Making decisions (part 1)

Computational Cognitive Science 2014

Dan Navarro

Lecture outline

- What is decision-making?
- The classical approach: expected utility
- Calculating probabilities isn’t easy
- Assigning utilities is tricky

“Decision making” is a broad term

Which of these is darker?

Which of these is darker?

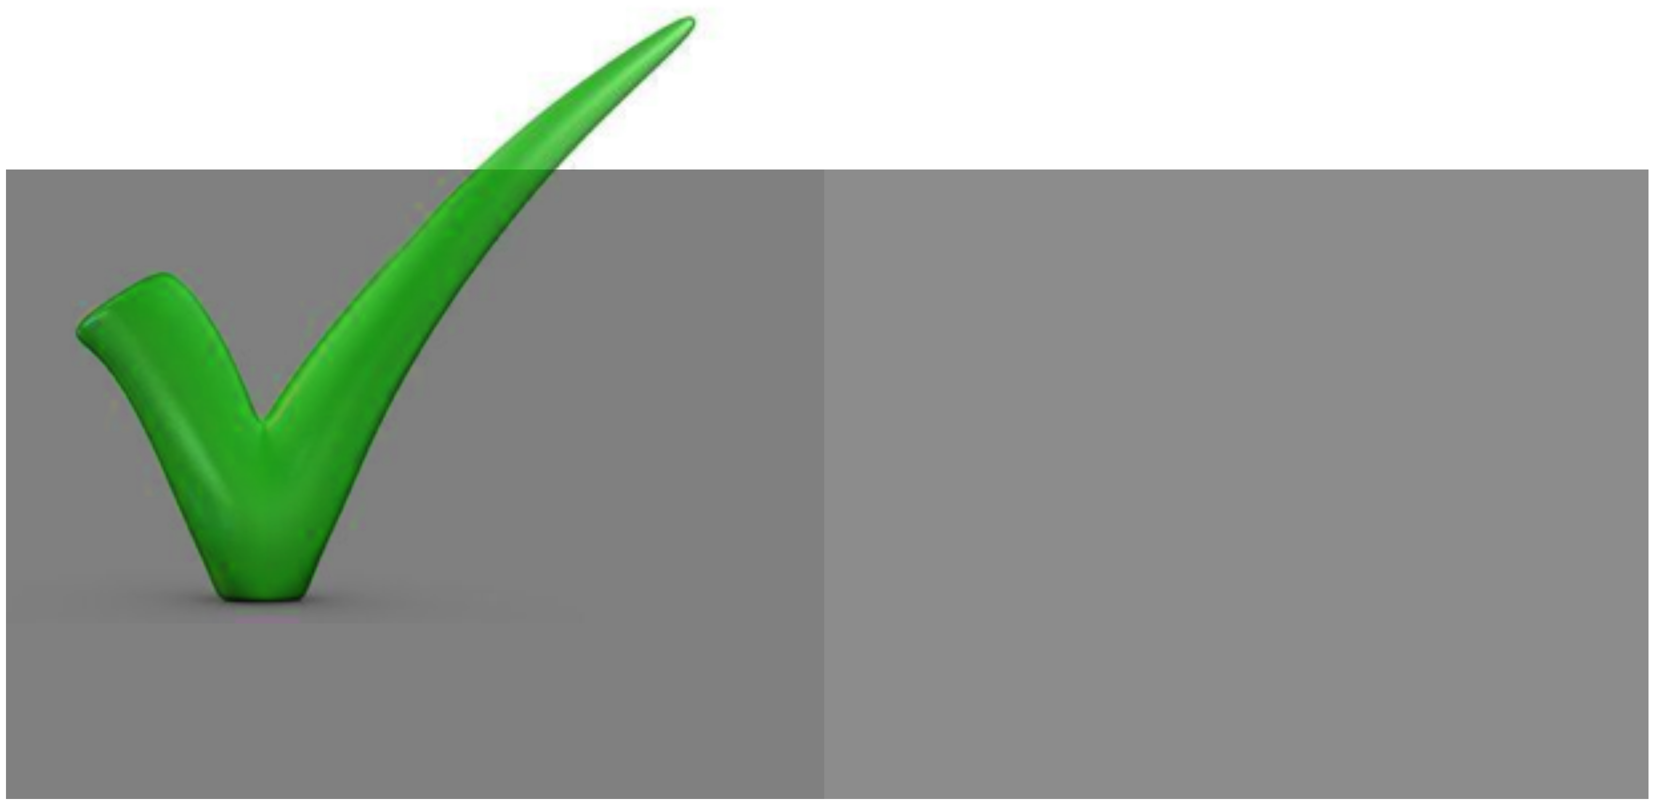

Should I eat the beef or the fish?

Should I eat the beef or the fish?

Vacation in Rome or in Venice?

Vacation in Rome or in Venice?

Cut the left wire or the right wire?

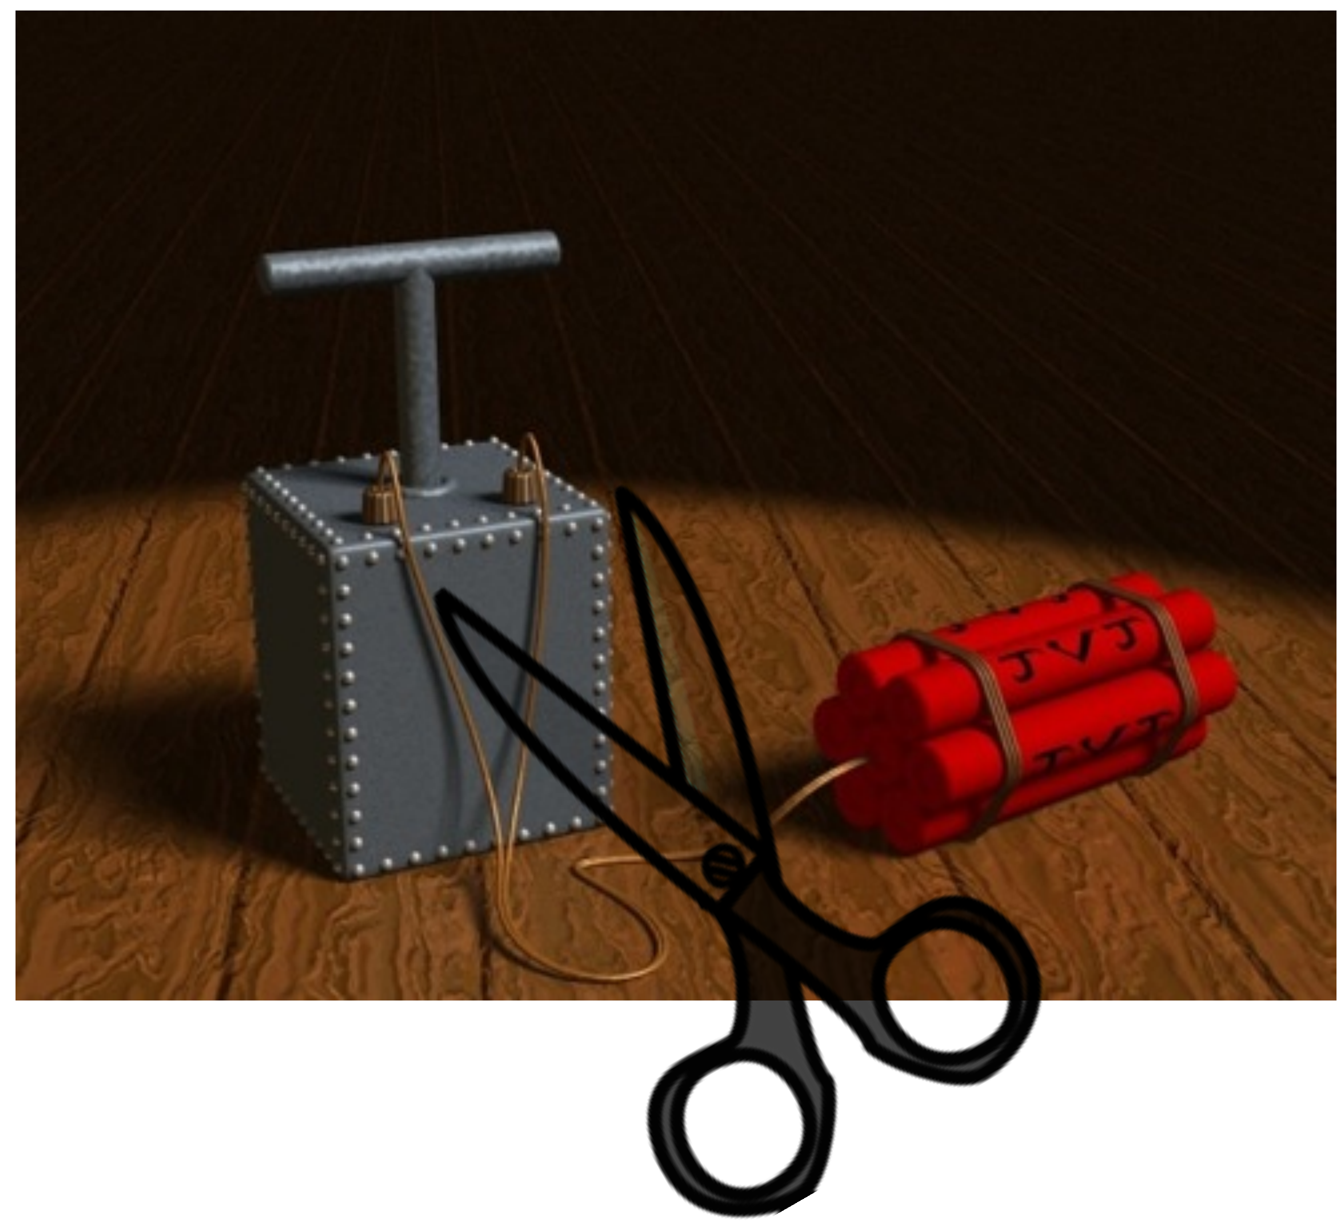

Cut the left wire or the right wire?

