Fachbereich II – Mathematik - Physik - Chemie

BEUTH HOCHSCHULE FÜR TECHNIK BERLIN
University of Applied Sciences

02/2019

Ulrike Grömping

The Sleeping Beauty Problem Demystified

Das Dornröschen-Problem enträtselt (englischsprachig)

Reports in Mathematics, Physics and Chemistry

Berichte aus der Mathematik, Physik und Chemie

ISSN (print): 2190-3913

Reports in Mathematics, Physics and Chemistry

Berichte aus der Mathematik, Physik und Chemie

The reports are freely available via the Internet:
http://www1.beuth-hochschule.de/FB_II/reports/welcome.htm

02/2019, March 2019

The Sleeping Beauty Problem Demystified
Das Dornröschen-Problem enträtselt (englischsprachig)

Editorial notice / Impressum

Published by / Herausgeber:
Fachbereich II
Beuth Hochschule für Technik Berlin
Luxemburger Str. 10
D-13353 Berlin
Internet: http://public.beuth-hochschule.de/FB_II/
E-Mail: fbiireports@beuth-hochschule.de

Responsibility for the content rests with the author(s) of the reports.
Die inhaltliche Verantwortung liegt bei den Autor/inn/en der Berichte.

ISSN (print): 2190-3913

The Sleeping Beauty Problem Demystified

Ulrike Grömping

23 March 2019

Abstract

A perfectly rational Sleeping Beauty is put to sleep on Sunday night, and intermittently awoken either on Monday only (if a fair coin shows Heads) or on both Monday and Tuesday (if the fair coin shows Tails). According to experimental conditions, she neither remembers previous awakenings nor knows about passage of time, but she is aware of experimental conditions. Her credence for the coin having come up Heads is at the center of a scientific (mostly philosophical) controversy: halfers insist it should be 1/2, while thirders claim it should be 1/3. Winkler (2017) reviewed the literature, in terminology accessible to statisticians. This paper presents a comprehensive treatment of the probability side of the problem, from the perspective of a frequentist applied statistician. Perceived contradictions are resolved (in line with Groisman 2008) or traced back to violation of model assumptions. Modeling was at the heart of the confusion and is also at the heart of the solution.

MSC2000: Primary: 62A01. Secondary: 62C99, 60A05
Keywords: Probability Dilemma, Probability Paradox, Conditional Probability, Decision Theory, Monte Carlo simulation, Sleeping Beauty Problem.

1 Introduction: The Sleeping Beauty Problem

Elga (2000) popularized a decision-theoretic problem that sparked a division of philosophers and probability theorists into "halfers" and "thirders", and a growing number of dualists. The latter claim that both camps are correct, because they interpret the problem in different ways that can both be justified. Groisman (2008) introduced the dualist perspective, titling "The end of Sleeping Beauty's nightmares". Beauty's sleep, whether with or without nightmares, did not end with his work, but continued to be widely discussed. This paper extends Groisman's take on the problem, using a frequentist applied statistics perspective. Contrary to Groisman, however, the author is a convinced halfer (but without missionary zeal).

The Sleeping Beauty problem is a probability dilemma which, after more than a decade of scientific discussion, still does not have a universally agreed solution, contrary to the well-known Monty Hall problem[1] for which all serious scientists agree on the correct solution (if they agree on the model assumptions). The philosophical literature uses both problems as a test ground for various theories (see e.g. Mann and Aarnio 2018 for a recent contribution on both the Monty Hall and the Sleeping Beauty problem). The present paper resolves the probability aspects of the Sleeping Beauty problem; the problem will no doubt remain of interest to philosophers.

Recently, Winkler (2017) gave an overview of the discussion of the Sleeping Beauty problem that is – contrary to many other works in the field – accessible to mathematicians or statisticians who are not deeply involved in philosophical reasoning; readers are referred to this source for a review. This paper will only work with a small selection of references. Winkler (2017) stated the problem as follows:

> Sleeping Beauty agrees to the following experiment. On Sunday, she is put to sleep, and a fair coin is flipped. If it comes up Heads, she is awoken on Monday morning; if Tails, she is awoken on Monday morning and again on Tuesday morning. In all cases, she is not told the day of the week, is put back to sleep shortly after, and will have no memory of any Monday

[1] The Monty Hall problem was first introduced by Selvin 1975, and was popularized in 1990 in the journal *Parade* by columnist Marylin vos Savant who answered a reader question. In the standard version of the problem, the game show host Monty Hall shows three closed doors to a contestant, and tells the contestant that there is a luxury car behind one of the doors and a goat behind each of the other two doors. The contestant is asked to pick a door, and picks one. Now, Monty Hall opens one of the other two doors and presents a goat to the contestant. Subsequently, the contestant is offered the choice to switch to the other closed door or to stick with the original pick. It is advantageous to switch, because the probability for the already picked door to contain the car remains 1/3, while the new information increases the probability for the car to be behind the other door to 2/3 (provided that Monty Hall never spoils the game by opening the picked door or the door that hides the car).

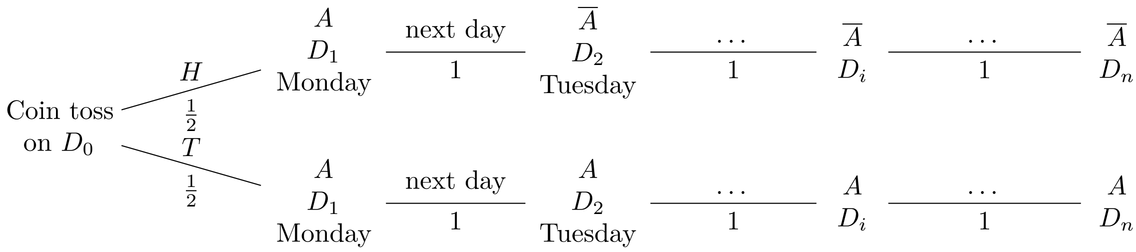

Figure 1: Visualization of the Sleeping Beauty problem with n laboratory days $D_1, \ldots, D_n$ and A for awakening, $\overline{A}$ for not awakening. The coin is tossed before the beginning of the experiment.

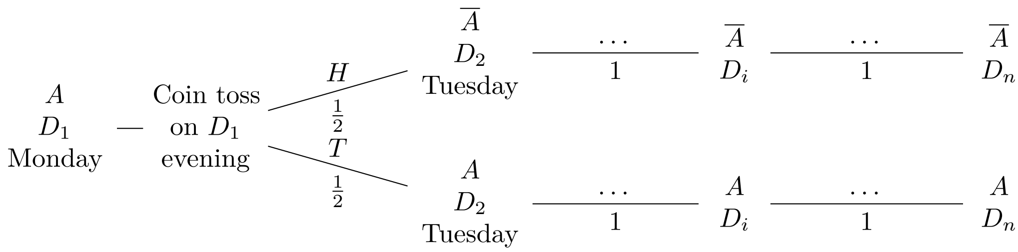

Figure 2: Visualization of the Sleeping Beauty problem with n laboratory days $D_1, \ldots, D_n$ and A for awakening, $\overline{A}$ for not awakening. The coin is tossed after the certain awakening on D_1.

