Sleeping Beauty goes to the lab: The psychology of self-locating evidence

Matteo Colombo

Tilburg Center for Logic, Ethics, and Philosophy of Science

Tilburg University

Jun Lai

Tilburg Center for Logic, Ethics, and Philosophy of Science

Tilburg University

Vincenzo Crupi

Center for Logic, Language, and Cognition

University of Turin

Sleeping Beauty goes to the lab: The psychology of self-locating evidence

Abstract. Analyses of the Sleeping Beauty Problem are polarised between those that advocate the "1/2 view" and those that advocate the "1/3 view." One source of disagreement between advocates of different views concerns the evidential relevance of self-locating information. Unlike halfers, thirders regard self-locating information as evidentially relevant in the Sleeping Beauty Problem. The present study advances the debate, providing a more nuanced and empirically grounded account of the evidential impact of self-locating information. By systematically manipulating the kind of information available in different formulations of the Sleeping Beauty Problem, we show that human reasoners acknowledge self-locating evidence as relevant, but discount its weight. This indicates that patterns of judgment on different formulations of the Sleeping Beauty Problem do not fit either the "1/2 view" or the "1/3 view." Our results suggest that an adequate explication of the evidential relevance of self-locating information should take into account that self-locating information may trigger more cautious judgments of confirmation than familiar kinds of statistical evidence.

Keywords: sleeping beauty problem; probability; reasoning; self-locating evidence.

Sleeping Beauty goes to the lab: The psychology of self-locating evidence

Introduction

The Sleeping Beauty Problem (SBP) is a challenging puzzle in probabilistic reasoning. It raises questions of unsuspected theoretical relevance for the foundations of probabilistic reasoning, belief update, decision-making, and beyond (Titelbaum, 2013).

In its standard formulation, the problem goes as follows:

> On a Sunday, some researchers are going to put you to sleep. During the two days that your sleep will last, they will briefly wake you up either once or twice, depending on the toss of a fair coin (Heads: once; Tails: twice). After each waking, they will put you back to sleep with a drug that makes you forget that waking. When you are awakened, to what degree ought you believe that the outcome of the coin toss is Heads? (Elga, 2000; see Piccione & Rubinstein, 1997 Example 5 for the original formulation)

Opinions on this puzzle are split between two camps. For so-called *halfers*, your credence in Heads should be 1/2 on the probability scale. At the outset, you know all the details of the experiment, including that the coin is fair and that you will lose your memory of an earlier awakening. When you wake up, all new information you have is the self-locating information 'I am awake now.' Self-locating information (aka 'centred information') concerns one's spatio-temporal location in the world or one's identity—such as the information 'Today is Monday or Tuesday,' or 'I am Jun.' Instead, non-self-locating information (aka 'uncentred information') — such as the information 'The coin landed Tails' or 'Jun was born in Nanjing' — concerns what the world is like. Self-locating information bears no relevant connection to the outcome of the coin flip, according to the halfer. Since before the experiment you know that the coin is fair, you should then retain a credence of P(Heads) = P(Tails) = 1/2 (Lewis, 2001; Arntzenius, 2002; Cozic, 2011; Hawley, 2013).

So-called *thirders* disagree and submit that your credence in heads should be 1⁄3. The pair {It is now Monday} and {It is now Tuesday} partitions your space of centred possibilities. When you wake up, the pieces of information on which you can rely are 'I am awake and it is now Monday' and 'I am awake and it is now Tuesday.' Since these pieces of information are jointly exhaustive and mutually exclusive, you may have three distinct self-locating beliefs:

H1: It is now Monday and the fair coin landed Heads.

T1: It is now Monday and the fair coin landed Tails.

T2: It is now Tuesday and the fair coin landed Tails.

By the law of total probability, it follows that P(Heads) = P(Heads | It is now Monday) x P(It is now Monday) + P(Heads | It is now Tuesday) x P(It is now Tuesday) = 1/2 x 2/3 + 0 x 1/3 = 1/3. Your credences should then be P(Heads) = 1/3 and P(Tails) = 2/3 on any particular awakening (Elga, 2000; Dorr, 2002; Weintraub, 2004; Titelbaum, 2008).

