

*Truth and (Sleeping) Beauty: all ye need to know.*

Beauty is a rational agent participating in an experiment in which a coin is tossed on Sunday night. She is awakened Monday morning, asked her credence in *heads* and told the outcome. If the coin landed heads, that's the end of the experiment. If *tails*, Beauty is given a drug that puts her to sleep for another 24 hours and erases all memory of her Monday awakening. Then on Tuesday morning she is again awakened, asked her credence in *heads*, told the outcome, and that's the end of the experiment. The problem is what Beauty's personal credence in *heads* should be on Monday morning. A *halfer* says one-half. A *thirder* says one-third. I maintain that Beauty's credences are underspecified, and will defend versions of *both* solutions.

1. *The naive solution: thirding.*

I will begin this section with a positive argument for thirding, then address some recently raised doubts about thirding. The positive argument is intended to be merely representative, though it may have no exact extant counterpart. Most halfers, including Lewis [6], have maintained that Beauty gains no information from Sunday night to Monday morning. Many thirders, for example Horgan [4], think that Beauty does gain information, and that this is the basis for her change in credence. I hold with those who maintain that Beauty loses (*de se*) information.<sup>1</sup>

A reflection principle attributed to Van Fraassen states that if at some future time (after conditioning on some partition, to be more accurate—more about this later) one has expected credence  $r$  in event  $E$ , then they should have credence  $r$  in  $E$  *now*. I'll be employing for my argument an inverse reflection principle which says that if one has *lost* information (again subject to some "partition" condition) potentially relevant to  $E$ , and knows their prior credence in  $E$ , they should revert to the prior credence.

The argument runs as follows. Let  $x$  be Beauty's credence in *heads* during an awakening. I will assume that Beauty accepts Elga's "restricted principle of indifference" [1], which says that  $P(\text{Monday tails})$  and  $P(\text{Tuesday tails})$  share a common value, here  $\frac{1}{2}(1 - x)$ . She knows she has lost information regarding when the coin was tossed, and suspects that this information is potentially relevant to *heads*. By the inverse reflection principle, her credences should revert to what they were prior to the loss. If it is Monday, to those of Sunday night—namely  $\frac{1}{2}$ , and if it is Tuesday, to those of (on the naive view) Monday night—namely zero. If Beauty doesn't know what day it is, she should use the weighted average  $\frac{1}{2} \cdot P(\text{Monday}) + 0 \cdot P(\text{Tuesday}) = \frac{1}{2}(\frac{1}{2}(1 + x)) = \frac{1}{4} + \frac{1}{4}x$ . On the other hand her credence in *heads* is  $x$ , so we get  $x = \frac{1}{4} + \frac{1}{4}x$ , which has solution  $x = \frac{1}{3}$ .

---

<sup>1</sup>On Sunday night Beauty knows that the coin was flipped in the previous 24 hours. On Monday morning, she doesn't. At first blush the information lost appears to be at least potentially relevant to *heads*, for conditioned on the experiment still going on, credence in *heads* clearly depends on whether or not the coin was flipped in the previous 24 hours.

So much for the positive argument; I move now to the defense. Ross in [8] claims that the one-third solution is at odds with the principle that rational credences should be countably additive. I believe his argument is essentially correct, but will advise that the conclusion has more limited scope than has been acknowledged. At nomically accessible worlds, in particular, rational thirders should not feel threatened.

Ross begins by defining a general “Sleeping Beauty problem” to be “a problem in which a fully rational agent, Beauty, will undergo one or more mutually indistinguishable awakenings...” where the number of such awakenings is a function of a discrete random variable into a set  $S$  of hypotheses. His argument then proceeds by means of a pathological Sleeping Beauty problem in which the expected number of awakenings is infinite, starting with the claim that thirders are committed to the following “indifference principle”:

*Finitistic Sleeping Beauty Indifference (FSBI).* In any Sleeping Beauty problem, for any hypothesis  $h$  in  $S$ , if the number of times Beauty awakens conditional on  $h$  is finite, then upon first awakening, Beauty should have equal credence in each of the awakening possibilities associated with  $h$ .