> or Tuesday awakenings. When Sleeping Beauty is awoken on Monday or Tuesday, what – to her – is the probability that the coin came up Heads?
>
> — Winkler (2017)

Bostrom (2007) generalized the number of laboratory days from two (Monday and Tuesday) to $n \geq 2$, in order to demonstrate perceived implausibilities, when increasing n. In this extended setting, Beauty is still awoken only on the first experimental day if the coin comes up Heads and on all n experimental days if the coin comes up Tails. All other experimental conditions also apply without change: in particular the "no memory" condition applies to all awakenings. Figure 1 illustrates the extended experiment from Winkler's quote, Figure 2 a slight modification thereof in which the coin toss does not happen before but after the Monday (D_1) awakening, which should be possible and inconsequential because the actions on Monday (D_1) do not depend on the outcome of the coin toss (but beware that this modification may interfere with what assumptions may appear plausible to Beauty).

In many accounts of the Sleeping Beauty problem, two questions are considered:

- **Question 1** (already part of the quoted setup by Winkler). On any instance of an awakening, which probability (or credence) should Beauty assign to the coin having come up Heads?
- **Question 2.** On a Monday (or D_1) awakening, which probability (or credence) should Beauty assign to the coin having come up Heads, if she gains the additional (and credible) information that today is Monday (or D_1)?

Winkler omitted Question 2 from the definition of the Sleeping Beauty problem but also considered the corresponding probability in a part of his discussion.

This paper equates credence to probability, assuming that a canonical Bayesian would request the credence to be chosen as the probability, if a probability can be derived. We assume that Beauty is perfectly rational in that sense.

The division into thirders (e.g. Elga 2000) and halfers (e.g. Lewis 2001) is based on Question 1: thirders answer that question with $1/3$ for $n = 2$, and $1/(n + 1)$ in general, while halfers generally answer it with $1/2$. The main argument of halfers is: Beauty does not get new information from knowing that she is awake now, because she knew all along that she would be woken during the experiment; thus prior information (fair coin) should be upheld. Winkler (2017) summarized various arguments that have been brought forward as justification of the thirders' view, the most prominent one being an indifference argument among awakenings, derived from symmetry considerations. Halfers and thirders also give different answers to Question 2. As was mentioned before, Bostrom (2007) introduced the extension of the Sleeping Beauty problem to n instead of only two experimental days; he found fault with both Elga's and Lewis' lines of reasoning and argued that the thirders' answer $1/(n+1)$ to Question 1 as well as the halfers' answer $n/(n+1)$ to Question 2 (see Section 5) yields very implausible results for large values of n; he proposed a hybrid model – in a terminology that is inaccessible to this author – which answers $1/2$ to both questions, regardless of n. He is thus a so-called double-halfer.

Luna (forthcoming) concisely summarized the most frequently emphasized aspect of the conflicting positions on Question 1:

> The strong law of large numbers and considerations concerning additional information strongly suggest that Beauty upon awakening has probability $1/3$ to be in a heads-awakening but should still believe the probability that the coin landed heads in the Sunday toss to be $1/2$.
>
> — Luna (forthcoming, quote from preprint)

He employed the philosophical distinction between ordinary and centered worlds for resolving the apparent contradiction.

The present paper constrains itself to ordinary probability theory and applied statistics, without employing centered worlds or other philosophical concepts not easily accessible to an applied statistician. The goals are

- to provide a few simple probabilistic models that explain each of the proposed answers to Questions 1 and 2 and relate to Beauty's potential assumptions on what exactly she knows. The simplest model is deterministic for everything except the coin toss. Two further models employ different randomness assumptions on how Beauty's current awakening day arises.
- to clarify the different perspectives on Question 1 that one can take, sharpening the contribution of Groisman (2008) who argued that halfers and thirders give correct and valid answers to different questions, so that their points of view are not contradictory.
- to explicitly relate the classical arguments by Elga (2000) and Lewis (2001) to the probability models of this paper, and to illustrate the causes for the long-standing confusion around the Sleeping Beauty problem.

Various sources tried to justify the correctness of their camp using betting or gambling; however, this brings along specifications, depending on which events trigger which outcomes. For example, if Beauty wants to be correct often (e.g. because of a per-awakening reward), she should pick T all the time; this is a setting for which the thirders' answer (under the "outcome population" perspective, see Section 4) would suit her needs. Mutalik (2016a) emphasized that proponents of all camps agreed on correct solutions, as soon as such concretizations were added to the problem. We will therefore not consider the betting approach.

The problem is interesting for the statistical community because it shows that modeling assumptions are extremely important and that using applied statistics for problems like this can help untangle arguments that are much harder to resolve solely by probability calculus or philosophical discourse. The route to adequate modeling is supported by thinking about appropriate Monte Carlo simulation of the problem. Therefore, simulation of the models is always considered along with the models themselves. This is a two-way street: simulation makes it easier for an applied statistician like this author to understand the Sleeping Beauty problem, and the performance of naïve simulation approaches can serve as a caveat on paying attention to the experimental setup in simulations for real applications.

Draper (2017), referring to the Oscar Seminar (2008, "An objectivist argument for thirdism"), asked for more objectivist involvement in answering the Sleeping Beauty problem, and expressed the expectation that the problem would remain difficult even for objectivists. Cisewski et al. (2016) proposed a very general model that is able to accomodate arbitrary knowledge Beauty may bring into her credence (like aches, pains or bodily functions at the current awakening) at the expense of making it hard to directly connect it to previous works. The present paper can be seen as a more pragmatic objectivist attempt by an applied statistician (based on frequentist probability theory) at explaining all aspects of the Sleeping Beauty problem, without allowing information from outside the direct experimental context to enter the picture. Section 2 introduces notation and sample space considerations. Sections 3 and 4 concentrate on Question 1: Section 3 discusses Beauty's potential assumptions and relates them to models, and Section 4 discusses which different probabilities Beauty can target (sharpening the Groisman (2008) point). Section 5 discusses answers for Question 2 and explicitly relates the models of this paper to the classsical reasoning by Elga (2000) and Lewis (2001). Section 6 summarizes the assessments of the three possible pairs of answers to Questions 1 and 2, and the final section remarks on the steps taken to resolve the confusion about the Sleeping Beauty problem, aspects that make the problem difficult and the role of applied statistics and simulation in relation to the Sleeping Beauty problem.

2 Notation, sample space, and events for capturing Beauty's knowledge

$\wedge$ denotes conjunction, $\vee$ disjunction, and $\mathbb{1}_{\text{condition}}$ denotes an indicator that yields 1 if the condition is true and 0 otherwise. This paper writes events as sets, and thus uses set notation for combining events ($\cap$ for intersection, $\cup$ for union, and $\times$ for Cartesian product). A conditional probability of event A given event B is denoted as $P(\mathsf{A}|\mathsf{B})$.

A parsimonious view on the sample space is deterministic, apart from the coin toss outcome:

$$\Omega_1 = \Omega_{coin} = \{H, T\}. \tag{1}$$

Both elements of this sample space have probability 1/2, and they have a one-to-one correspondence to the corresponding experimental trajectories that are depicted in Figures 1 and 2. This sample space is sufficient without any further assumptions, as long as one is only interested in the probability of the initial coin toss being Heads, and as long as Beauty believes that she has not gained additional information by being awake, because she knew in advance that she would be awoken. Days and individual awakenings are included in Ω_1 through the deterministic trajectories. Credences involving days can be derived from applying an indifference principle to a reasonable selection of cases of interest (see also Sections 3.2 and 5.2).