Disagreement on SBP persists, partly because it is unclear whether or not self-locating information is *evidentially relevant* to credences about uncentred hypotheses. Thirders believe that learning new self-locating information in the SBP impacts your credences about the outcome of the coin flip (e.g., Horgan 2004; Weintraub 2004; Titelbaum 2008; Draper 2013). Halfers believe that self-locating information is evidentially irrelevant to uncentred hypotheses; and so it has no evidential impact on your credences about the outcome of the coinflip (e.g., Lewis 2001; White 2006; Bradley 2012; Hawley 2013).

Epistemologists and philosophers of science have developed several probabilistic explications of the concept of *evidential relevance* (Fitelson 1999; Crupi 2015). One way of evaluating the adequacy of these alternative explications is to examine their degree of similarity to ordinary usage and judgment (Carnap 1950; Kemeny & Oppenheim 1952). To the extent that an explication will only illuminate a concept if it fits central cases of ordinary usage and judgment, empirical results from the psychology of reasoning bear on philosophical questions about the evidential relevance of self-locating information in the SBP (cf., Schupbach 2015; Colombo 2016; Tentori et al. 2007). Psychological results will provide philosophers with data helpful to discover and assess instances of concept pluralism underlying an explication; they can help philosophers to identify the explicandum's central features and their relation with other concepts; and they can point to sources of bias affecting philosophers' judgments about such an unusual case as the SBP (Shepherd & Justus 2014).

With the studies reported here, we bring Sleeping Beauty to the lab for the first time, so as to gain a more nuanced understanding of the evidential impact of self-locating information. Our results indicate that an adequate explication of the *evidential relevance of self-locating information* should take into account that self-locating information may trigger more cautious judgments of confirmation than familiar kinds of statistical evidence.

Overview of the experimental scenarios

For our study, we adapted the standard SBP to make it as transparent as possible to naïve participants, as follows.

BASIC version

On a Sunday, you will be administered one of two pills, depending on the toss of a fair coin (HEADS: regular pill; TAILS: strong pill). You will not be told the outcome of the coin toss, and the two pills look identical. However, you know the following.

If the coin landed HEADS:

– the pill you're given on Sunday (regular pill) will make you sleep for one day;

– then you will wake up (on Monday).

If the coin landed TAILS:

– the pill you're given on Sunday (strong pill) will make you sleep for one day;

– then you will wake up a first time (on Monday), and shortly afterwards fall back asleep for another day, forgetting that you just woke up;

– then you will finally wake up a second time (on Tuesday).

Imagine you've just woken up. You don't know which day it is, and you do not know whether or not you have already woken up before. You are now asked to express your belief about the outcome of the coin toss that was made on Sunday.

As anticipated, the halfer's and thirder's predictions diverge critically in this basic version. According the thirder, one should judge P(Heads) = 1/3 and P(Tails) = 2/3. According to halfer, instead, the correct answer here is P(Heads) = P(Tails) = ½, just as in the following *No Evidence* version, which we employed as a control condition.

NO EVIDENCE version

[same introductory paragraph as above]

If the coin landed HEADS:

– the pill you're given on Sunday (regular pill) will make you sleep for one day;

– then you will wake up (on Monday).

If the coin landed TAILS:

– the pill you're given on Sunday (strong pill) will make you sleep for two days;

– then you will wake up (on Tuesday).
Imagine you've just woken up. You don't know which day it is. You are now asked to express your belief about the outcome of the coin toss that was made on Sunday.

As *P*(Heads) = *P*(Tails) = ½ is uncontroversially the correct response in this version, participants' judgments should differ between the *Basic* vs. *No Evidence* condition in case they reason as thirders and self-locating information has impact on their credences. In order to gain a more fine grained understanding of the evidential impact of self-locating information, and further disentangle halfers' and thirders' predictions, we relied on yet another benchmark variant, where ordinary (non-self-locating) and evidentially relevant information was involved.