*FSBI*, together with some additional premises, leads to the following:

*Generalized Thirder Principle (GTP).* In any Sleeping Beauty problem, upon first awakening, Beauty’s credence in any given hypothesis in  $S$  must be proportional to the product of the hypothesis’ objective chance and the number of times Beauty will awaken conditional on this hypothesis.

The pathological example to follow shows that *GTP* is in conflict with:

*Countable Additivity (CA).* For any set of countably many centered or uncentered propositions, any two of which are incompatible, rationality requires that one’s credences in the propositions in this set sum to one’s credence in their disjunction.

Here is the example.<sup>2</sup>

*Sleeping Beauty in St. Petersburg (SBSP).* Let  $S = \mathbf{N}$  and suppose that Beauty awakens  $2^X$  times, where  $X$  is a random variable with  $P(X = n) = 2^{-n}$ ,  $n \in \mathbf{N}$ .

If Beauty subscribes to *GTP*, then in *SBSP* it would appear that she must assign equal credences to the exhaustive and mutually exclusive assertions  $X = n$ , which violates *CA*.

It’s easy to see how this argument could become popular—it’s counterintuitive and seem-

---

<sup>2</sup>For those who know of such things, the original problem is basically a positive recurrent Markov chain, the one-third solution deriving from its stationary probability measure, while *SBSP* is a null recurrent chain, for which all stationary distributions are infinite. That null recurrent chains are essentially useless as models for nomic phenomena is generally accepted; see e.g. [2], in particular Example 6.1.1 and Section 6.9 (chapter summary). What saves thirders is that this limitation is not shared by positive recurrent chains.

ingly plausible at worlds where a literal reading of *FSBI*, indeed an apparent premise of thirders, can be rationally defended. At our world, however, thirders can't rationally sanction a literal interpretation of *FSBI*, as they know that with non-zero probability, Beauty will die between tails awakenings. That thirders have failed to be explicit about mortal contingencies (which would rightly be perceived as tedious) is, I think, innocent. Ignoring such remote possibilities simplifies the thirder model, and of course their effect on Beauty's credence function vanishes with increasing remoteness. Ignoring mortality in the class of examples to which *SBSP* belongs, however, surrenders one to pathology. Indeed, *SBSP* is so sensitive to conditions of implementation that acknowledging any prospect for mortality whatsoever (that Beauty might be transformed into a marble bust of Pascal by unfortunate quantum effects will do) eliminates any perceived conflict with *CA*.

Some details: when Beauty is explicit about mortality, she employs not *FSBI* but:

*Sleeping Beauty Partiality (SBP)*. If the number of times Beauty awakens is  $M$ , then for any hypothesis  $h$  in  $S$ , upon first awakening, Beauty's credence in the  $k$ th awakening associated with  $h$  should be proportional to  $Ch(M \geq k|h)$ .

Here  $Ch(\cdot)$  is objective chance. *SBP*, together with other plausible hypotheses (first night mortality rates independent of  $h$  and credences in first  $h$  awakenings proportional to  $Ch(h)$ ), implies that Beauty's absolute credence in the  $k$ th awakening associated with  $h$  should be

$$P(h \wedge k) = \frac{Ch(h) \cdot Ch(M \geq k|h)}{\sum_{j \in S, l \in \mathbf{N}} Ch(j) \cdot Ch(M \geq l|j)} = \frac{Ch(h) \cdot Ch(M \geq k|h)}{E(M)}.$$

