

# Personalized dosing decisions using Bayesian multistate modeling

Bayesian Biostatistics 2025, Leiden

Juho Timonen, Eric Novik, Jacqueline Buross

# Personalized dosing

# The problem


- **Event risks:** Risk of death and other adverse events depends on drug dose
- **Patient preferences:** Patients may rate the harmfulness of each event differently
- How to find the optimal dose for a given patient?

# Event risks

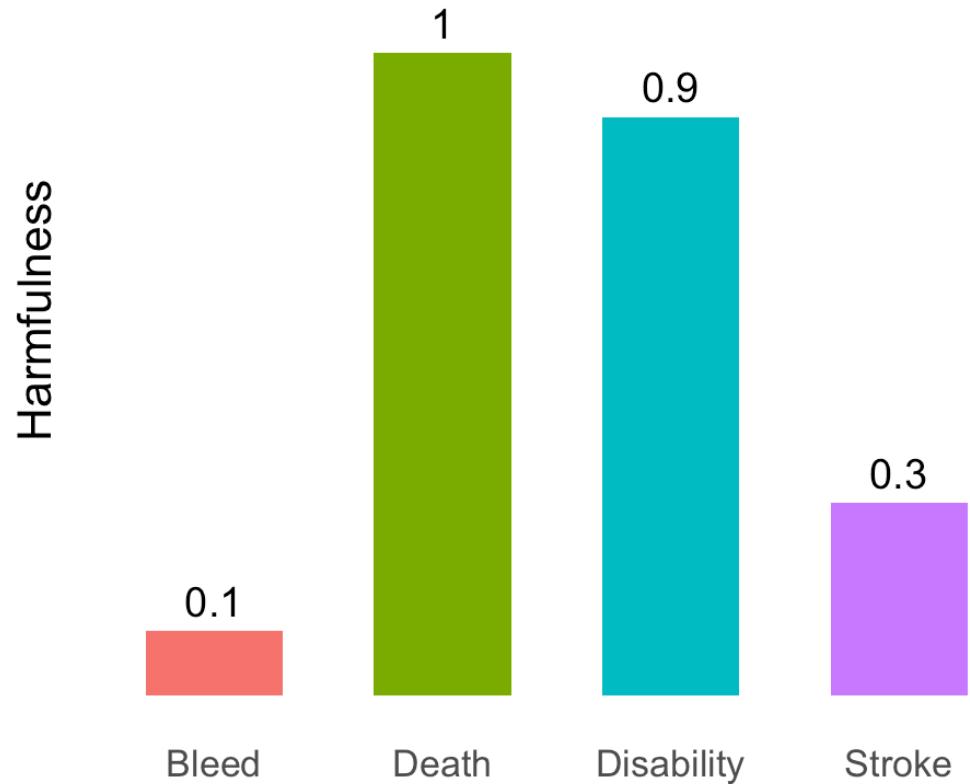
- Motivating example: anticoagulant drug
- Higher dose → higher bleeding risk
- Lower dose → higher risk of ischemic stroke