MARBLE version

On a Sunday, five small, empty, and closed boxes are placed in front of you; and you will then be administered one of two pills, depending on the toss of a fair coin (HEADS: regular pill; TAILS: strong pill). You will not be told the outcome of the coin toss, and the two pills look identical. However, you know the following.
If the coin landed HEADS:
– the pill you're given on Sunday (regular pill) will make you sleep for one day;
– meanwhile, one of the five boxes will be filled with a marble, then closed again (the other four boxes remain closed and empty);
– then you will wake up (on Monday), and open one of the five boxes at random.
If the coin landed TAILS:
– the pill you're given on Sunday (strong pill) will make you sleep for two days;
– meanwhile, all five boxes will be filled with five marbles (one each), then closed again;
– then you will wake up (on Tuesday), and open one of the five boxes at random.
Imagine you've just woken up. You don't know which day it is. You open one of the five boxes at random: you find a marble. You are now asked to express your belief about the outcome of the coin toss that was made on Sunday.

Like in the *No Evidence* version, the halfer's and the thirder's analyses converge in this case, because *P*(Tails) = 5/6 (computed as (1*1/2) / [(1*1/2) + (1/5 * 1/2)]). Instead, they critically diverge as concerns the comparison with the following counterpart version:

PLUS version

On a Sunday, you will be administered one of two pills, depending on the toss of a fair coin (HEADS: regular pill; TAILS: strong pill). You will not be told the outcome of the coin toss, and the two pills look identical. However, you know the following.

If the coin landed HEADS:

– the pill you're given on Sunday (regular pill) will make you sleep for one day;

– then you will wake up (on Monday).

If the coin landed TAILS:

– the pill you're given on Sunday (strong pill) will make you sleep for one day;

– then you will wake up a first time (on Monday), and shortly afterwards fall back asleep for another day, forgetting that you just woke up;

– the same will happen on each of the following days, until you finally wake up a fifth time (on Friday).

Imagine you've just woken up. You don't know which day it is, and you do not know whether or not you have already woken up any time before. You are now asked to express your belief about the outcome of the coin toss that was made on Sunday.

The *Plus* version tries out the halfer's intuition further. According to the halfer, the switch from two to five awakenings (or ten, for that matters) would still leave the self-locating evidence irrelevant, so that judgments in the *Plus* condition are expected to differ from the *Marble* condition but not from the *Basic* (and *No Evidence*) condition. An opposite prediction arises from the thirder's analysis: responses in the *Plus* version should line up with those in the *Marble* version and differ from both the *Basic* and the *No Evidence* variants.

Table 1. Comparison between the Halfers' and the Thirders' judgments about the probability that the coin landed TAILS in different versions of the SBP.

	No Evidence	Basic	Plus	Marble
Halfers	½	1/2	1/2	5/6
Thirders	½	2/3	5/6	5/6

In summary, we experimentally manipulated the SBP's formulation as a function of the kind of evidence available to human reasoners (Table 1). We thereby addressed three questions with our study: whether human reasoners acknowledge self-locating information as evidentially relevant in the SBP, whether the impact of self-locating information differs from the impact of objective statistical information like finding a marble in a randomly chosen box, and whether the standard halfer and thirder accounts are empirically adequate.

Study 1

***Method*.** Two hundred and forty-one participants (Mean age, 35, SD = 10, male 137, female 104) were recruited using Amazon MTurk. We only allowed MTurk workers with an approval rate > 95% and with a number of HITs approved > 5000 to participate in our study. Instructions and material were presented in English on the Qualtrics Survey Software. Participants were randomly assigned to one of four experimental groups, and none took part in more than one experiment.

As pointed out above, four conditions were sufficient to disentangle relevant predictions from standard halfer and thirder accounts, thus putting them to empirical test. Halfers and thirders agree on their predictions that *P*(Tails) = 1/2 in the *No Evidence* condition, and that *P*(Tails) = 5/6 in the *Marble* condition. For the *Basic* and *Plus* conditions, instead, halfers and thirders disagree (Cisewski *et al*., 2016; see also Ross 2010). Halfers predict that *P* (Tails) = 1/2 in both the *Basic* and the *Plus* condition, since they claim that self-locating information bears no relevance relation with the outcome of the coin toss. Instead, thirders predict that *P* (Tails) = 2/3 in the *Basic* condition, and that *P* (Tails) = 5/6 in the *Plus* condition as a function of the partition of centred possibilities.