Summing over  $k \in \mathbf{N}$ , Beauty's credence in  $h$  should be

$$P(h) = \frac{Ch(h) \cdot E(M|h)}{E(M)}.$$

Define the *fidelity* of an implementation to be  $Ch(N = M)$ , where  $N$  is the number of times Beauty is told she will awaken and  $M$  is the number of times Beauty does awaken. The *variation distance* between two discrete credence functions  $R$  and  $Q$  on a set  $S$  is the quantity

$$v(R, Q) = \frac{1}{2} \sum_{h \in S} |R(h) - Q(h)|.$$

Under suitable hypotheses, Beauty's credence in  $h$  is  $P(h) = \frac{Ch(h) \cdot E(M|h)}{E(M)}$ . Therefore, if  $E(N) < \infty$  then as fidelity approaches 1 Beauty's credences converge in variation to the distribution  $Q(h) = \frac{Ch(h) \cdot E(N|h)}{E(N)}$ . On the other hand if  $E(N) = \infty$  then as fidelity approaches 1 Beauty's credence in  $h$  approaches zero for every  $h$ , and her credences diverge in variation. In the former case, the distribution  $Q$  constitutes a stable solution to the problem. In the latter case, there is no stable solution, meaning that individual agents cannot avoid mortality estimates in establishing or even approximating their credences.

This is consistent with *CA* and recovers the one-third solution, modulo agreement that thirder Beauty may assign credence  $\frac{1}{3} + \epsilon$  to heads for a smallish  $\epsilon$ . Diehard thirders who insist on  $\epsilon = 0$  meanwhile may indeed run afoul of *CA* per [8]. Where rational, however, such thirders inhabit nomically (at least) inaccessible worlds.<sup>3</sup>

## 2. The theoretical solution: halving.

In this section I'll be explicating the halving scheme of Lewis [6], which I think is attractive and plausible but suffers from underdevelopment and somewhat mysterious motivation. As I explained in Section 1, Beauty loses *de se* information which, on the naive view, is relevant to *heads*. Lewis is not satisfied with the naive view, as for him Sleeping Beauty scenarios involve what might be termed a "possible world sampling bias" (in the original problem, for example, the *tails* world is oversampled). Employing this language, thirders adopt sample (*de se*) proportions as their personal credences; Lewis prefers to adopt population (objective chance) proportions. To get at these, one must correct the bias.

What I think is mysterious about the way Lewis sets things up concerns his use of the following premiss: (L1) Only new relevant evidence, centred or uncentred, produces a change in credence; and the evidence ( $H1 \vee T1 \vee T2$ ) (sic) is not relevant to HEADS versus TAILS. Here of course *H1* is *Monday heads*, *T1* is *Monday tails* and *T2* is *Tuesday tails*. This premiss makes no provision for *loss* of (*de se* or "centred") information, or that such information could be relevant to *heads*. In other words, it makes no provision for the naive view. A move like that requires strong theoretical support, of which Lewis offers none.

One way of giving such support (while shedding light on some of the issues) is to make use of the logarithmic scoring rule,<sup>4</sup> where the agent with estimate  $p$  for  $P(\text{heads})$  scores the negative of surprisal, here  $\log p$  if *heads* and  $\log(1 - p)$  if *tails*. Assuming a fair coin, minimum expected surprisal occurs at  $p = \frac{1}{2}$  (a calculus exercise). Of course, Beauty is *surprised twice* if the coin lands tails, and some (namely thirders) will contend that her surprisal is actually  $2 \log(1 - p)$  in this event, with minimum expectation at  $p = \frac{1}{3}$ . The reason that both methods of accounting are defensible is that logarithmic scoring measures distance (in bits) from omniscience, and there are (if one believes [5]) two legitimate readings of *omniscient*, according to whether one admits *de se* information. Accounting

---

<sup>3</sup>At worlds supporting faithful implementations of *SBSP* and the like, the first thing Beauty will try in her attempt to rectify *GTP* and *CA* will be to modify *GTP* (and *FSBI*) with the language "In any Sleeping Beauty problem of *finite expected duration*...." However, this still runs up against *CA* modulo the following principle of: *Completeness*: Any rational agent must assign probability zero to any event  $h$  having the property that for any  $x > 0$  there are events  $k$  and  $E$  containing  $h$  such that  $P(k|E)$  exists and is less than  $x$ . (In *SBSP* take  $k = h = (X = c)$  and  $E = (X \leq n)$  for a largish  $n$  to conclude that  $P(X = c) = 0$ .)