Whenever Beauty wants to incorporate additional information regarding the day of her awakening, or wants to work out proper *probabilities* about events that relate to awakening days, a more detailed sample space is asked for:

$$\begin{aligned} \Omega_2 &= \Omega_{coin} \times \Omega_{\text{day}} = \{H, T\} \times \{D_1, \ldots, D_n\} \\ &= \{(H, D_1), \ldots, (H, D_n), (T, D_1), \ldots, (T, D_n)\} \\ &= \{(\omega_1, \omega_2) : \omega_1 \in \Omega_{coin} \wedge \omega_2 \in \Omega_{\text{day}}\}. \end{aligned} \tag{2}$$

Furthermore, according to the experimental rules, awakening functions A and $A_{\text{experiment}}$ are defined on this sample space as follows:

$$\begin{aligned} A : \Omega_2 &\to \{0, 1\} \\ (\omega_1, \omega_2) &\mapsto \mathbb{1}_{\{\omega_1 = T \,\vee\, \omega_2 = D_1\}}, \end{aligned} \tag{3}$$

$$\begin{aligned} A_{\text{experiment}} : \Omega_2 &\to \{0, 1\} \\ (\omega_1, \omega_2) &\mapsto 1. \end{aligned} \tag{4}$$

Note the rationale for function $A_{\text{experiment}}$: it yields the value 1 if and only if Beauty is awoken at least once in the experiment; it treats every sampled element as a representative of its trajectory; since both potential trajectories contain an awakening with certainty, $A_{\text{experiment}}$ assigns the value 1 to every element of the sample space. The purpose of this function is to make the sample space compatible with Beauty's knowledge about the overall experiment, in spite of sampling a single day only from the experimental trajectory.

For use in models that are based on Ω_2, the following events are defined:

$$\begin{aligned} \mathsf{H} &= \{\omega \in \Omega_2 : \omega_1 = H\}, \\ \mathsf{T} &= \{\omega \in \Omega_2 : \omega_1 = T\}, \\ \mathsf{D}_i &= \{\omega \in \Omega_2 : \omega_2 = D_i\}, \\ \mathsf{A}_{\text{day}} &= \{\omega \in \Omega_2 : A(\omega) = 1\} = \{\omega \in \Omega_2 : \omega_1 = T \vee \omega_2 = D_1\}, \\ \mathsf{A}_i &= \{\omega \in \Omega_2 : A(\omega) = 1 \wedge \omega_2 = D_i\}, \\ \mathsf{A}_L &= \{\omega \in \Omega_2 : A_{\text{experiment}}(\omega) = 1\} = \Omega_2. \end{aligned} \tag{5}$$

H and T, as well as D_i have the obvious meanings. A_L is the event that Beauty is awoken in the laboratory during the experiment, which happens with probability 1. A_{day} is the event that Beauty is awoken *today*, i.e. she is in an awakening day of the experiment; however, under the experimental conditions of memory loss and ignorance of which day she is in, *today* does not convey information over and above the fact that Beauty is in a combination of day and coin toss that implies an awakening. While A_L is something that Beauty knew from the start, her knowledge can be seen to be A_{day}, when actually being awake during the experiment. Should she, by any chance, learn that the current day is D_i, while being awake, her knowledge would become A_i. We can also handle Ω_1 as isomorphic to Ω_2 by employing the σ-algebra $\{\emptyset, \mathsf{H}, \mathsf{T}, \Omega_2\}$ on Ω_2, and we can therefore also use events H and T under Ω_1.

When considering a series of N experiments, this paper will use the shortened expressions "Heads experiment" or "Tails experiment" for experiments in the series for which the coin toss outcome was Heads or Tails, respectively. Furthermore, an awakening that occurs during a Heads experiment is called a Heads awakening.

3 Beauty's assumptions, and implied models

The one doubtlessly random experiment in the Sleeping Beauty problem is the toss of the fair coin (on D_0 or after the D_1 awakening), which implies

$$P(\mathsf{H}) = P(\mathsf{T}) = 0.5. \tag{6}$$

Given the result of this random experiment, Beauty's awakening(s) occur(s) according to deterministic rules, and D_1 to D_n occur deterministically one after the other, as was shown in Figures 1 and 2.

3.1 Beauty's assessment of her own knowledge

If Beauty assesses her knowledge to only consist of A_L (she has been woken in the experiment at least once), no update of her prior belief is warranted, because A_L is the certain event, regardless how the sample space is modelled.

If Beauty believes that her knowledge is A_{day}, i.e. she is currently in an awakening day of the experiment, implications of that knowledge depend on whether or not Beauty incorporates the experimental setup (which she is aware of in all accounts of the Sleeping Beauty problem) into her assumptions:

If she does, she has to assign unequal probabilities to the elements of Ω_2 (as proposed by Lewis 2001), either because of hierarchical indifference assumptions within Model M_1 or as outlined in Model M_2 below. This yields the halfers' answer.

Otherwise, she can apply indifference over awakening days (outside of probability theory within Model M_1, as proposed by Elga 2000) or use a probability model that ignores the experimental setup (see Model M_3). This yields the thirders' answer.

This author works on experimental design; thus, in her view, a fully rational approach has to incorporate the experimental setup.

3.2 Days remain deterministic (Model M_1)

Model M_1 assumes an equiprobable pick from Ω_1 (i.e. toss of the fair coin) only, with the corresponding experimental trajectory arising in due course. The simulation of N_{simul} trajectories under the deterministic assumption (M_1) proceeds as follows: Simulate N_{simul} coin toss outcomes and record the corresponding trajectories.

If Beauty believes that all she knows is that she was awoken during the experiment, standard probability theory can provide her probability: $P(\mathsf{H}|\mathsf{A}_L) = 1/2$, because A_L is the certain event.

If Beauty pictures herself in one of the $n + 1$ possible awakening settings (those that would be part of A_{day} when considering the sample space Ω_2) and applies an intuitive indifference principle, she can assess the probability of picking a Heads awakening from the population of awakenings produced by a series of experiments, as $1/(n+1)$ (see Section 5.2 for Elga's symmetriy argument for this indifference).

Alternatively, Beauty can picture herself in the experiment, consider herself with probability 1/2 in the Heads branch and therefore with probability 1/2 in the single possible Heads awakening (so far, it's probability theory), whereas she optionally applies a non-probabilistic indifference argument to the Tails awakenings for each of which she assesses the chance $1/n$ given Tails, and thus $1/(2n)$ overall (see also Model M_2 for a probability-based version of this aproach). Thus, she can assess the probability for H, given an awakening day as 1/2, even when solely using probability theory, without any formal conditioning (no new information obtained).