Participants read one version of the SBP, and were asked to express their belief about the outcome of the coin toss in the situation described in one vignette. Participants' beliefs were collected on a 7-point Likert-scale ranging from 'Certain that it was Heads and not Tails' to 'Certain that it was Tails and not Heads.' Finally, participants were asked to indicate their age, sex, and level of education.

***Results*.** A Kruskal-Wallis test showed that each of our four manipulations had a significant effect on participants' judgment, χ^2 (3) = 28.76, *p* = .000. Across conditions, we also found significant differences concerning the degree of certainty that the outcome of the coin toss was Tails (ranging in 4-7, i.e. from "equally likely" to "certain"), χ^2 (3) =60.16, *p* = .000. A Dunn's test was further performed to test all possible pairwise comparisons between the four different conditions. The results showed that difference of scores between *Basic* & *Marble*, *Marble* & *No Evidence*, *Marble* & *Plus*, and *No Evidence* & *Plus* were significant ($p<0.01$). We found no effect of age, sex, or education.

Study 2

***Method*.** Study 1 revealed that participants' judgments of the SBP depended on the type of evidence available. In particular, its results are consistent with the idea that naïve reasoners acknowledge self-locating information as relevant in the SBP (*Plus* condition). Study 2 examined whether these results may have been affected by a focus on the coin mechanism in the question participants were asked.

A new sample of two hundred and thirty-seven participants (Mean age 34, SD = 10, male 137, female 100) was recruited from Amazon MTurk. As in Study 1, we only allowed MTurk workers with an approval rate > 95% and with a number of HITs approved > 5000 to participate in our study. Instructions and material were presented in English on the Qualtrics Survey Software. Participants were randomly assigned to one of four experimental groups, and none took part in more than one experiment.

Unlike in Study 1, participants did not express their belief about the outcome of the coin toss. Instead, participants expressed their belief about the pill they were administered in the situation they were asked to consider. Otherwise, the versions of the SBP used in this second study were identical to the versions we used in Study 1. Responses were again collected on a 7-point Likert-scale and participants again provided their age, sex, and level of education.

***Results*.** A Kruskal-Wallis test showed that all groups differed significantly in their answers, χ^2 (3) =17.43, p = .001. All groups differed significantly in their certainty of a Tails outcome, χ^2 (3) =29.18, p = .000. A Dunn's test was further performed to test all possible pairwise comparisons between the four different conditions. The results showed that difference of scores between *Basic* & *Marble*, and *Marble* & *No Evidence* were significant ($p<0.01$). Like in Study 1, we found no effect of age, sex, and education.

A Mann-Whitney Test showed that there was no significant difference between the answers of the participants of this study (M = 4.31) and the answers of participants from Study 1 (M = 4.24), p =

.83. Aggregating data from both studies, the difference between the *Basic* and the *No Evidence* condition did not reach significance, *p* = .42. However, a significant difference was found between the *Plus* and the *Basic* condition, *p* = .03, *d* = .29, and between the *Marble* and the *Plus* condition, *p* = .03, *d* = .25 (Table 2 and Figure 1).

Table 2. Mean scores for each group (scores ranging from 1 to 7).

Conditions	**No Evidence**		**Basic**		**Plus**		**Marble**	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Coin (Study 1, n = 241)	3.83	.78	3.98	1.05	4.40	1.15	4.87	1.60
Pill (Study 2, n = 237)	3.92	1.04	3.90	1.28	4.23	1.71	4.82	1.54
Combined (1&2, n = 478)	3.87	.91	3.94	1.17	4.32	1.46	4.85	1.57

Figure 1. Mean scores for each group (scores ranging from 1 to 7)