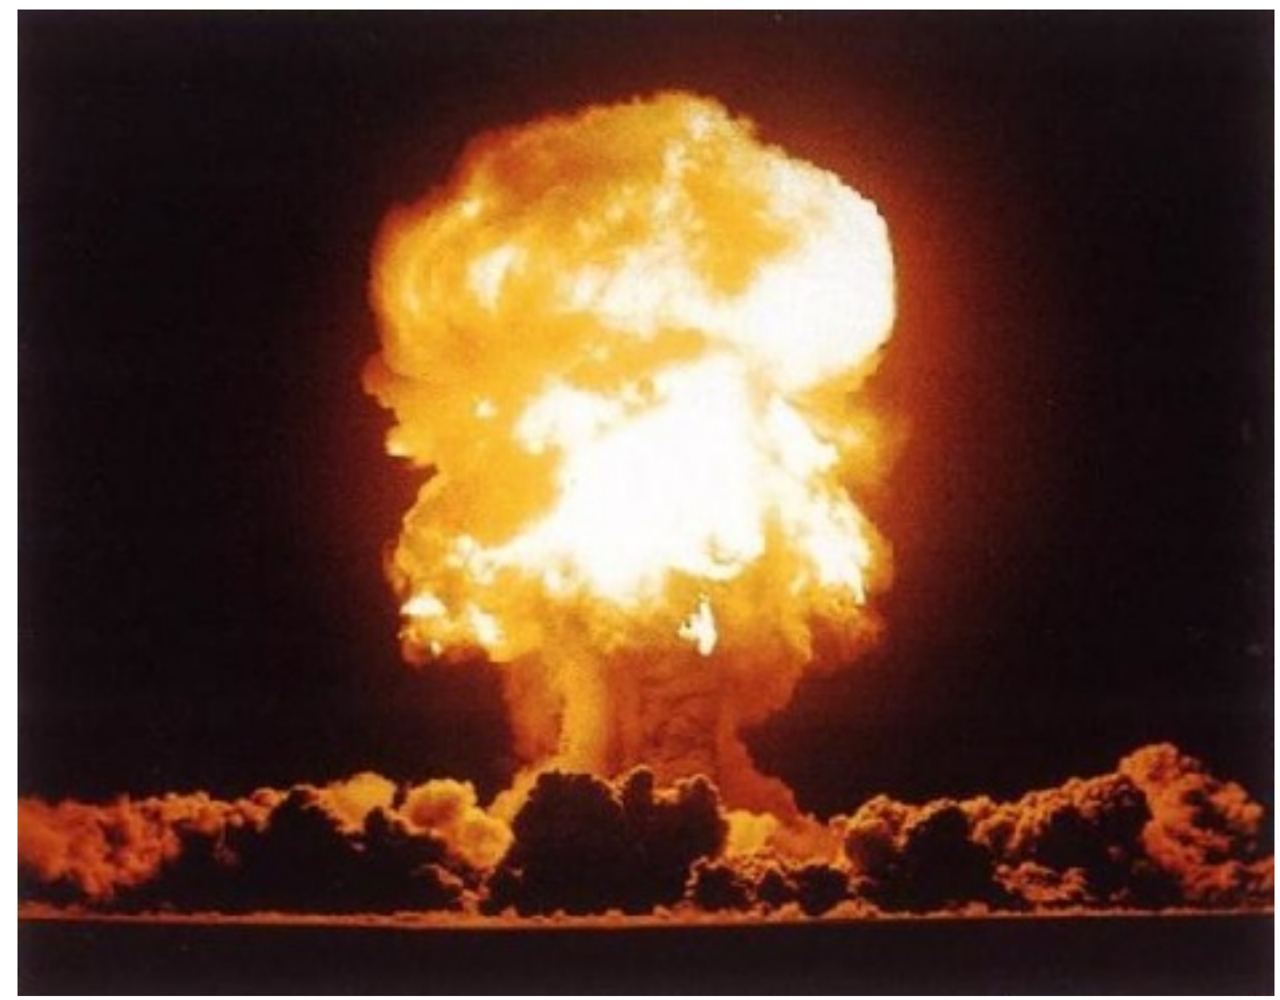

A set of possible actions

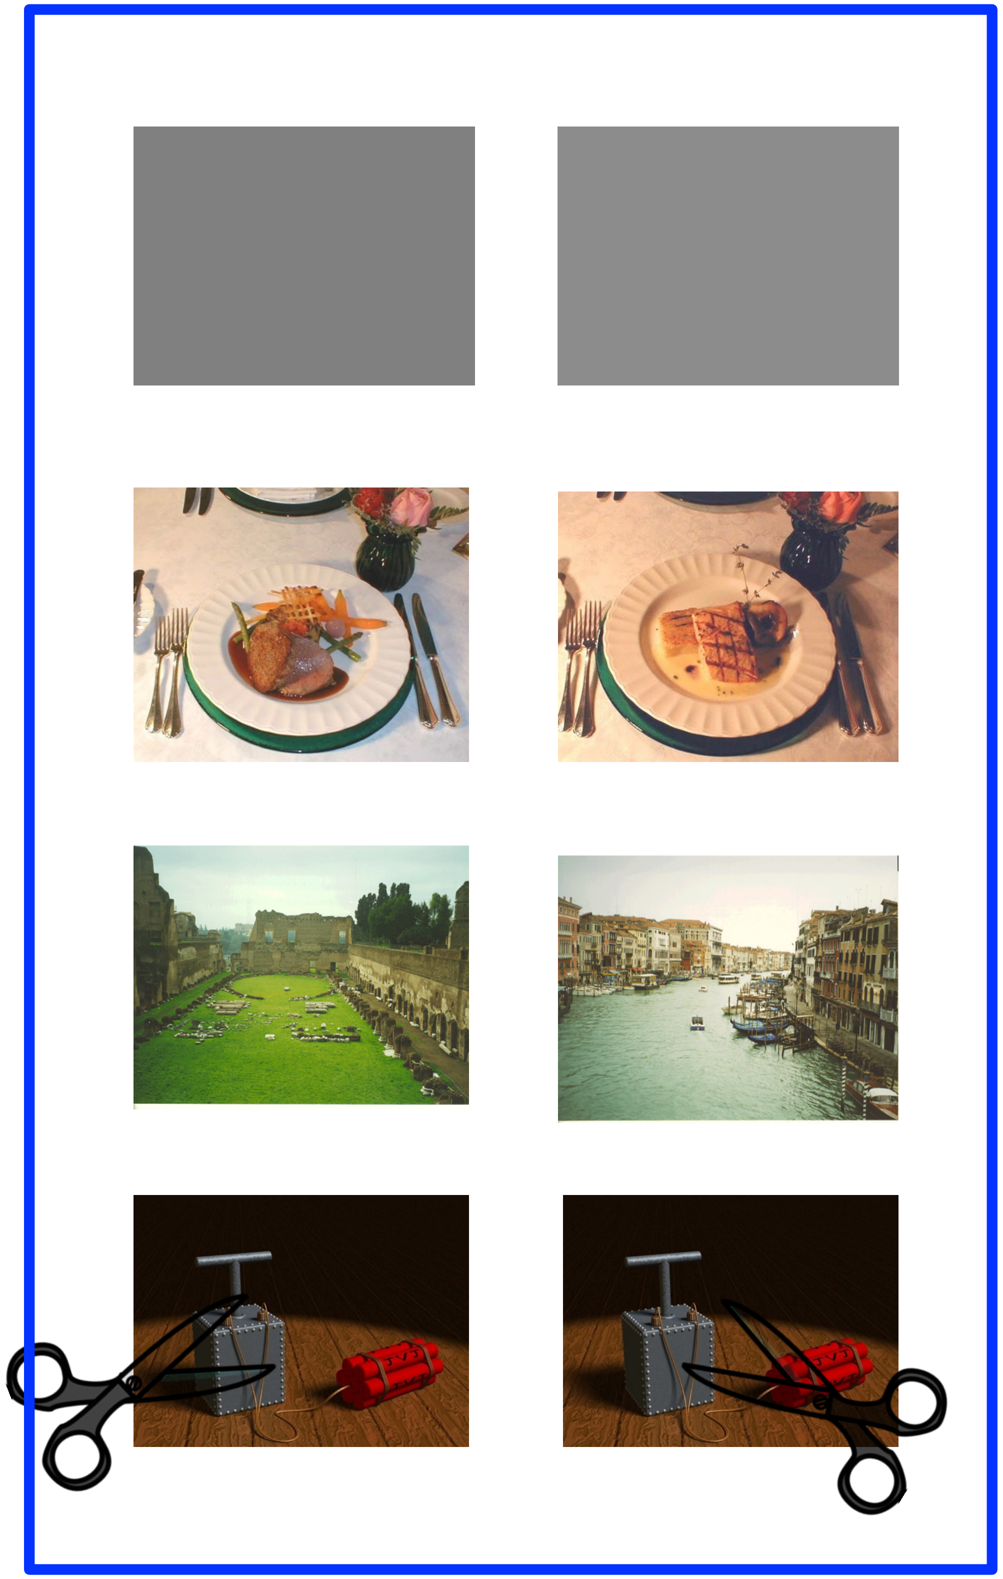

Outcomes associated with actions

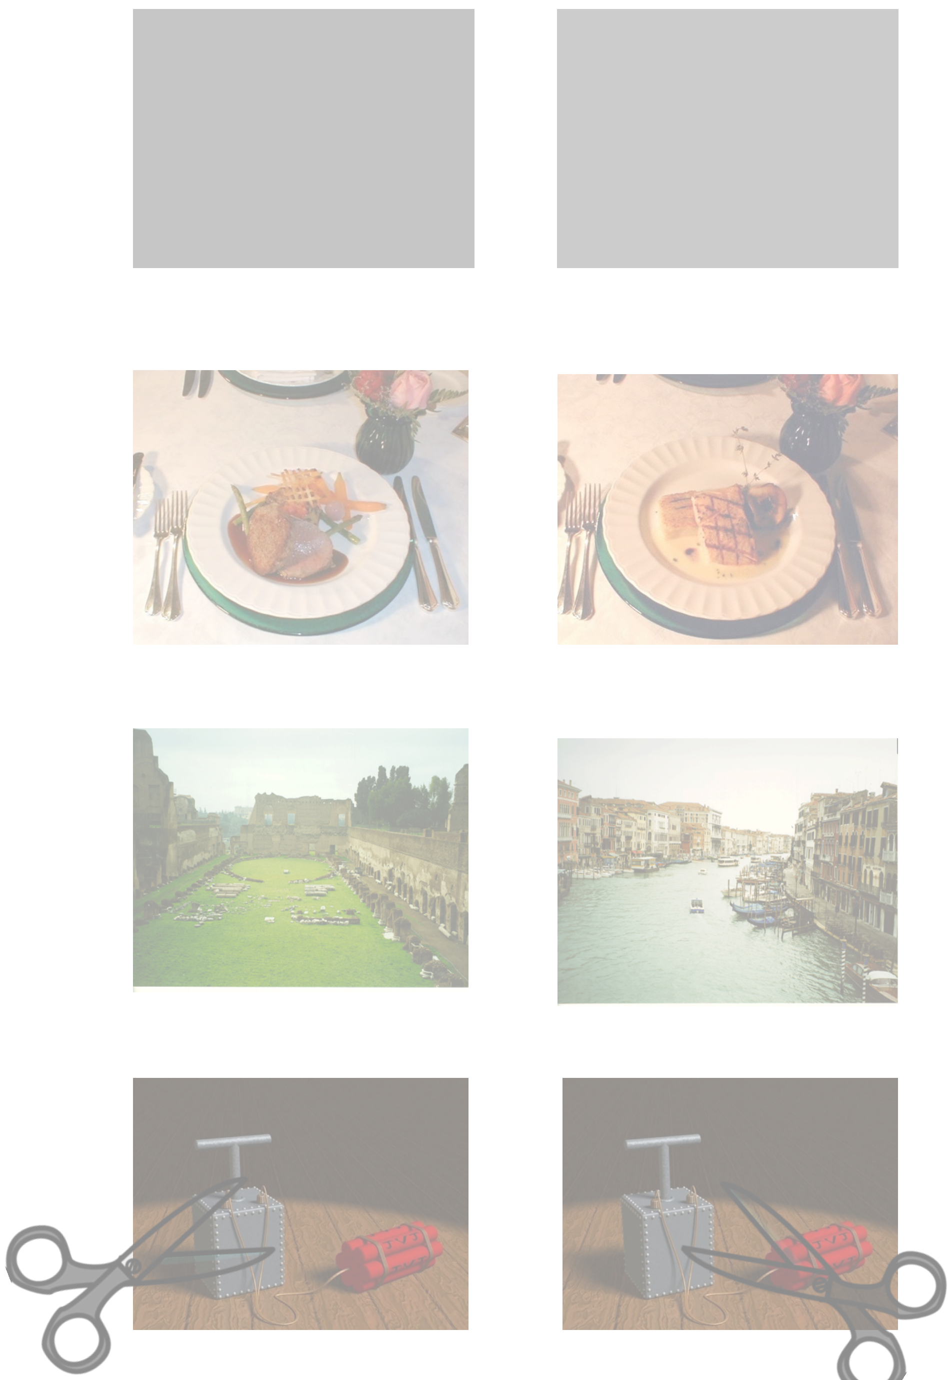

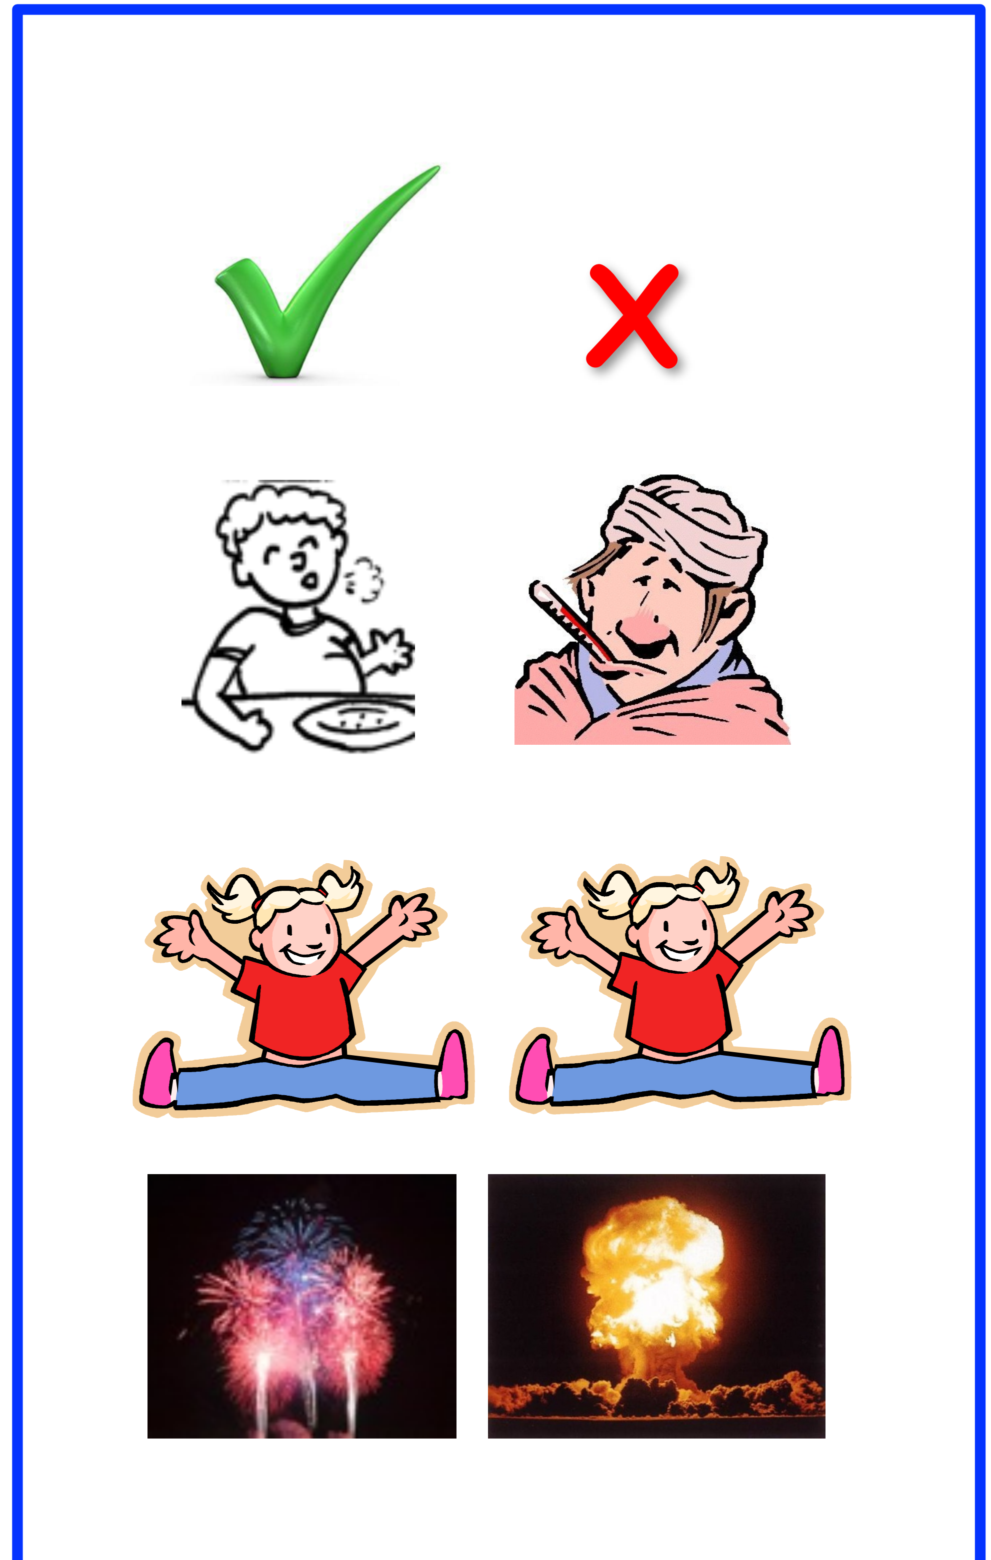

Utilities associated with outcomes?

