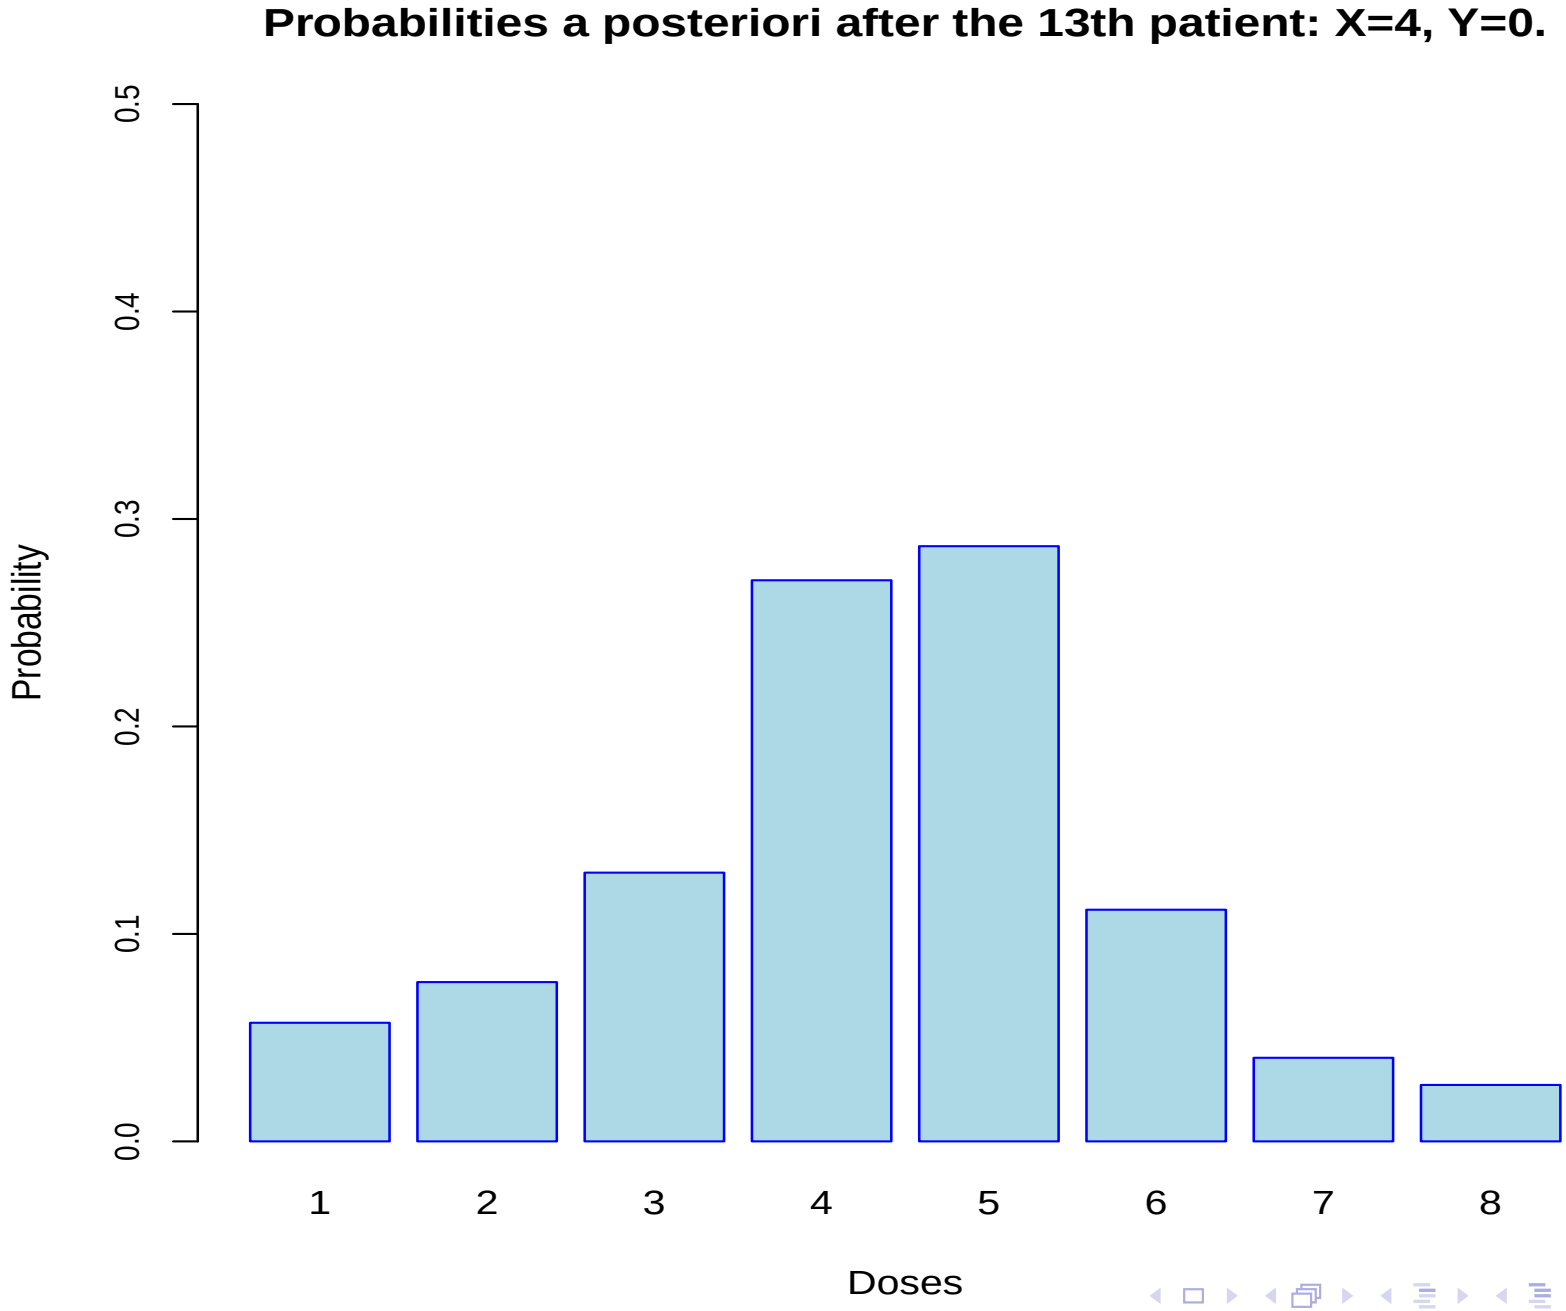
Probabilities a posteriori after the 13th patient: X=4, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

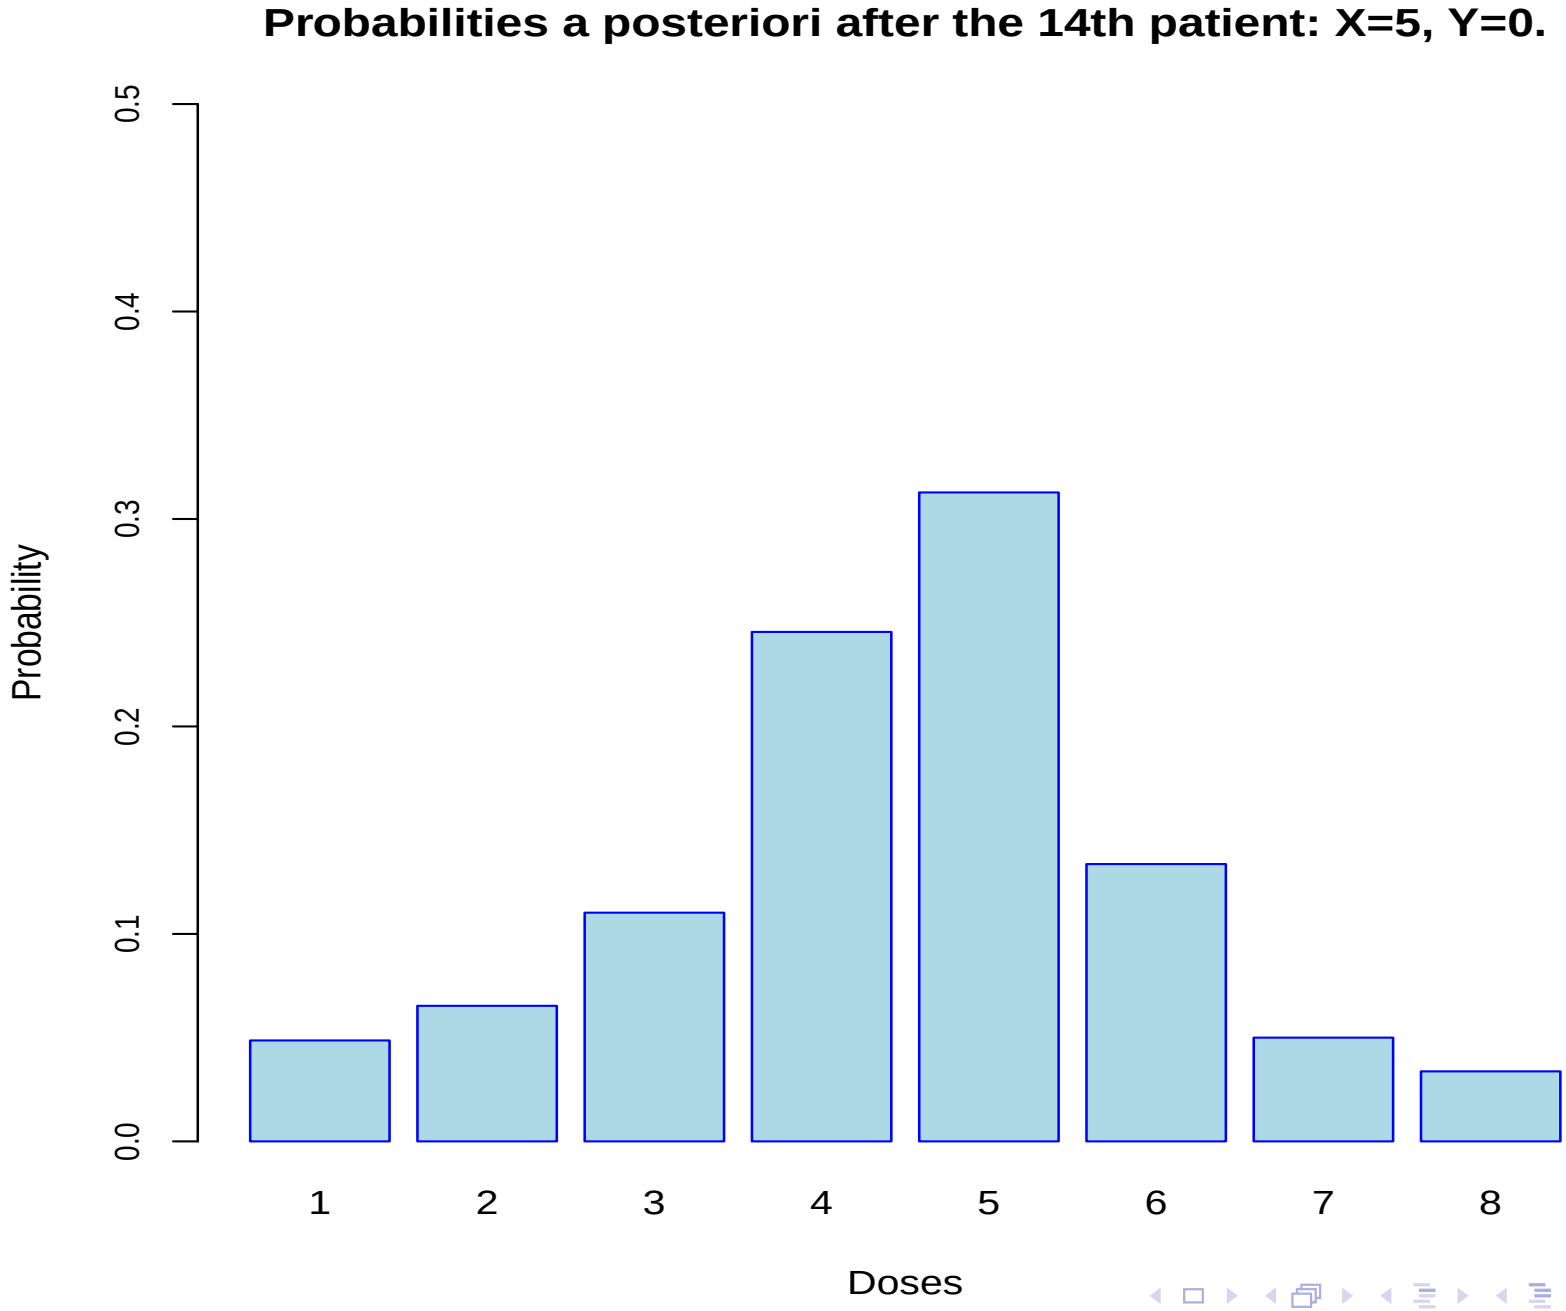
Probabilities a posteriori after the 14th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