3.3 Unequal probabilities on $\Omega_2 = \Omega_{coin} \times \Omega_{\mathbf{day}}$ (Model M_2)

Beauty pictures herself in each of the two equiprobable experimental trajectories; given the trajectory (≙coin toss outcome), she pictures herself in a random pick of possible awakening days. Of course, given Heads, she is certain that it must be D_1, given Tails she makes a uniform random assumption over $D_1, \ldots, D_n$. The justification of this assumption is more plausible if the coin is tossed before Beauty's current awakening, thus Figure 2 should not be assumed for this variant. Simulation of N_{simul} elements of Ω_2 under this conditional uniform distribution of awakening days *given coin toss* (Model M_2) proceeds as follows:

- Simulate N_{simul} coin toss outcomes,
- and pick the awakening day as D_1 for each H toss and as a random element from Ω_{day} for each T toss.

The conditional sampling of awakening days implies that $P(\mathsf{A}_{\text{day}}) = 1$. The elements of Ω_2 have the following probabilities:

$$\begin{aligned} P(\mathsf{A}_i \cap \mathsf{T}) &= P(\{(T, D_i)\}) = P(\mathsf{T} \cap \mathsf{D}_i) = 1/(2n) \quad \text{for } i = 1, \ldots, n, \\ P(\mathsf{A}_1 \cap \mathsf{H}) &= P(\{(H, D_1)\}) = P(\mathsf{H} \cap \mathsf{D}_1) = 1/2, \\ P(\mathsf{A}_i \cap \mathsf{H}) &= P(\{(H, D_i)\}) = P(\mathsf{H} \cap \mathsf{D}_i) = 0 \quad \text{for } i = 2, \ldots, n. \end{aligned} \tag{7}$$

The probabilities in (7) coincide with those stated by Lewis for $n = 2$. Probability calculus trivially yields $P(\mathsf{H}|\mathsf{A}_L) = P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/2$. As an aside, within this model (which implicitly conditions on awakening), marginal probabilities for specific awakening days are

$$P(\mathsf{A}_1) = P(\mathsf{D}_1) = \frac{n+1}{2n}, \qquad\qquad P(\mathsf{A}_i) = P(\mathsf{D}_i) = \frac{1}{2n}, \quad i = 2, \ldots, n,$$

i.e. $P(\mathsf{A}_1) = 3/4$ and $P(\mathsf{A}_2) = 1/4$ for $n = 2$.

3.4 Equal probabilities within $\Omega_2 = \Omega_{coin} \times \Omega_{\mathbf{day}}$ (Model M_3)

Beauty pictures herself in every laboratory day (not restricted to awakening days) with the same probability:

$$P(\mathsf{D}_i) = 1/n, i = 1 \ldots n. \tag{8}$$

Every day is independently combined with the experimental coin toss, which implies equiprobable elements of Ω_2:

$$P(\omega) = 1/(2n) \quad \text{for each} \quad \omega \in \Omega_2. \tag{9}$$

Simulations of N_{simul} elements of Ω_2 under the unconditional uniform assumption on the distribution of laboratory days proceeds as follows: Independently simulate

- N_{simul} coin toss outcomes uniformly from Ω_{coin},
- N_{simul} days uniformly from Ω_{day}.

Note: For larger n, N_{simul} must be made quite large in order to achieve enough samples from individual laboratory days (needed for Question 2).

In Model M_3, considering the entire experiment at once, each sampled element of Ω_2 has to be treated as a representative of the entire trajectory; this is why function $A_{\text{experiment}}$ has been defined as the constant function returning 1, so that A_L becomes the certain event. For day-wise considerations, the sample space elements stand for themselves only. In this model, A_{day} is not the certain event, but $P(\mathsf{A}_{\text{day}}) = (n+1)/(2n)$, with $P(\mathsf{A}_1) = 1/n$ and $P(\mathsf{A}_2) = \cdots = P(\mathsf{A}_n) = 1/(2n)$.

Clearly, probability calculus yields

$$\begin{aligned} P(\mathsf{H}|\mathsf{A}_L) &= P(\mathsf{H}) = 1/2, \\ P(\mathsf{H}|\mathsf{A}_{\text{day}}) &= P(\{(H, D_1)\}|\mathsf{A}_{\text{day}}) = P(\mathsf{H} \cap \mathsf{D}_1)/P(\mathsf{A}_{\text{day}}) = 1/(n+1). \end{aligned} \tag{10}$$

For $n = 2$, $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/3$ (the thirders' perspective).

Model M_3 directly implies Elga's indifference reasoning on awakening days: (9) implies that each element of A_{day} receives the same probability in the conditional distribution, i.e.

$$P(\mathsf{H} \cap \mathsf{D}_1|\mathsf{A}_{\text{day}}) = P(\mathsf{T} \cap \mathsf{D}_1|\mathsf{A}_{\text{day}}) = \cdots = P(\mathsf{T} \cap \mathsf{D}_n|\mathsf{A}_{\text{day}}) = 1/(n+1). \tag{11}$$

Thus, Model M_3 can serve as a probabilistic justification for uniform sampling from awakening days (let's call it Model M_{3a}). This support for the thirders' view arises because Model M_3 refrains from conditioning on the known experimental setup. In fact, the model would apply without change to a modified experiment with a daily coin toss also on D_2 to D_n for deciding independently for each day whether or not Beauty is awoken (yes if Tails, no if Heads).

4 Interpretation of Question 1

For the thirders' view, the strong law of large numbers is often cited, presumably with repeated experiments in mind. The quote from Luna (forthcoming) shown in the introduction illustrates the conflict: in a sequence of N experiments with $n = 2$, a third of the awakenings are expected to be Heads awakenings, while half of the coin tosses are expected to come out Heads – trivial, really. Groisman (2008) emphasized that it makes a difference whether Beauty wants to form a credence about the coin toss showing Heads (related to the entire experiment), equivalent to *the current awakening having been produced by a Heads coin toss* but different from *an awakening sampled from the population of awakenings being a Heads awakening*. Exposing Beauty to repeated experiments supports the perception of this distinction, because it highlights the fact that T coin toss outcomes are overrepresented among awakenings, i.e. in the population of awakening outcomes, in spite of occurring with the same probability as H coin toss outcomes. Therefore, this section will work with repeated experiments.

Winkler's definition of the Sleeping Beauty problem asks "what – to her – is the probability that the coin came up Heads". We will explain the three perspectives one can take on this question in this section, first for the unmodified Sleeping Beauty problem, then illustrated by introducing a commemorative scarf, and finally using a general coin that might also be unfair.

4.1 Three perspectives

I Heads coin toss
: Beauty may want to obtain a probability for a Heads outcome of the coin toss *at the beginning of the experiment she is currently in.* Conditioning on A_L within any model (certain event, i.e., no new information), or conditioning on A_{day} within Model M_2 (also certain event, i.e. no new information), she will obtain probability 1/2. This perspective is equivalent to perspective II below.

II Heads awakening, "in experiment" perspective
: If Beauty's focus is on a probability for the *current awakening being* a Heads awakening, i.e. a probability for the current experiment having produced a Heads awakening, she can again employ Model M_2 and will obtain $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/2$ (see also I and Section 3.3).

III Heads awakening, "outcome population" perspective

If Beauty wants to obtain a probability for *the current awakening being* a Heads awakening, viewing it as part of the population of awakenings, she can decide to ignore the experimental setup and to be guided by the composition of the population of awakenings created by a series of experiments. She can use Model M_3 (or employ an indifference argument among awakenings in Model M_1). Thus, $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/(n+1)$ is obtained (1/3 for $n = 2$, thirders' perspective).