# Patient preferences

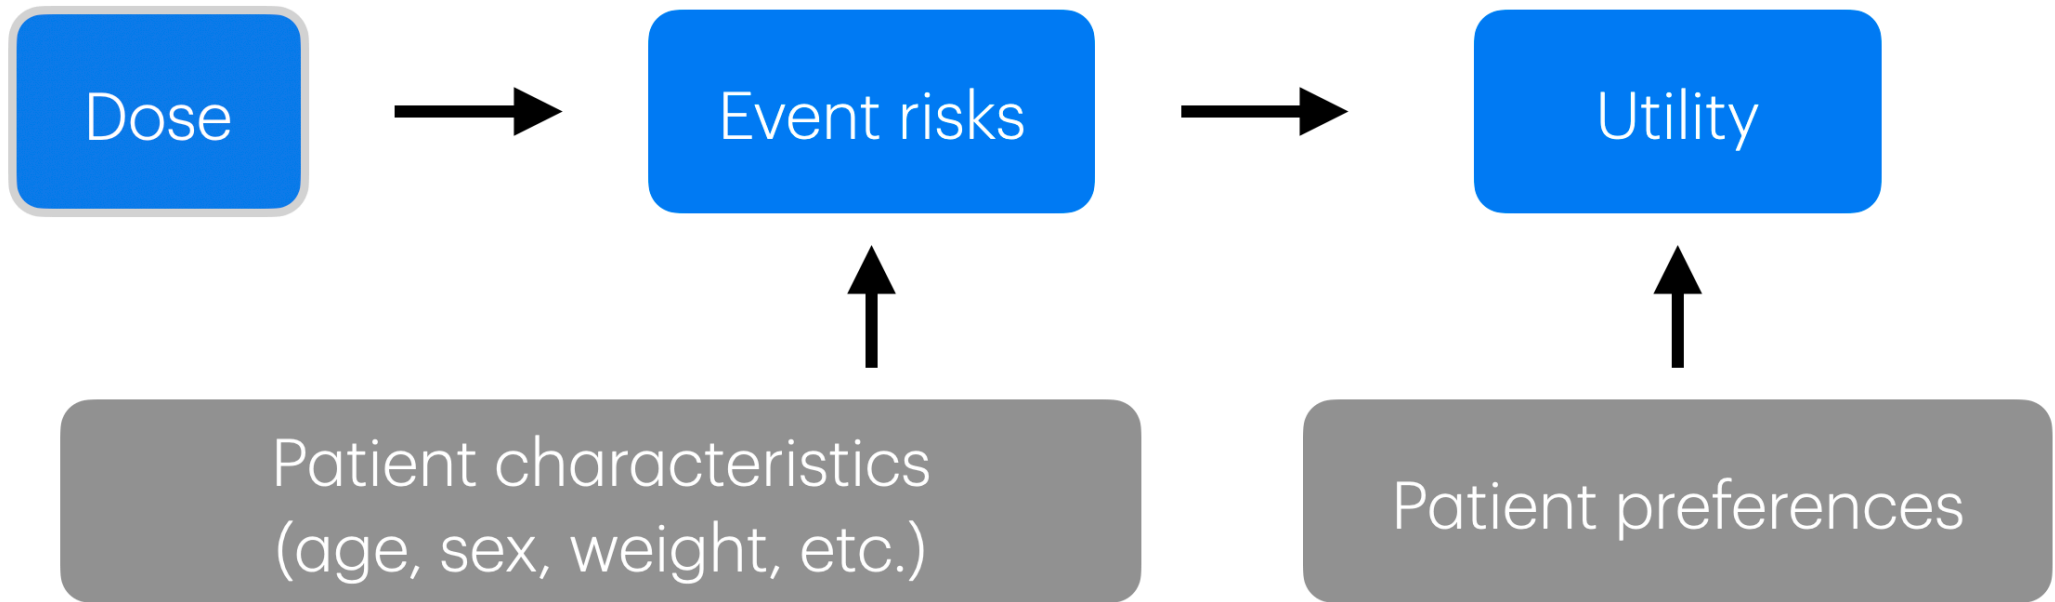
1) Avoid death



2) Other events matter



# Approach



- Select the dose that maximizes the utility

# Utility for patient $i$ given dose $D$

$$U_i(D) = - \sum_{e=1}^E P_i(e | D) H_i(e)$$

- Patient preferences captured by  $H_i(e) \in [0, 1]$ 
  - how harmful patient  $i$  rates event  $e$
- Event risks captured by  $P_i(e | D)$ 
  - probability that patient  $i$  experiences event  $e$  (in the next  $n$  years) if given dose  $D$  with regular dosing schedule

# Event risk model

# Event risk prediction

- We never observe the actual event risks (probabilities), just what events occurred and when
- We use time-to-event modeling to fit hazard functions
- The fitted hazard model can then be used to estimate event probabilities