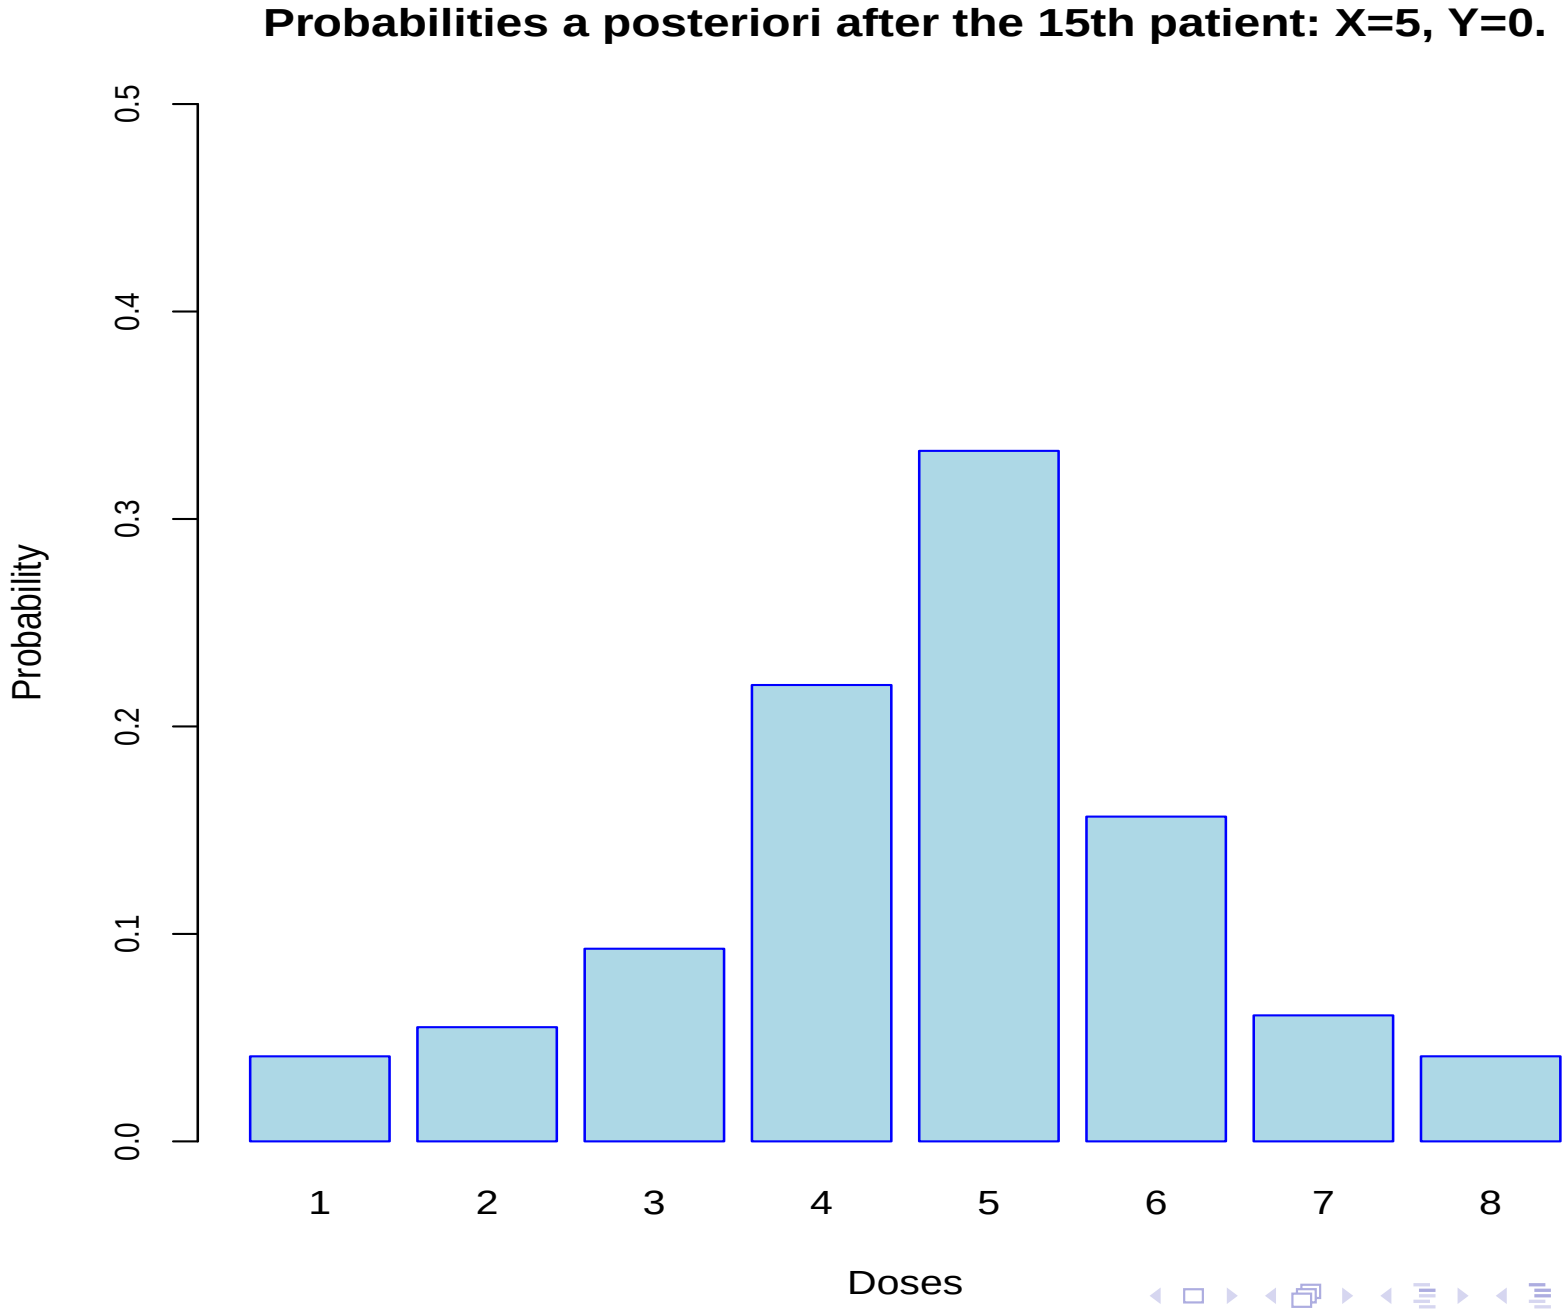
Probabilities a posteriori after the 15th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

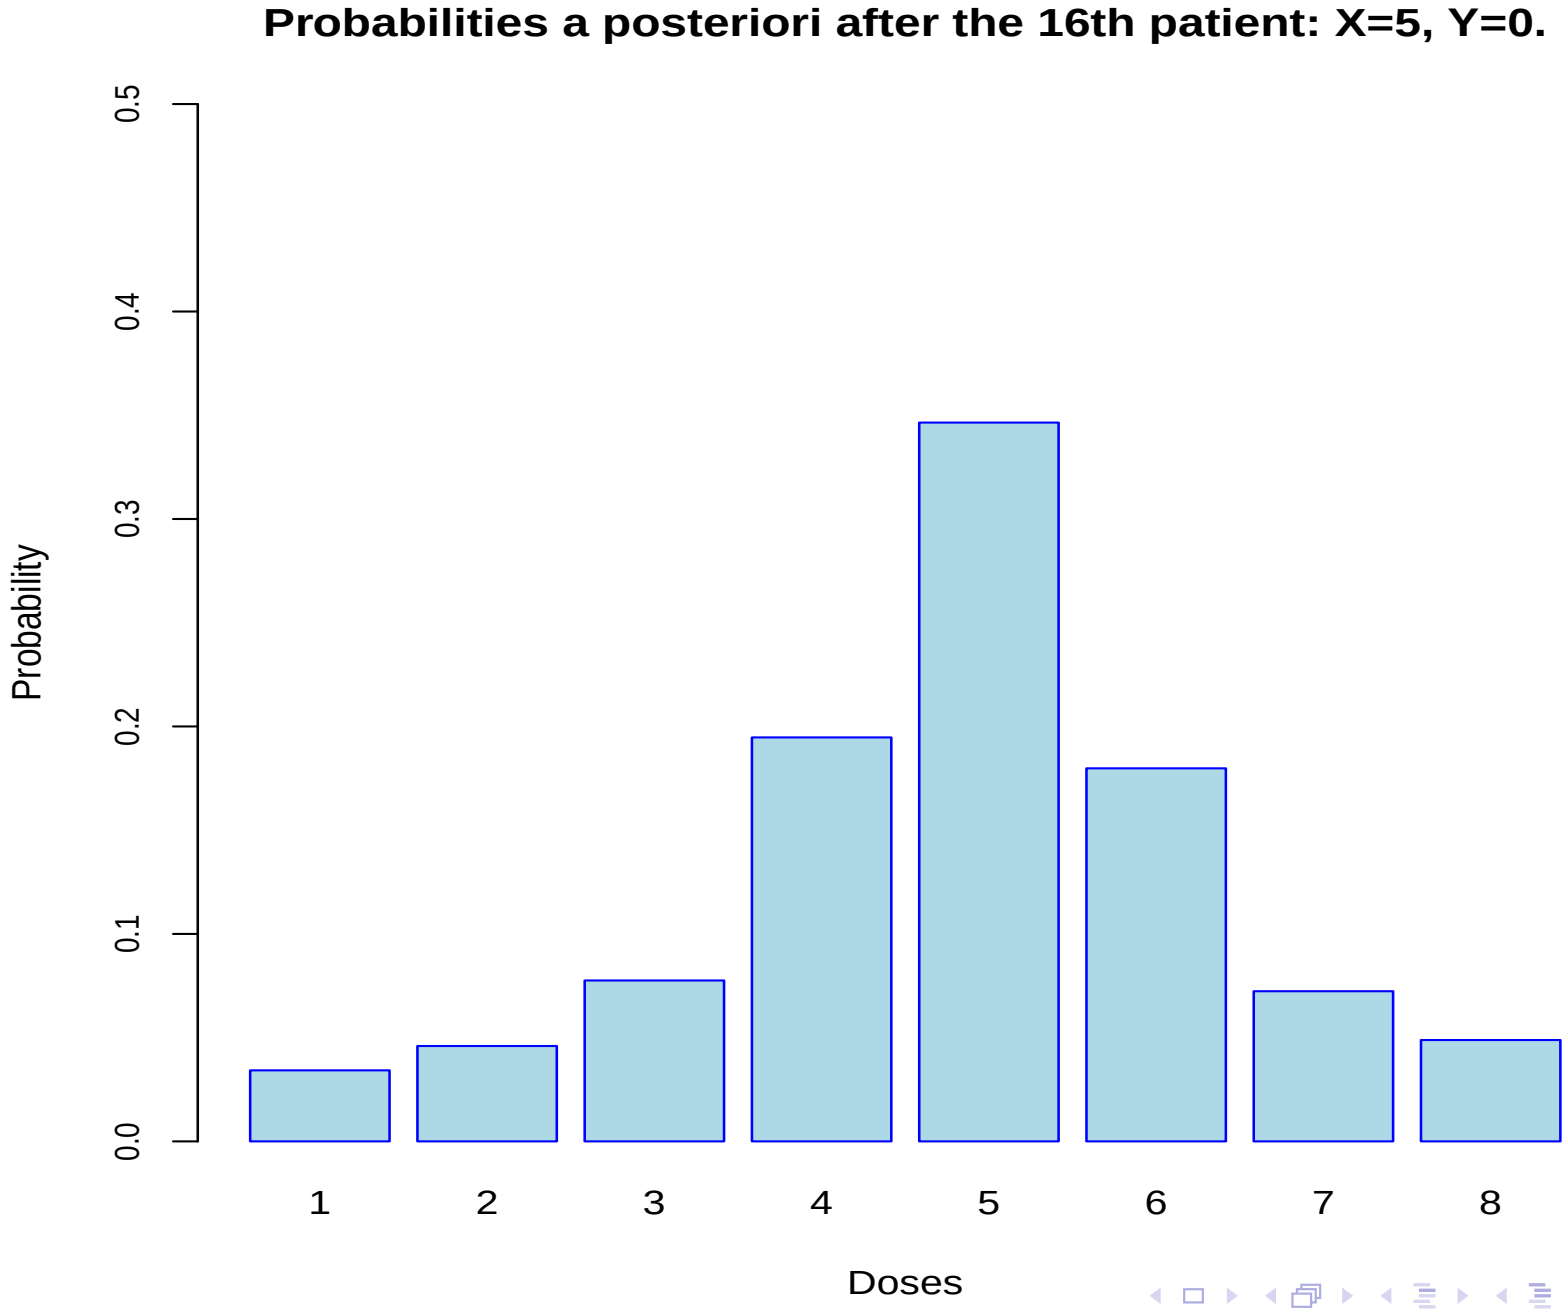
Probabilities a posteriori after the 16th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

