Sleeping Beauty and Other Probability Conundrums

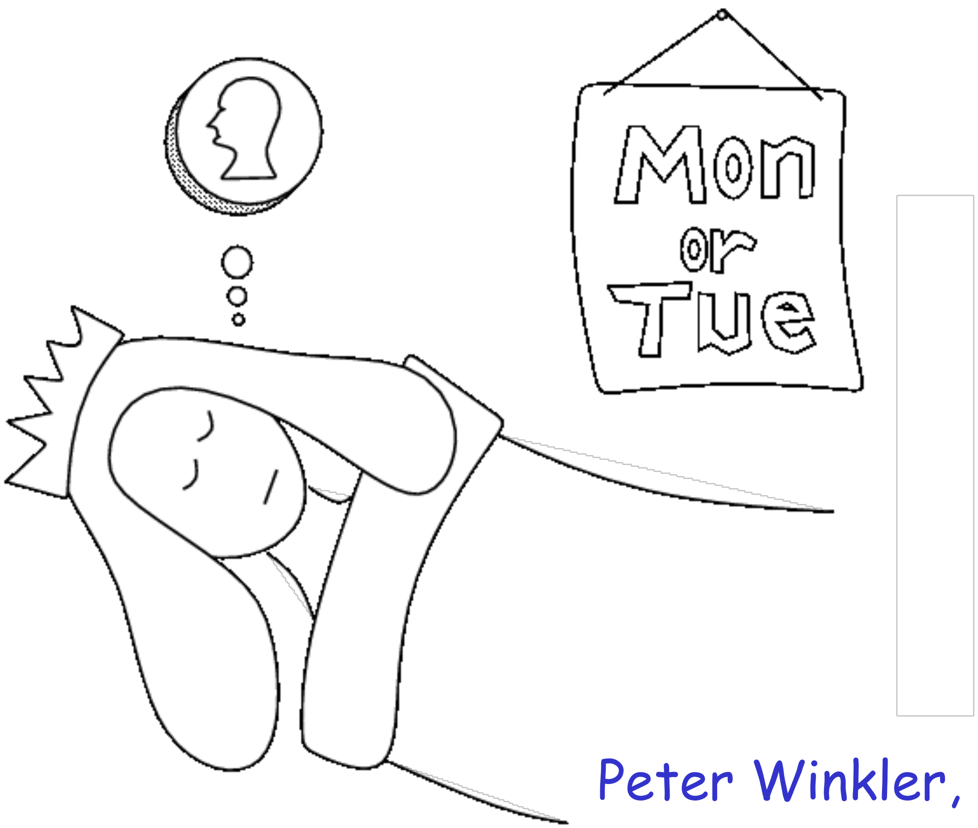

Peter Winkler, Dartmouth

Is it obvious how to define probability?

Three academics are debating the likelihood that there is intelligent life elsewhere in the universe.

"1000 to 1 in favor," says one. "New planets are discovered daily."

"1000 to 1 *against*," says the second. "It takes an exponentially unlikely sequence of conditions and events to produce intelligence. The universe is finite. Don't be fooled by anthropocentrism!"

"About even," says the third. "We have no idea, so it would be foolish to assign an extreme probability either way."

During the discussion one professor accidently rubs an ancient lamp, and a genie appears. "In return for releasing me," says the genie, "I'll tell you anything you want to know."

The genie is stymied

The three academics consult with one another.

"OK," says one to the genie. "We want to know the probability that there is intelligent life elsewhere in the universe."

"You want to know if there are other intelligent beings out there? No problem! The answer is…"

"NO NO NO!!" shout the professors. "That wouldn't settle anything! We want to know the true ***probability*** that there is intelligent life elsewhere in the universe."

The "conjunction fallacy" (Kahneman & Tversky)

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

Which is more probable?
1. Linda is a bank teller.
2. Linda is a bank teller and is active in the feminist movement.

I am particularly fond of this example [the Linda problem] because I know that the [conjoint] statement is least probable, yet a little homunculus in my head continues to jump up and down, shouting at me—"but she can't just be a bank teller; read the description."

---Stephen J. Gould

How did you get the information?

There are 100 persons who fit the description above (that is, Linda's). How many of them are:

Bank tellers? __ of 100

Bank tellers and active in the feminist movement? __ of 100

Compare:

Mrs. Smith has two children. One of them is a boy. What is the probability that the other one is a boy?

A coin is drawn from an urn containing a two-headed coin and a fair coin. The coin is flipped and comes up "heads." What is the probability that the other side is also a head?