I and III highlight exactly the conflict that Luna (forthcoming) referred to. II is a variant of I that makes it appear much more similar to III: both II and III refer to the current awakening as the entity, on which Beauty is required to derive a probability for Heads. However, the perspective of looking at the awakening is different, as was explained by Groisman (2008) with an example of putting red and green balls into a box or taking them out again: he called the halfer variant “under the setup of coin tossing” and the thirder variant “under the setup of awakening”, which in this author's opinion does not fully capture the subtlety of the point. In the next subsection, a commemorative scarf is introduced as an auxiliary tool for illustrating the distinction between the “in experiment” perspective (I/II) and the “outcome population” perspective III. For simplifying communication, we will henceforth denote the probability of a Heads coin toss in the experiment (perspective I or II) as p_H, as opposed to the probability of sampling a heads awakening from the overall population of awakenings *from a series of experiments* (III), which we will denote as p_S.

In general, one might be interested in one or the other probability, depending on context. In the Sleeping Beauty problem, this author considers it rational for Beauty to answer Question 1 with p_S (and not with p_H), if she is in a series of experiments for which she cannot (or at least does not) focus on which experiment she is currently in but considers the series of awakenings only. Nevertheless, even in a series of experiments, as she knows she must be in a particular experiment, she might still apply the “in experiment” perspective.

4.2 A commemorative scarf

Consider a small inconsequential experimental condition (inspired both by Groisman's (2008) green and red balls and by Mutalik's (2016b) brass plaque): At the beginning of a series of N experiments, it is explained to Beauty that outside of her sleeping room, a commemorative patchwork scarf will be knitted for her to take home afterwards: the scarf will be extended by a square patch (say 10cm x 10cm, including a thin gold-colored rim) at each awakening, with the patch's interior colored green at a Heads awakening or red at a Tails awakening. The scarf will tell Beauty about the sequence of awakenings and (by deduction) experiments she went through; we will use the term scarf even if it only consists of one or two patches, e.g. when $N = 1$ and $n = 2$. Obviously, the existence of the commemorative scarf does not change the information available to Beauty during the experiment, because she will not get to see it until after the experiment. However, it does make clear the distinction between II and III, and can also be related to I by forming an expectation about the scarf's length. Figure 3 illustrates the three perspectives that Beauty can take, when considering the commemorative scarf.

I Heads coin toss

For a single experiment ($N = 1$) Beauty can think about the expected length of her scarf initially as $p_H + n(1 - p_H) = (n+1)/2$. For updating this expectation (even though the update should not modify her expectation), she needs to update her probability p_H for a Heads coin toss at the beginning of the experiment. (Beauty could apply the analogous reasoning for the expected number of patches contributed to the scarf by the current experiment in a series of $N > 1$ experiments.)

II Heads awakening, "in experiment" perspective

Beauty can derive a probability about whether the patch that is currently added to her scarf is green (equivalent to I, targeting probability p_H).

III Heads awakening, "outcome population" perspective

Beauty can derive a probability about hitting a green patch when randomly picking a patch out of the finished scarf (targeting probability p_S).

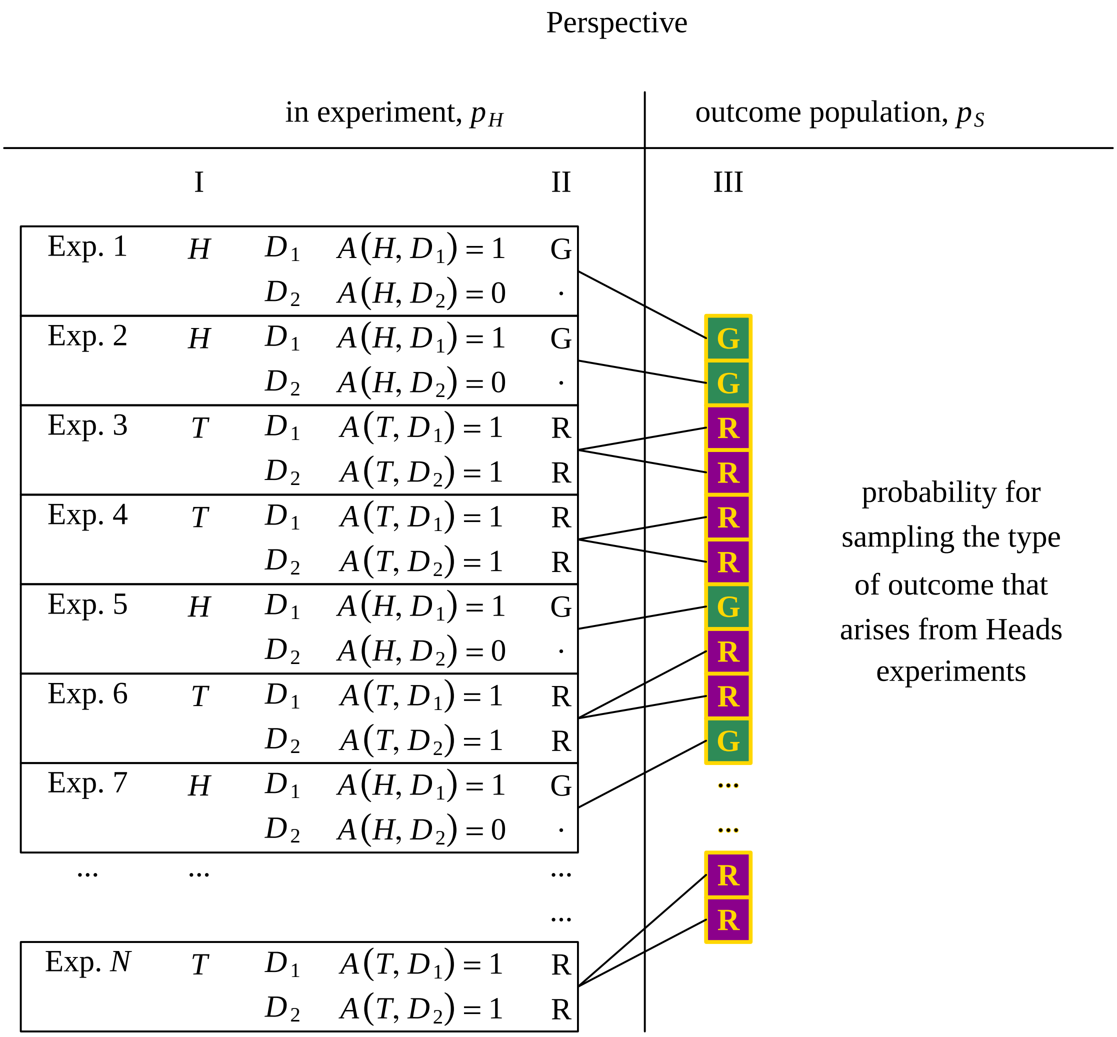

Figure 3: The three perspectives, illustrated with a sample series of experiments. G stands for a green patch, R for a red patch of the commemorative scarf.