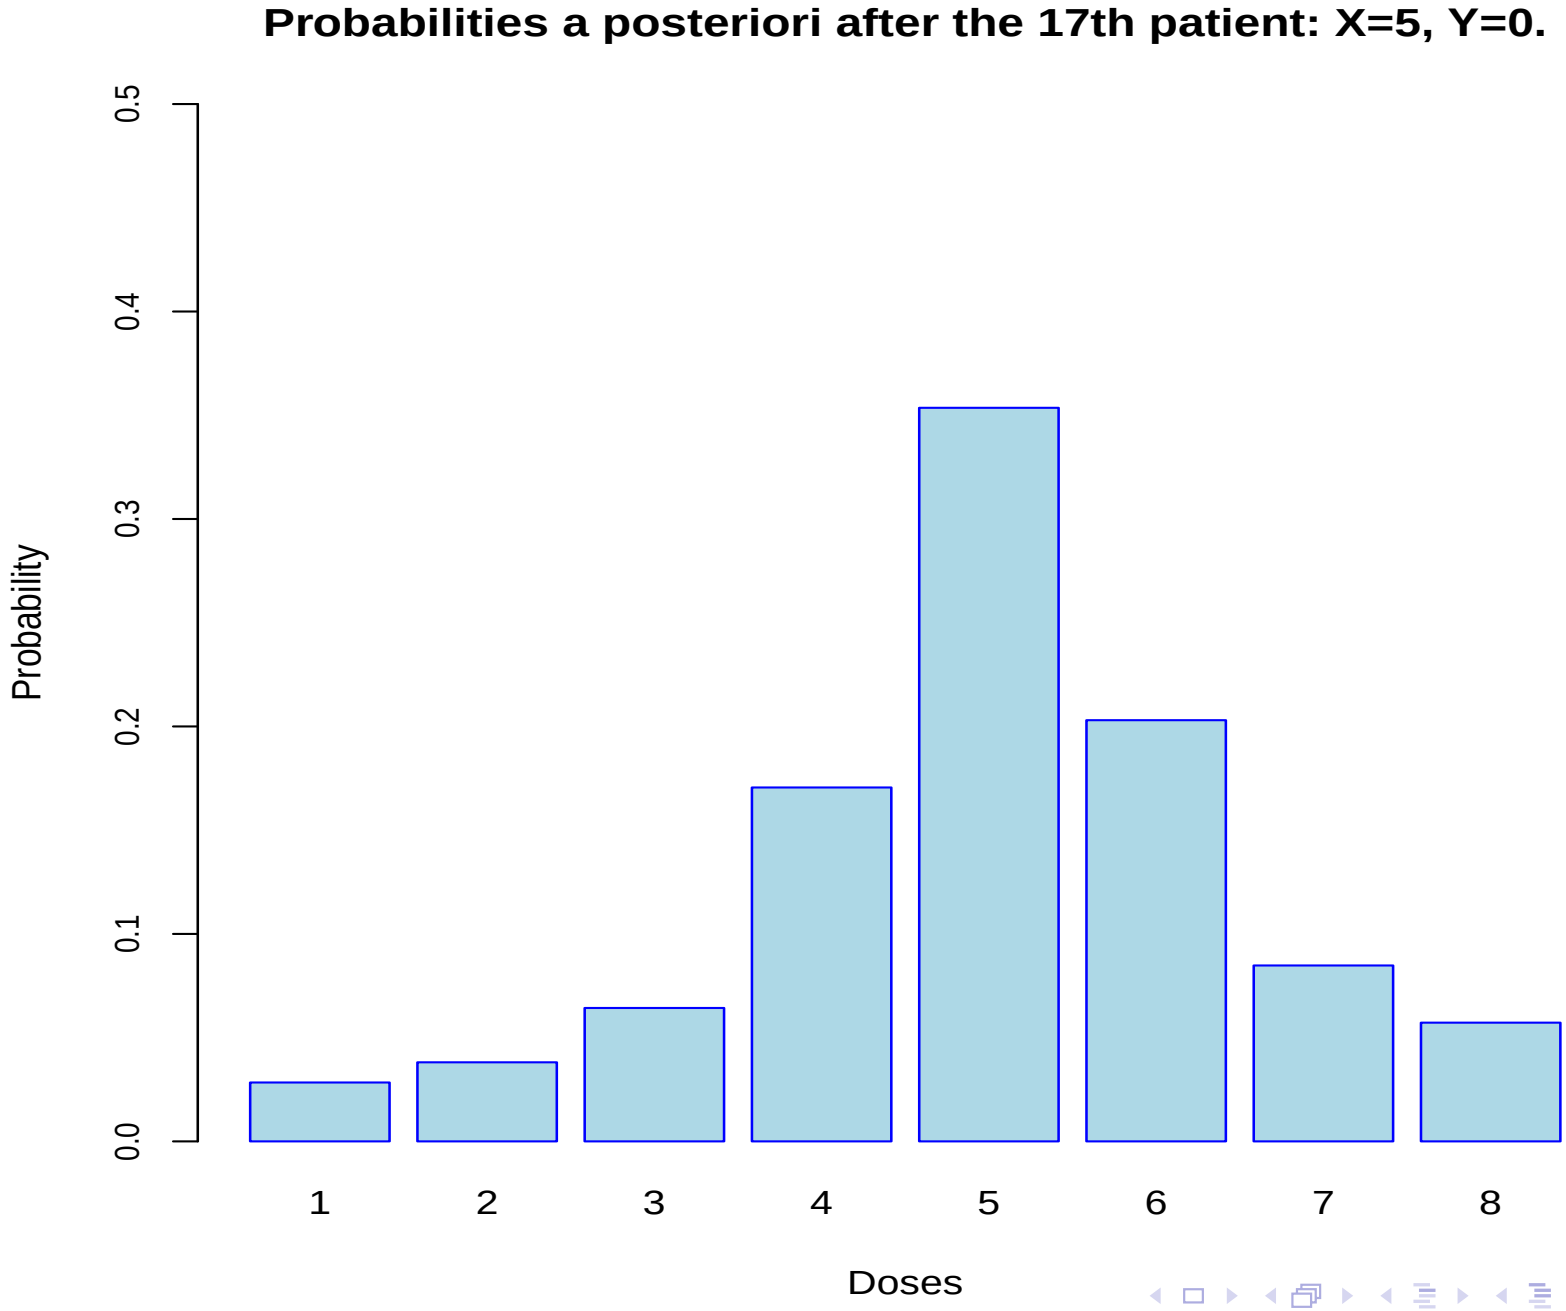
Probabilities a posteriori after the 17th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