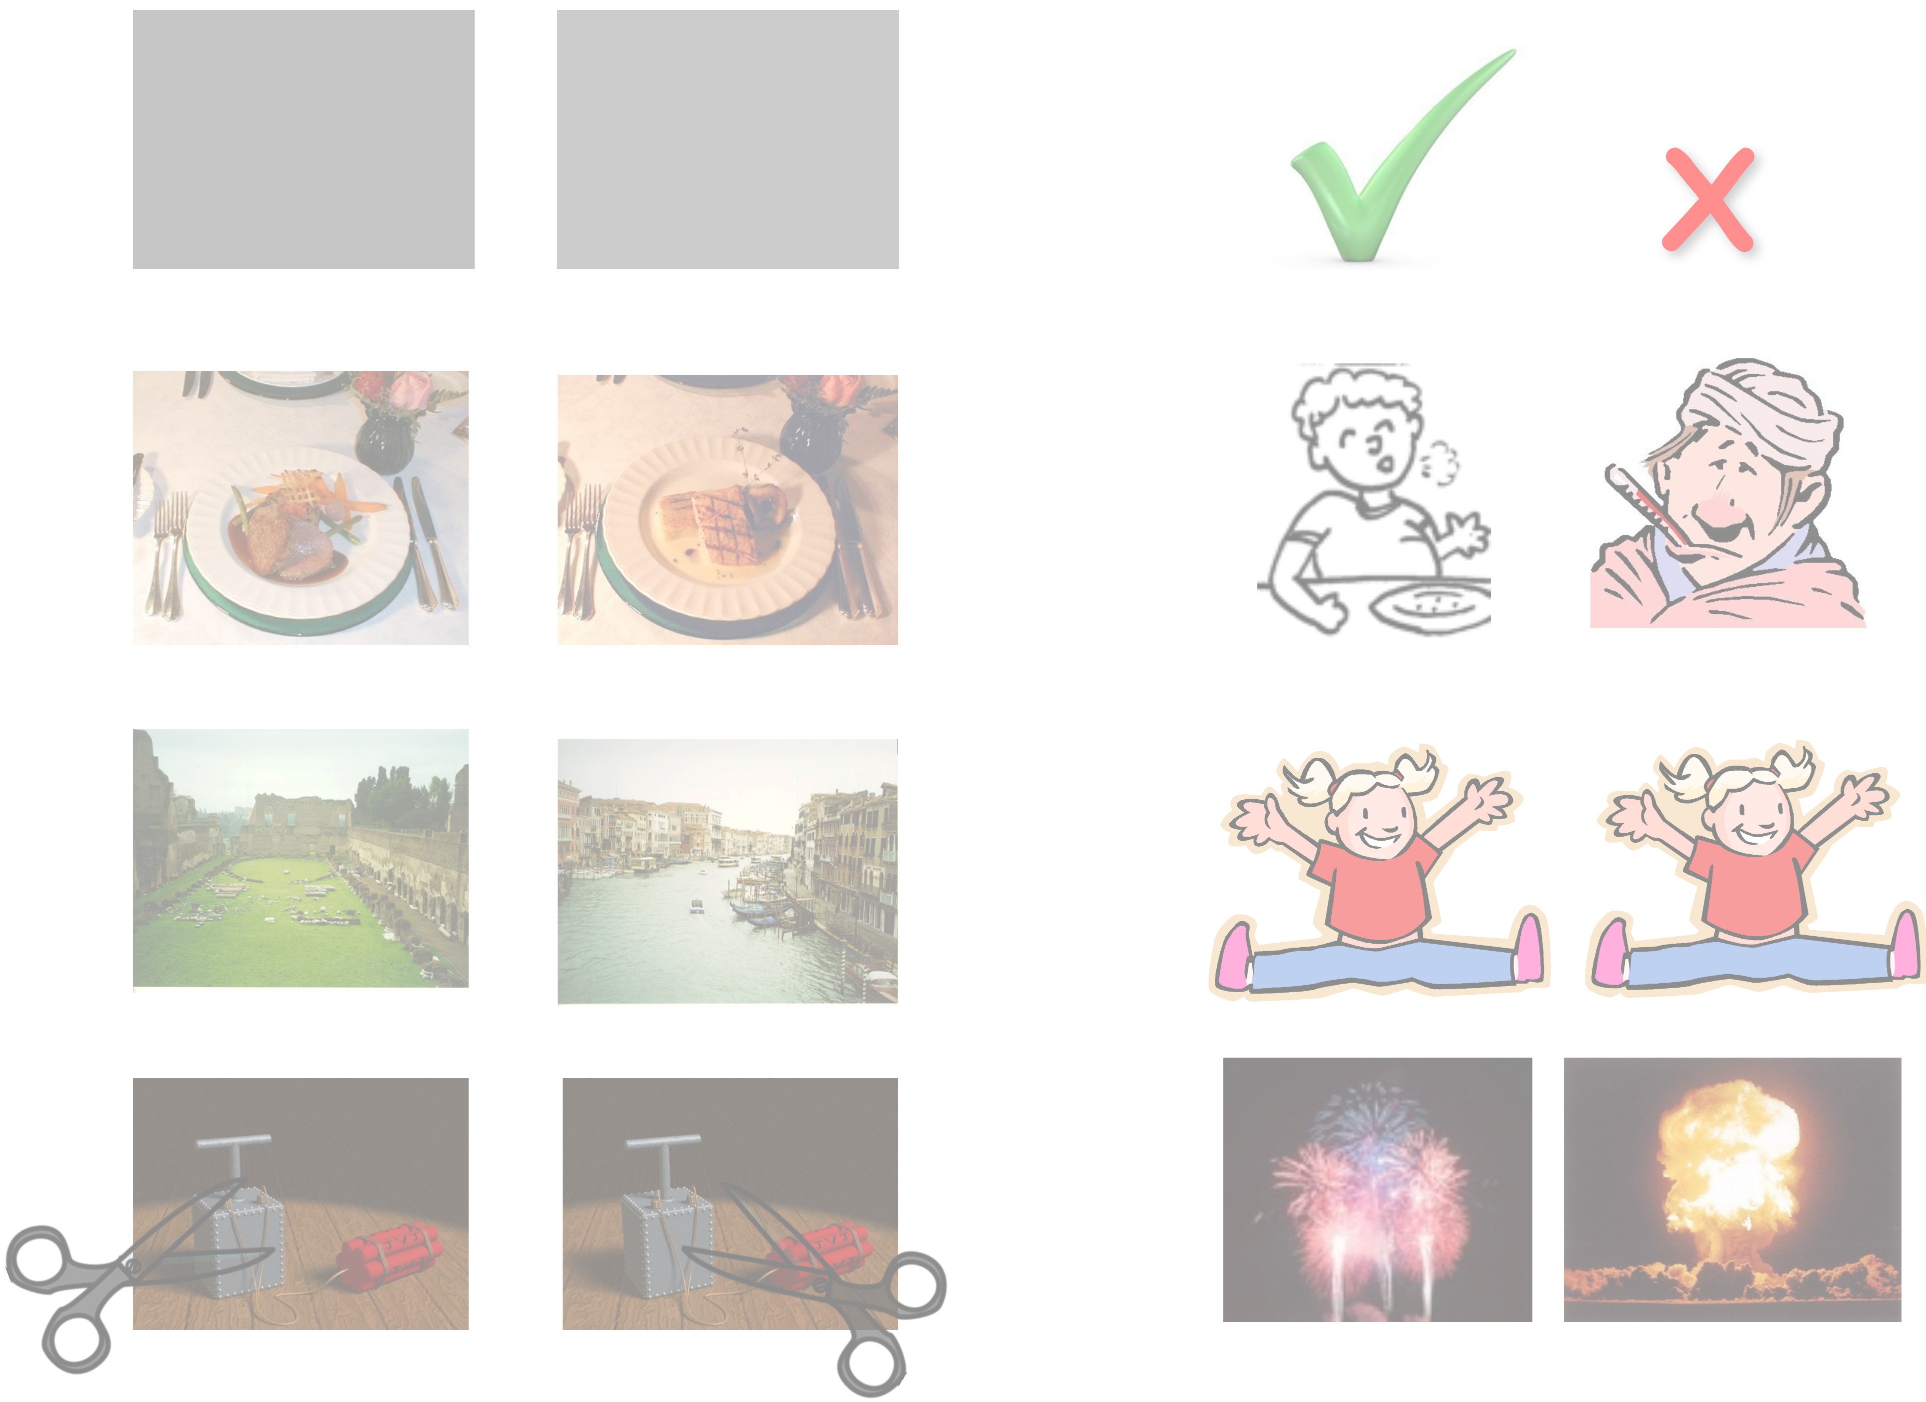

+1	-1
+5	-50
+100	+101
+10	-10000

Decision problems vary

- This is a perceptual problem
- It uses only one “kind” of information (brightness)
- The problem is familiar to people
- The utilities are symmetric (+1 & -1)
- The problem isn’t important (utilities near 0)

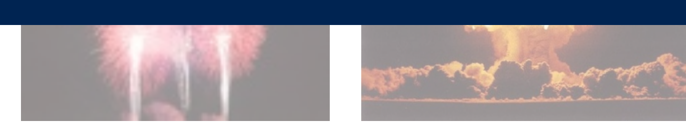

Decision problems vary

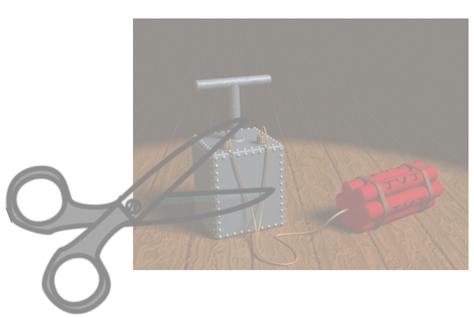

- This is both perceptual and cognitive
- Many "kinds" of information
- The problem is familiar to people
- The utilities are asymmetric (+5 & -50)
- The problem is not usually important

Decision problems vary

- Mostly cognitive, though some perception
- Many “kinds” of information
- Not a common problem, but moderately familiar
- The utilities are symmetric (+100 & +101)
- The problem is important (is the decision?)

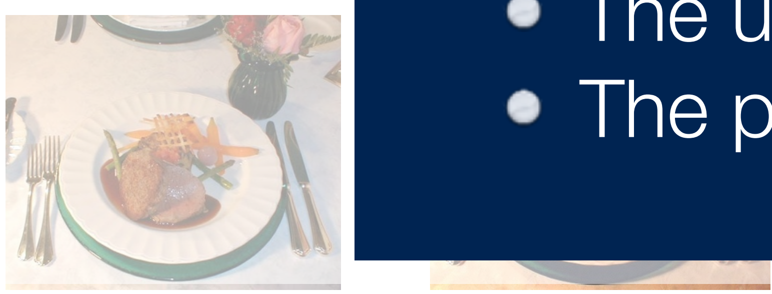

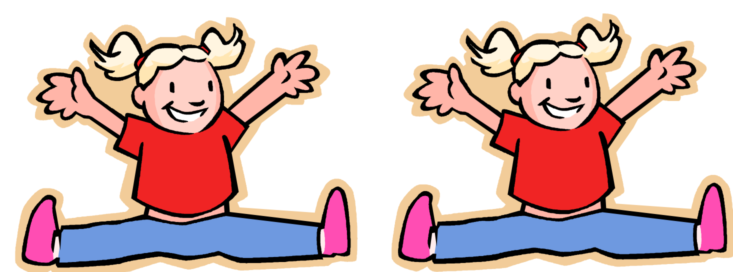

+100 +101

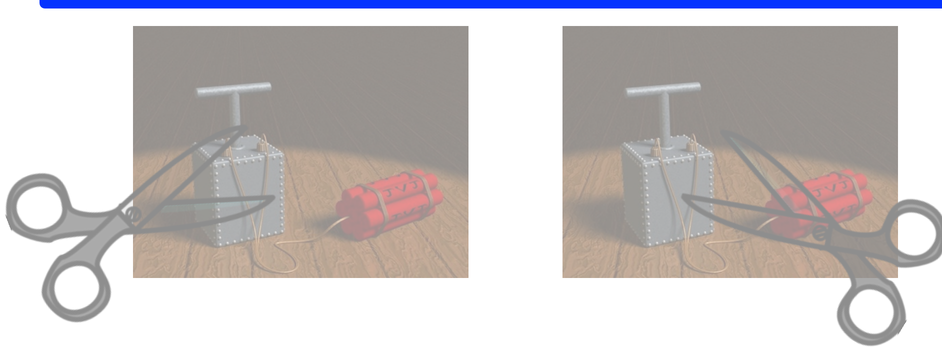

+10 -10000

Decision problems vary

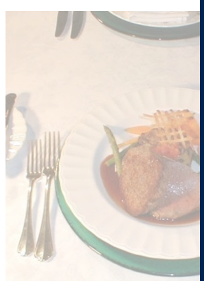

- Mostly cognitive, though some perception
- Many “kinds” of information
- Very unfamiliar to almost everyone
- The utilities are very asymmetric
- The problem is very important

-1

-50

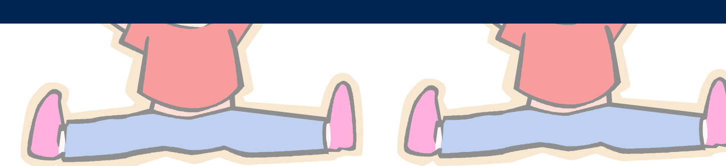

+101

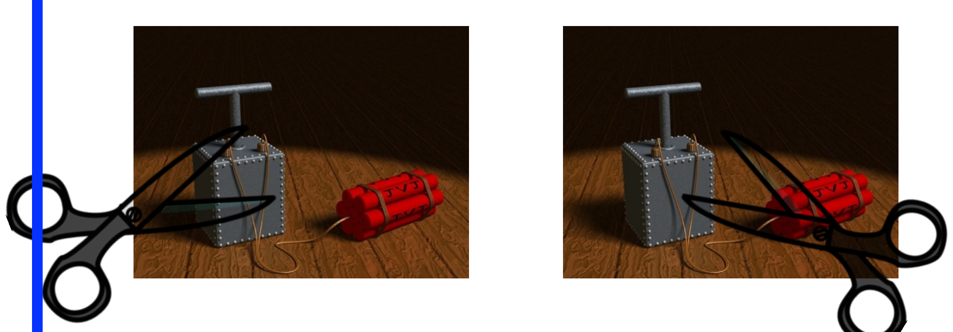

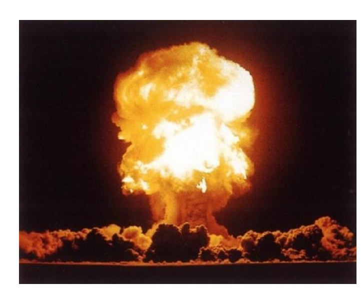

+10 -10000

Optimal decision making when the agent has full information

Choose the maximum utility option

When all the information about the outcomes is known, a rational actor should select the action that produces the most utility

$$a^* = \arg\max_{a_i} u(a_i)$$

The best possible action is chosen

u() is a function assigning utilities to actions

Soft rationality

$$P(a_i) = f(u(a_i))$$

When all the information about the outcomes is known, a sensible actor should be **more likely** to select the action that produces the most utility

Probability of choosing a particular action

Monotonically increasing function f() of the utility of that action

Triage decisions

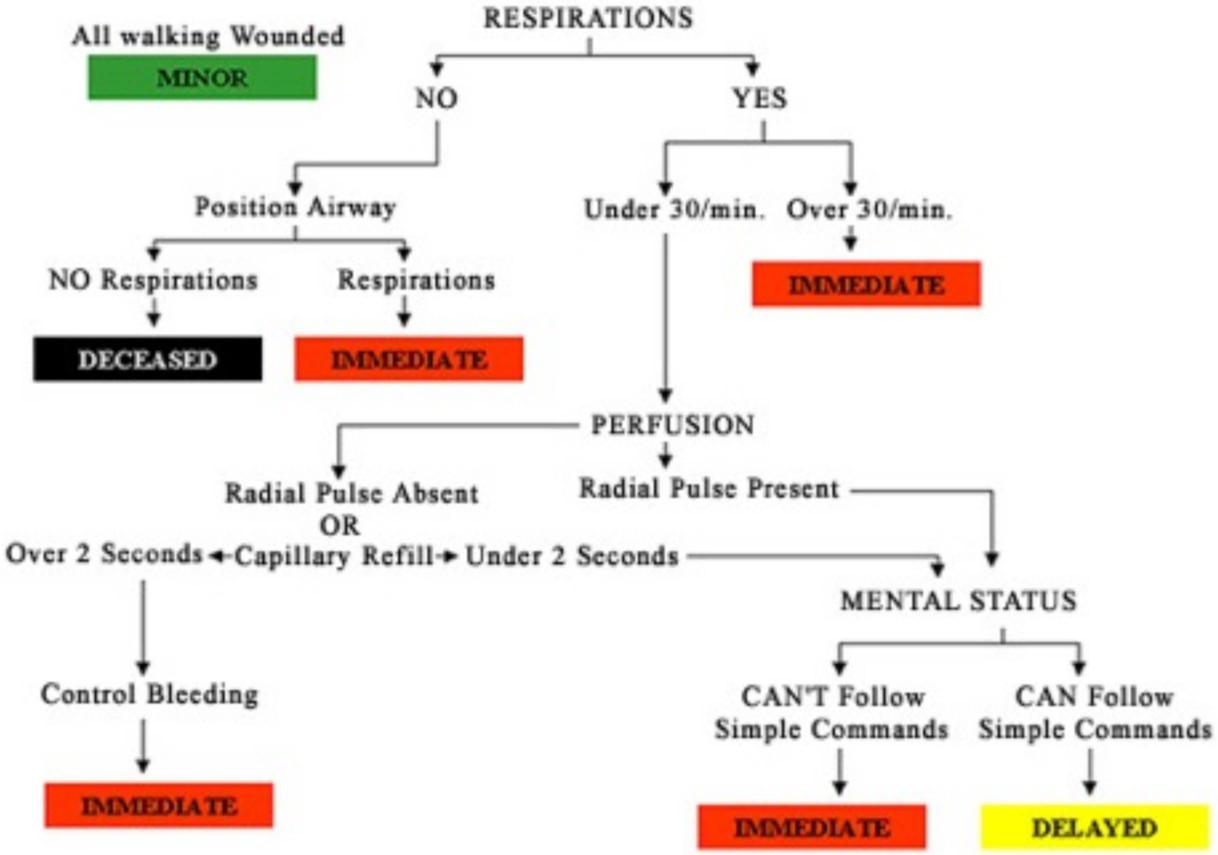

An unrealistic triage problem

- Triage:
  - Three critical patients arrive at the hospital at the same time.
  - Only one doctor is free to assist: will treat patients in order, passing over people if they have died, or if they die mid-treatment
- Information:
  - Patient A: Death in 5 minutes, takes 4 minutes to treat.
  - Patient B: Death in 11 minutes, takes 6 minutes to treat.
  - Patient C: Death in 35 minutes, takes 3 minutes to treat.
- Possible actions:
  - Set of 6 possible treatment orders to choose between.
  - ABC, ACB, BAC, BCA, CAB, CBA
  - Which is best?

Expected utility formulation

Each treatment order counts as a single action by the triage nurse

What happens?

Utility is a count of the number of survivors?

label	action	outcome	utility
a1	ABC	all live	+3
a2	ACB	B dies	+2
a3	BAC	A dies	+2
a4	BCA	A dies	+2
a5	CAB	A dies	+2
a6	CBA	A dies	+2

Optimal decision making in probabilistic environments

Real world triage only supplies partial information to the agent

You don't know exactly how long a patient has to live, nor do you know exactly how long it will take the doctor to stabilise the patient (or if that's even possible)

Maximise your expected utility

$$\mathrm{EU}(a) = \sum_z u(z)P(z|a)$$

Expected utility of option a

Utility of outcome z

Probability with which (you think) outcome z occurs given option a

The previous triage problem…

The **set of possible outcomes** is the set of survivor lists

	ABC	AB	AC	BC	A	B	C	-
ABC								
ACB								
BAC								
BCA								
CAB								
CBA								

The **action set** is the set of possible treatment orders

The previous triage problem…

	ABC	AB	AC	BC	A	B	C	-
ABC	**1**	0	0	0	0	0	0	0
ACB	0	0	**1**	0	0	0	0	0
BAC	0	0	0	**1**	0	0	0	0
BCA	0	0	0	**1**	0	0	0	0
CAB	0	0	0	**1**	0	0	0	0
CBA	0	0	0	**1**	0	0	0	0

Each possible action (row) is a probability distribution over possible outcomes (columns)

let's call this probability matrix P

The previous triage problem…

	ABC	AB	AC	BC	A	B	C	-
ABC	**1**	0	0	0	0	0	0	0
ACB	0	0	**1**	0	0	0	0	0
BAC	0	0	0	**1**	0	0	0	0
BCA	0	0	0	**1**	0	0	0	0
CAB	0	0	0	**1**	0	0	0	0
CBA	0	0	0	**1**	0	0	0	0

Refer to this matrix of probability distributions as **P**

The previous triage problem…

	+3	+2	+2	+2	+1	+1	+1	0
	ABC	**AB**	**AC**	**BC**	**A**	**B**	**C**	**-**
ABC	**1**	0	0	0	0	0	0	0
ACB	0	0	**1**	0	0	0	0	0
BAC	0	0	0	**1**	0	0	0	0
BCA	0	0	0	**1**	0	0	0	0
CAB	0	0	0	**1**	0	0	0	0
CBA	0	0	0	**1**	0	0	0	0

We have a vector that assigns utilities to outcomes **u**

The previous triage problem…

	+3	+2	+2	+2	+1	+1	+1	0	
	ABC	**AB**	**AC**	**BC**	**A**	**B**	**C**	**-**	
ABC	**1**	0	0	0	0	0	0	0	+3
ACB	0	0	**1**	0	0	0	0	0	+2
BAC	0	0	0	**1**	0	0	0	0	+2
BCA	0	0	0	**1**	0	0	0	0	+2
CAB	0	0	0	**1**	0	0	0	0	+2
CBA	0	0	0	**1**	0	0	0	0	+2

vector of expected utilities **e** is calculated by multiplying probability by utility and summing

… has a simple matrix formulation

$$\mathbf{u}' = \left[\begin{array}{cccccccc} 3 & 2 & 2 & 2 & 1 & 1 & 1 & 0 \end{array}\right]$$

$$\mathbf{P} = \left[\begin{array}{cccccccc} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{array}\right] \qquad \mathbf{e} = \left[\begin{array}{c} 3 \\ 2 \\ 2 \\ 2 \\ 2 \\ 2 \end{array}\right]$$

If **P** is the matrix of conditional probabilities, **u** is a column vector of outcome utilities, and **e** is a column vector of expected utilities of actions, then we're just doing the matrix operation **e** = **P u**

$$\mathbf{e} = \mathbf{Pu}$$

$$= \begin{bmatrix} .5 & .5 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & .25 & .25 & .25 & .25 & 0 & 0 & 0 \\ 0 & 0 & .33 & .34 & .33 & 0 & 0 & 0 \\ 0 & 0 & 0 & .5 & .5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ .125 & .125 & .125 & .125 & .125 & .125 & .125 & .125 \end{bmatrix} \times \begin{bmatrix} 3 \\ 2 \\ 2 \\ 2 \\ 1 \\ 1 \\ 1 \\ 0 \end{bmatrix}$$

$$= \begin{bmatrix} 2.50 \\ 1.75 \\ 1.67 \\ 1.50 \\ 2.00 \\ 1.50 \end{bmatrix}$$

Expected utility calculations are straightforward linear algebra even in the probabilistic case.