<sup>4</sup>Only the logarithmic rule is linear with respect to surprisal, which is the quantity rational agents trafficking in information seek to minimize.

propositionally, if Beauty knows everything except *tails*, she is one bit shy of omniscience, and is only charged one bit (logs are base 2). Under *de se* accounting practices, a halving Beauty gets charged two bits in case of *tails*, as her surprisals are doxastically independent.

If this is right, then even thirders might adopt a charitable attitude regarding efforts to make sense of a one-half solution. I'll be canvassing a few proposals in the following paragraphs and checking them for soundness. The surprisal-minimization argument (together with the sampling bias paradigm) can be taken as a guide; since only one *tails* surprisal is to count, a policy of *disenfranchisement* is instituted, according to which *Monday tails* and *Tuesday tails* versions of Beauty, between them, get counted only once, whether by stipulation, lots or weighted average. Apportionment by equally weighted lots or averages amounts to acceptance of Elga's indifference principle, and is the more common choice.

Unequal lots or asymmetric stipulation is championed in Hawley [3], where it is maintained that Beauty should assign credence 1 to *Monday* (and thus in particular to *Monday* conditioned on *tails*). Strangers to disenfranchisement may find it worrisome that Hawley will be updating credence in *heads* to 1 upon learning that *Monday tails* does not obtain, but what troubles me more is that the choice seems arbitrary—could not one just as easily assign credence 1 to *Tuesday* conditioned on *tails*? Hawley writes: “the best compromise...might well be to believe to degree 1 that it is Monday whenever she awakens. She will be surely be right on Monday, and possibly never be wrong about the day during the experiment.” What about *Muesday*, though, which is like Monday whenever the coin lies heads and like Tuesday whenever the coin lies tails? In favor of this grueish day one might counter: “the best compromise...might well be to believe to degree 1 that it is *Muesday* whenever she awakens. She will be surely be right on *Muesday*, and possibly never be wrong about the day during the experiment.” Halfers are on thin ice already; arbitrariness is pushing it.

Respect for symmetry, then, seems to advise that halfers accept Elga's principle. Halfers that do include Lewis, Meacham [7] and probably White [11], albeit not explicitly. Even among halfers agreeing on this much, however, there are differences in how propositional credences are updated in light of *de se* evidence. No method comes cheap; Lewis's policies lead famously to Beauty counterintuitively (it is claimed) updating credence in *heads* to  $\frac{2}{3}$  when she learns *Monday* during an awakening. Meacham meanwhile (White too, I think) abandons the multiplication rule  $P(A \wedge B) = P(A)P(B|A)$ . The latter concession is more costly, as I now illustrate by a brief sidequest through well-trodden probability lore.

Suppose that a **big prize** is hidden behind one of three doors, each with equal objective chance. The hypothesis *Door i* corresponds to the state of affairs in which the **big prize** is behind Door *i*. If *Door 1*, then Beauty will have a single awakening, on Monday. If *Door 2*, Beauty will have a single awakening, on Tuesday. And, if *Door 3*, Beauty will have two awakenings, on Monday and Tuesday. Halfers of course assign each of the alternatives credence  $\frac{1}{3}$  upon awakening.