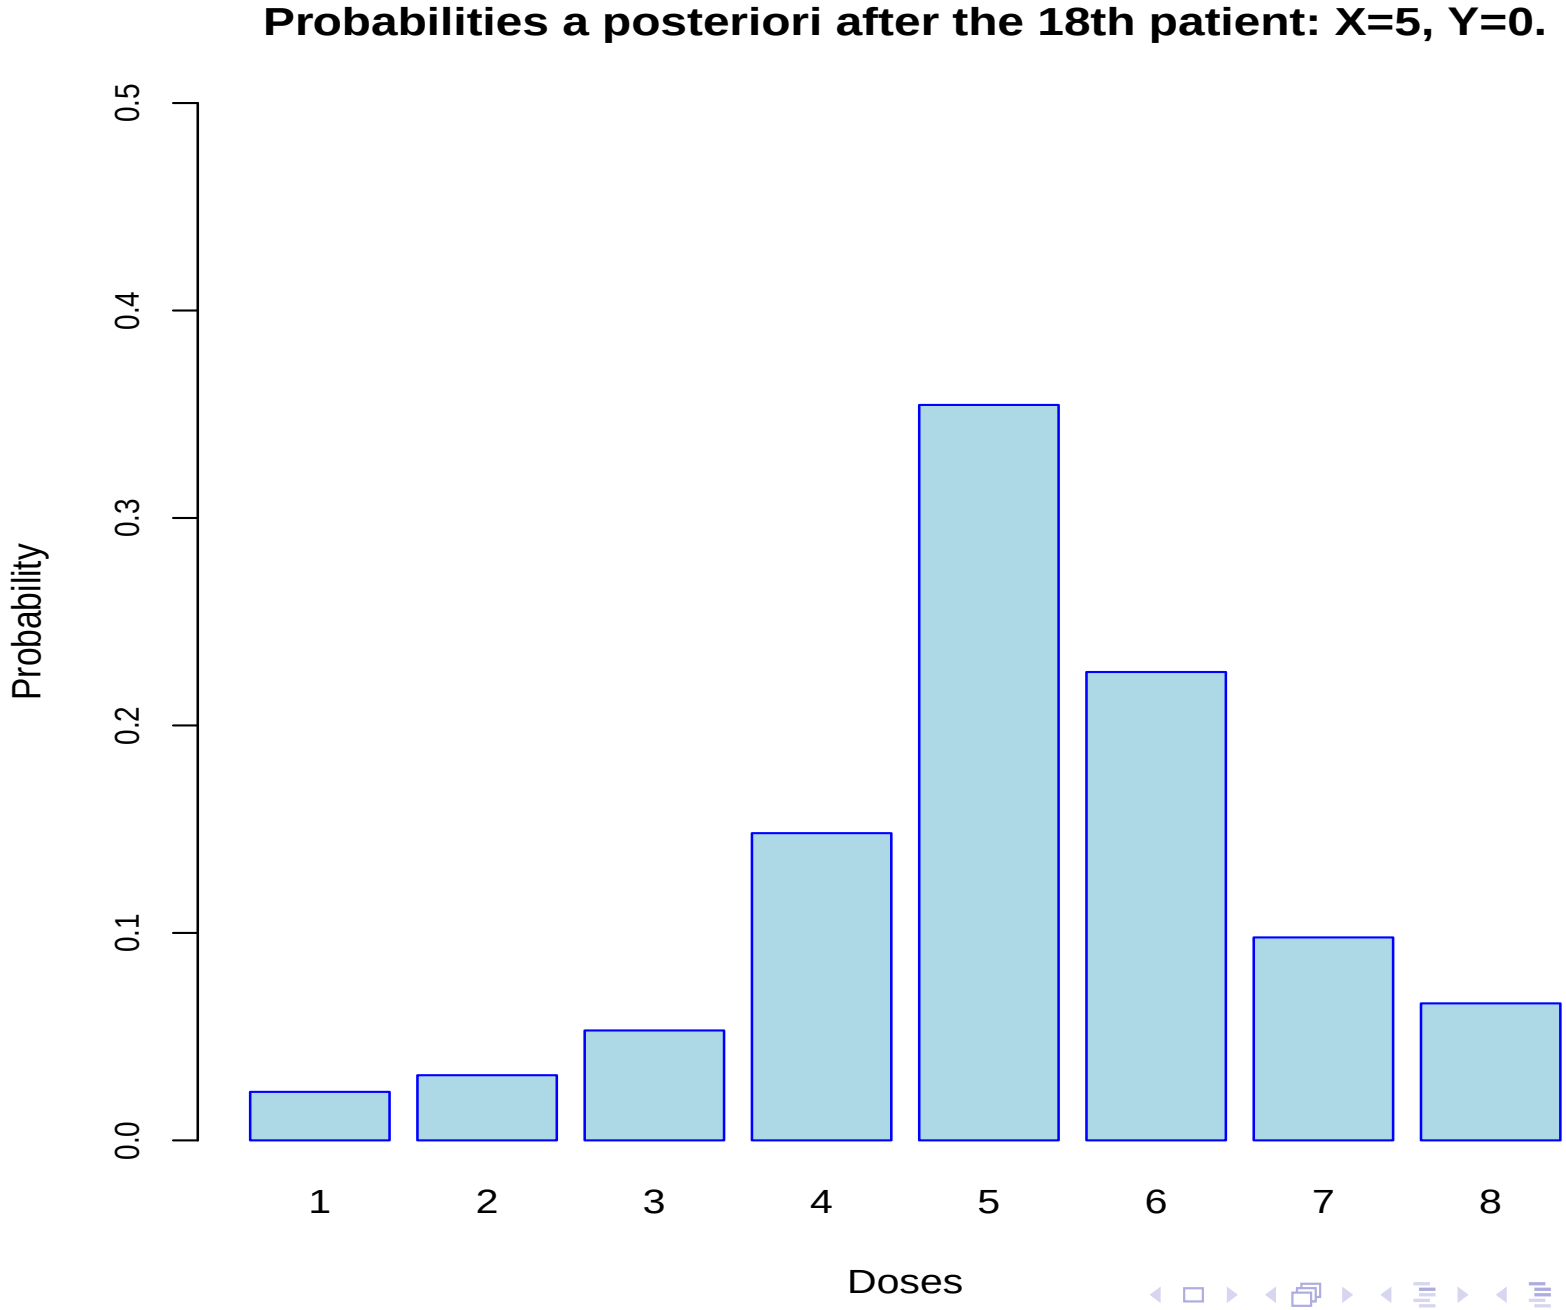
Probabilities a posteriori after the 18th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

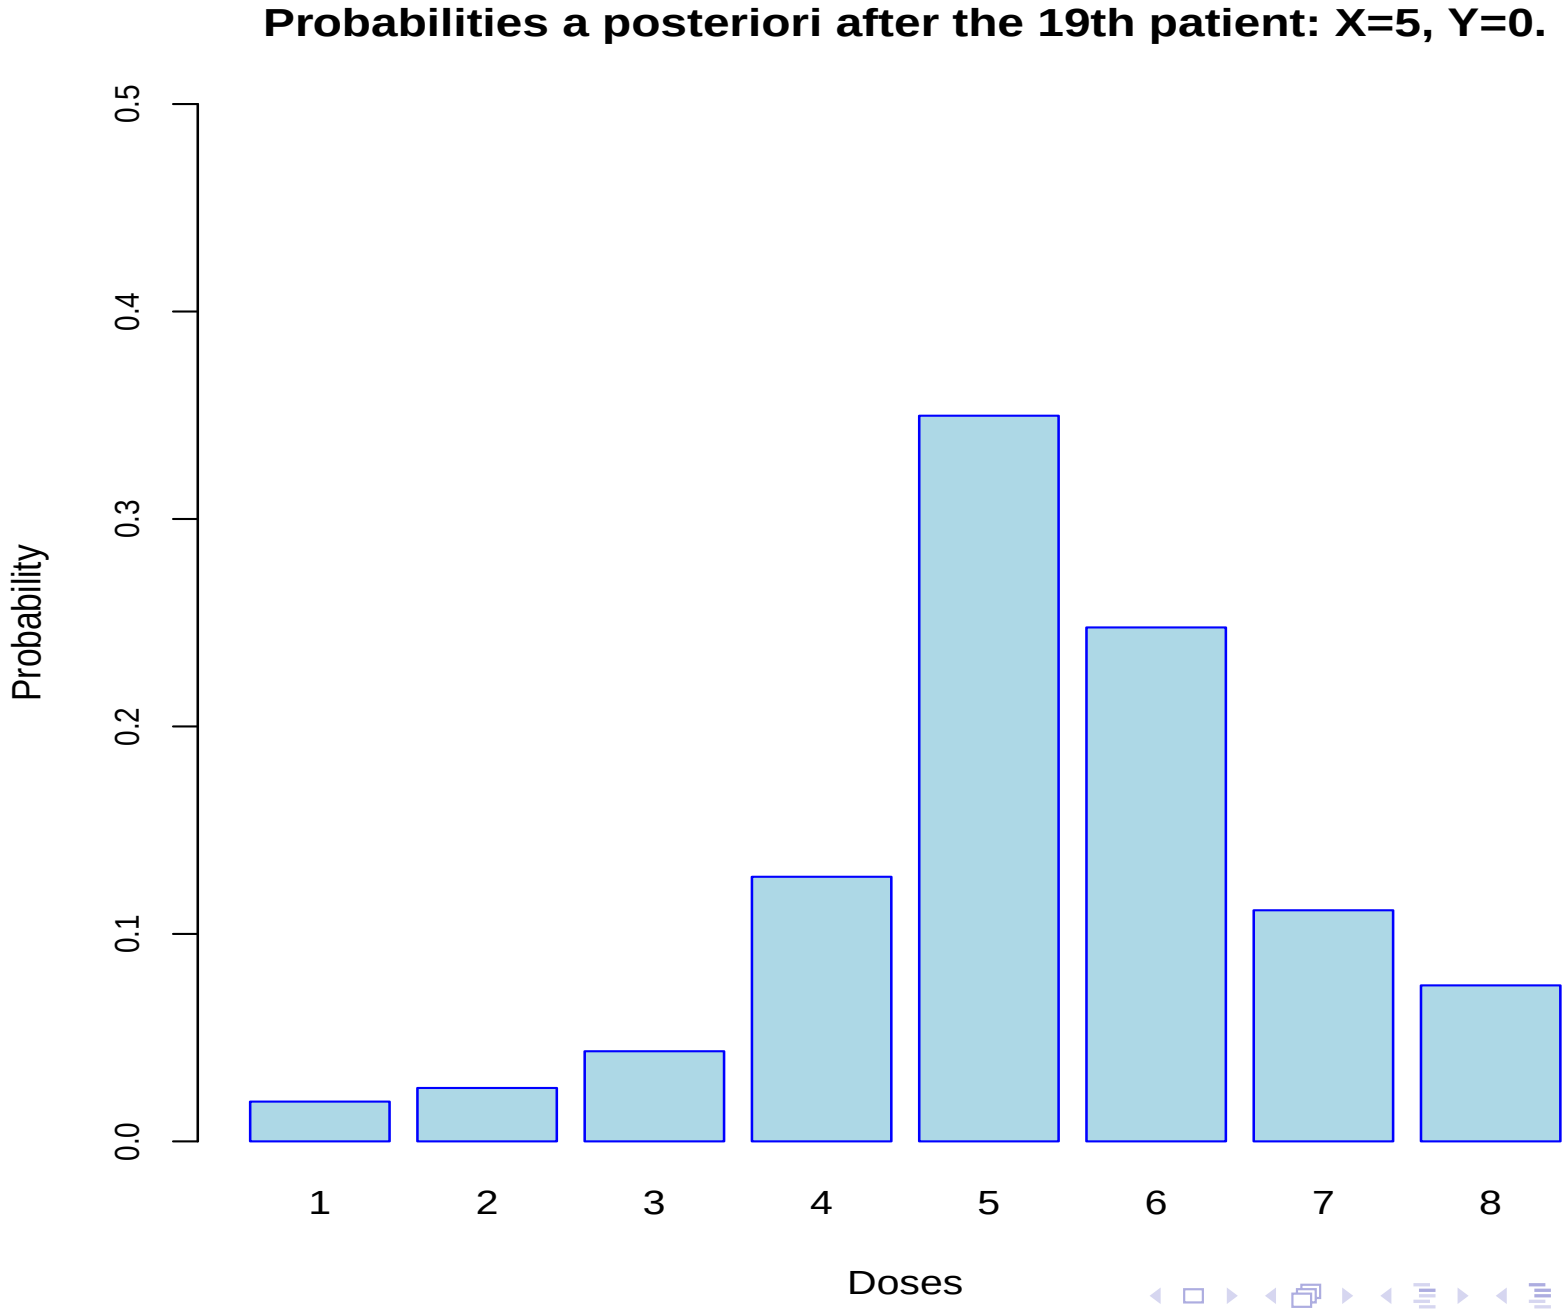
Probabilities a posteriori after the 19th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