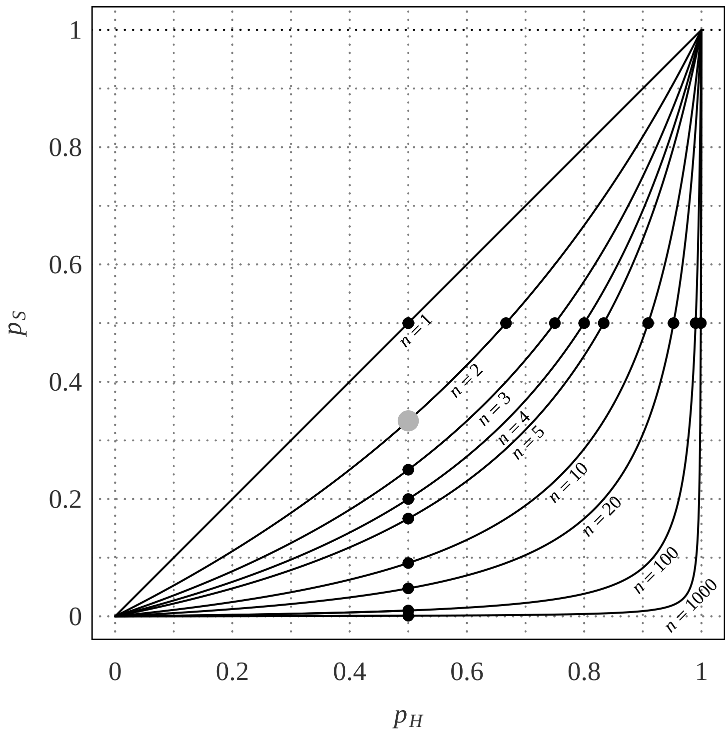

Figure 4: p_S values as functions of p_H values for different values of n. Large gray dot: the standard Sleeping Beauty problem with $p_H = 1/2$, $n = 2$, $p_S = 1/3$; black dots for $p_H = 1/2$: fair coin implies $p_S = 1/(n+1)$; black dots for $p_S = 1/2$: $p_H = n/(n+1)$ implies equally probable red and green patches in the commemorative scarf.

4.3 Sleeping beauty with a potentially unfair coin

This section stresses the general relation between the probability p_H of the coin to have shown Heads (I and II) and the probability p_S of an element sampled from the population of experimental outcomes to stem from a Heads experiment (i.e. a green patch, when sampling a random patch from the commemorative scarf, III). We now allow $p_H \in (0, 1)$, i.e. the coin may be arbitrarily unfair but not deterministic. The expected length of a scarf from N experiments is $N(p_H + n(1 - p_H))$, the expected number of green patches is Np_H, and the large sample proportion of outcomes from Heads experiments (i.e. Heads awakenings or green patches in the commemorative scarf) is

$$p_S = \frac{p_H}{p_H + n(1 - p_H)}. \tag{12}$$

Equation (12) can be solved for n: the duration of the experiment that transforms the Heads probability p_H into the sampling of Heads probability p_S is

$$n = \frac{p_H/(1 - p_H)}{p_S/(1 - p_S)}, \tag{13}$$

i.e., the ratio of the two odds. Thus, one would obtain $p_S = 1/2$ for $n = p_H/(1-p_H) \iff p_H = n/(n+1)$.

Figure 4 depicts the p_S values (“outcome population” perspective) as functions of the p_H values (“in experiment” perspective) for a selection of n values. Of course, $p_S = p_H$ is obtained for $n = 1$, while $n > 1$ implies $p_S < p_H$. $n = 2$ includes the classical Sleeping Beauty problem with $p_H = 1/2$ and $p_S = 1/3$, as shown by the large gray dot; for a fair coin ($p_H = 1/2$) and general n we obtain $p_S = 1/(n+1)$. $p_S = 1/2$ can be obtained for $p_H = n/(n + 1)$, as was pointed out above; for example, for $n = 2$, $p_H = 2/3$ would lead to equally probable red and green patches in the commemorative scarf. Note that p_S is *not* an update of p_H from gaining new knowledge because of an awakening, but the answer to a different version of Question 1.

All models can be generalized to an unfair coin: In Model M_1, the two trajectories have the potentially unequal probabilities p_H and $1 - p_H$; a hierarchical indifference reasoning within trajectories can still be applied. Because of the unequal probabilities for trajectories, an indifference reasoning between all $n + 1$ awakening variants is no longer justified for $p_H \neq p_S$. This also holds in Model M_3, for which the coin toss is combined with an independent pick of a day:

$$\begin{aligned} P(\mathsf{A}_i \cap \mathsf{T}) = P(\{(T, D_i)\}) = P(\mathsf{T} \cap \mathsf{D}_i) = (1 - p_H)/n \quad \text{for } i = 1, \ldots, n, \\ P(\mathsf{A}_i \cap \mathsf{H}) = P(\{(H, D_i)\}) = P(\mathsf{H} \cap \mathsf{D}_i) = p_H/n \quad \text{for } i = 1, \ldots, n, \\ P(\mathsf{A}_{\text{day}}) = 1 - p_H + p_H/n. \end{aligned} \tag{14}$$

Thus, Model M_3 yields $P(\mathsf{H}|\mathsf{A}_L) = p_H$ (trivial) and $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = P(\mathsf{H} \cap \mathsf{A}_{\text{day}})/P(\mathsf{A}_{\text{day}}) = p_H/(p_H + n(1 - p_H)) = p_S$. In a Bayesian framework, p_S can be considered as an update of p_H in Model M_3; within the Sleeping Beauty problem, however, p_S must be considered to be a change in perspective rather than an update from receiving new information, because Model M_3 neglects the experimental setup.

In Model M_2, an awakening is picked randomly, conditional on the coin toss result, which yields

$$\begin{aligned} P(\mathsf{A}_i \cap \mathsf{T}) = P(\{(T, D_i)\}) = P(\mathsf{T} \cap \mathsf{D}_i) = (1 - p_H)/n \quad \text{for } i = 1, \ldots, n, \\ P(\mathsf{A}_1 \cap \mathsf{H}) = P(\{(H, D_1)\}) = P(\mathsf{H} \cap \mathsf{D}_1) = p_H, \\ P(\mathsf{A}_i \cap \mathsf{H}) = P(\{(H, D_i)\}) = P(\mathsf{H} \cap \mathsf{D}_i) = 0 \quad \text{for } i = 2, \ldots, n. \end{aligned} \tag{15}$$

Consequently, $P(\mathsf{H}|\mathsf{A}_L) = P(\mathsf{H}|\mathsf{A}_{\text{day}}) = p_H$ remains correct for an unfair coin.

Thus, in perfect analogy to the fair coin, Model M_3 answers Question 1 under the "outcome population" perspective (III), while Model M_2 covers Question 1 under the "in experiment" perspective (I/II).

5 Treatment of Question 2 and relation to classical arguments

5.1 Treatment of Question 2

For Question 2, conditioning on $\mathsf{A}_1 = \mathsf{D}_1$ removes all doubts on what Beauty knows, and $\mathsf{A}_L \cap \mathsf{D}_1 = \mathsf{A}_{\text{day}} \cap \mathsf{D}_1 = \mathsf{D}_1$. It remains a valid question, whether Beauty wants to target p_H or p_S; however, the "in experiment" or "outcome population" perspective coincide for this special case, because this can be considered to be the $n = 1$ case (see Figure 4 in Section 4.3) for which $p_H = p_S$. These deliberations hold if it is known to Beauty that this question will be asked on D_1 of every experiment (as stated in various instances of the Sleeping Beauty problem).