R code: `e <- P %*% u`

Computing expected utilities isn't easy: A slightly more difficult triage problem

Probabilistic triage problems

- Keep it simple:
  - 3 patients
  - 1 doctor

- Robot doctor:
  - Treats patients in order
  - If a patient dies before or during treatment, moves immediately to the next one

How long do the patients have to live?

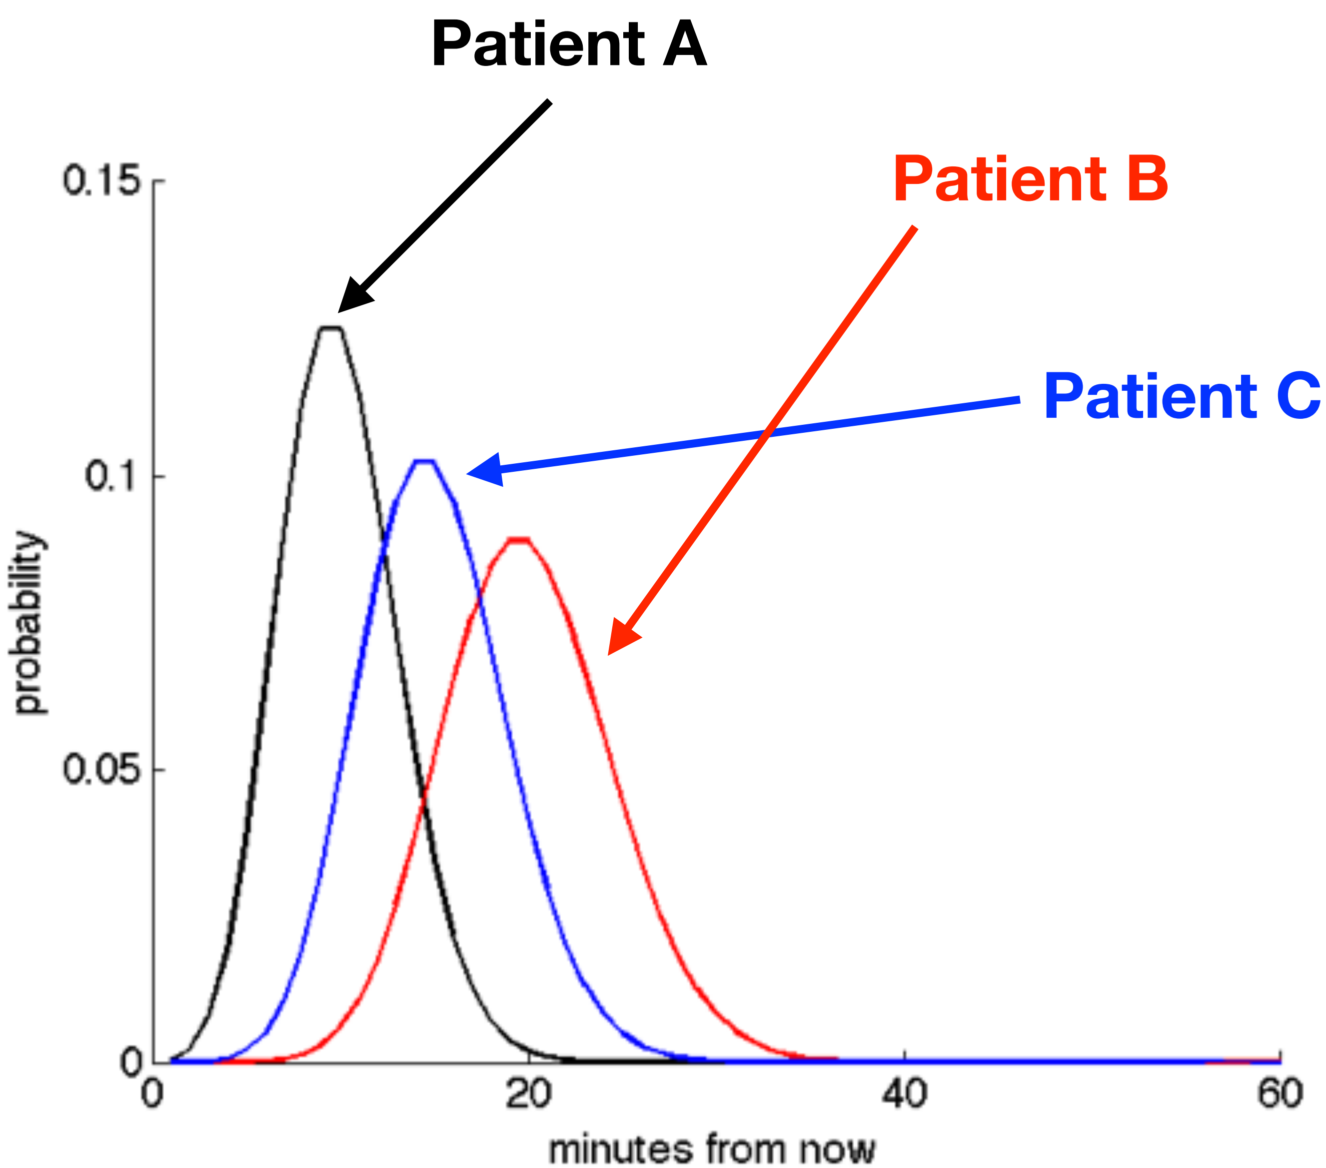

Nurse beliefs about patient death probabilities… the probability that the patient dies on exactly the t-th minute is Poisson(λ):

$$P(t|\lambda) = \frac{\lambda^t \exp(-\lambda)}{t!}$$

Patient A: λ = 10
Patient B: λ = 20
Patient C: λ = 15

How long does treatment take?

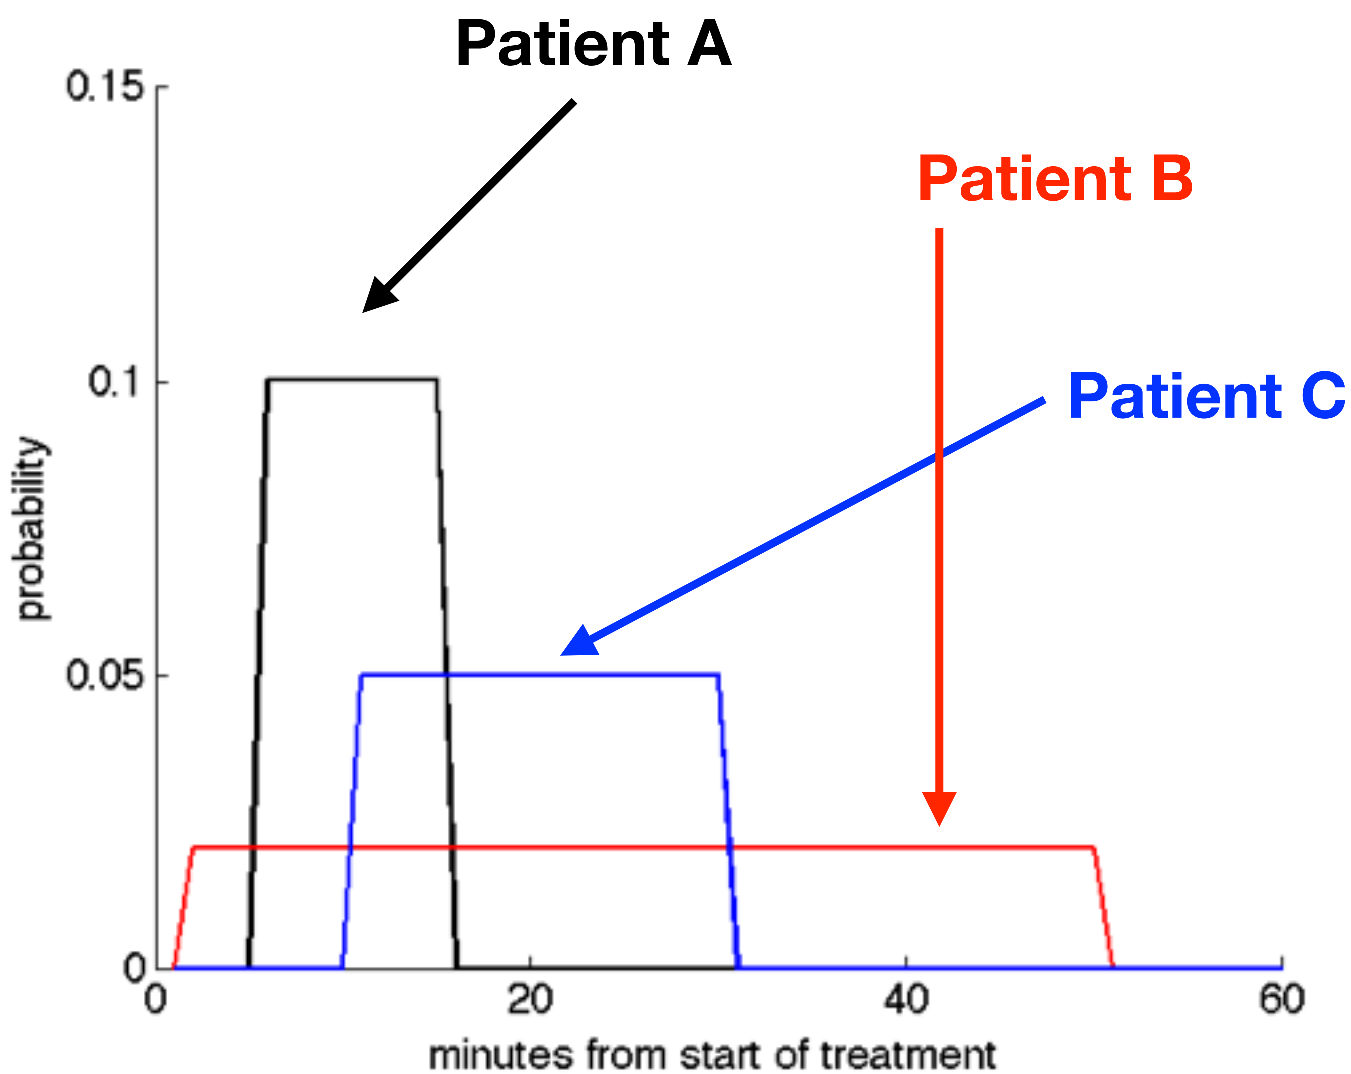

Amount of time taken to save the patient follows a uniform distribution:

Patient A: 6-15 minutes

Patient B: 2-50 minutes

Patient C: 11-30 minutes

(Although the patient might die while you're treating them!)