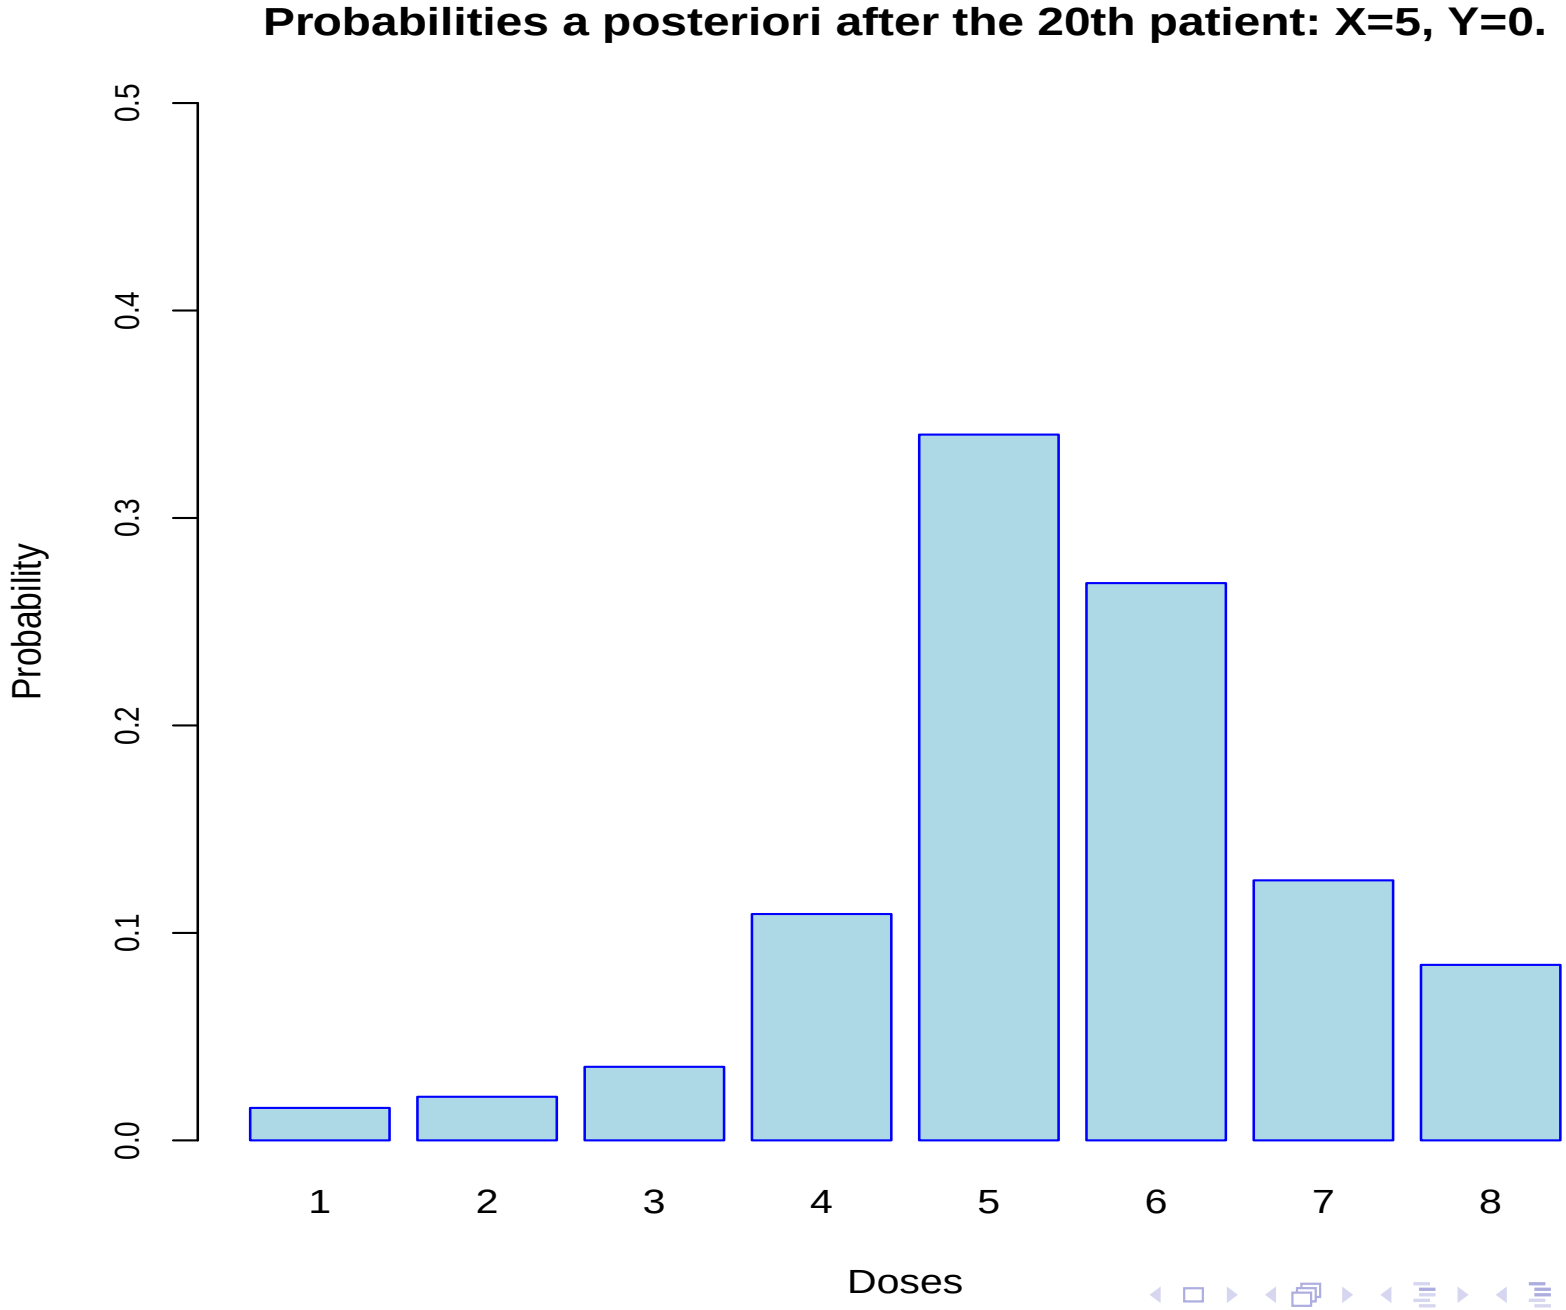
Probabilities a posteriori after the 20th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

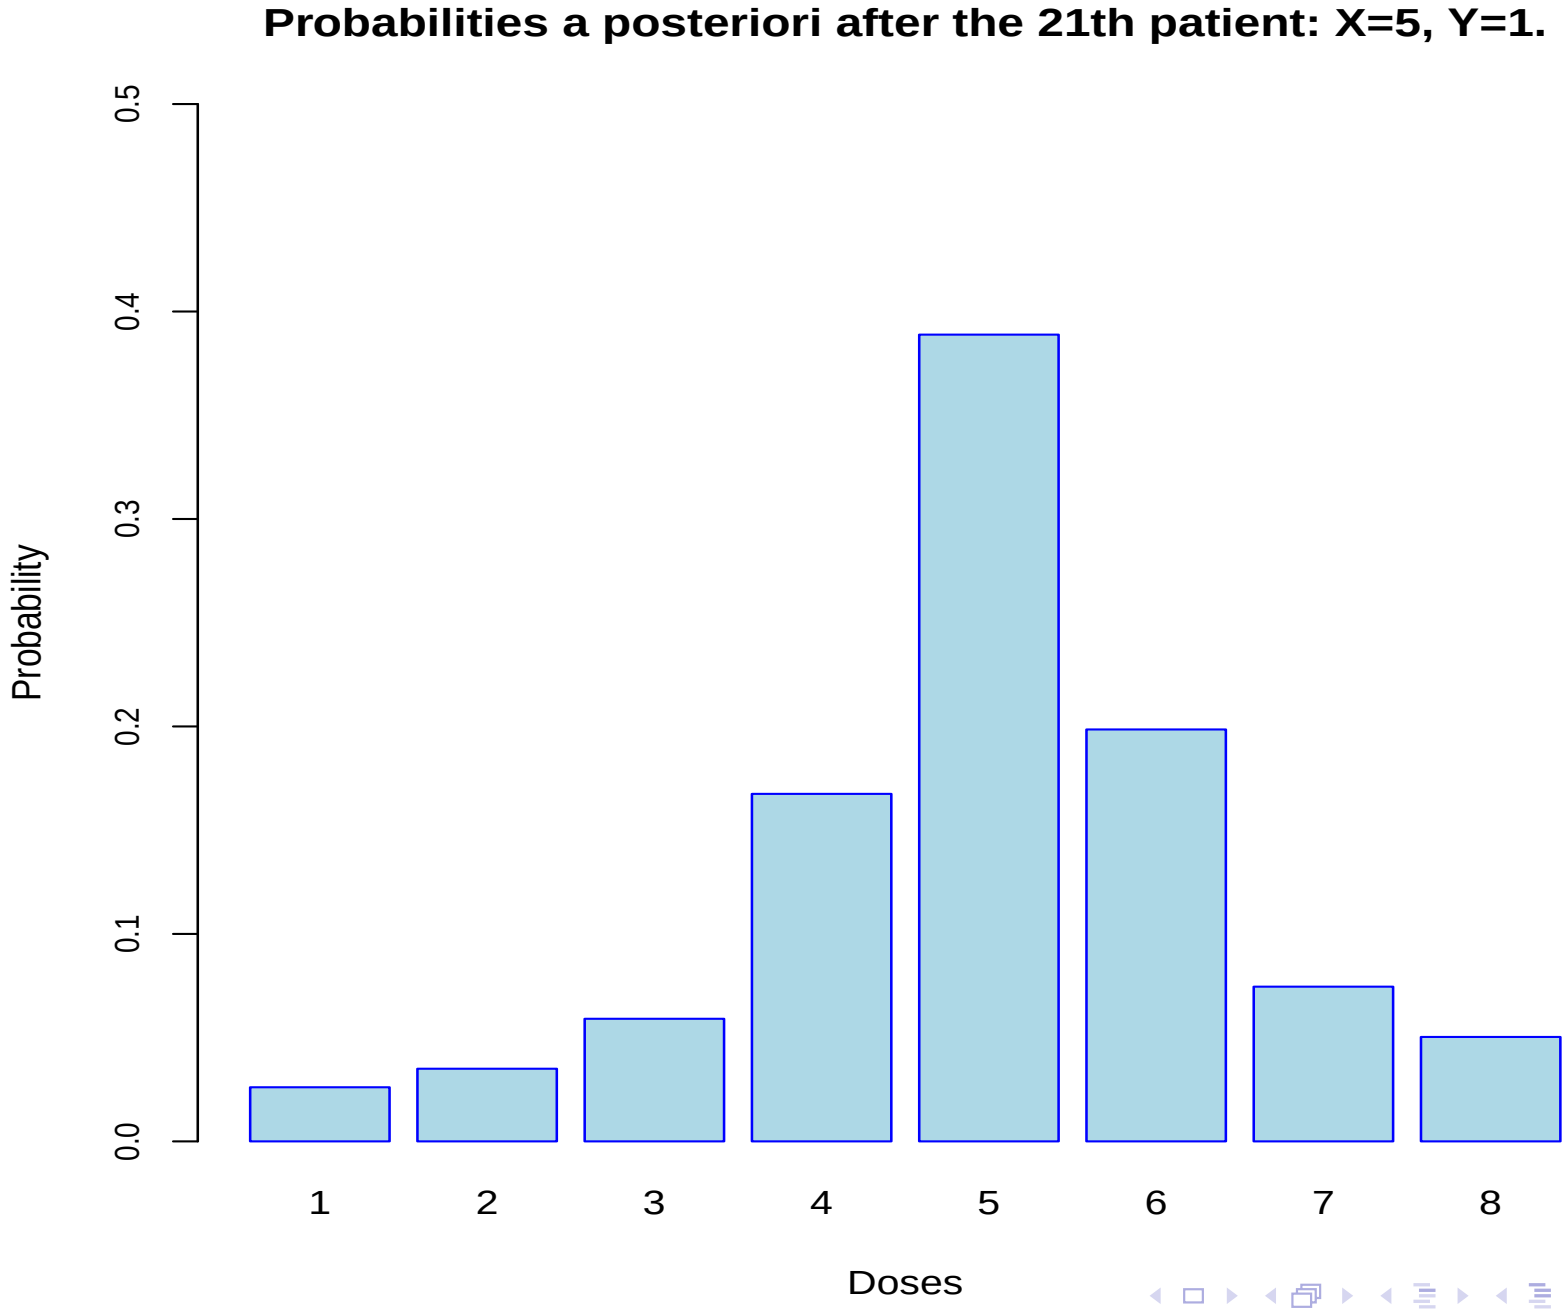
Probabilities a posteriori after the 21th patient: X=5, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