The Vietnam Draft Lottery

Below, Rep. Alexander Pirnie, R-NY, draws the first capsule in the lottery drawing held on Dec. 1, 1969. The capsule contained the date Sept. 14.

Should you be suspicious if some month got six dates drawn before some other month got any?

The actual lottery data

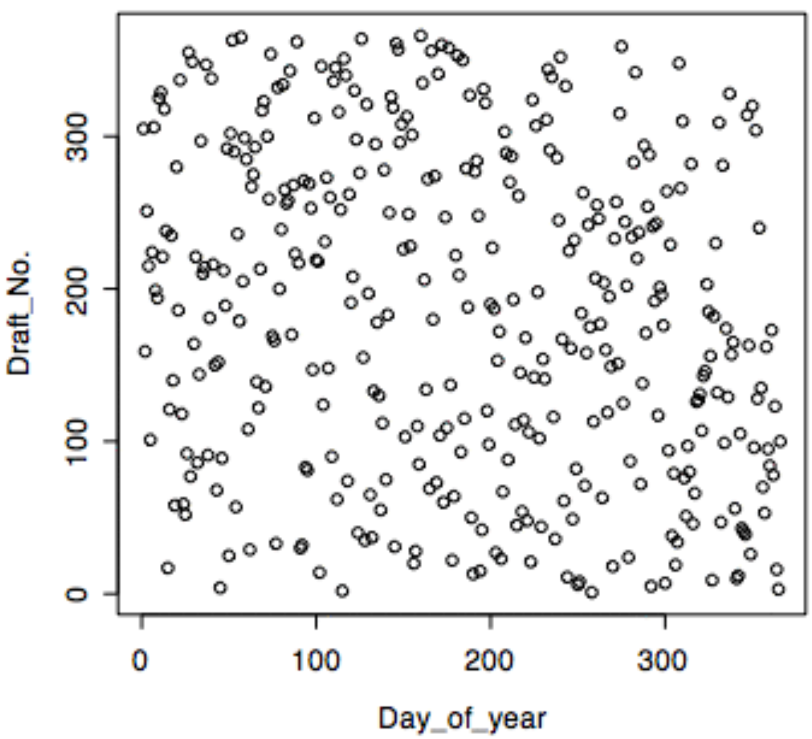

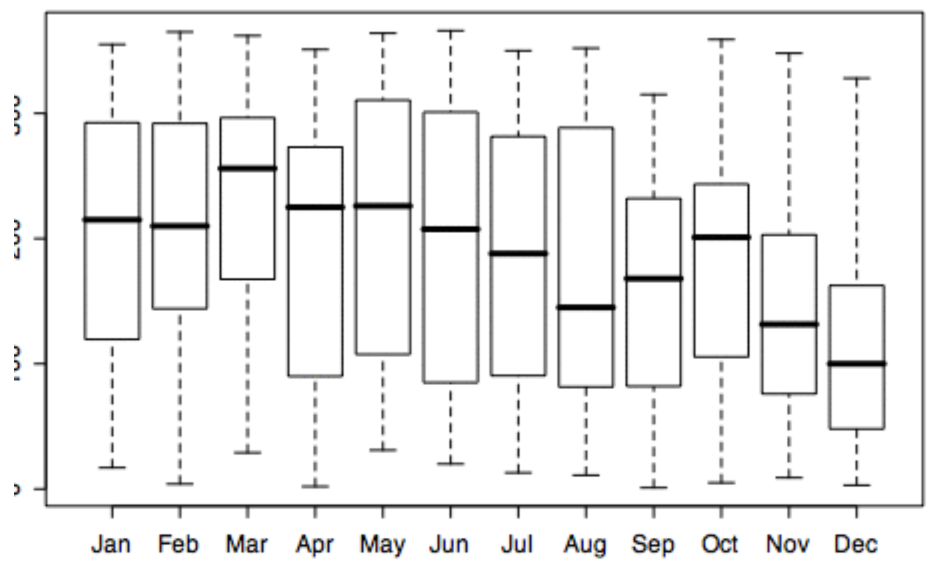

Causal vs. Evidential Decision Theory

Newcomb's Paradox:

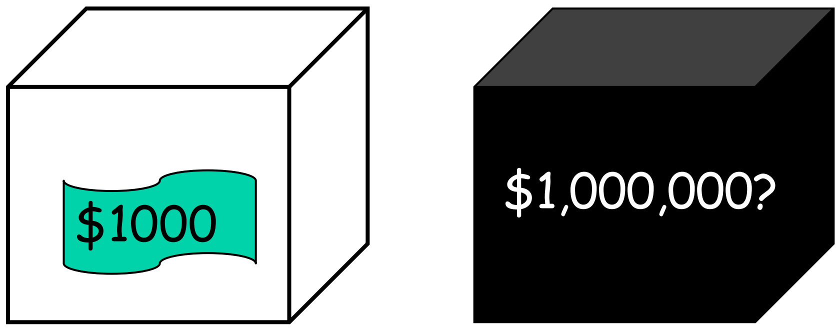

You get to take either the black box, which may or may not contain $1,000,000, or both boxes!