# Cox proportional hazards (PH) survival model

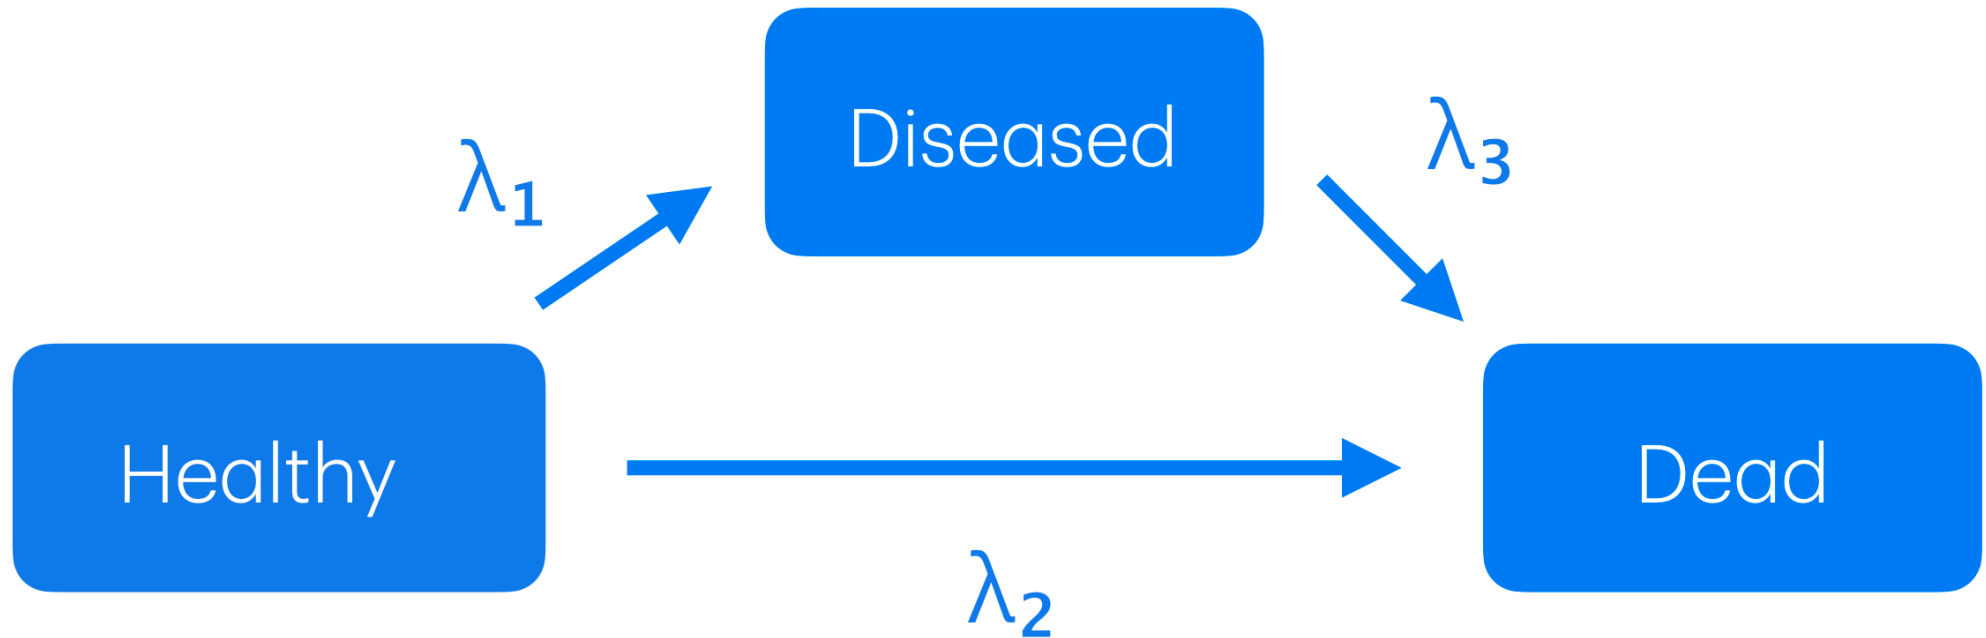
- Hazard function with time-independent covariates  $\mathbf{x}$  is