Suppose now that a halfer asks what day it is, and after hearing the answer is asked for

her updated credence in *Door 3*. Note: if the answer is *Monday*, this rules out *Door 1*. If the answer is *Tuesday*, this rules out *Door 2*. *Door 3* cannot be eliminated. Recall that our halfer has prior credence  $\frac{1}{3}$  in *Door i* for each *i* and, if she accepts Elga's principle, *Monday* and *Tuesday* are equally likely conditioned on *Door 3*. Suppose our halfer learns that *Monday* obtains. Then her predicament derives from a protocol (see [9]) isomorphic to that of the Monty Hall problem, and corresponds to the situation in which the contestant has initially chosen *Door 3* and seen the hypothesis *Door 1* eliminated.

Accordingly, any halfer who updates credences by conditioning on *not Door 1* is committing the well-known fallacy of those who answer  $\frac{1}{2}$  in the Monty Hall problem. Namely, conditioning on the proposition learned instead of on the fact that it, among other candidate propositions, was learned. On the contrary, Beauty's credence in *Door 3* must remain  $\frac{1}{3}$ . White and Meacham, who update propositional credences in response to *de se* evidence by conditioning on the largest propositional event consistent with that evidence, fall for the fallacy; Lewis gets it right.

How, though, are we to reconcile Lewis's counterintuitive consequences?<sup>5</sup> Most people, even thirders, acknowledge that the the objective chance of heads is still  $\frac{1}{2}$ , as they might imagine advising someone outside of the experiment, but when it comes to advising  $\frac{2}{3}$  after learning *Monday*, they balk. I believe that, taken together, these intuitions are wrong, and wrong for the same reason that naive Monty Hall intuitions are wrong. Namely, they ignore the role of protocol—in this case, the protocol by which the communication with the person outside of the experiment was arranged. If the person outside knows this protocol, he won't condition on Beauty's evidence *per se* but on the fact that it came to him. And, as I'll elaborate on below (in a betting context), on the most natural protocols for which he would take Beauty's initial advice of  $\frac{1}{2}$  at face value, he would take  $\frac{2}{3}$  at face value when Beauty had learned that it was the first day of the experiment.

There are, moreover, powerful statistical precedents for the coherence of Lewisian updating, at least at the level of mathematics. For an artificial example, Lewis's policies would be vindicated were we to force Beauty to subject one of her *tails* versions to electoral disenfranchisement by requiring her to mail herself a postcard during each awakening and to then read her credence in *heads* off of a single (randomly selected if multiple) postcard Wednesday morning. Less artificially, they agree in spirit and letter with standard methods of accounting for biased sampling.

Suppose for example that a lake contains catfish and sunfish. If you go fishing, two-thirds of the fish you catch are sunfish, but you know that individual catfish, being bottom feeders, are only half as likely to get hooked as individual sunfish. Asked for the correct proportion

---

<sup>5</sup>If it isn't counterintuitive enough that he assigns credence  $\frac{2}{3}$  to *heads* conditioned on *Monday*, consider that he sticks by this even when the coin is not to be flipped until Monday night, or that by iterating the erasure procedures over shorter and shorter time frames on Monday, one can drive credence in *heads* arbitrarily close to 1.

of catfish, you may wonder whether “sample proportion” or “population proportion” is intended. If the latter, you will correct the sampling bias (typically using *sample weights*, where in this case individual catfish turning up in samples are given twice the weight of individual sunfish) and answer  $\frac{1}{2}$ . Suppose moreover that half the sunfish are warmouth and half are bluegill. Equal numbers of catfish, warmouth and bluegill appear in catch samples, but if a fish is selected from the lake uniformly at random and is found not to be a bluegill, the probability that it’s a catfish will be  $\frac{2}{3}$ . What’s going on with Lewis appears to be strongly analogous.<sup>6</sup> So, if one takes divergence of *de se* credences from objective chance in Sleeping Beauty scenarios to be an undesirable artifact of *tails* world oversampling, it’s plausible that one should do exactly what Lewis does: correct the bias in formulating what might be called “post-theoretic” personal credences.