But the $1,000,000 is there only if X predicted you would take only one box.

Identical Twins Prisoner's Dilemma

The Absent-Minded Driver

Piccione & Rubinstein '97 (and an entire issue of *Games & Economics*):

An absent-minded driver wants to drive home via Exit 2 but can't tell Exit 2 from Exit 1 and never remembers if he's passed an exit.

If he gets off at Exit 1 he's in a bad neighborhood (payoff: 0); if he misses both exits he's a long way from home (payoff 1); if he misses Exit 1 and gets off at Exit 2, he's home (payoff 4).

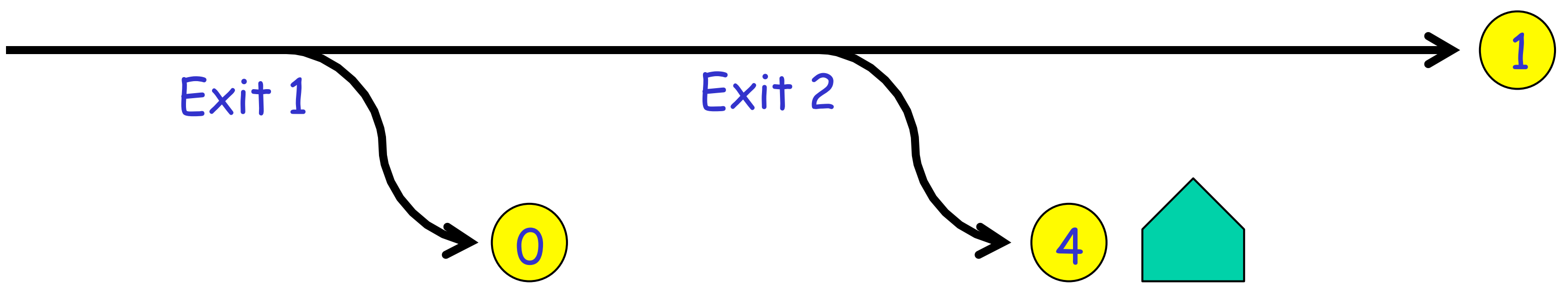

If he exits with probability p, his expected payoff is $p^2 + 4p(1-p)$ which is maximized at $p = 2/3$. So that's his plan.

But when he gets to an exit he figures he's at Exit 1 with some probability q, and now has expected payoff $q(p^2 + 4p(1-p)) + (1-q)(p + 4(1-p))$.

Adam Elga's puzzle (2001):

Sleeping Beauty ("SB") agrees to the following experiment.

SB goes to sleep on Sunday and then a fair coin is flipped.

If it comes up Heads, SB is awakened briefly on Monday, then sleeps again until Wednesday when the experiment is over.

If the coin comes up Tails, SB is awakened briefly on Monday *and again on Tuesday*, then sleeps until Wednesday.

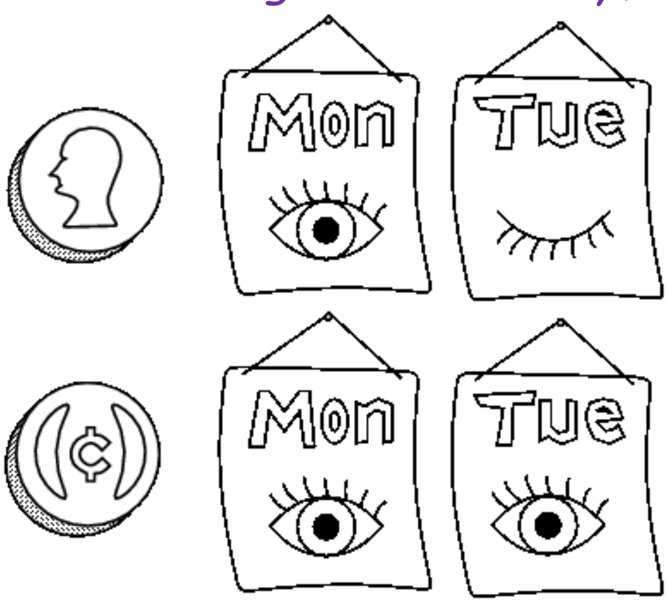

SB will have no memory of any awakenings, nor will she be told the current day of the week or the state of the coin.