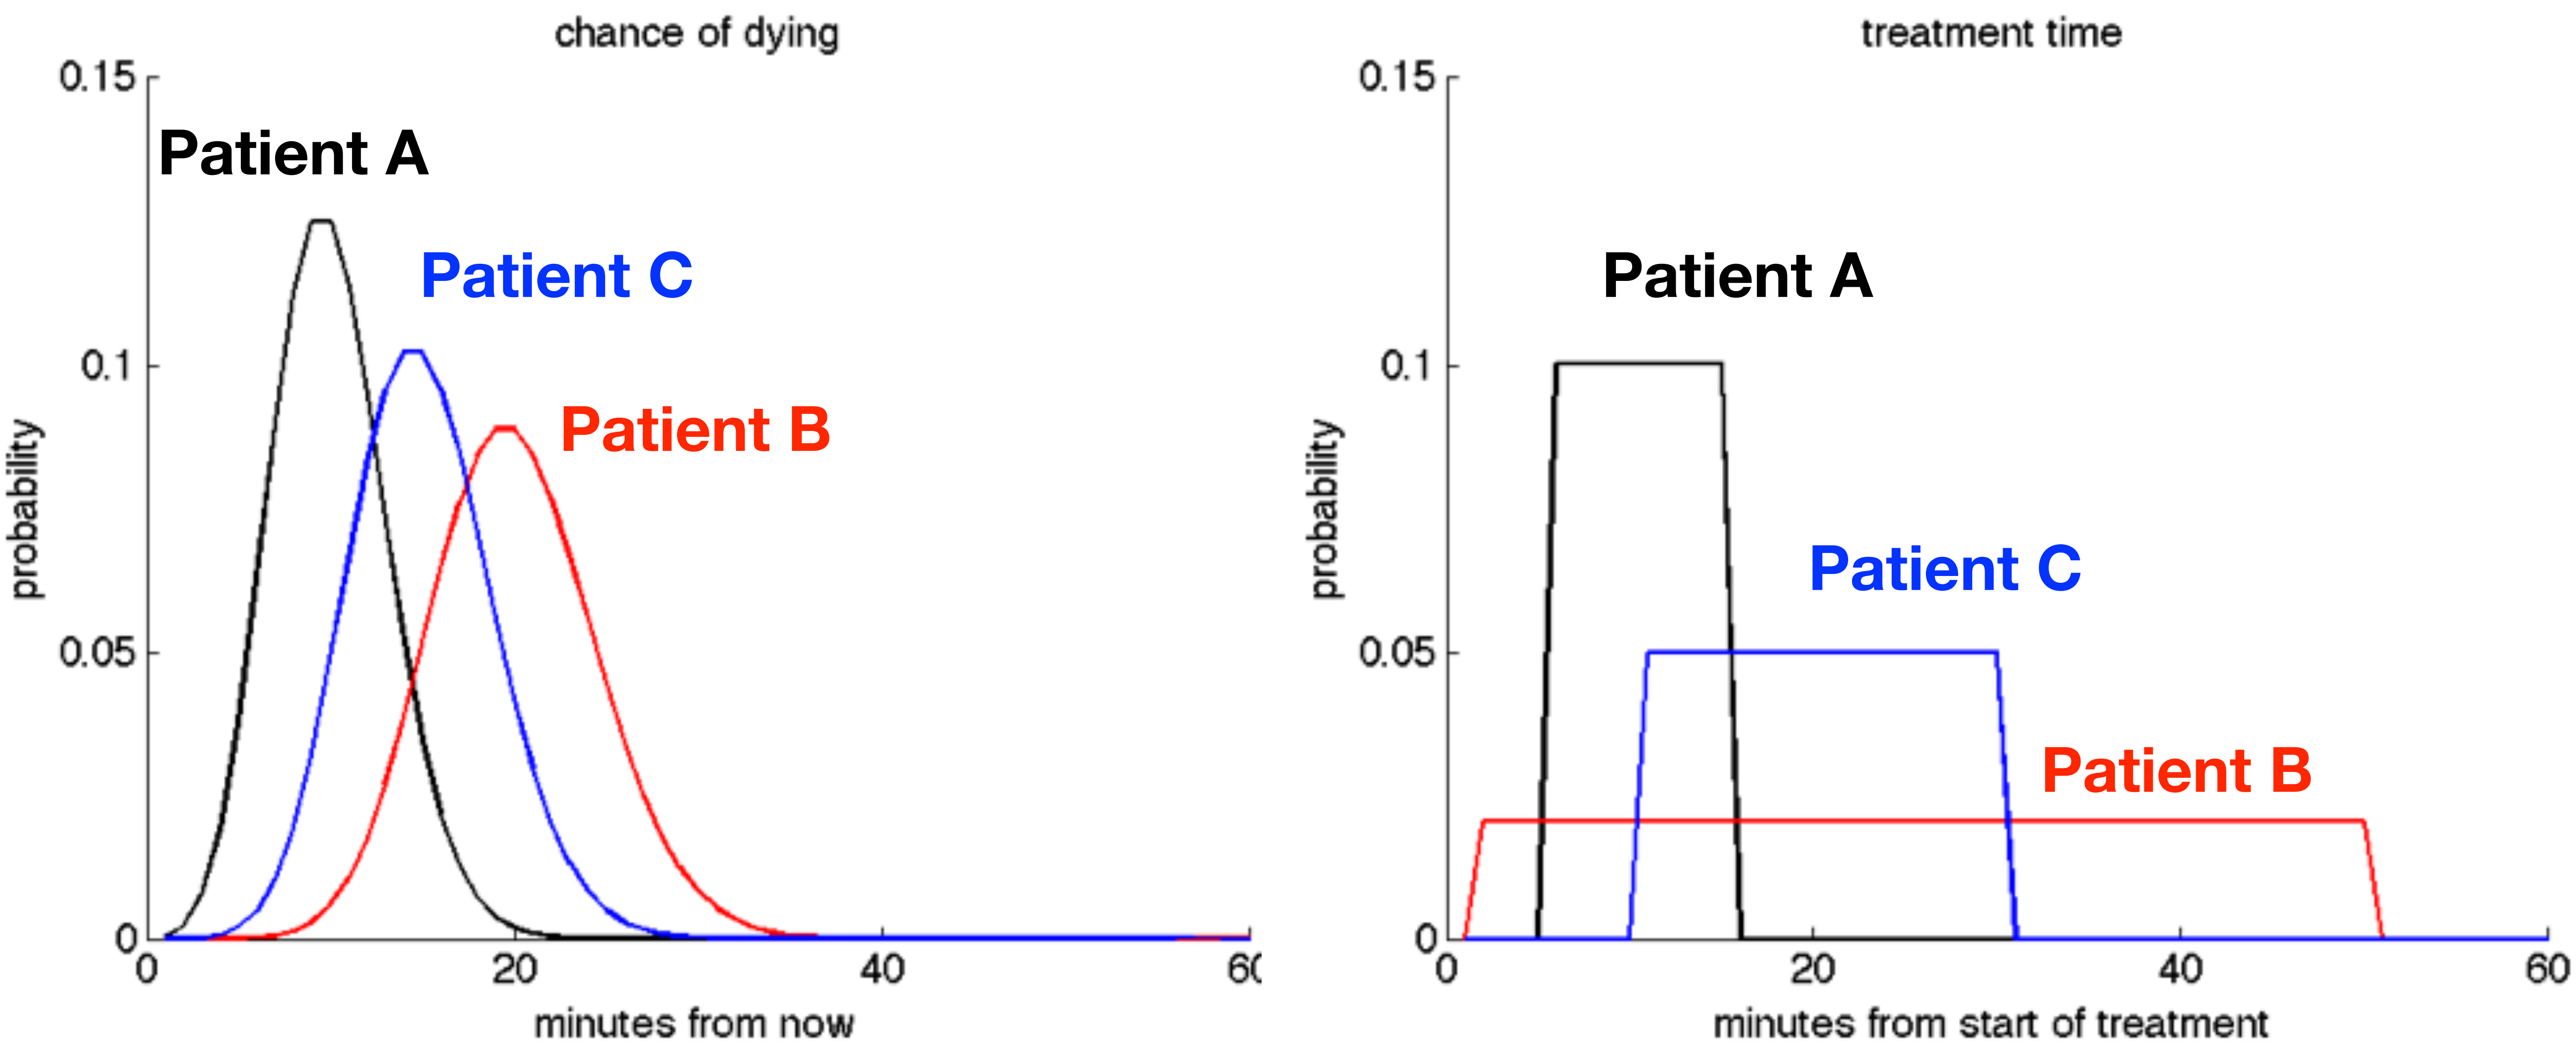

It's your first shift on the ER.
You have 20 seconds to make this decision.

Suppose it turned out like this…

- Treatment times:
  - Patient A takes **10** minutes to treat
  - Patient B takes **30** minutes to treat
  - Patient C takes **13** minutes to treat
- Death times:
  - Patient A dies in the **12th** minute
  - Patient B dies in the **20th** minute
  - Patient C dies in the **15th** minute
- What happens?

And suppose you picked ACB

- Treatment times:
  - Patient A takes **10** minutes to treat
  - Patient B takes **30** minutes to treat
  - Patient C takes **13** minutes to treat
- Death times:
  - Patient A dies in the **12th** minute
  - Patient B dies in the **20th** minute
  - Patient C dies in the **15th** minute
- What happens?

- Only patient A is saved.
  - Starts on A at minute 1.
  - A saved on minute 10.
  - Starts on C at minute 11.
  - C dies on minute 15.
  - Starts on B at minute 16.
  - B dies on minute 20.

What about the other actions?

- Treatment times:
  - Patient A takes **10** minutes to treat
  - Patient B takes **30** minutes to treat
  - Patient C takes **13** minutes to treat
- Death times:
  - Patient A dies in the **12th** minute
  - Patient B dies in the **20th** minute
  - Patient C dies in the **15th** minute
- What happens?

- More generally…
  - If you start with patient B, all three people die.
  - If you start with patient A, only A survives
  - If you start with patient C, only C survives

But you **<u>don't</u>** know when people will die or how long the treatment takes!