$$\lambda(t) = b(t) \exp(\beta^\top \mathbf{x})$$

where  $b(t)$  is a baseline hazard function

# Survival model as a multistate model

# Illness-death multistate model



- Hazards of transitions starting from same state are competing

# Transition hazards

$$\lambda_h(t) = b_h(t) \exp(\beta_h^\top \mathbf{x})$$

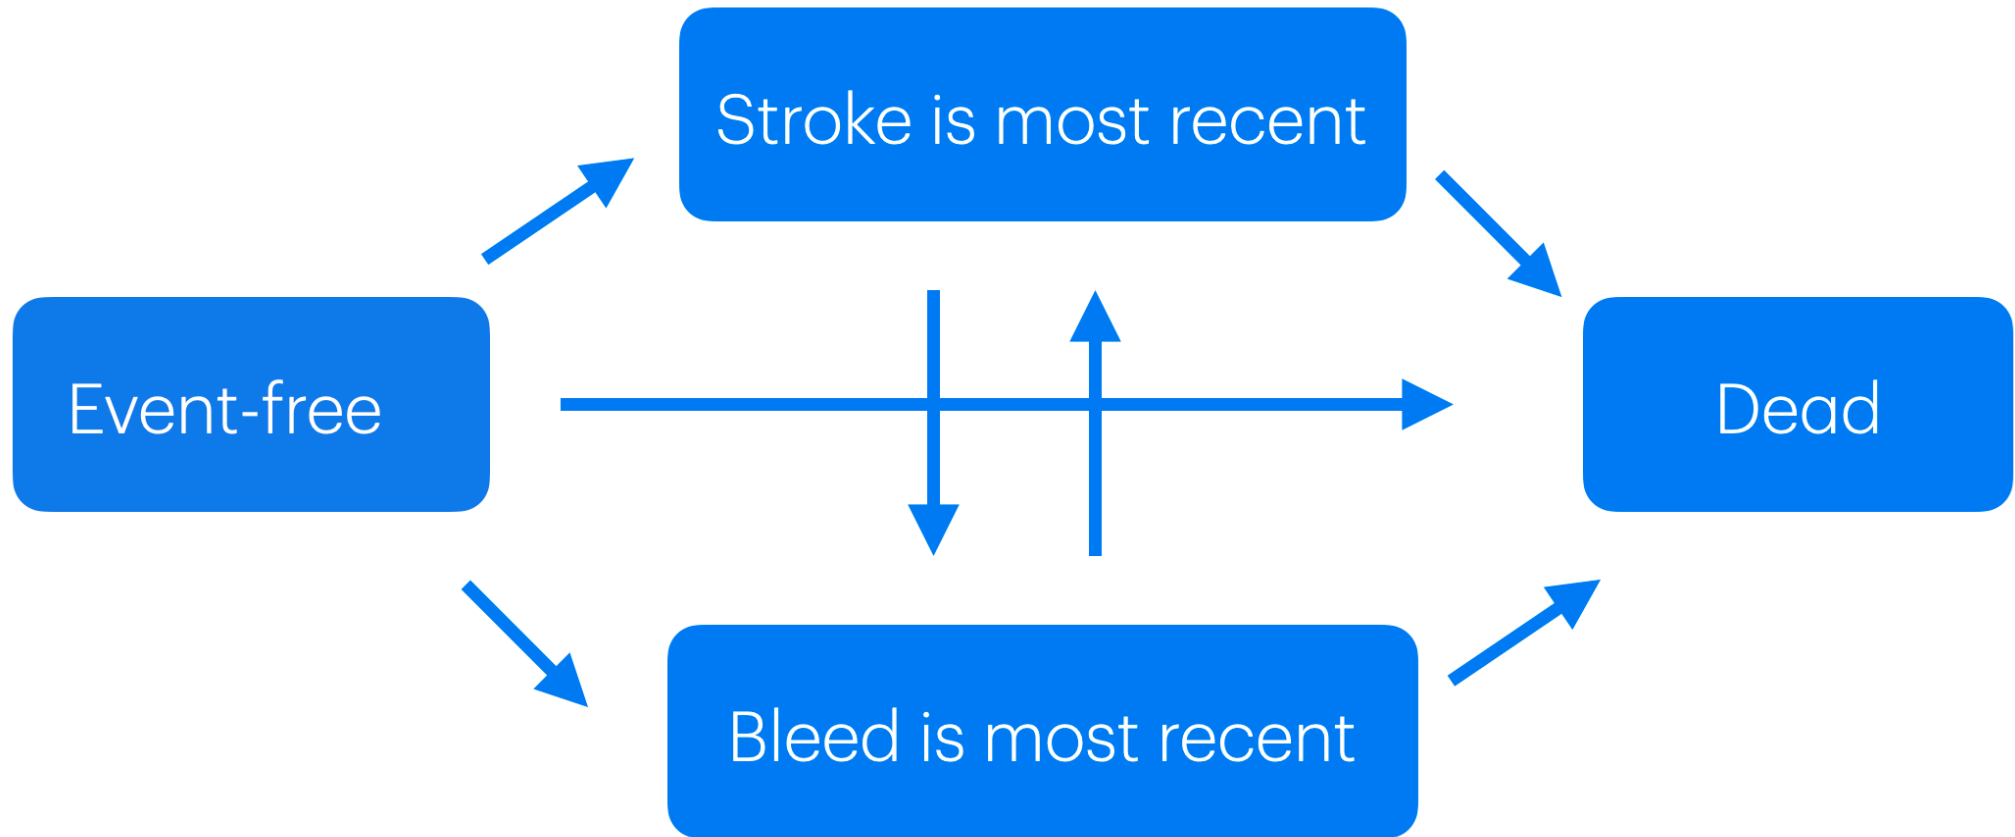
- We model  $b_h(t)$  using splines
- Unknown parameters are  $\beta_h$  and spline weights
- Bayesian inference using Stan

# Markov vs. semi-Markov

$$\lambda_h(t) = b_h(t) \exp(\beta_h^\top \mathbf{x})$$

- Markov:  $t$  is time from start of study (*we use this*)
- semi-Markov:  $t$  is time since entering current state (*alternative*)


# Larger multistate model



- Multistate model with 4 states and 7 transitions
- State is defined based on most recent event

# State paths

- The current state determines the hazards of moving to each state



# Fitting a multistate model

- Data is the observed event sequence for each patient
- First we translate these into state paths
- This splits the timeline into intervals
- On each interval we can tell what state the patient was in and what transitions were at risk

# Likelihood

- Log likelihood of observing a given state path
- Sum over  $K$  intervals and  $H$  transitions

$$\log p(\text{PATH} \mid \theta) = \sum_{k=1}^K \sum_{h=1}^H \left[ \delta_{k,h} \log(\lambda_h(t_k \mid \theta)) - \mathbf{I}_{k,h} \right]$$

- Binary variables  $\mathbf{I}_{k,h}$  and  $\delta_{k,h}$  indicates which transitions occurred and which were at risk

# Event risk prediction with multistate models

# Continuous markov models (no loops in model graph)

- A generator (transition intensity) matrix  $\mathbf{Q}(t)$  of size  $S \times S$ , where,  $S$  is the number of states
- The corresponding transition probability matrix  $\mathbf{P}(t)$  can be solved from

$$\frac{d}{dt}\mathbf{P}(t) = \mathbf{P}(t)\mathbf{Q}(t)$$

using the edge condition  $\mathbf{P}(0) = I$ .