- Model M_1, by restricting both trajectories to D_1, yields the unambiguous answer 1/2.
- Model M_3 yields the same answer, as $p_H = p_S = 1/2$.
- Model M_2 is not applicable to Question 2 in its usual form, in which the calendar day is disclosed to Beauty on D_1 of *every experiment* (see the discussion below).

When simply applying probability calculus within Model M_2,

$$P(\mathsf{H}|\mathsf{D}_1) = P(\mathsf{H} \cap \mathsf{D}_1)/P(\mathsf{D}_1) = \frac{n/(2n)}{(n+1)/(2n)} = \frac{n}{n+1}. \tag{16}$$

This probability (2/3 for $n = 2$) has been brought forward by some halfers, e.g. Lewis (2001). This result would be appropriate under the assumption that D_1 has been the outcome of a random pick of an awakening day (as is assumed in Model M_2) on which to disclose the calendar day to Beauty; such a random pick would have probability 1 in case of Heads and $1/n$ only in case of Tails. Bayes' Theorem then plausibly yields the posterior probability $n/(n + 1)$ for Heads as shown in Equation (16). However, if the calendar day is disclosed to Beauty on D_1 of *every experiment*, and if this is known to Beauty, she must abandon her model M_2 of a random pick of awakening day for answering Question 2. This does of course not invalidate the model for Question 1.

5.2 Relation to Elga (2000) and Lewis (2001)

In the notation of this paper, Elga's symmetry argument, as reported by Winkler (2017), is as follows:

- $P(\mathsf{H}|\mathsf{A}_1) = P(\mathsf{T}|\mathsf{A}_1)$, as we just agreed (answer to Question 2).
- By multiplying with $P(\mathsf{A}_1)$, one gets

$$P(\mathsf{H} \cap \mathsf{A}_1) = P(\mathsf{H}|\mathsf{A}_1)P(\mathsf{A}_1) = P(\mathsf{T}|\mathsf{A}_1)P(\mathsf{A}_1) = P(\mathsf{T} \cap \mathsf{A}_1).$$

- Furthermore, it is doubtlessly plausible that

$$P(\mathsf{T} \cap \mathsf{A}_1) = \cdots = P(\mathsf{T} \cap \mathsf{A}_n).$$

- The preceding two bullets can be combined into the indifference equation (11) of Model 3a, i.e. into the thirders' perspective.
- Elga concluded that each of the $n+1$ possible awakening combinations must receive probability $1/(n+1)$.

Elga's entire reasoning implicitly conditions on A_{day}. The second bullet is compatible with Model M_3 only, for which $P(\mathsf{A}_1) = 2 \cdot P(\mathsf{A}_i), \quad i > 1$. As was mentioned before, Model M_3 answers Question 1 from the "outcome population" perspective (which is the point made by Groisman 2008), i.e. provides the answer for Question 1 in terms of the "outcome population" probability $p_S = 1/3$.

Recently, the thirders' perspective seems to have become the dominant view, according to Winkler (2017). However, given the experimental setup, the double halfers position does have considerable clout, and we present a slightly shifted Lewis (2001) reasoning below, which is based on Model M_2:

- $P(\mathsf{H}) = P(\mathsf{T}) = 1/2$, $P(\mathsf{A}_{\text{day}}) = 1$, so that (of course) $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/2$ (answer to Question 1).
- $P(\mathsf{A}_1|\mathsf{H}) = 1$, which implies $P(\mathsf{A}_1 \cap \mathsf{H}) = 1/2$.
- $\sum_{i=1}^{n} P(\mathsf{A}_i|\mathsf{T}) = 1$, which implies $\sum_{i=1}^{n} P(\mathsf{A}_i \cap \mathsf{T}) = 1/2$; assuming conditional uniformity for awakening days, $P(\mathsf{A}_1|\mathsf{T}) = 1/n$, which implies $P(\mathsf{A}_1 \cap \mathsf{T}) = 1/(2n)$ ($1/4$ for $n = 2$).
- Consequently, $P(\mathsf{A}_1) = (n+1)/(2n)$ ($3/4$ for $n = 2$).
- Hence, $P(\mathsf{H}|\mathsf{A}_1) = n/(n+1)$ ($2/3$ for $n = 2$).

As was pointed out in the previous subsection, the last bullet is plausible only, if the awakening day is a random pick from the possible awakening days. *If this assumption is plausible to Beauty*, a rational Beauty should answer Question 2 with $n/(n+1)$. However, as was pointed out before, Question 2 of the Sleeping Beauty problem is usually assumed to be asked on every D_1 instance, and this assumption is also included by Lewis. If Beauty is aware of that fact, she should answer Question 2 with the probability $1/2$, e.g. by applying Model M_1.

6 Wrap-up

This section summarizes the relations of different solution patterns for the Sleeping Beauty problem to the models, assumptions and perspectives discussed in this paper.

Answer $1/2$ (or p_H) to both Question 1 and Question 2
: This arises under all models if Beauty assumes her knowledge to be A_L only. If she assumes her knowledge to be A_{day}, the answer to Question 1 arises by applying Model M_2 (or the corresponding hierarchical indifference reasoning within M_1, which does not require that she uses arguments outside of probability theory, see Section 3.2). For Question 2, Beauty has to realize that the assumptions for Model M_2 (random pick of awakening day, given coin toss) cannot be upheld; she can apply Model M_1 in a purely probabilistic fashion with trajectory shortened to D_1 (Monday) instead. Beauty remains within the "in experiment" perspective (I or II). This double-halfer pair of solutions appears the most plausible to this author.

Answer $1/2$ (or p_H) to Question 1 and $n/(n+1)$ ($2/3$ for $n = 2$, p_S in general) to Question 2
: This pair arises under Model M_2 (or the corresponding hierarchical indifference reasoning within M_1, which now requires the use of non-probabilistic indifference reasoning among Tails awakenings) if Beauty applies the model to both questions. This may be considered rational if she believes the day of Question 2 to be a random pick of awakening day, which is incompatible with most accounts of the Sleeping Beauty problem which involve Question 2. Beauty remains within the "in experiment" perspective (I or II).

Answer $1/(n+1)$ (1/3 for $n = 2$ or p_S in general) to Question 1 and 1/2 to Question 2
: This pair arises from Model M_3 (or the corresponding non-hierarchical indifference reasoning among awakenings within Model M_1). It refers to the "outcome population" perspective (III). Within Model M_3, the probability appears to arise from a Bayesian update. However, for the Sleeping Beauty problem, since the known experimental setup is ignored in the model assumptions, the approach is not an update of the (in experiment) probability under the Bayes theorem, but a switch in perspective (as explained in Section 4). Note that the answer to Question 2 arises in both perspectives, because $n = 1$ implies $p_H = p_S$.

Note that the thirders' answer could also be justified by Model M_3 in a modified Sleeping Beauty problem with a daily coin toss (see the end of Section 3.4), for which a rational Beauty could use conventional Bayesian conditioning for obtaining the probability $P(\mathsf{H}_i|\mathsf{A}_{\text{day}}) = 1/(n+1)$, with H_i denoting Heads on D_i. For $n = 2$, and with the coin toss occurring after the D_1 awakening (Figure 2), the daily coin toss may be what some thirders have in mind.