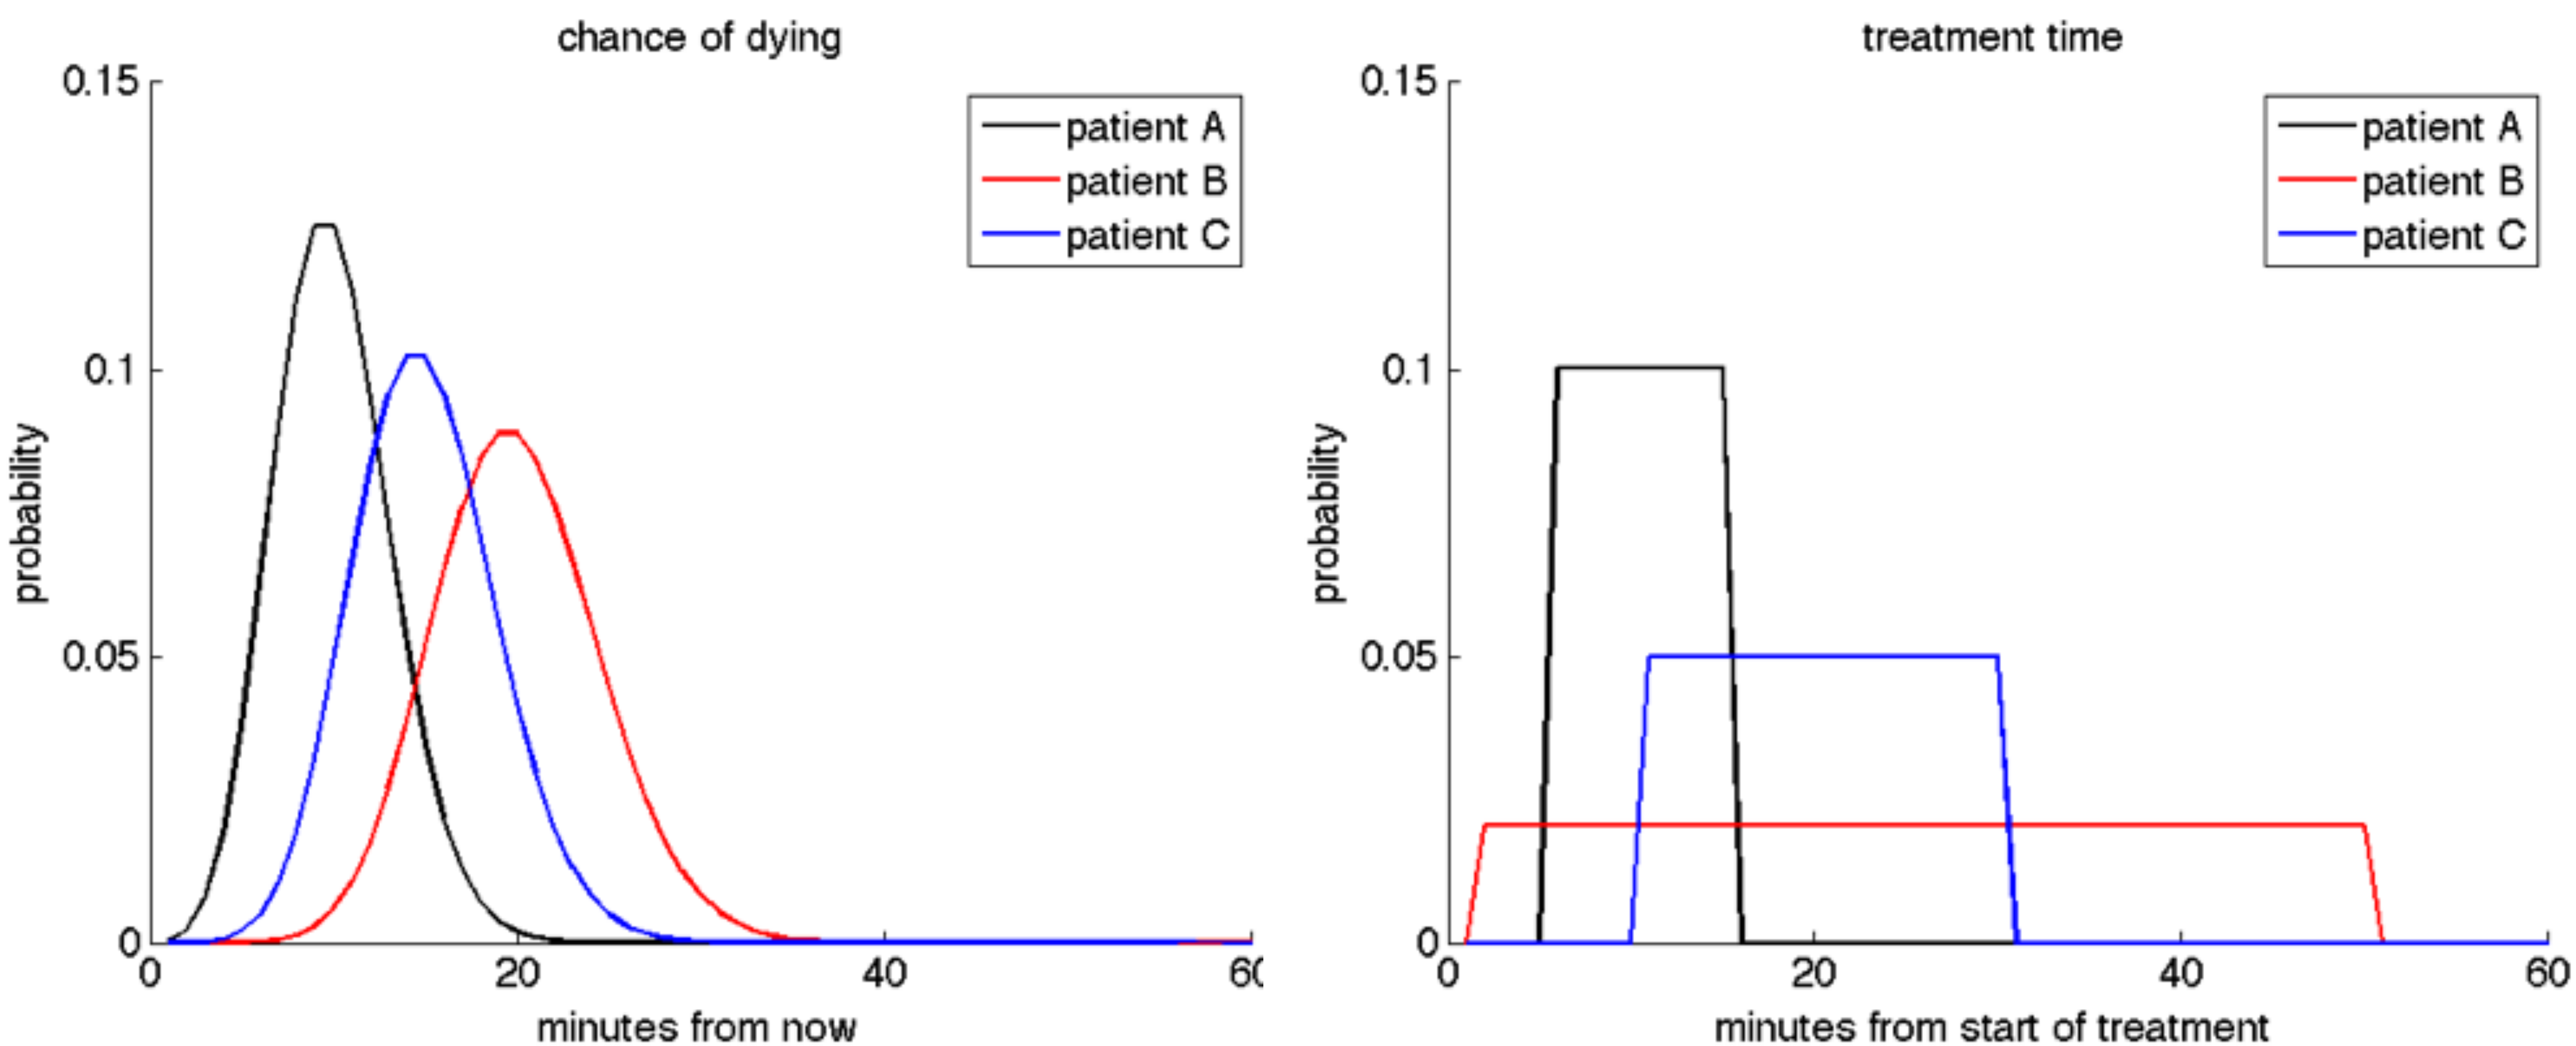

The answer

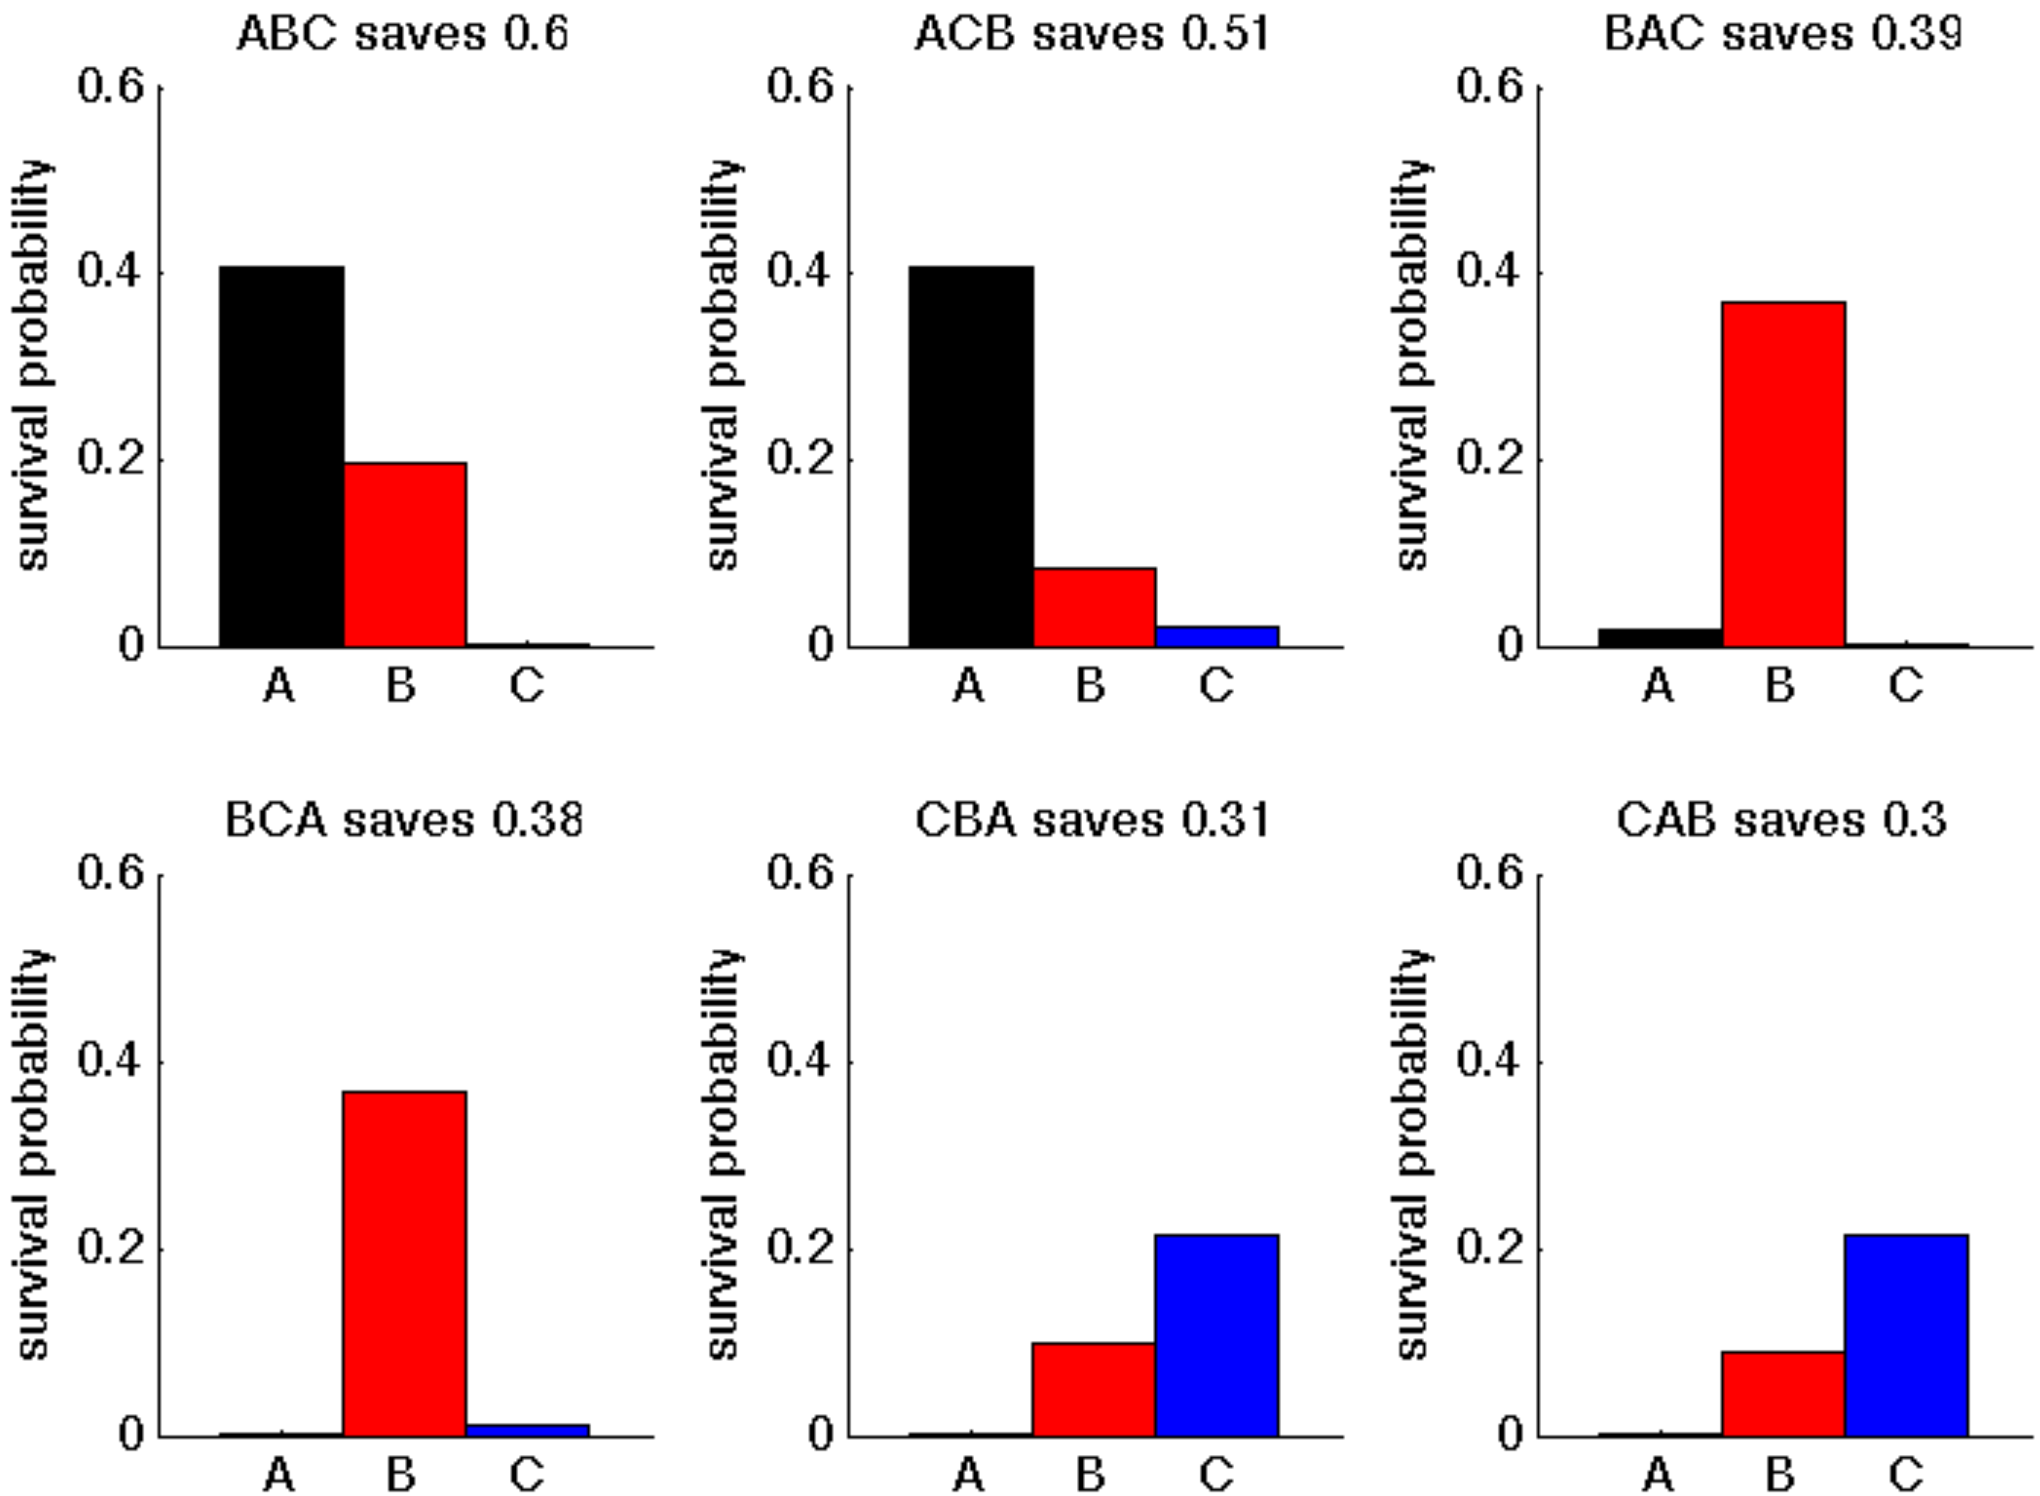
ABC saves 0.6
ACB saves 0.51
BAC saves 0.39
BCA saves 0.38
CBA saves 0.31
CAB saves 0.3
survival probability
0
0.2
0.4
0.6
A
B
C

ABC maximises expected utility

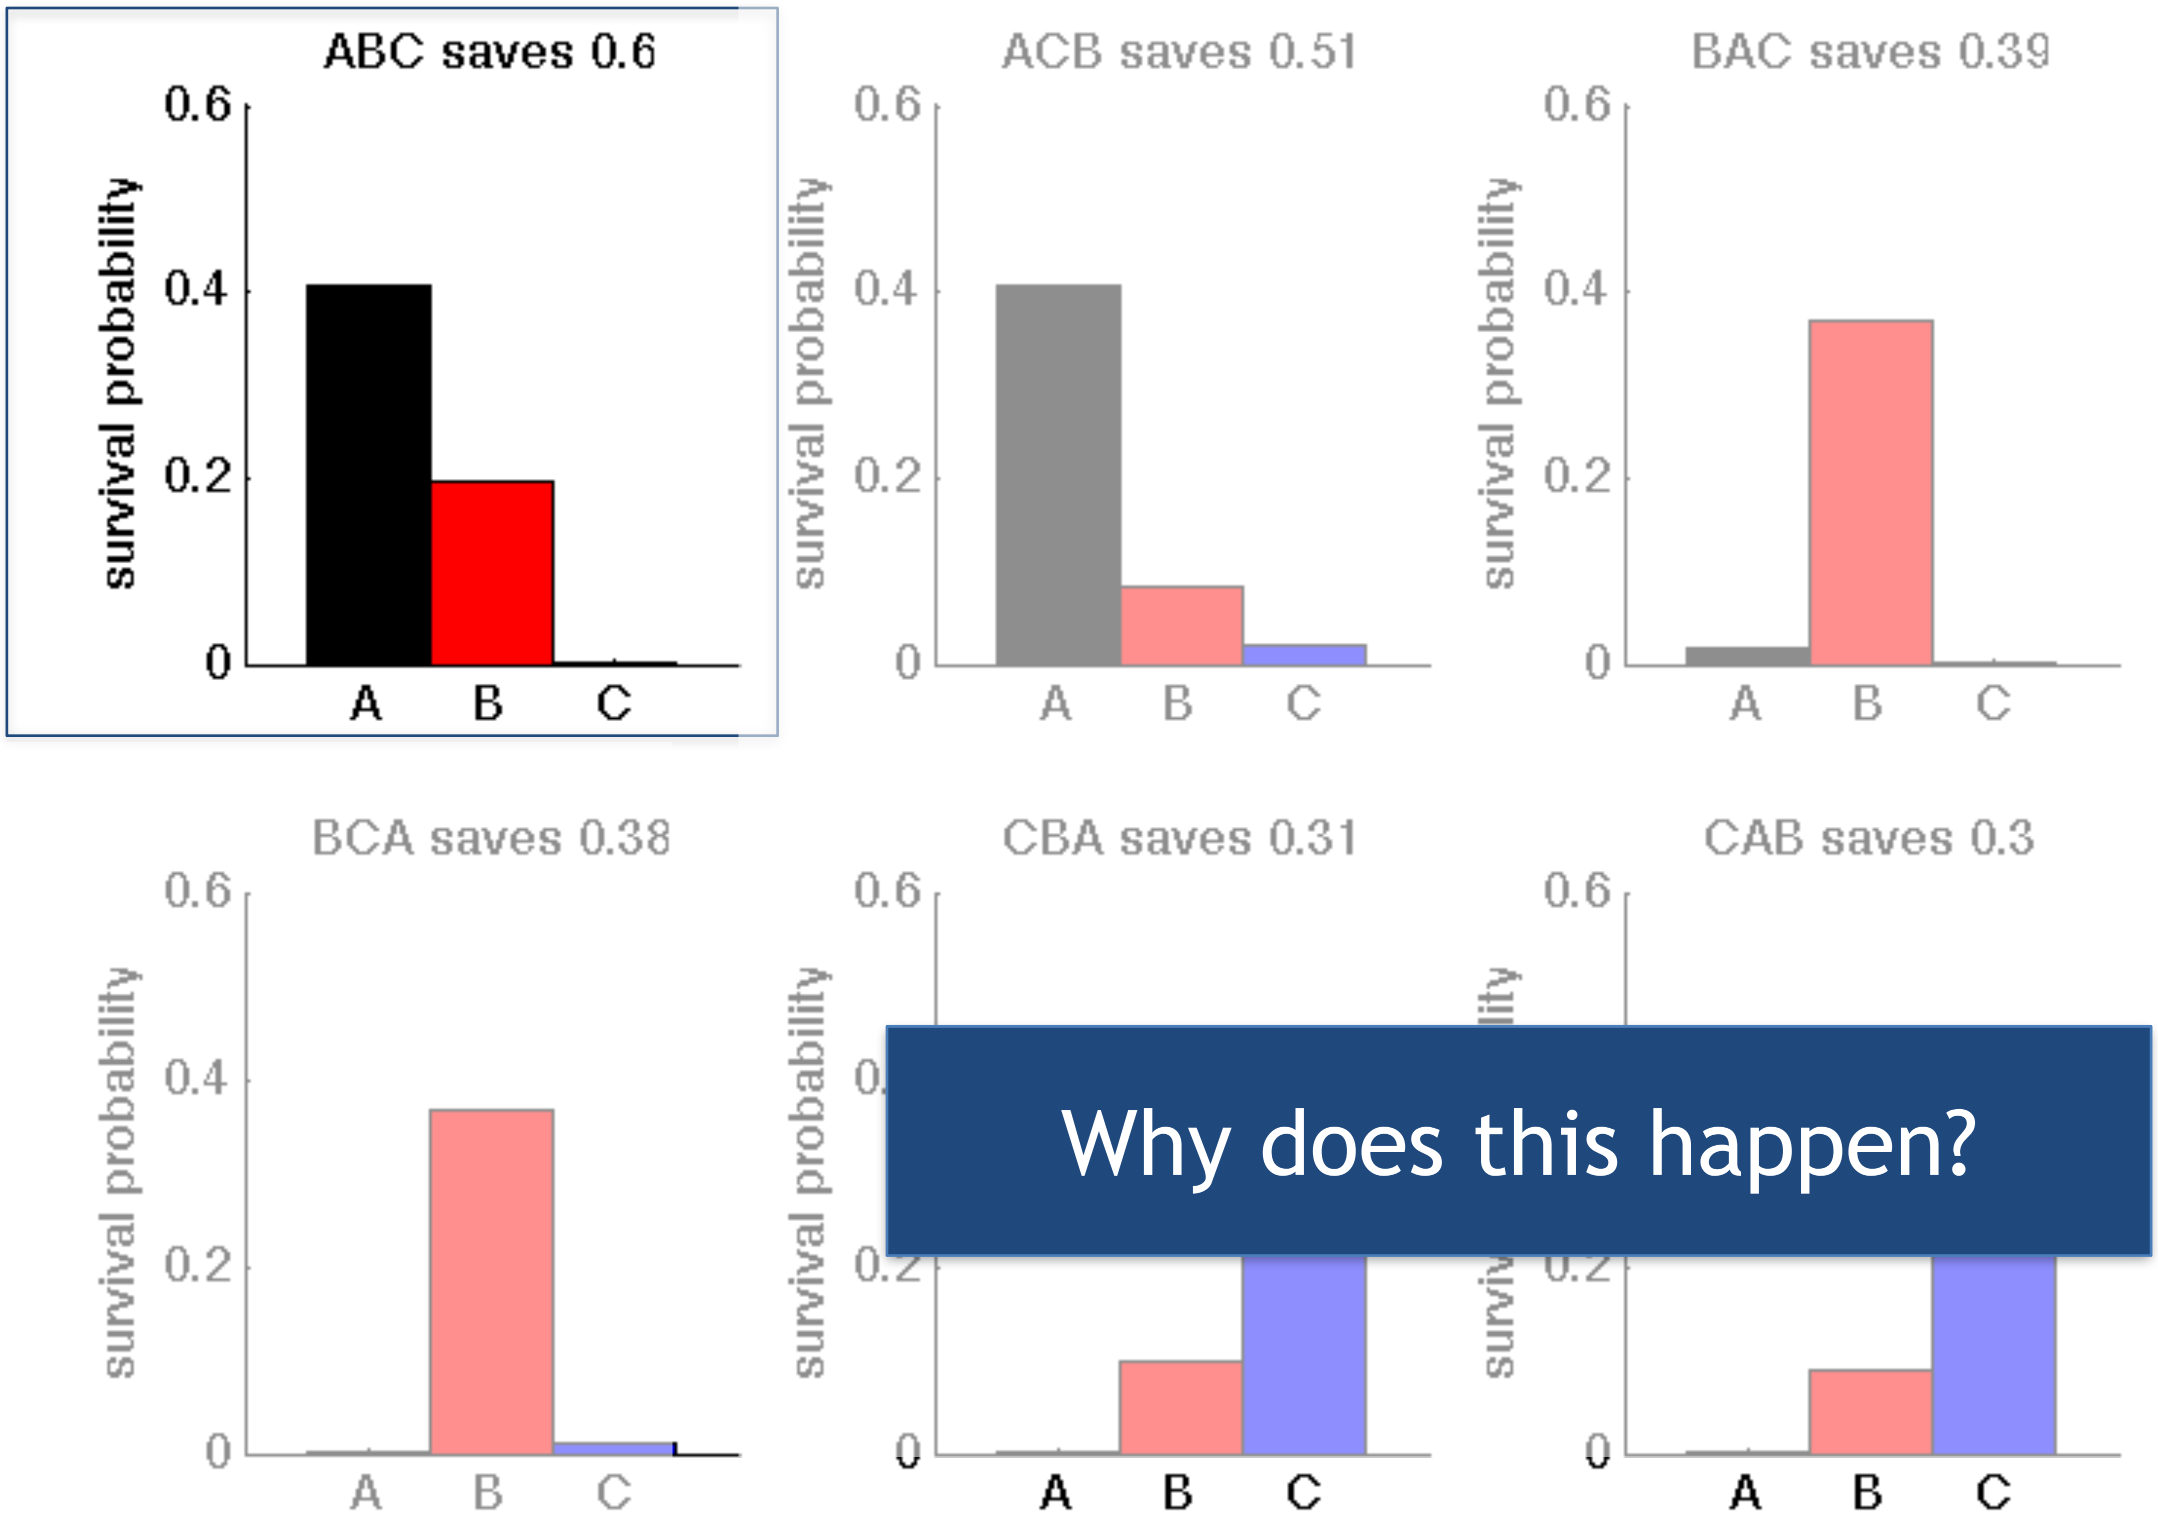

Why does this happen?