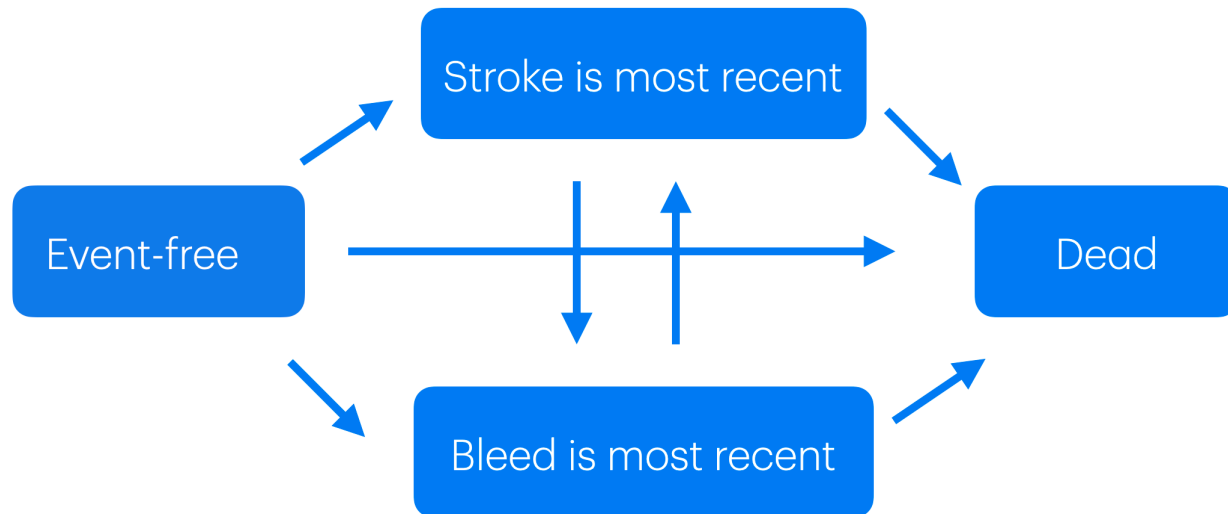
# Generating paths

- We can simulate paths according to the assumed data-generating process
- We can estimate event risks based on the proportion of paths where the event occurred
- Requires a lot of computation and is random, but more generally applicable

# Simulation study

# Setup

- Data simulated using a multistate (MS) model
- Age and dose affect each transition hazard differently
- The baseline hazard is higher for example for Stroke  $\rightarrow$  Dead than Event-free  $\rightarrow$  Dead




# Questions of interest

- Can we fit a MS model that recovers the true death risks?
- How well does a standard survival model (Cox PH) compare?

# 3-year death risk prediction


**a**



- Starting from  $t = 0$  and event-free state
- Here we compare with the true event risks used to simulate the data

# Starting from $t = 1$ year

State at  $t = 1$  year    ● Bleed    ▲ Event-free    ■ Stroke



- Death risk prediction for the remaining 2 years

# Concordance of risk predictions

Concordance indices of predicted 3-year event probabilities

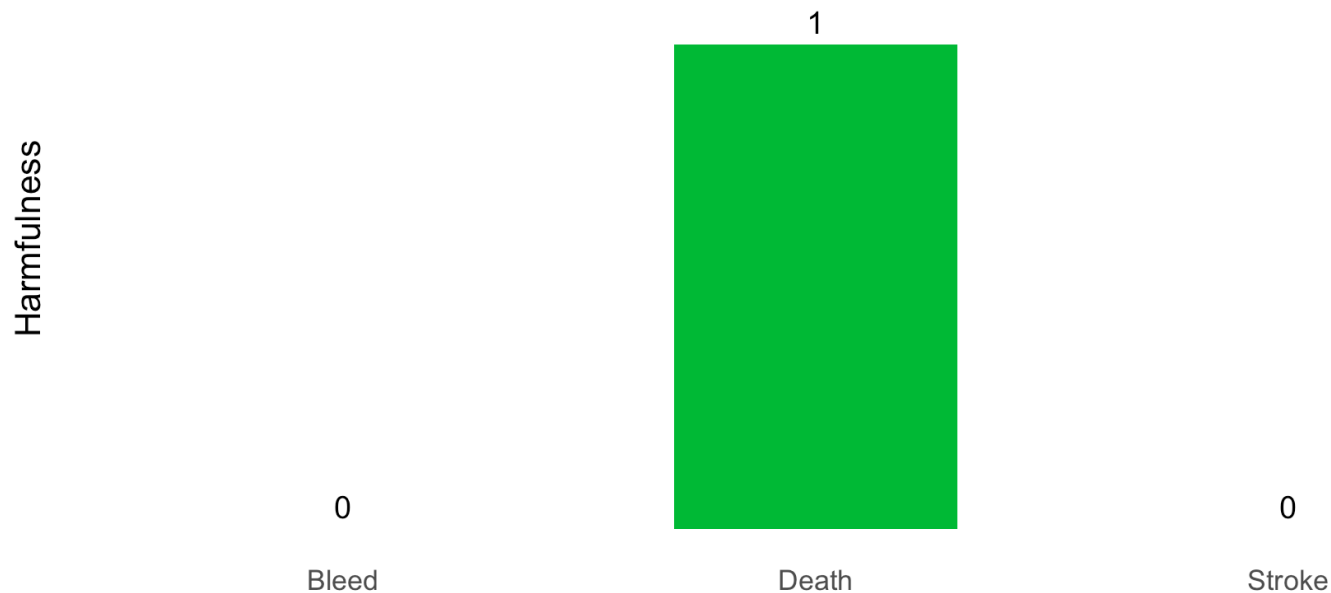
Start time	Model		
	MS.FIT	COXPH.FIT	GROUND TRUTH
$t_0 = 0$	0.572	0.547	0.571
$t_0 = 1$ year	0.681	0.521	0.683