In light of these considerations, a rather idiosyncratic picture emerges. One version of *tails Beauty*<sup>7</sup>...either Monday’s version or Tuesday’s, is viewed as a redundant copy, and counts for nothing. Eccentric as it is, though, any argument for thirding that doesn’t address it in detail risks begging the question. For example, Weintraub in [10] (Horgan [4] is another example, but with fewer awakenings) considers a scenario in which Beauty is to have three awakenings in either case and finds out during an awakening that (more or less) it is either *Monday heads*, *Monday tails* or *Tuesday tails*. She argues that the correct credence in *heads* here is  $\frac{1}{3}$ , and that the original problem is analogous. But by Lewis’s lights the cases are different, and in fact he too assigns *heads* credence  $\frac{1}{3}$  in hers.

Another argument against Lewis is that he appears to violate reflection. Indeed, suppose that Beauty awakens at 9:00, is told what day it is at 9:05 and is informed about the outcome of the flip at 9:10. Then for Lewis her expected credence in *heads* at 9:07 Monday is  $\frac{2}{3}$ , and she knows this on Sunday night. By reflection, ought not she update to  $\frac{2}{3}$  then? The answer is of course no, but why? We can answer this by analogy with an apparent violation of reflection in the example of the prisoner (see [0]). Suppose a prisoner is waiting on a possible stay of his execution from a governor, which he expects to receive with probability  $\frac{1}{2}$ . There is a light in his cell that a meticulous guard has agreed to turn off at midnight if the stay is granted. There is no clock in the cell. As we advance toward midnight, the prisoner sees his credence in *stay* drop from  $\frac{1}{2}$  to some lower value  $L$ . Let’s say the expected value of  $L$  is  $\alpha$ . (There is a heuristic argument suggesting that  $\alpha \approx \frac{1}{3}$ , but nothing turns on this.) The prisoner knows all of this; does reflection require him to update his credence in *stay* to  $\alpha$  now, at 6 P.M.? Of course no again...but why?

The resolution, hinted at earlier, is that reflection applies not to expected credences at

---

<sup>6</sup>If you think this example is disanalogous, imagine instead that all of the sunfish are changelings, spending half of their lives as warmouth and half as bluegill; in a uniform sample now,  $\frac{2}{3}$  of the non-bluegill are catfish.

<sup>7</sup>Selected at random. Or half of each, as one might imagine doing in versions of the problem where Beauty’s body is replicated, but Beauty’s soul is divided, in case of *tails*.

future times but to expected credences upon conditionalization on partitions.<sup>8</sup> That’s not what happens here. First, conditioning on the state  $S$  of his internal clock alone (as if marking time in a blindfold) does nothing to his credence in *stay*, and what reflection implies about a subsequent partitioning on there being light or not is that

$$P(\textit{stay}) = P(\textit{stay}|S) = P(\textit{on}|S)P(\textit{stay}|S \wedge \textit{on}) + P(\textit{off}|S)P(\textit{stay}|S \wedge \textit{off});$$

it does not imply that  $P(\textit{stay}|S) = P(\textit{stay}|S \wedge \textit{on})$ . The prisoner’s credence drops in response to conditionalization on evidence, then, but not in response to conditionalization on a partition. The same is true for a Lewisian Beauty’s  $\frac{2}{3}$  credence in *heads* at 9:07 on Monday.<sup>9</sup> What reflection implies about partitioning on the current day is that

$$P(\textit{heads}) = P(\textit{Monday})P(\textit{heads}|\textit{Monday}) + P(\textit{Tuesday})P(\textit{heads}|\textit{Tuesday});$$

it does not imply that  $P(\textit{heads}) = P(\textit{heads}|\textit{Monday})$ .<sup>10</sup>