CALCULATE SURVIVAL PROBABILITY P_{XO}, FOR ALL PATIENTS X AND ALL TREATMENT ORDERS O:

- for M iterations
  - for $X = [A,B,C]$
    - generate a random **time of death** d_X for patient X
    - generate a random **length of treatment** t_X for patient X
  - for all possible treatment orders, O
    - simulate **doctor's behaviour**, and determine which patients survive
    - for $X = [A,B,C]$
      - if patient X survives, increment count: $N_{OX} = N_{OX}+1$
- for $X = [A,B,C]$
  - for all possible treatment orders, O
    - $P_{XO}=N_{XO} / M$

GENERATE TIME OF DEATH, $d \sim \text{Poisson}(\lambda)$

set $L=\exp(-\lambda)$; $d=0$; $p=1$;

do while $p > L$

  $d = d+1;$

  generate $u \sim \text{Uniform}([0,1])$

  $p = p^{*}u;$

$d = d\text{-}1;$

GENERATE LENGTH OF TREATMENT, t

I'm assuming that this one is obvious… randomly select t from the set of treatment times (e.g., 6,7,8,…15 for patient A)

SIMULATE THE DOCTOR'S BEHAVIOUR

input: treatment order o, time of death d, time to treat t

set minutes elapsed $m = 0$

do until all patients treated or dead

look up next patient, denoted X

if $t_X + m < d_X$

patient X lives

$m = m + t_X$

else

patient X dies

$m = \max(m, d_X)$

output: list of survivors

demonstration code: triage.R

Summary

- What did we calculate here?
  - Computation of P(z|a), the probability of an outcome given an action
  - e.g., P("A lives, BC die" | "order is ABC")
- The point:
  - Decision making depends on the correct evaluation of the probabilities
  - This is often hard to do.

Utilities are not straightforward

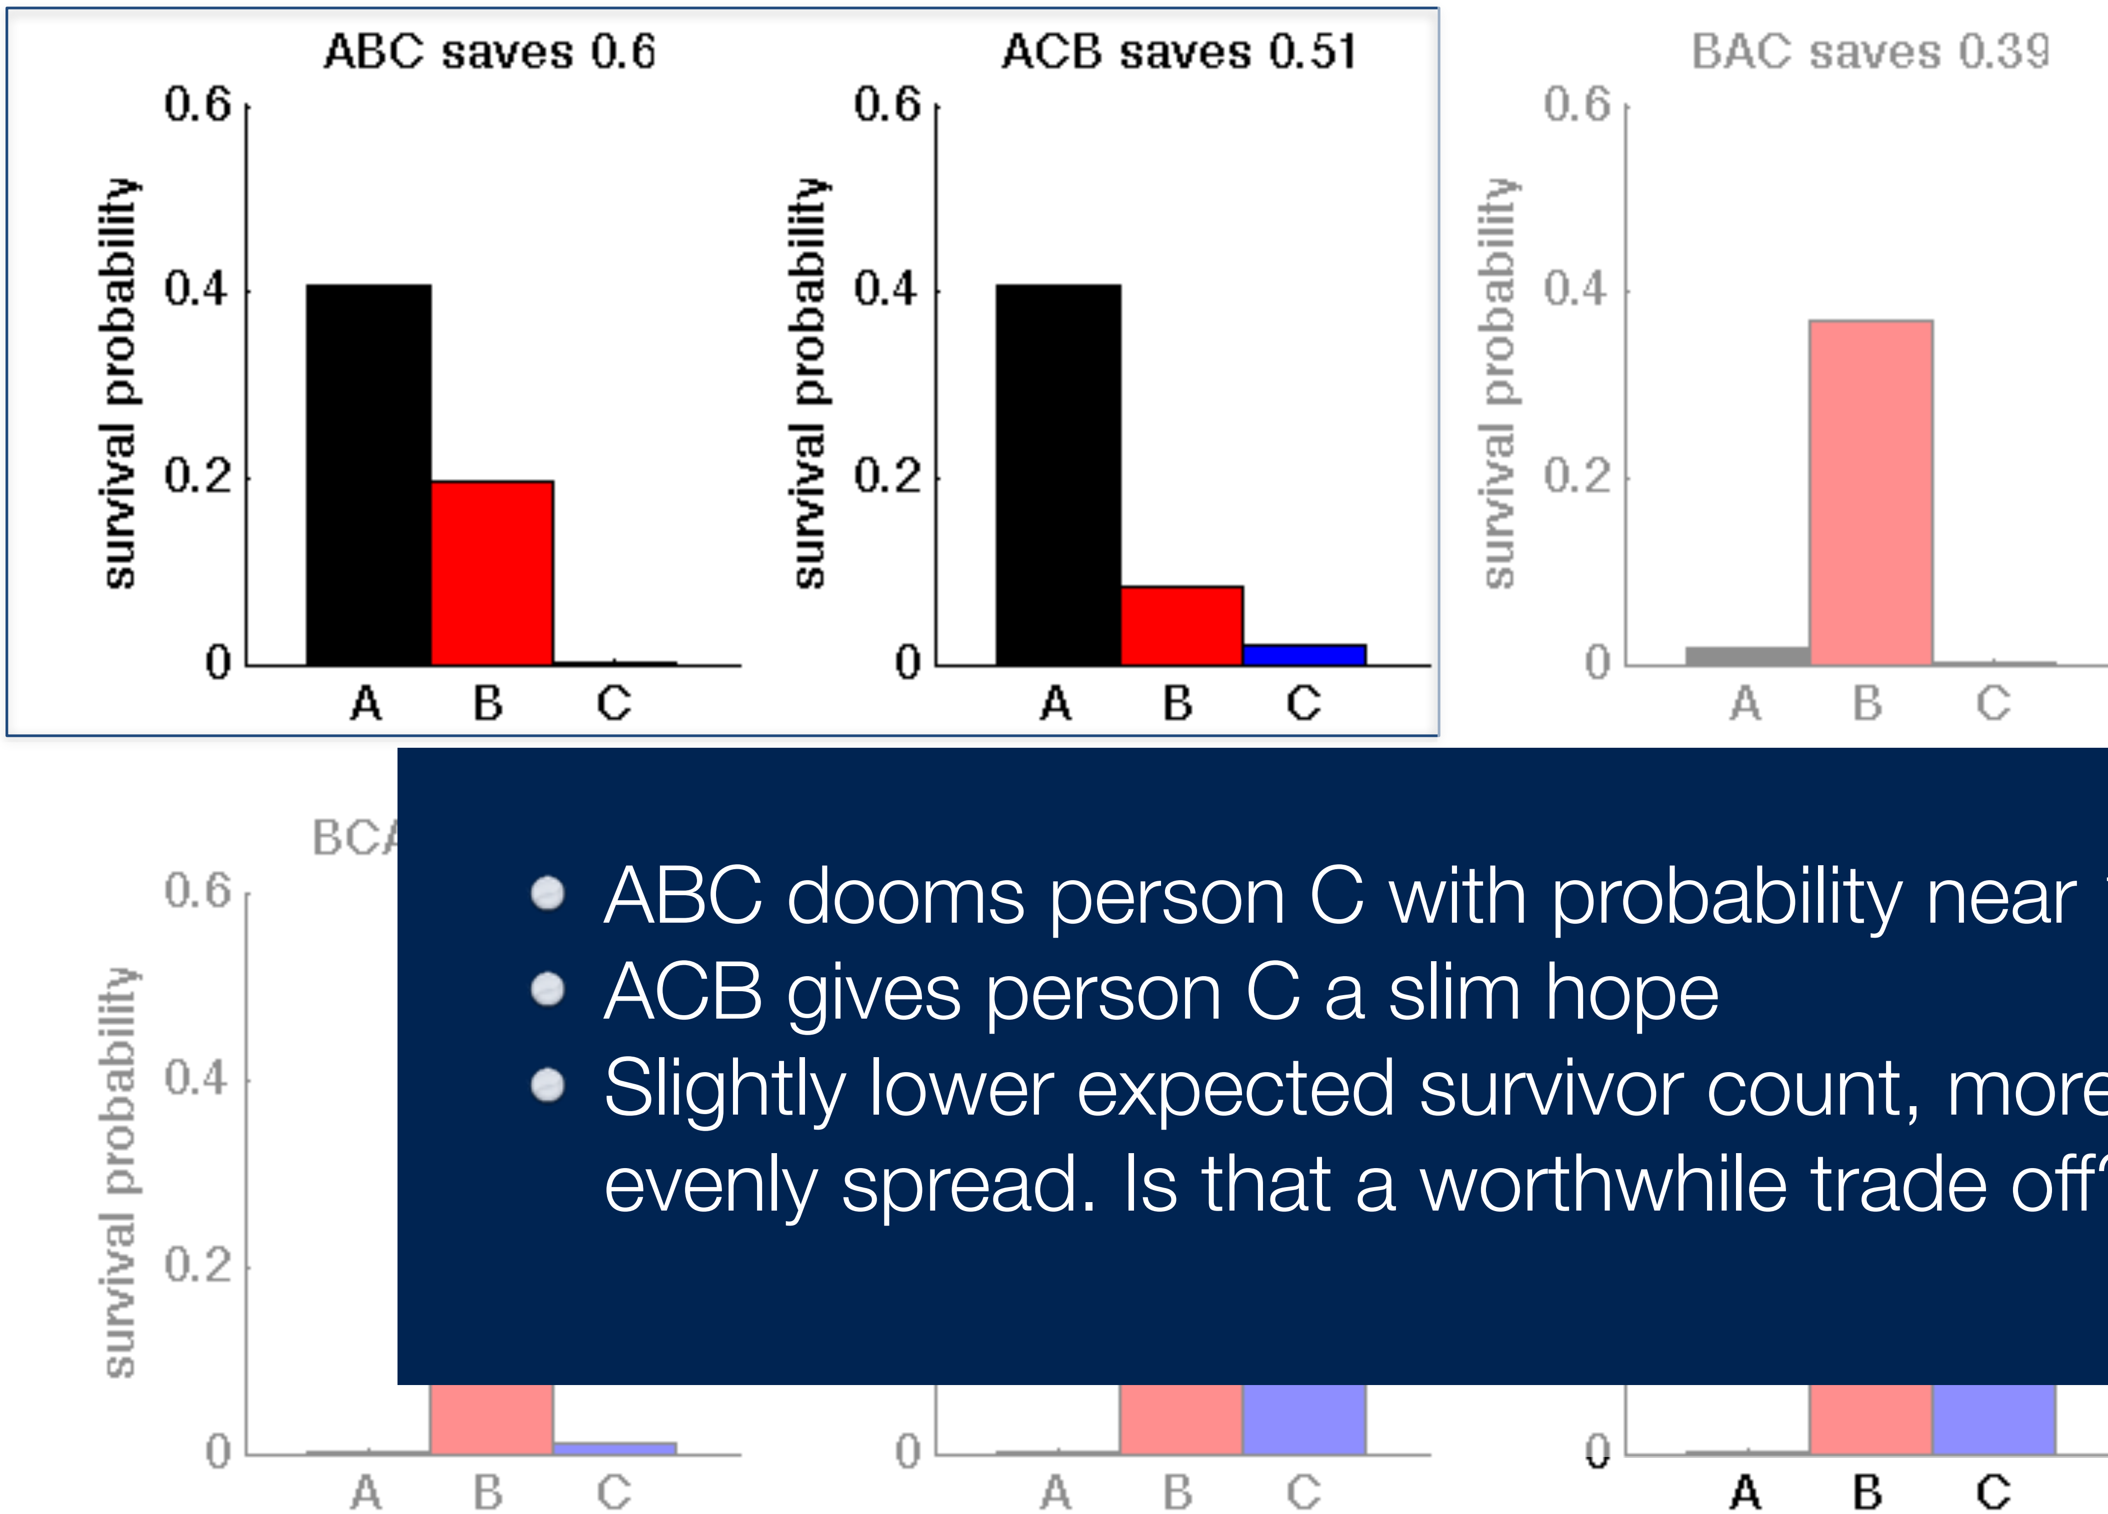

- ABC dooms person C with probability near 1
- ACB gives person C a slim hope
- Slightly lower expected survivor count, more evenly spread. Is that a worthwhile trade off??

It's not purely a computational problem: The St Petersburg paradox

Flipping coins until you get a tail

outcome:	**T**	**HT**	**HHT**	**HHHT**	etc
payout:	$1	$2	$4	$8	etc

What is a fair price to pay in exchange for the opportunity to play this game?

The game is infinitely valuable?

$$\begin{aligned} \mathrm{EU} &= \sum_{\mathrm{outcome}} u(\mathrm{outcome}) P(\mathrm{outcome}) \\ &= u(\mathtt{T})P(\mathtt{T}) + u(\mathtt{HT})P(\mathtt{HT}) + u(\mathtt{HHT})P(\mathtt{HHT}) + \ldots \\ &= \left(1 \times \frac{1}{2}\right) + \left(2 \times \frac{1}{4}\right) + \left(4 \times \frac{1}{8}\right) + \ldots \\ &= \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \ldots \\ &= \infty \end{aligned}$$

Most people are only willing to pay a few dollars. Are they being irrational?

Solution #1

- Money doesn't equal happiness
  - Utility doesn't scale linearly with dollar value
  - Specifically, it has **decreasing marginal utility**

- Daniel Bernoulli (1788):
  - "The determination of the value of an item must not be based on the price, but rather on the utility it yields…. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount."