7 Final remarks

The Sleeping Beauty problem has been extensively discussed, especially in the philosophical literature. The author considers the objective probability side of the problem resolved by the present paper, building on simple probabilistic models in combination with the idea brought forward by Groisman (2008). The philosophical discussions will certainly continue, however: even the Monty Hall problem whose probability assessment has been resolved for a long time, continues to serve as an example setting for testing philosophical ideas.

The previous section summarized the understanding gained in this paper regarding the objective probability aspects of the Sleeping Beauty problem. The key steps in sorting out assumptions, models, and perspectives were

- to view the experiment as a series of events, as depicted in Figures 1 and 2; this understanding is supported by considering a general n instead of the default $n = 2$.
- to acknowledge Groisman's (2008) distinction between considering *a Heads coin toss on the experiment level* (I) equivalent to *being in a Heads awakening* in terms of its *production* (II, "in experiment" perspective) on the one hand, or *being in a Heads awakening* viewed from the population of awakenings (III, "outcome population" perspective), on the other hand. This distinction is best understood for repeated experiments and was illustrated by introducing the commemorative scarf.
- to explicitly incorporate sampling *of awakening days* conditional on the outcome of the initial coin toss (Model M_2, definitely applicable under Figure 1, but only with a stretch applicable under Figure 2) or *of laboratory days* independently of the coin toss, ignoring the experimental setup (Model M_3, equivalent to sampling from the outcome population).
- to explicitly distinguish between p_H for the "in experiment" perspective and p_S for the "outcome population" perspective, which is supported by also considering an unfair coin (Section 4.3, and especially Figure 4).

While probability calculus involves application of Bayes' theorem, the only instance of an actual update of probability because of new information *within the Sleeping Beauty problem* occurs for $P(\mathsf{H}|\mathsf{A}_1) = n/(n+1)$ (2/3 for $n = 2$), if one assumes a random pick of awakening day; however, as was discussed before, this assumption of Model M_2 is not compatible with most accounts of Question 2. $P(\mathsf{H}|\mathsf{A}_{\text{day}}) = 1/(n+1)$ according to Model M_3 formally also arises from applying Bayes' theorem. Within the Sleeping Beauty problem with a single coin toss at the beginning of the experiment, however, it is not an update because of new information of $P(\mathsf{H}) = p_H$ within the "in experiment" perspective, because Model M_3 neglects the information on the experimental setup; rather, the formal application of Bayes' theorem in Model M_3 must be seen as a change in perspective from p_H to $p_S = p_H/(p_H + n(1 - p_H))$ (see Section 4.3).

Various factors make it difficult to obtain unambiguous answers to the Sleeping Beauty problem: the natural Model M_1 does not provide a probability framework for all aspects of interest, and neither Model M_2 nor Model M_3 perfectly captures all aspects of the Sleeping Beauty problem. Things are further muddled by a lack of separation between awakening days and calendar days in many accounts. The disagreement between the "in experiment" and the "outcome population" perspectives for $n > 1$ laboratory days (Question 1), combined with the agreement between the two perspectives for $n = 1$ laboratory days (and thus for Question 2) also adds difficulty to the problem.

Many publications quoted the strong law of large numbers in support of the thirders' answer to Question 1, presumably with the repeated experiments for the "outcome population" perspective in mind. The considerations of this paper should make it obvious that the strong law of large numbers is not partial for a particular answer. If one correctly applies it under a set of assumptions, it yields the answer that is compatible with those assumptions. If the strong law of large numbers is identified with a large sample from Model M_1, subsequently calculating the proportion of Heads experiments among awakenings, the thirders' answer is indeed obtained, because this answer refers to the "outcome population" perspective (picking a patch from the finished commemorative scarf). From the "in experiment" perspective, however, *hierarchical* application of indifference arguments is asked for, which yields 1/2 and coincides with the model that conditions on the coin toss (see Model M_2). The fact that Beauty is currently in an awakening does not provide new information on the probability p_H of the coin toss of the current experiment having been Heads, even if one takes A_{day} as the given information (as opposed to A_L only). The thirders' answer can be justified by shifting attention to the different question of finding the probability p_S for picking a Heads produced element from the outcome population.

The simulation considerations of this paper highlight the well-known and important fact that a simulation's results crucially depend on the assumptions that underlie the model. For example, in Model M_3, the author's initial (mistaken) approach in the simulation was to discard all non-awakening days; this is fine, as long as one wants to estimate the probability for Heads underlying the current awakening from the "outcome population" perspective; for assessing the probability for Heads underlying the experiment (i.e., perspective I) from the same simulation, however, discarding the non-awakening elements of the sample space will (of course) create a massive selection bias against H. Simulation assumptions are of course crucial in practical problems as well, and not only when simulating for a decision theoretic problem. Thus, this work on the interesting and much-discussed but very theoretical decision-theoretic Sleeping Beauty problem can also serve as a reminder of the importance of assumptions for applied statisticians who use simulation in their work.

8 Acknowledgment

Discussions with my colleague Tim Downie were very helpful, and his comments on earlier drafts led to substantial improvements of this paper.

9 References

Bostrom, N. (2007). Sleeping Beauty and Self-location: A Hybrid Model. *Synthese* **157**, 59-78.

Bradley, D.J. (2011). Self-location is no problem for conditionalization. *Synthese* **182**, 393-411.

Cisewski, J., Kadane, J. B., Schervish, M. J., Seidenfeld, T. and Stern, R. (2016). Sleeping Beauty's credences, *Philosophy of Science* **83**, 324-347.

Draper, K. (2017). Even for Objectivists, Sleeping Beauty Isn't So Simple. *Analysis* **77**, 29-37.

Elga, A. (2000). Self-Locating Belief and the Sleeping Beauty Problem, *Analysis* **60**, 143-147.

Groisman, B. (2008). The end of Sleeping Beauty's nightmare. *British Journal for the Philosophy of Science* **59**, 409-416.

Lewis, D. (2001). Sleeping Beauty: Reply to Elga, *Analysis* **61**(3), 171-176.

Luna, L. (forthcoming). Sleeping Beauty: Exploring a Neglected Solution. *British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axy045.

Mann, A.L. and Aarnio, V. (2018). A Logical Analysis of Monty Hall and Sleeping Beauty. *Studia Logica* **106**, 1123-1162.

Mutalik, P. (2016a). Solution: 'Sleeping Beauty's Dilemma'. *QUANTA* (January 29, 2016). https://www.quantamagazine.org/solution-sleeping-beautys-dilemma-20160129

Mutalik, P. (2016b). Why Sleeping Beauty is lost in time. *QUANTA* (March 31, 2016). http://quantamagazine.org/20160331-why-sleeping-beauty-is-lost-in-time.

Oscar Seminar (2008). An objectivist argument for thirdism. *Analysis* **68**, 149-155. Selvin, S. (1975). A Problem in Probability (Letter to the Editor). *The American Statistician* **29**, 67.

Winkler, Peter (2017). The Sleeping Beauty Controversy. *American Mathematical Monthly* **124**, 579-587.