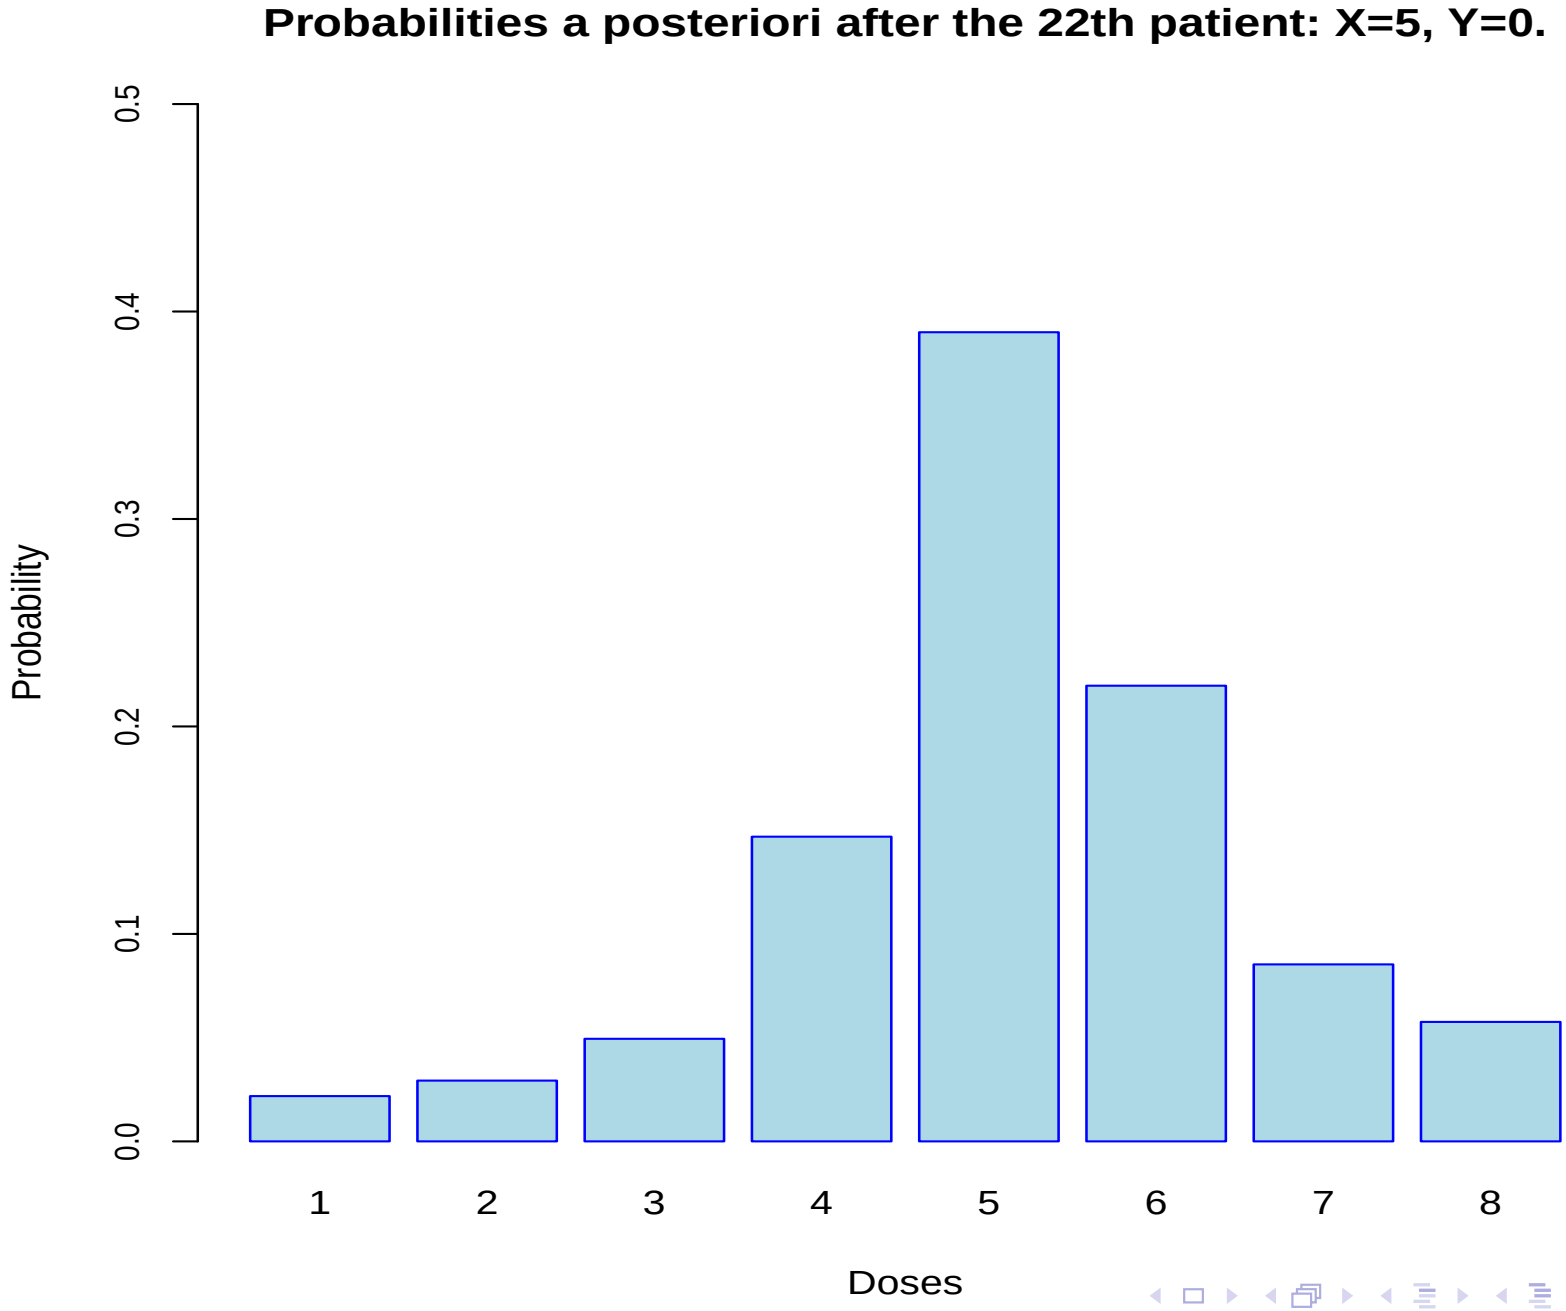
Probabilities a posteriori after the 22th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

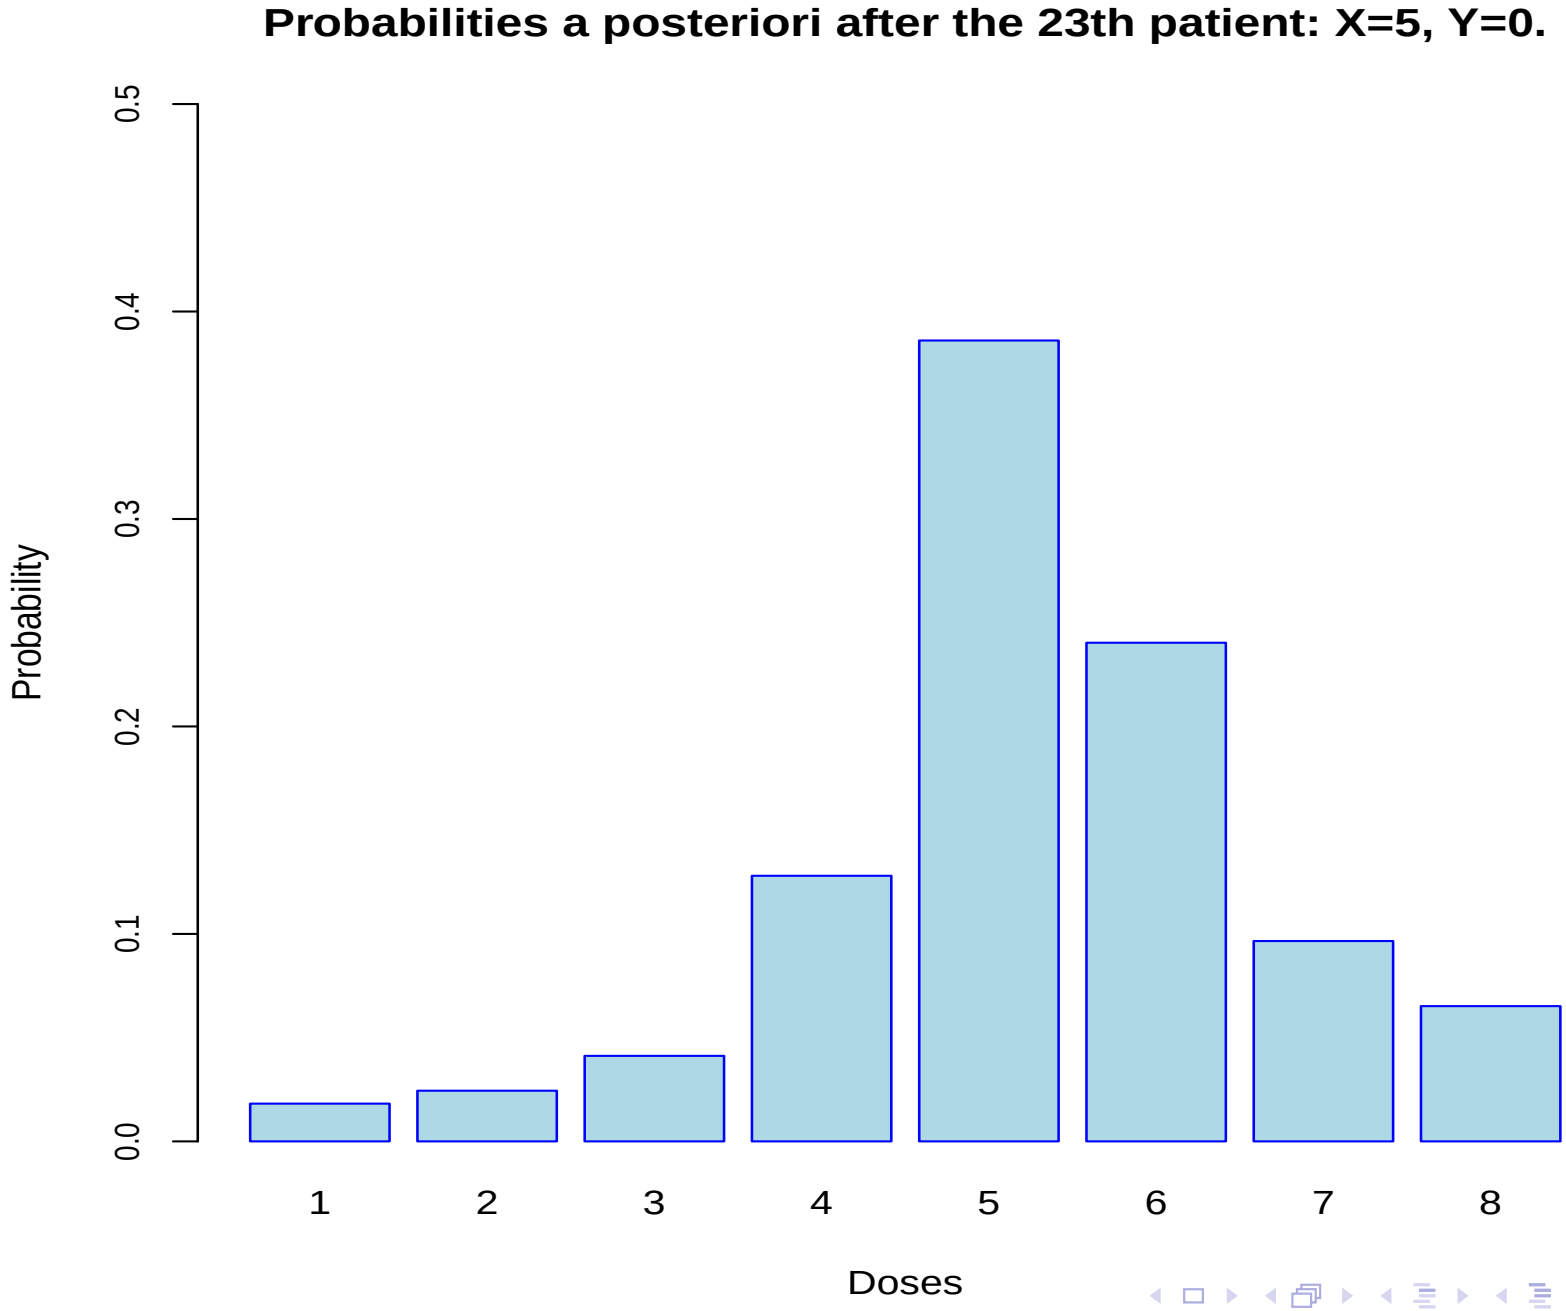
Probabilities a posteriori after the 23th patient: X=5, Y=0.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