- In reality we have to rely on for example concordance index
- Even knowing the true event risks does not give perfect concordance

# Dose selection

- Here, all patients have the same preference
- Only predicting risk of death matters

Patient preferences



# Dose selection

For each model, percentage of subjects for which each dose was selected as the best one

Dose (mg)	Model		
	1. GROUND TRUTH	2. MS FIT	3. COXPH FIT
15	0.00%	0.00%	0.00%
30	99.77%	92.37%	0.00%
60	0.23%	7.63%	100.00%

- Here we know the correct dosing decision for each patient

# Challenges

# Challenge: large number of states

- The more possible states and transitions we include in the model, the more likely it is that we get transitions that were never observed
- Should one combine e.g. major bleed and minor bleed into same state?

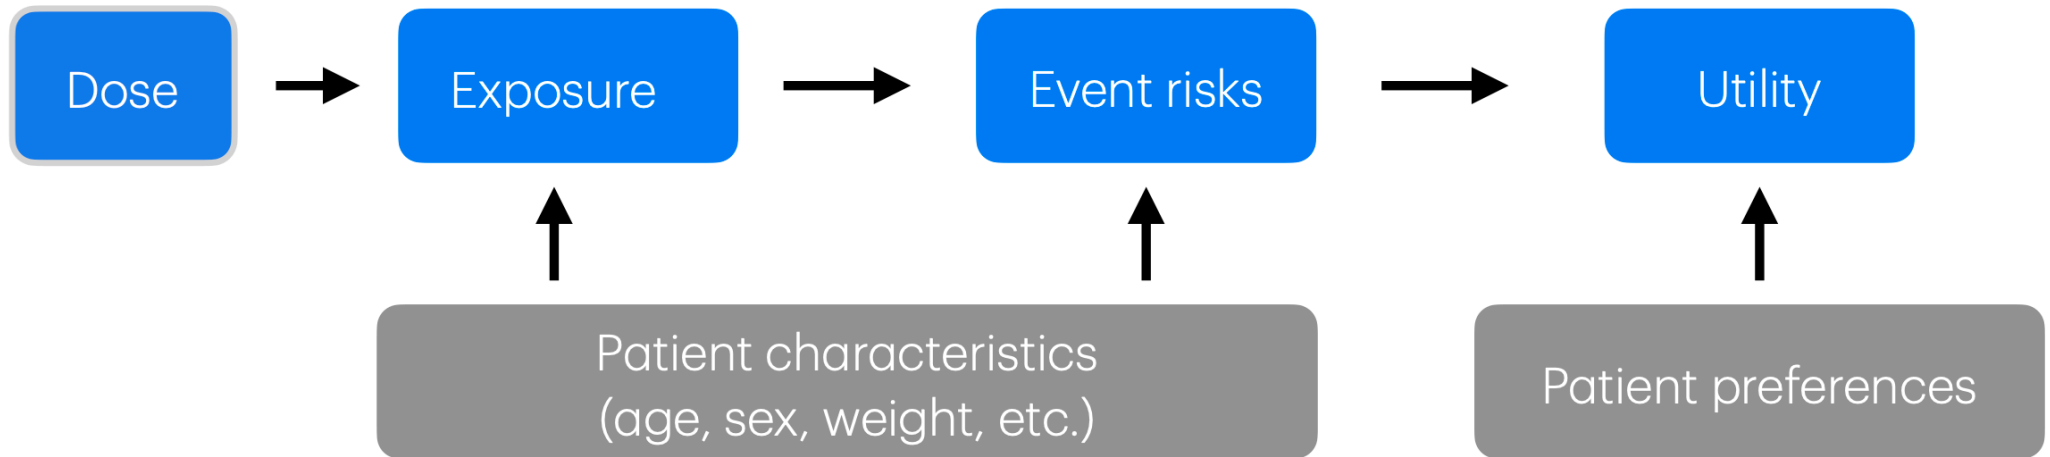
# Challenge: including relevant history

# Challenge: going off dose

- We are assuming the same constant dose
- In reality patients may have gone off dose during the trial

# Exposure-hazard model

# Exposure-hazard model



- Exposure depends on how much of the drug enters the body, how long it stays there, and at what concentration

# Exposure models

- Simplest model: Exposure = Dose
- More complex: Pharmacokinetic (PK) modeling of the drug concentration
  - needs concentration measurement data to fit
  - can account for patient characteristics that affect drug absorption and clearance