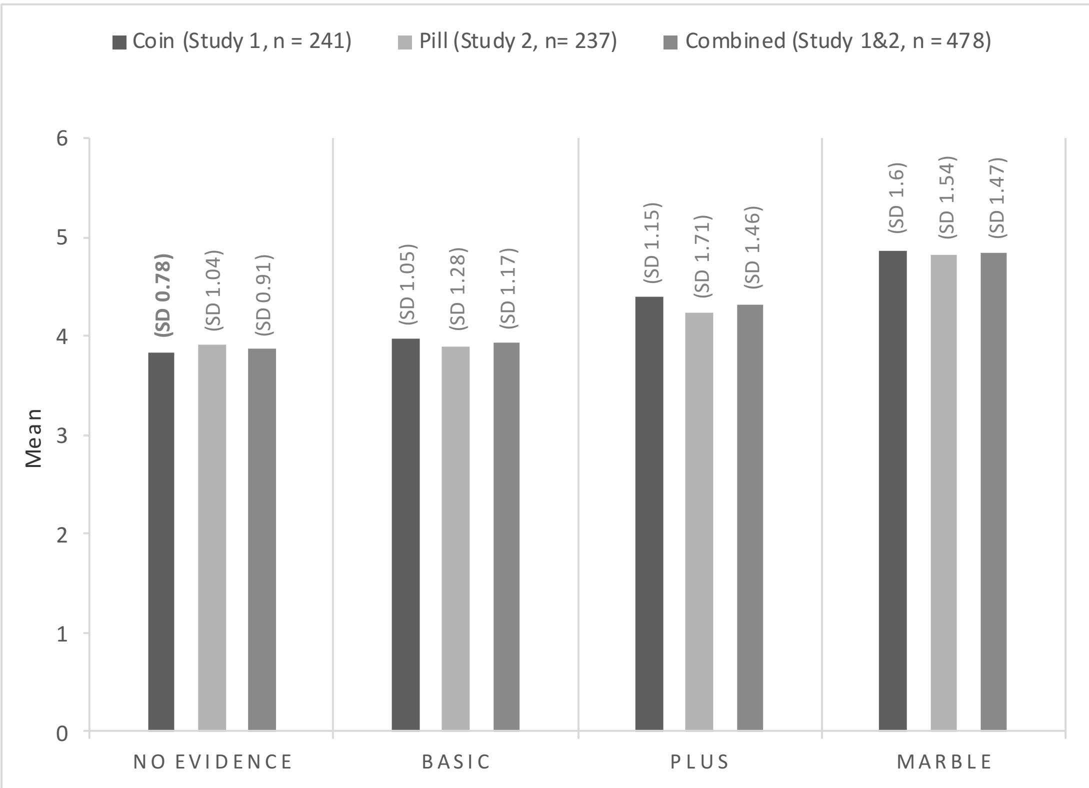

Discussion

Our results show that experimentally observed reasoning in the SBP does not simply fit either the halfer or the thirder analyses. The halfer's analysis is consistent with the lack of a significant difference between the *Basic* and *No Evidence* conditions, but is at odds with the finding that the *Plus* and *Basic* conditions reliably differed (recall that for the halfer one should have *P*(Tails) = 1/2 in both cases). The thirder's analysis, on the other hand, is consistent with the latter result, but is inconsistent with the finding that the probability of Tails is reliably judged to be higher in the *Marble* than in the *Plus* condition (for the thirder, one should have *P*(Tails) = 5/6 in both cases).

Because in some of the conditions the variance was relatively high, we examined whether we could distinguish a group of halfers and a group of thirders among the participants. If our participants actually included these two groups, then the pattern of responses across critical conditions like *Basic* and *Plus* should be bimodally distributed. However, we did not find indications of a bimodal distribution in any of the conditions.

Furthermore, our results seem to rule out the possibility that differences in how people reason about SB do not depend on differences in how they treat self-locating information, but instead depend on different people having differing models for the type of evidence that SB will observe (cf. Cisewski et al. 2016). If the significant difference that we found between the *Basic* and *Plus* conditions only depended on the presence of one group of halfers and one group of thirders among our participants, who would have differing models for the type of evidence that SP will observe, then the variance of the responses in the *Plus* condition should be higher than the variance in the *Basic* condition. After all, halfers should respond 1/2 in both conditions, while thirders should respond 2/3 in *Basic* and 5/6 in *Plus*. However, if there are actually two groups among our participants, then the variance in the *Plus* condition should also be systematically higher than the variance of the responses in the *Marble* condition, because, in this latter condition, both halfers and thirders agree on the 5/6 response. Yet again, this is not what we found: the variance of participants' responses did not systematically increase from the *Marble* to the *Plus* conditions. So, the hypothesis that our participants included a group of halfers and a group of thirders does not seem to account for our results.