Well, we’ve seen thirding and several forms of halving. Theory aside, whether or not these schemes are any good depends on the existence or non-existence of reasonably natural applications. So suppose that, during an awakening, Beauty gets a call from her bookie, Rango Maleficent. Rango explains that there is an opportunity to bet on the coin toss, and he needs to know her credence in *heads*. How should she respond, if she wishes to be invulnerable to a Dutch book? How should she respond if she knows that it’s Monday (or Tuesday)? As previously noted, her appropriate responses are underspecified, because they depend on the protocol by which Mal’s call was put through. Some possible protocols, in decreasing order of “naturalness”: (1) call to come in each day of the experiment—Beauty adopts thirder credences. (2) exactly one call to come in, Monday and Tuesday equally likely conditioned on *tails*—Beauty adopts Lewisian credences. (3) single call to come in on the first (respectively last) day of the experiment—Beauty engages in Hawley-style (respectively Tuesday-style) halving.

---

<sup>8</sup>Somewhat more generally, reflection is valid at almost surely finite future times depending only on evidence gathered up to that time, i.e. *stopping times*; reflection in this context follows from the bounded martingale stopping theorem. Where time and evidence are discrete, this is equivalent to conditionalization on countable partitions, which is simpler.

<sup>9</sup>Nor is 9:07 Monday a stopping time for a half-half disenfranchiser; in fact it fails to be finite almost surely, as it is never reached conditioned on *Tuesday*, i.e. never reached by the half that is disenfranchised on Monday.

<sup>10</sup>Disenfranchisement also allows halvers to dodge the apparent violation of inverse reflection entailed by the contrapositive of the argument given for thirding in Section 1; if *Monday tails* is disenfranchised, then on Tuesday Beauty’s credences revert to those not of Monday (which has been disenfranchised), but of Sunday, while if halves of each are disenfranchised, *Monday* and *Tuesday* exist in parallel, not series (different halves are disenfranchised, if you will), so the credences of each revert to those of Sunday.



I've maintained that proper credences in Sleeping Beauty scenarios are underspecified. Thidding is the most natural choice and coheres, at nomically accessible worlds, with countable additivity of credences, but Lewisian halving, the frontrunner among halving schemes, has not been refuted by thirders. They aren't without resources; a fair case might be made that disenfranchisement is a just-so move that jurymanders personal credence into false accord with objective chance. I have conflicted feelings about this. For while thirders might have *Truth* on their side, surely Lewis has *Beauty* on his.

## References

- [0] F. Arntzenius. 2003. Some problems for conditionalization and reflection. *The Journal of Philosophy* 100:356-370.
- [1] Elga, Adam. 2000. Self-locating belief and the Sleeping Beauty problem. *Analysis* 60:143-147.
- [2] Gallager, Robert G. 2011. *Stochastic Process: Theory for Applications* (draft). Available at <http://www.rle.mit.edu/rgallager/notes.htm>
- [3] Hawley, Patrick. 2012. Inertia, Optimism and Beauty. *Nous*. To appear. Available at <http://philsci-archive.pitt.edu/5319/1/iob.pdf>
- [4] Horgan, Terry. 2004. Sleeping Beauty awakened: new odds at the dawn of the new day. *Analysis* 63: 10-21.
- [5] Lewis, David. 1979. Attitudes *De Dicto* and *De Se*. *The Philosophical Review* 88: 513-543.
- [6] Lewis, David. 2001. Sleeping Beauty: Reply to Elga. *Analysis* 61:171-176.
- [7] Meacham, Christopher. 2008. Sleeping Beauty and the Dynamics of *De Se* Beliefs. *Philosophical Studies* 138: 245-69.
- [8] Ross, Jacob. 2010. Sleeping Beauty, countable additivity, and rational dilemmas. *The Philosophical Review* 119: 411-447.
- [9] G. Shafer, Conditional probability. *International Statistical Review* 53: 261-275.
- [10] Weintraub, Ruth. 2004. Sleeping Beauty: A Simple Solution. *Analysis* 64: 8-10.
- [11] White, Roger. 2006. The generalized Sleeping Beauty problem: a challenge for thirders. *Analysis* 66: 114-119.