# PK modeling

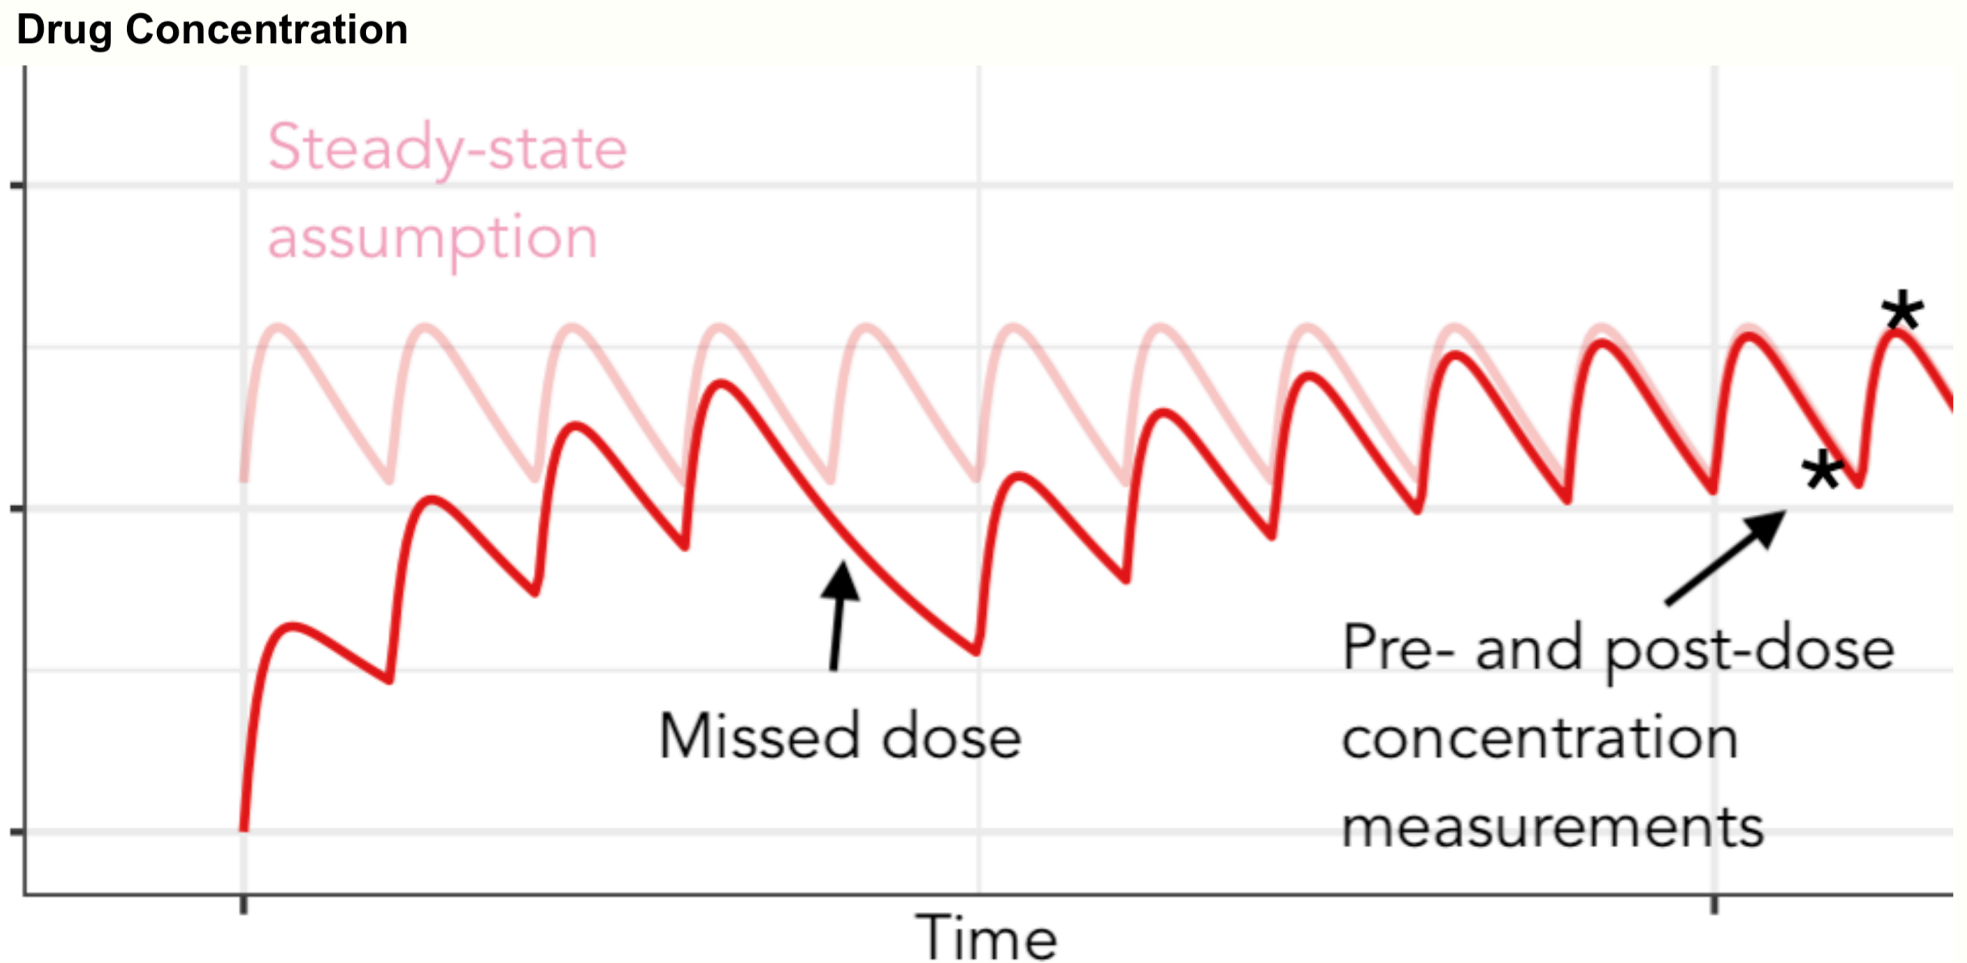
- Concentration after hundreds of doses is computationally demanding to evaluate
- However, with regular dosing schedule, a steady state is reached