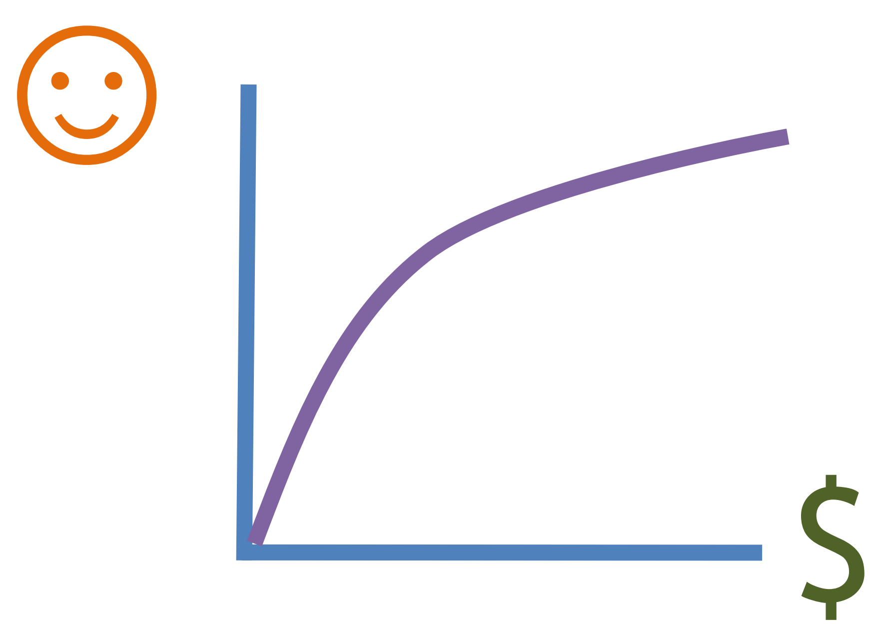

Solution #1

- It now depends on how much money you start with
  - Calculations are based on the utility of your bank balance
  - Compare the current utility to the expected utility after the game

- What happens now?
  - Say you start with a bank balance of *w* and the game costs *c* to play
  - Expected utility of bank balance scales logarithmically with $
  - So if you get *k* heads before the first tail the utility of the game is:

$$\ln(w - c + 2^k)$$

  - The utility of not playing is just the current utility of your bank balance:

$$\ln(w)$$

  - You should play if the the utility of playing is larger

Solution #1

- More generally:
  - Getting k heads before a tail has probability $2^{-(k+1)}$
  - Expected utility of the playing the game is computed by taking a probability-weighted average of all possible outcomes
  - So:

$$\text{EU}(\text{“play”}) = \sum_{k=0}^{\infty} \frac{\ln(w - c + 2^k)}{2^{k+1}}$$

$$\text{EU}(\text{“don't play”}) = \ln(w)$$

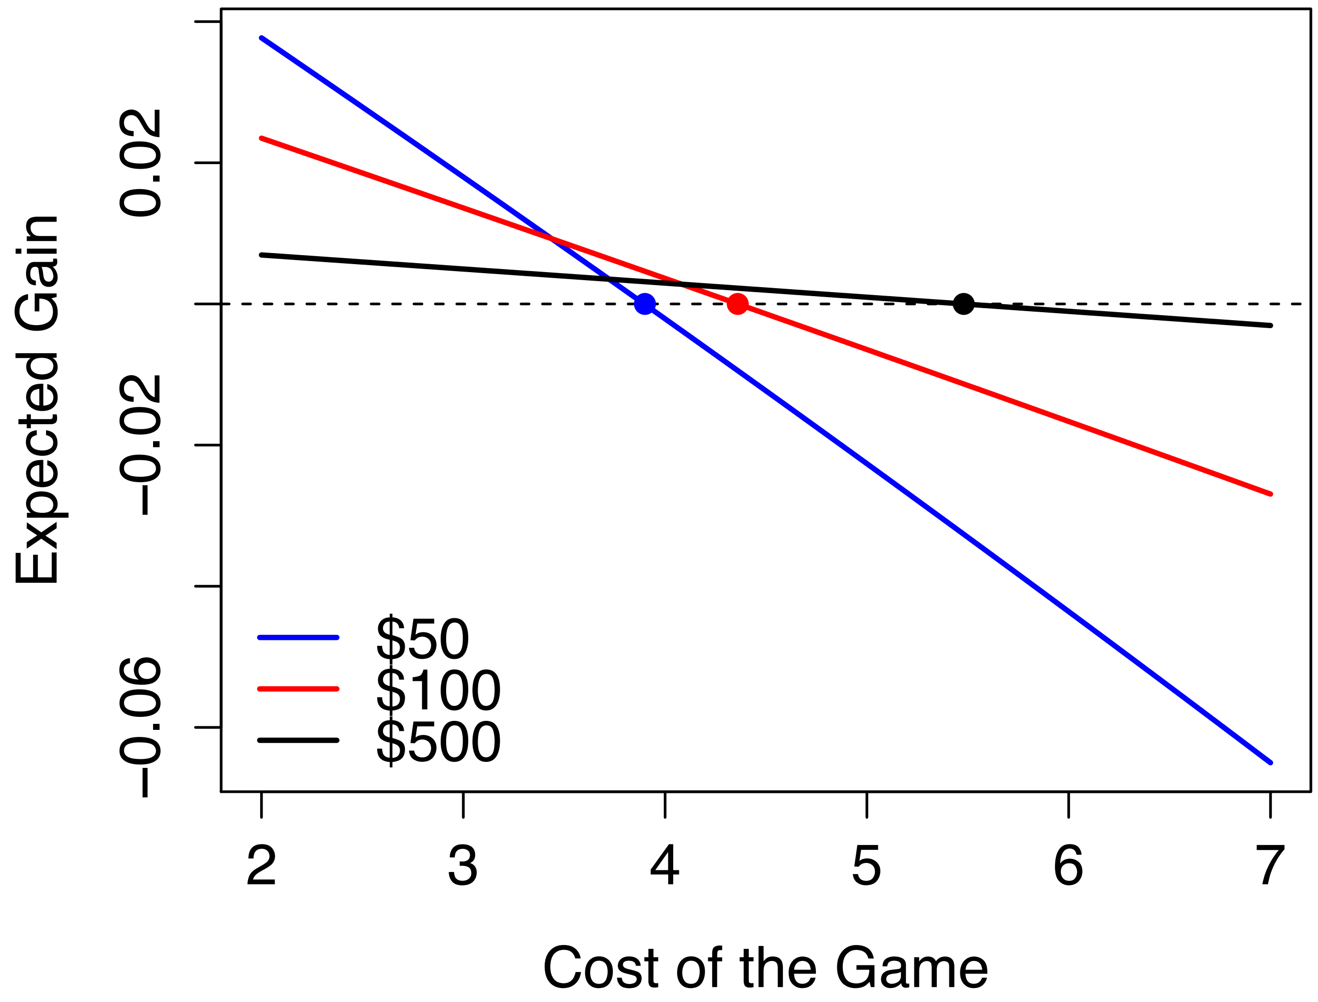

If the real utility of money scales logarithmically with nominal value, the game is only worth a few dollars

Solution #2

- Finite wealth of the casino
  - If the casino only has w dollars in their bank account…
  - Or is only insured up to a payout of w dollars…
  - This imposes an upper bound on your winnings
  - You hit this limit after L flips, where $L = 1 + \lfloor \log_2(W) \rfloor$

- Now the value of the game is:

$$\begin{aligned} E[u(\text{game})] &= \sum_{k=0}^{\infty} \frac{1}{2^{k+1}} \min(2^k, W) \\ &= \sum_{k=0}^{L-1} \frac{1}{2^{k+1}} 2^k + \sum_{k=L}^{\infty} \frac{1}{2^{k+1}} W \\ &= \frac{L}{2} + \frac{W}{2^L} \end{aligned}$$

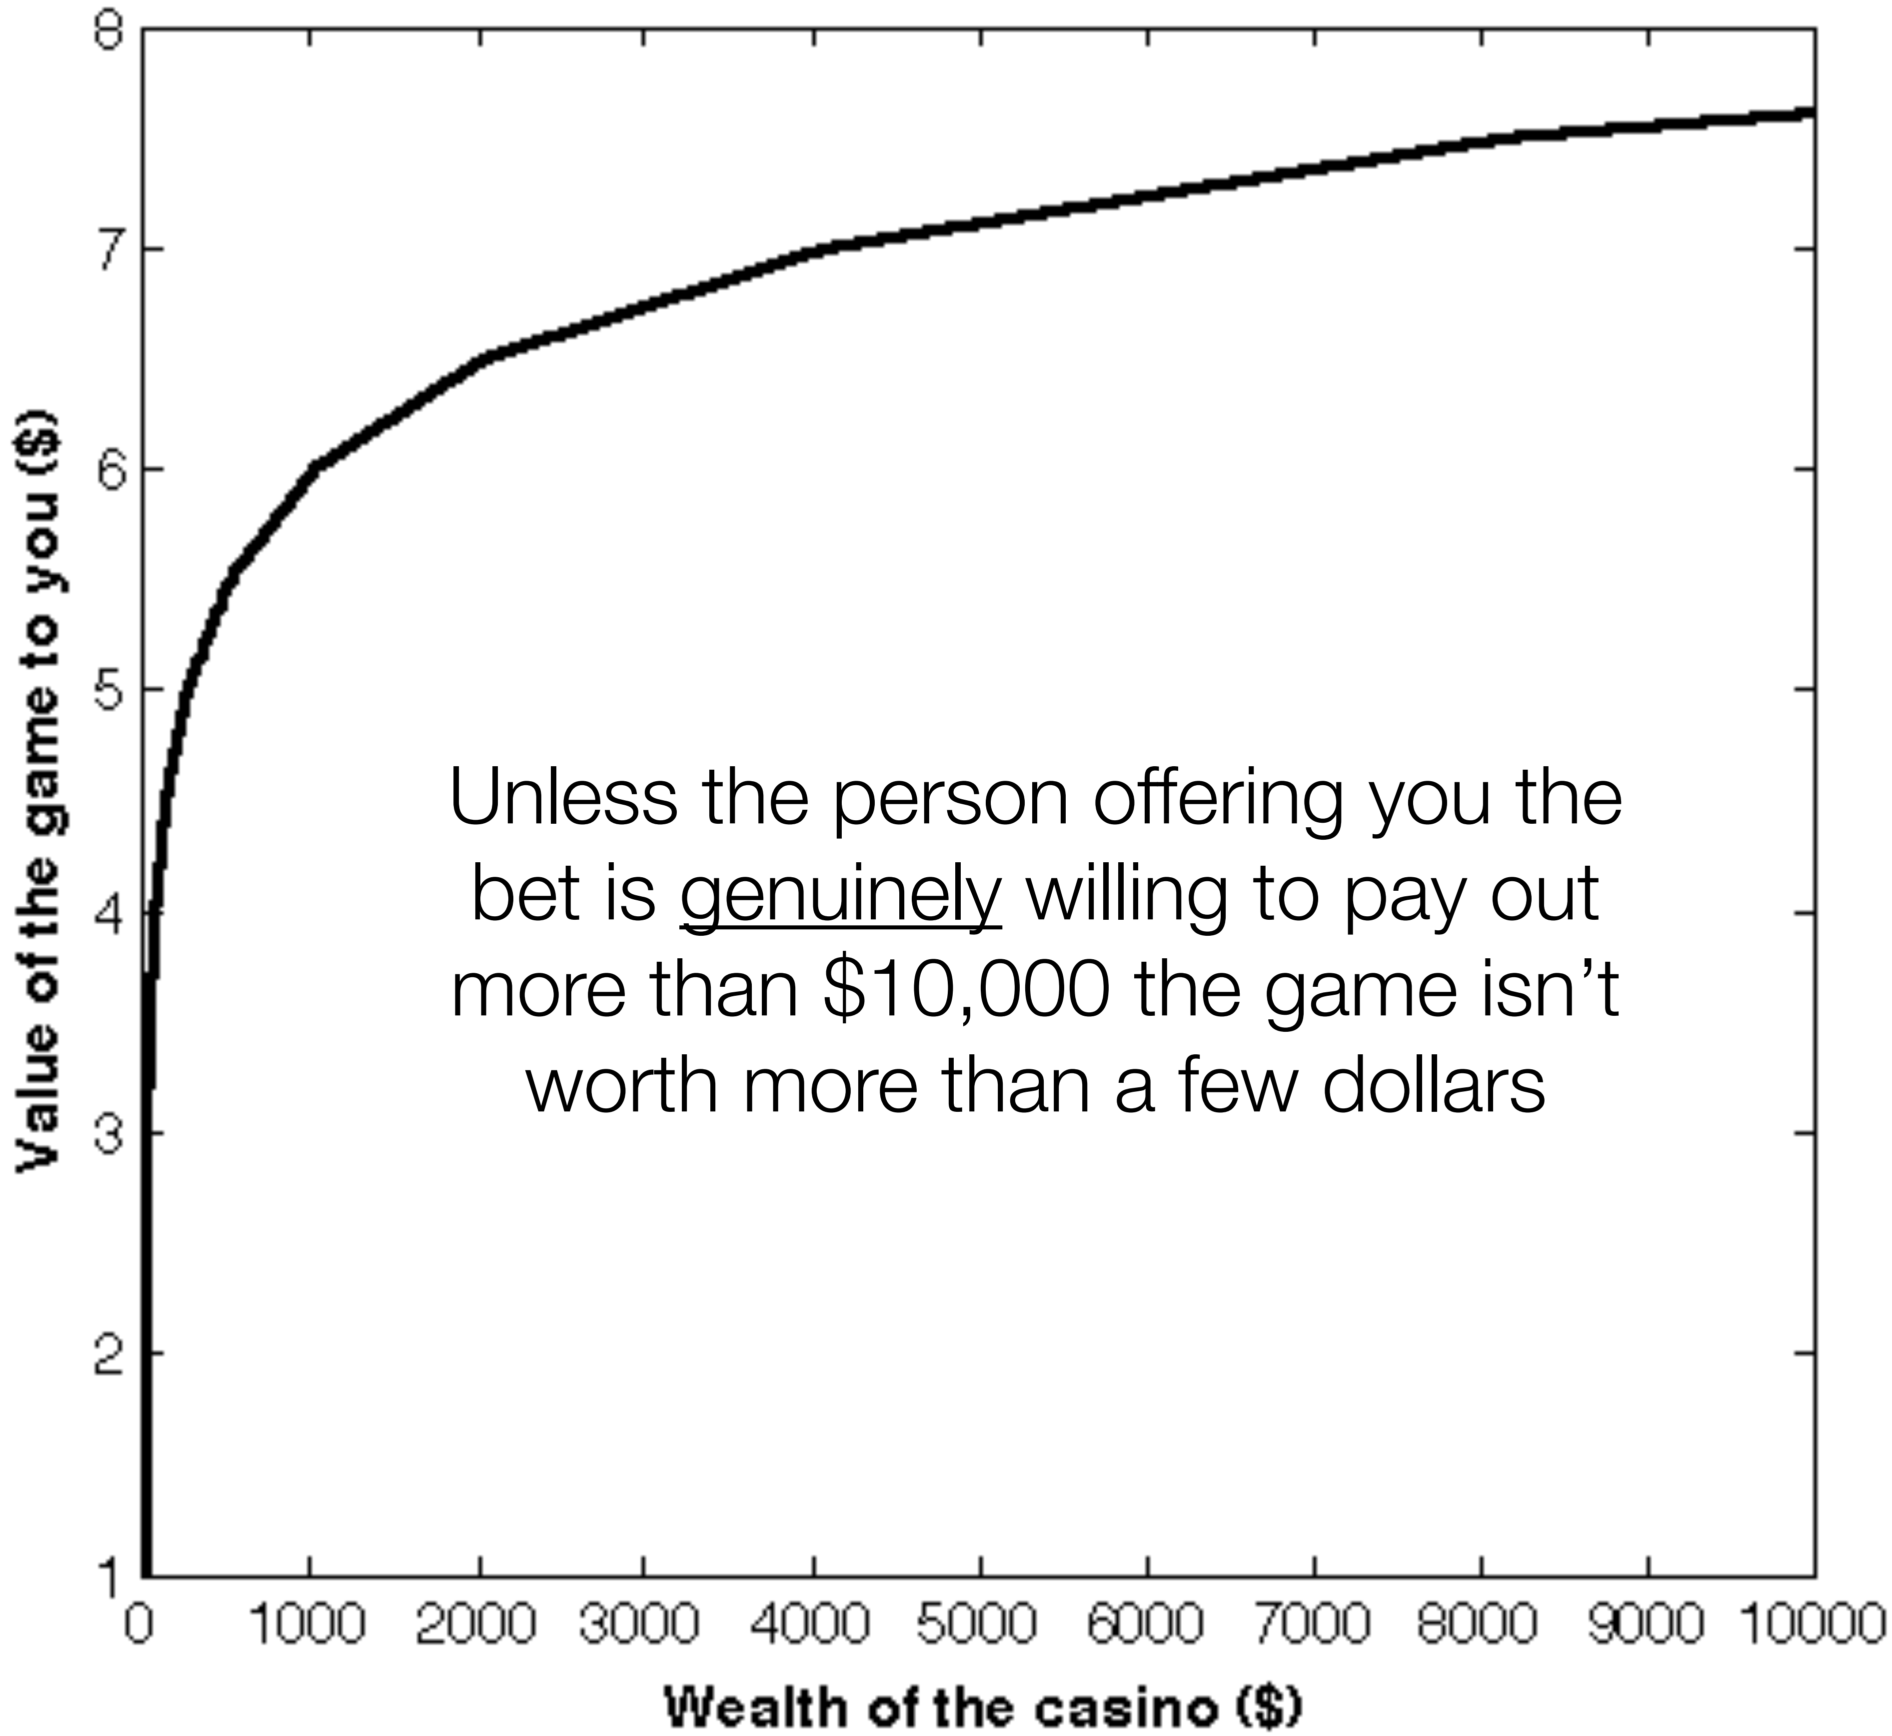
Unless the person offering you the bet is genuinely willing to pay out more than $10,000 the game isn't worth more than a few dollars
Value of the game to you ($)
Wealth of the casino ($)
1
2
3
4
5
6
7
8
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
10000

Some issues…

Utility is a difficult concept

- “Cost functions” and “utility functions”
  - They’re everywhere in machine learning and statistics
  - It’s hard to define the behaviour of an optimal agent without them
  - But they’re psychologically tricky

- It’s not easy to map between $10 and a utility value
- If we can’t do that, how will we assign utility to…
  - The first cup of coffee on a cold grey morning
  - An unexpected phone call from an old friend
  - Solving an annoying puzzle
  - Your first kiss

Utility is a difficult concept

- “Cost functions” and “utility functions”

Does it even make sense to try?
Is there really any such thing as a “utility scale” that allows you to compare these things?
Or are they truly incommensurate?

- If we can’t do that, how will we assign utility to…
  - The first cup of coffee on a cold grey morning
  - An unexpected phone call from an old friend
  - Solving an annoying puzzle
  - Your first kiss

Summary

- Expected utility theory
- Calculating action utilities can be hard
- Assigning outcome utilities can be hard

- Next lecture: improving on EU theory