When SB is awakened (on Monday or Tuesday), what---to her---is the probability that the coin has come up Heads?

Two (of many) persuasive arguments:

When SB goes to sleep on Sunday, she knows Pr(Heads) = 1/2. She knows she will be awakened, so when she is, she has no new information---thus, no basis upon which to change her mind.

So the answer is 1/2.

Suppose that SB is awakened on both Monday and Tuesday, regardless of the coinflip. But, if the coin comes up Heads, SB is told so 15 minutes after her *Tuesday* awakening. (She knows all this in advance.)

Then immediately after SB is awakened, she reasons that Monday-Heads, Tuesday-Heads, Monday-Tails and Tuesday-Tails are equally likely.

When 15 minutes have passed and she is not told that it is Tuesday-Heads, that possibility dies and the remaining possibilities remain equally likely.

So the answer is 1/3.

Camps:

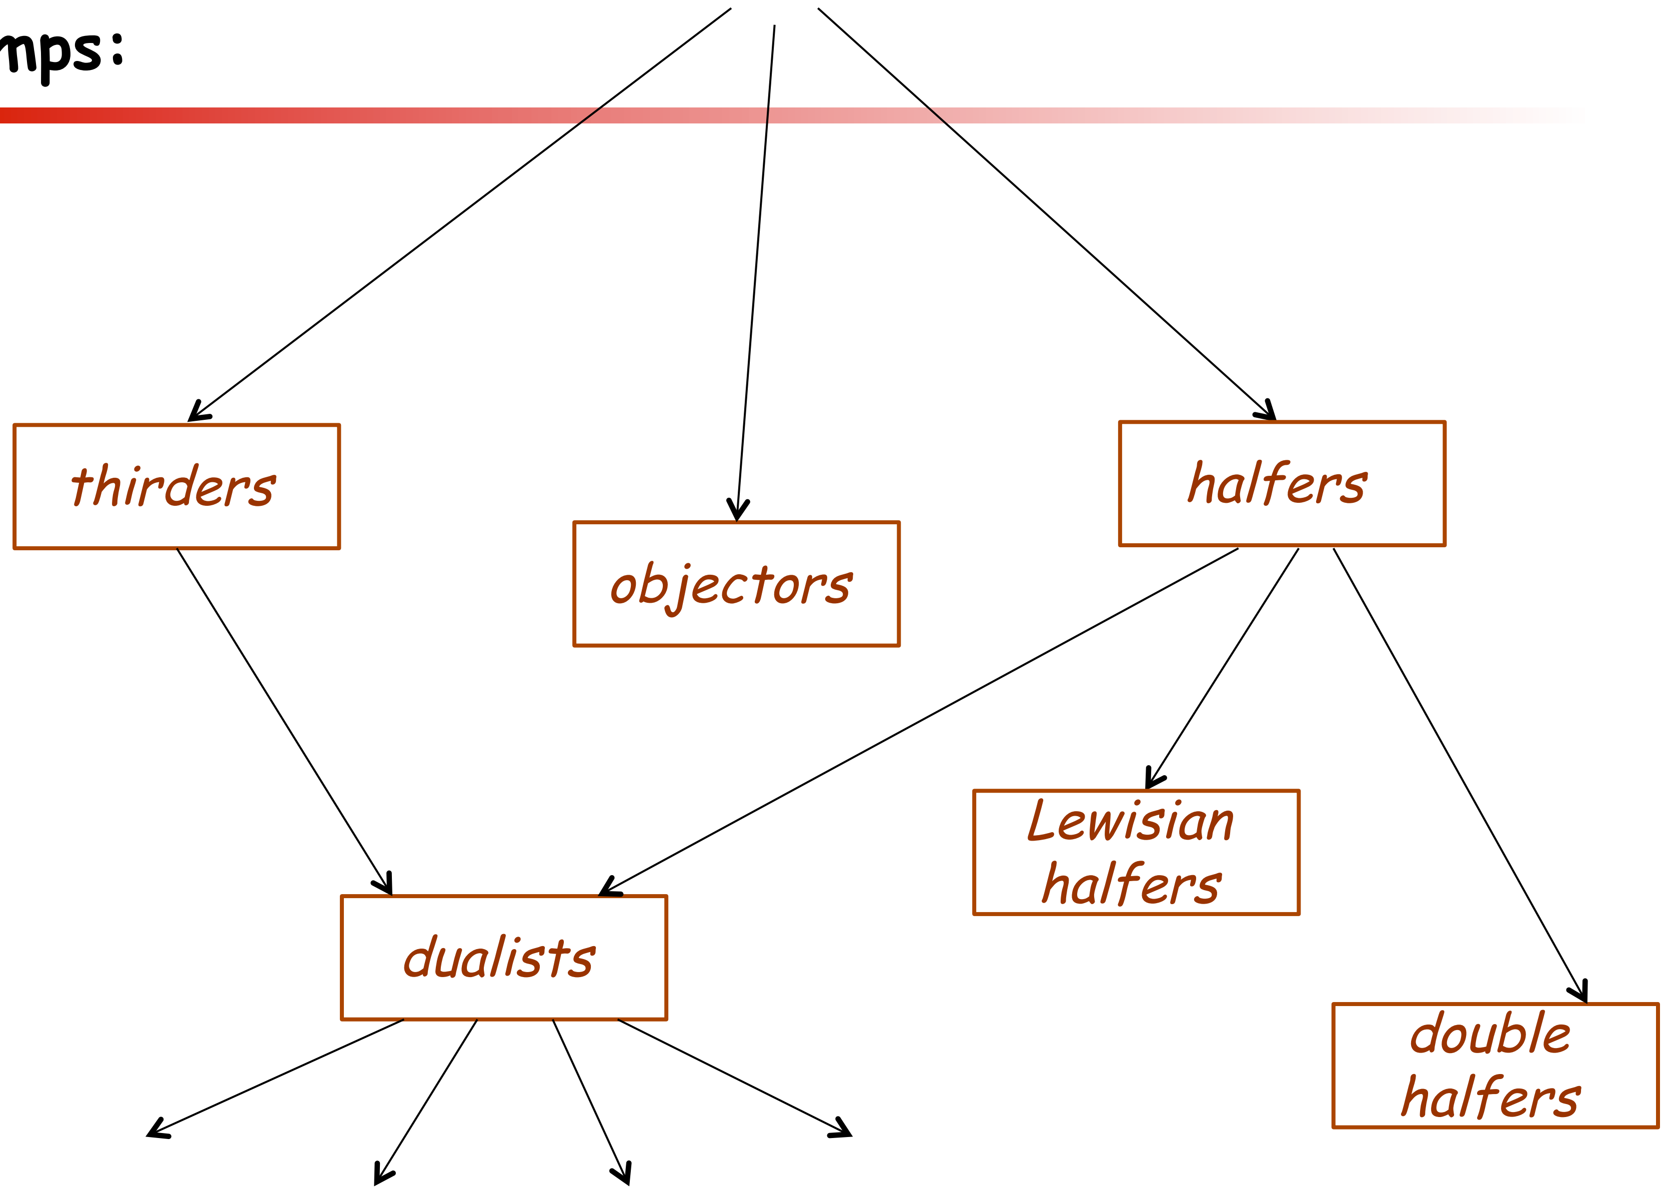
thirders
objectors
halfers
Lewisian halfers
dualists
double halfers

Whether your answer is 1/2 or 1/3 depends on:

whether you are an *evidential* or a *causal* decision theorist
(Briggs, rebutted by Conitzer);

whether you are a *many-world* or a *single-world* quantum
physicist (P. Lewis, rebutted by Peterson and others);

whether you believe in robust perspectivalism
(Pittard, who does);

whether you are asked about "the coin associated with this experiment"
or "the coin associated with this awakening" (Mutalik);

whether there is a time-measuring device in SB's room, even
if she can't read it (Meacham);

whether it is possible that SB has slight indigestion on Tuesday
(Cisewski et al., rebutted by me).

Winkler's Theory of Crackpots

You are suspicious of a supposedly random phenomenon when a low-probability event occurs *which has special significance.*

But *special significance* is sometimes a subjective judgment---especially when the subject is *you.*

Example: do dreams come true? Suppose you dream that the number 36,581,994 will win the New York State lottery, and it comes true.

Since there are 'way more than 36,581,994 people in the world, dreams with even less likelihood than yours will come true every day.

Nonetheless, your dream didn't happen to someone else, it happened to *you.* So if you assign *a priori* probability greater than 1/36,581,994 to the event that dreams come true, you are entitled to believe they do.

Thanks for listening!