Given that no standard theoretical analysis accounts for observed behavior, one might be tempted to complement a thirder framework with an appeal to cognitive limitations akin to those arising in other known puzzles of probabilistic reasoning. In particular, from a thirder's perspective, the SBP may seem structurally similar to the Monty Hall problem. And Monty Hall is known to invite 1/2 as a largely dominant response because of the representational and computational difficulty of the task for the unaided human mind (Krauss & Wang, 2003). Although appealing, this remark is not sufficient to explain our results either. After all, thirders could easily see that the *Marble* and the *Plus* variants have the same mathematical structure, with the number of marbles corresponding to the number of awakenings; and so, they will provide similar judgments in both variants. However, this is not what we found. Our participants' judgments differed significantly across the *Marble* and *Plus* condition, which suggests that our findings cannot simply be explained in terms of representational difficulty of the task like in the Monty Hall problem.

A thirder would also be unable to explain our results by relying solely on a general tendency to *conservatism* in probability updating, according to which people's belief updating is generally conservative relative to the predictions of Bayesian conditioning (Phillips & Edwards 1966; Edwards 1968; Fischoff & Beyth-Marom, 1983; Slovic & Lichtenstein, 1972). The provision of new evidence would then have less impact on people's belief updating than what Bayesian conditioning predicts. But this general tendency to conservativism in probability updating still does not explain the difference we found between judgments in the *Marble* and *Plus* conditions, which were structurally analogous from a general probabilistic point of view.

Consistent with our results is instead the explanation that people show a tendency to conservativism in specific settings, only with respect to specific kinds of evidence. The idea is that our participants extracted less certainty from the available self-locating evidence in comparison to the "observed" evidence of the marble in a randomly chosen box. This conservativism in belief updating with self-locating information can be further explained in a number of ways (cf., Edwards 1968). Participants could have mis-perceived the true evidential impact of self-locating information, discounting its weight more than "observed" evidence. Or, although the weight of self-locating information was perceived in a similar way as the weight of non-self-locating information, participants might have applied an adjusted combination rule to self-locating information when revising their beliefs.

Source reliability provides another factor that can explain our results in a way that is consistent with Bayesian belief update. As Corner, Harris & Hahn (2010) suggest, in situations where information is received from some source, people should not update their beliefs as much as if they had directly observed the evidence. This is because it is rare that people receive information from sources that are fully reliable; and it is even more rare that people produce themselves fully reliable information, on which they update their beliefs. Consequently, self-locating information produced by oneself could be diluted to some degree, in line with a Bayesian analysis of source reliability (Corner et al. 2010).

In summary, our results are consistent with a pattern of judgment that is qualitatively different from either the halfer or thirder analyses, where self-locating evidence is acknowledged as relevant but its quantitative impact is discounted significantly as compared to more standard statistical evidence. Other factors were previously shown to have such diluting effects on reasoning with evidence, such as second-order uncertainty about the values of a relevant statistical distribution (Tentori, Crupi, & Osherson, 2010). Although "mixed" models of the SBP exist (Bostrom, 2007; Meacham 2008), they do not take into account this particular diluting, conservative effect involved in reasoning with self-locating information. Our results will then contribute to put analyses of the SBP, and more generally of probabilistic reasoning with self-locating information, on plausible descriptive grounds.

References

Arntzenius, F. (2002). Reflections on Sleeping Beauty. *Analysis* 62(1): 53-62.

Bostrom, N. (2007). Sleeping Beauty and self-location: A hybrid model. *Synthese*, *157* (1), 59-78.

Bradley, D. (2012). Four problems about self-locating belief. *Philosophical Review* 121: 149–177.

Carnap, R. (1950). *Logical Foundations of Probability*. University of Chicago Press, Chicago.