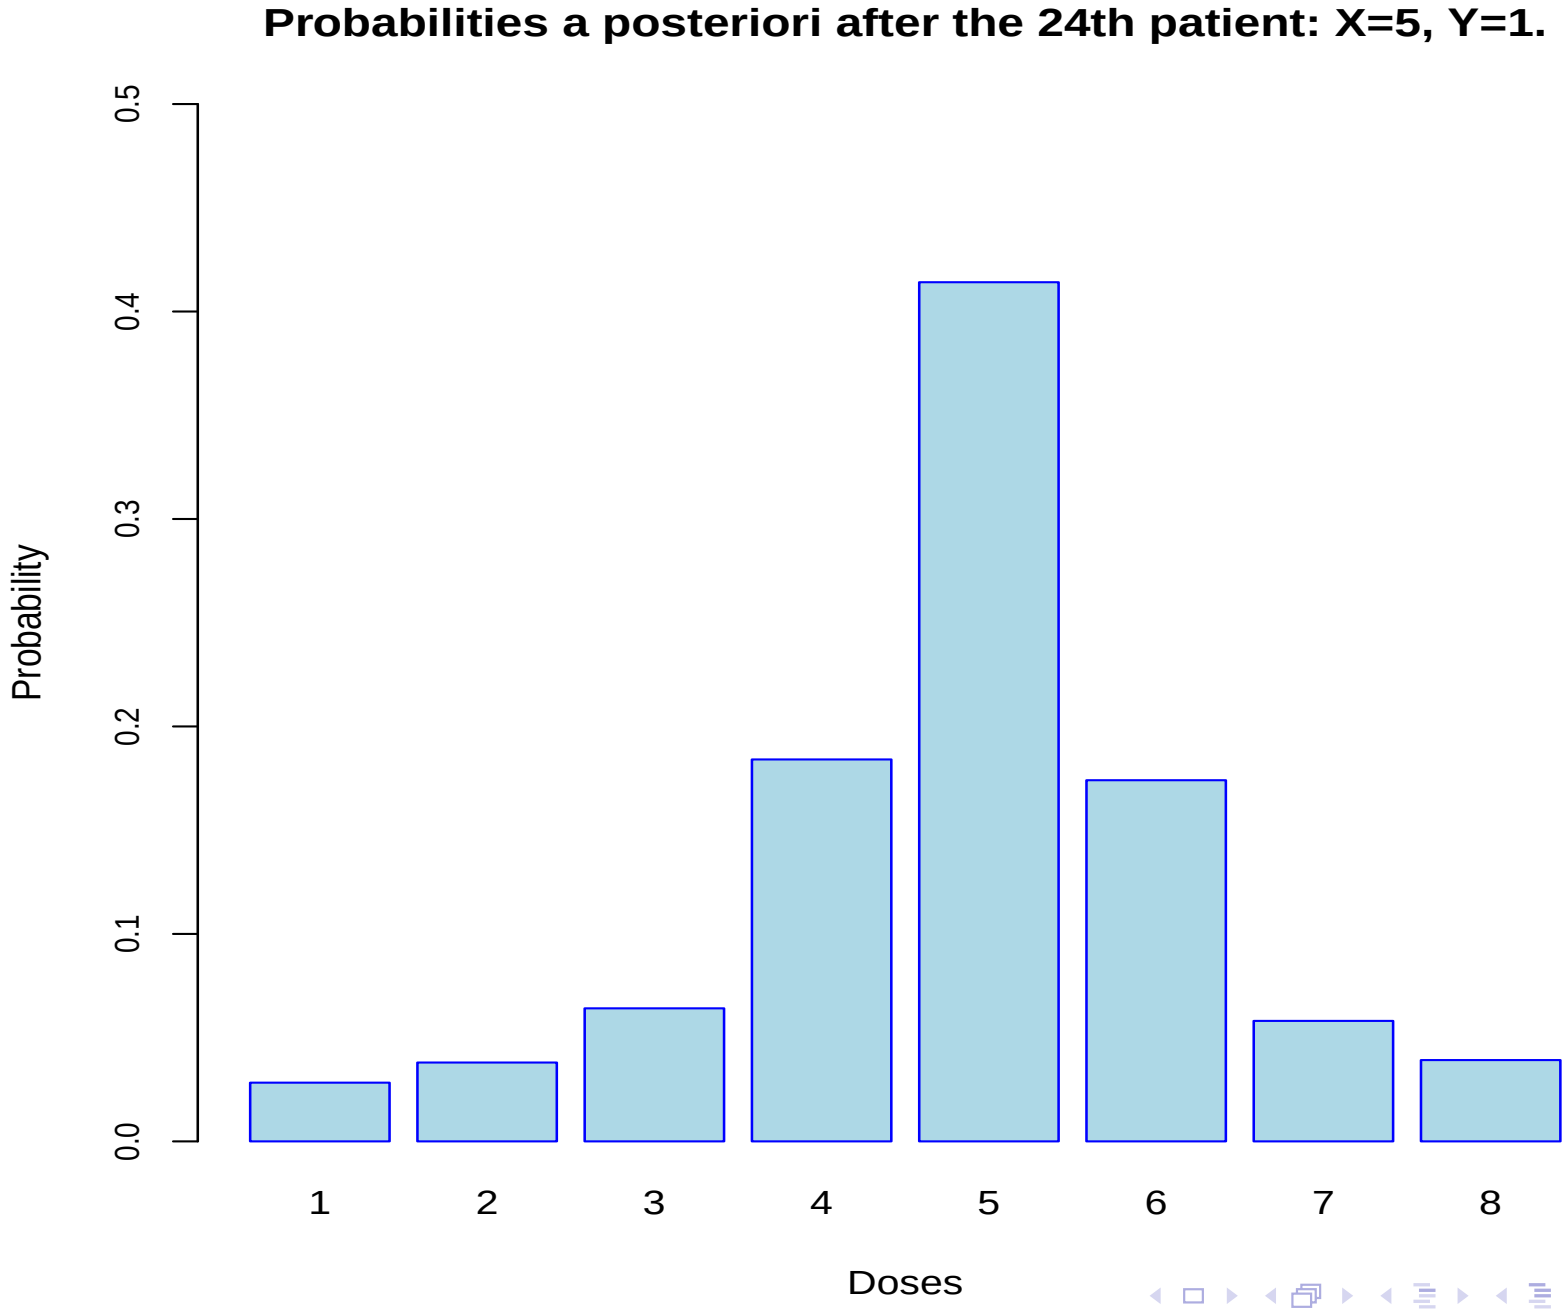
Probabilities a posteriori after the 24th patient: X=5, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