# Computationally efficient strategy

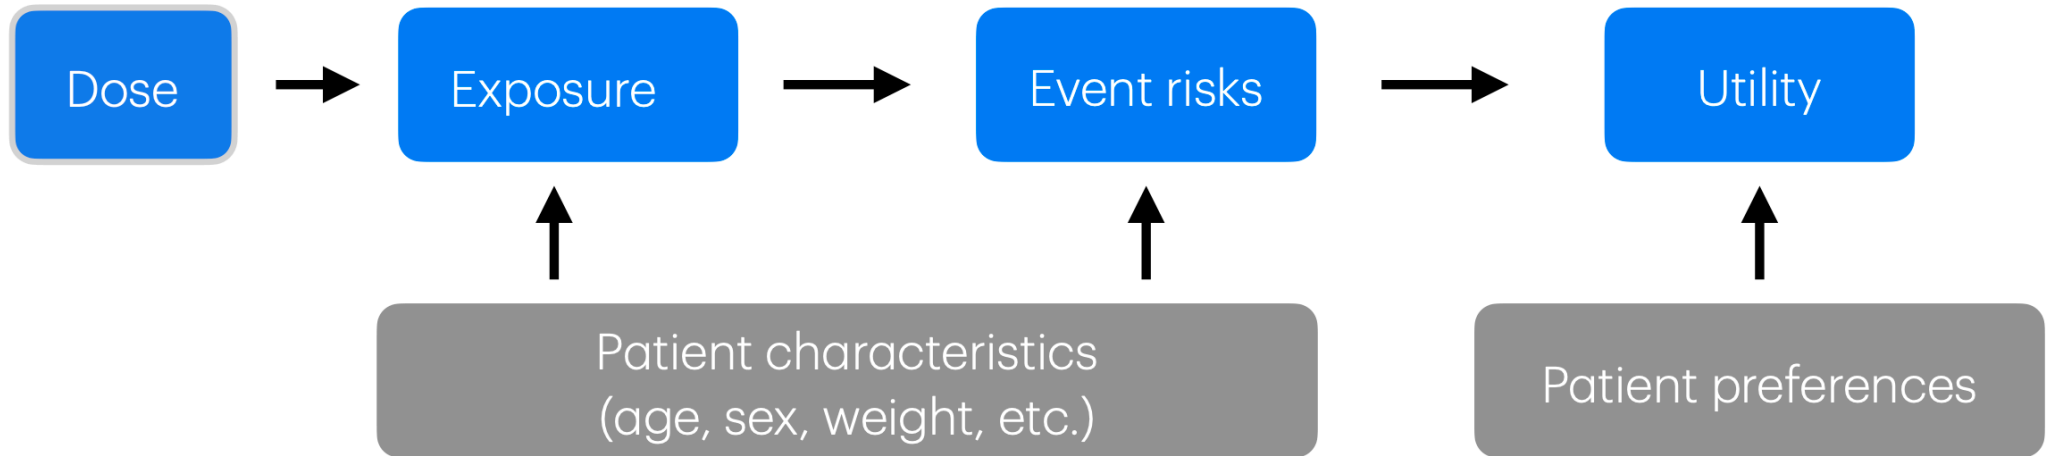
- Exposure = average concentration at steady-state (regular dosing schedule)
- Needs only a single pair of pre- and post-dose concentration measurements
- We can fit this jointly with the event risk prediction model

# Idea

- Assume steady state until last few dose occasions



# The end



- Select the dose that maximizes the utility