Cisewski, J., Kadane, J.B., Schervish, M.J., Seidenfeld, T., & Stern, R. (2016) Sleeping Beauty's Credences. *Philosophy of Science* 0; 0(ja). DOI: 10.1086/685741

Colombo, M. (2016). Experimental Philosophy of Explanation Rising: The case for a plurality of concepts of explanation. *Cognitive Science*. DOI: 10.1111/cogs.12340

Corner, A. J., Harris, A. J. L., & Hahn, U. (2010). Conservatism in belief revision and participant skepticism. In R. Camtrabone & S. Ohlsson (Eds.), *Proceedings of the 32nd Annual Meeting of the Cognitive Science Society* (pp. 1625–1630). Portland, OR: Cognitive Science Society.

Cozic, M. (2011). Imaging and Sleeping Beauty: A case for double-halfers. *International Journal of Approximate Reasoning*. 52: 137–143.

Crupi, V. (2013), "Confirmation", in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2015 Edition). Url: http://plato.stanford.edu/archives/fall2015/entries/confirmation/.

Crupi V., Tentori K., & Gonzalez M. (2007). On Bayesian measures of evidential support: Theoretical and empirical issues. *Philosophy of Science*, 74: 229-252

Dorr, C. (2002). Sleeping Beauty: In Defense of Elga. *Analysis* 62: 292-296

Draper, K. (2013). The evidential relevance of self-locating information. *Philosophical studies*, *166*(1), 185-202.

Edwards, W. (1968). Conservatism in Human Information Processing. In B. Kleinmuntz (Ed.), *Formal Representation of Human Judgment* (pp. 17–52). New York, NY: Wiley.

Elga, A. (2000). Self-locating Belief and the Sleeping Beauty problem. *Analysis, 60* (2): 143-147.

Fischoff, B & Beyth-Marom, R. (1983). Hypothesis Evaluation from a Bayesian Perspective. *Psychological Review* 90 (3) 239-260.

Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science*, 66: S362–S378.

Hawley, P. (2013) Inertia, Optimisim, and Beauty. *Noûs* 47:1 85-103.

Horgan, T. (2004). Sleeping Beauty awakened: new odds at the dawn of the new day. *Analysis*. 64: 10–21.

Kemeny, J. G. & Oppenheim, P. (1952). Degree of factual support. Philosophy of Science, 19: 307–324.

Krauss, S., & Wang, X.T. (2003). The psychology of the Monty Hall problem: discovering psychological mechanisms for solving a tenacious brain teaser. *Journal of Experimental Psychology: General*, *132* (1), 3-22.

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

Meacham, C. (2008). Sleeping Beauty and the Dynamics of *De se* Belief. *Philosophical Studies* 138:2, 245-269.

Phillips, L. D., & Edwards, W. (1966). Conservatism in a simple probability inference task. Journal of Experimental Psychology, 72, 346–354.Piccione, M., & Rubinstein, A. (1997). On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, *20* (1), 3-24.

Ross, J. (2010). Sleeping Beauty, countable additivity, and rational dilemmas. *Philosophical Review*, *119*(4), 411-447.

Schupbach, J. N. (2015). Experimental Explication. *Philosophy and Phenomenological Research*.

Slovic, P. & Lichtenstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgement. *Organizational Behavior & Human Processes* 6, 649-744.

Tentori, K., Crupi, V., Bonini, N., & Osherson, D. (2007). Comparison of confirmation measures. *Cognition*, *103*, 107-119.

Tentori, K., V. Crupi, & Osherson, D. (2010). Second-order Probability Affects Hypothesis Confirmation. *Psychonomic Bulletin & Review*, *17*, 129–34.

Titelbaum, M.G. (2013). Ten reasons to care about the Sleeping Beauty problem. *Philosophy Compass*, *8* (11), 1003-1017.

Titelbaum, M. G. (2008). The relevance of self-locating beliefs. *Philsophical Review* 117: 555–605.

Weintraub, R. (2004). Sleeping Beauty: A Simple Solution. *Analysis* 64(1): 8-10.

White, R. (2006). The generalized Sleeping Beauty problem: a challenge for thirders. *Analysis*, *66*(290), 114-119.