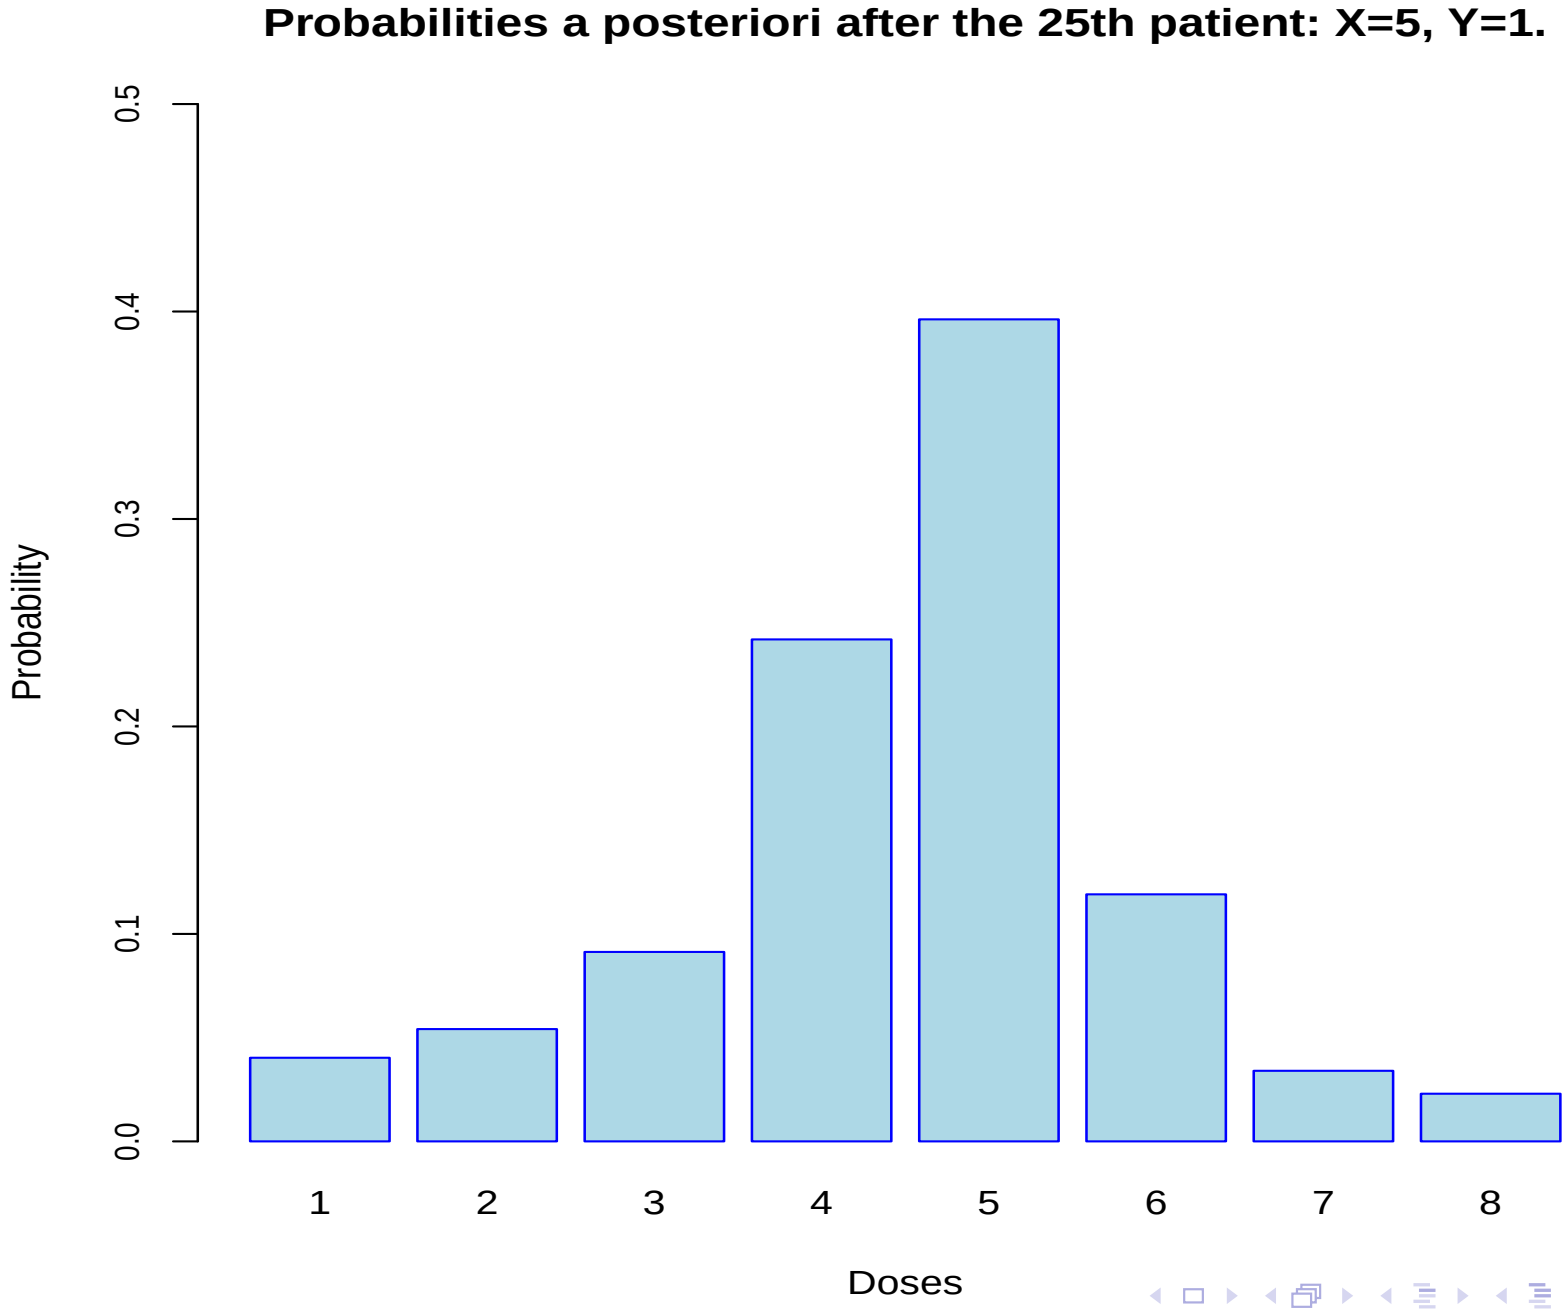
Probabilities a posteriori after the 25th patient: X=5, Y=1.
Probability
0.0
0.1
0.2
0.3
0.4
0.5
1
2
3
4
5
6
7
8
Doses

A trial under the scenario: $\beta = (0.02, 0.04, 0.18, 0.20, 0.28, 0.34, 0.40, 0.56)$.

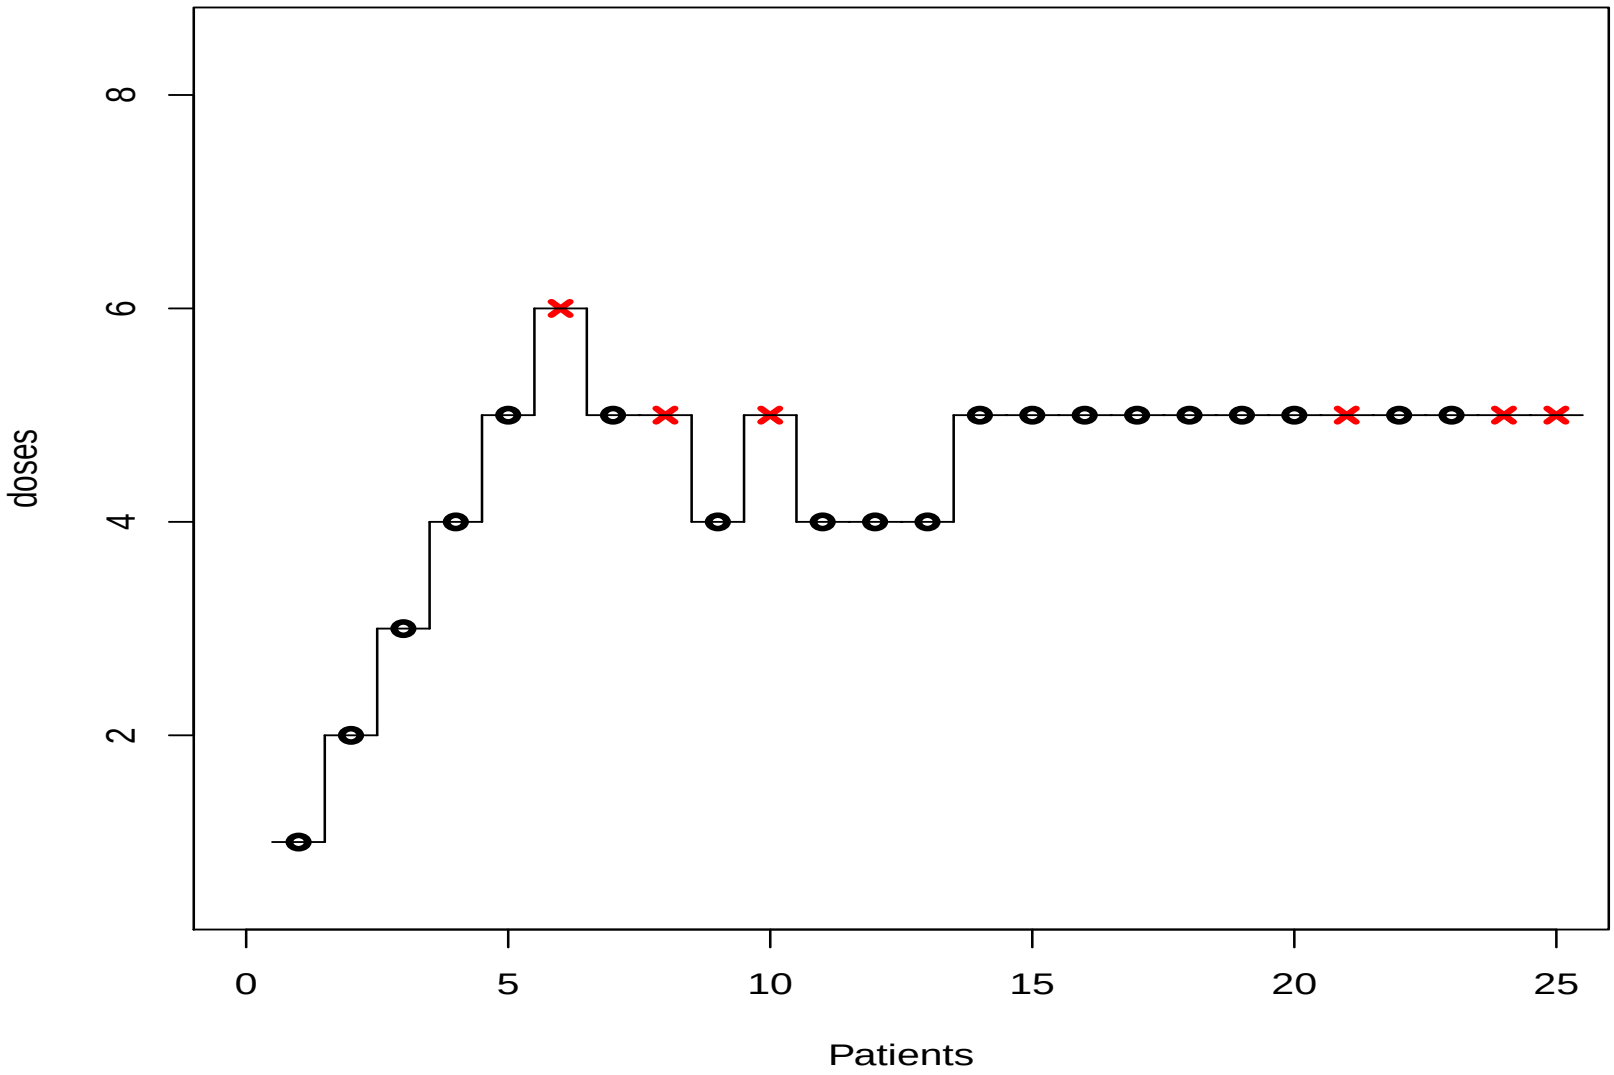

References


Azriel, D. and Mandel, M. and Rinott, Y.
The treatment versus experimentation dilemma in dose finding studies, *J. Statist. Plann. Inference*141-8:2759–2768, 2011.


Cheung, Y. K. and Chappell, R.
A Simple Technique to Evaluate Model Sensitivity in the Continual Reassessment Method, *Biometrics*, 58-3:1541-0420, 2002.


Cheung, Y. K.
Coherence principles in dose-finding studies, *Biometrika*, 92-4:863-873, 2005.


O'Quigley, J. and Pepe, M. and Fisher, L.
Continual reassessment method: a practical design for phase 1 clinical trials in cancer, *Biometrics*, 46-1:33–48, 1990.


O'Quigley, J. and Shen, L. Z.
Continual reassessment method: a likelihood approach, *Biometrics*, 46-1:673-684, 1996.

Coherent Behavior

Coherence Principle (Cheung)

A method $\mathcal{M}$ is coherent:

- In escalation, if $\mathbb{P}_{\mathcal{M}}(X_{n+1} > X_n | Y_n = 1) = 0$,
- In de-escalation, if $\mathbb{P}_{\mathcal{M}}(X_{n+1} < X_n | Y_n = 0) = 0$.

Coherent Behavior

Assumption

1. ***Structure:*** *The support S_θ of Λ_θ is:* $S_\theta = B^{\theta-1} \times I \times A^{m-\theta}$.
2. ***Independence:*** *Λ_θ is a product of unidimensional distributions at each dose:* $\Lambda_\theta = \Lambda_\theta^1 \times \ldots \times \Lambda_\theta^m$.
3. ***Stastistical ordering:*** *For any $n \in \mathbb{N}$, marginal posteriors on any dose d satisfy:* $\theta > \theta' \Rightarrow \Lambda^d_{\theta,n} \underset{st}{\preceq} \Lambda^d_{\theta',n}$.

Theorem

Under the preceding assumption, SPM is coherent.