Beauty 's Cautionary Tale

Susan Vineberg
Wayne State University

The Problem

Sleeping Beauty will participate in an experiment in which she will be put into a deep sleep on Sunday night. She will be awakened during the experiment either once or twice depending on the toss of a fair coin. In either case she will be awakened on Monday. She will be unaware upon awakening that it is Monday, but subsequently will be told that it is Monday and then given drugs to put her back to sleep. These drugs will erase her memory of having been awake. If the coin lands heads, Beauty will then sleep until the experiment ends on Wednesday. If it lands tails, she will be awakened briefly again on Tuesday, unaware of the Monday awakening. Beauty knows all of this, and will be able to distinguish waking up within the experiment from other awakenings.

Before being put to sleep, it is reasonable for Beauty to assign a probability of 1/2 to the proposition that the coin lands heads. Now, let us assume that the experiment has begun and Beauty has just awakened from a deep sleep. What should her probability be now for heads? It seems that surely her probability should remain at 1/2. Beauty knew in advance that she would be awakened at least once, and so upon awakening receives no new information that could count as evidence for or against heads. If however the experiment were to be repeated, in the long run there be would twice as many tails awakenings as heads awakenings, and this suggests that the answer should be 1/3. Additional arguments have been proposed for both answers. Since the discussion has been largely a response to Elga (Elga 2000), who defends the answer of 1/3, I shall begin with his argument.

Elga

Elga starts with the assumption that upon awakening, Beauty is in one of the following scenarios:

H_1 : Heads and its Monday.
T_1 : Tails and its Monday.
T_2: Tails and its Tuesday.

Let P_0 be Beauty's probability function before being put to sleep on Sunday, P_1 her probability function upon awakening, and P_2 her probability function after being told that it is Monday. Elga argues that being in T_1 is qualitatively like being in T_2, and hence that $P_1(T_1) = P_1(T_2)$. He then notes that the experimenters could either toss a coin before or after the Monday awakening. He assumes that the timing of the toss is irrelevant, but for definiteness assumes that it is to take place after Beauty is put back to sleep on Monday. He claims that upon learning that it is Monday, Beauty should update her probabilities by conditionalizing on this information, which can be represented as $H_1 \vee T_1$. Elga reasons that once Beauty is told that is Monday, she should attach a probability of 1/2 to Heads (after all the coin has not yet been tossed). Thus,

$P_2(\text{heads}) = P_1(\text{heads}_1/H_1 \vee T_1) = 1/2.$

From this it follows that $P_1(H_1) = P_1(T_1)$, and hence that $P_1(H_1) = P_1(T_1) = P_1(T_2) = 1/3$, and so $P_1(\text{heads}) = 1/3$.

Elga observes that T_1 and T_2 do not pick out different possible worlds, but different locations within a single possible world. Possibilities about which world is actual are associated with classes of uncentered worlds, whereas possibilities about where one is in the world are associated with classes of centered worlds. So here, H_1 ,T_1 and T_2 are to be associated with classes of centered worlds. Within this framework, Elga takes it that upon awakening Beauty does not receive uncentered evidence, but that she does receive the centered evidence that H_1 v T_1 v T_2. For Elga, the shift from P_0 to P_1 is attributed to the learning of the latter centered evidence, though on his view Beauty obtains no uncentered evidence.

It should be noted that if Elga is correct then Beauty violates Van Fraassen's Reflection Principle, which as a special case requires that if an agent is certain that she will attach probability x to S tomorrow, then she should attach that probability to S today. Further, at the outset of the experiment P_0(heads)= 1/2 and Beauty believes that it is Sunday, and so given Elga's representation must attach a probability of 0 to H_1 v T_1. This means that the probability shift in heads from 1/2 to 1/3 cannot take place by conditionalization. Accordingly, Elga regards the Sleeping Beauty Problem as raising an issue about how beliefs should change when considering centered, rather than uncentered, possibilities.

Lewis

In his reply to Elga (Lewis 2001), Lewis accepts Elga's framework and assumptions except for the claims that H_1 v T_1 v T_2 is relevant to heads and that P_2(heads) = 1/2. Indeed, from the premise that upon awakening Beauty learns nothing relevant to the probability of heads, Lewis argues that P_1(heads)=1/2, and perhaps more surprisingly that P_2(heads) = 2/3. This follows from the claim that P_1(heads)=1/2 , together with the shared assumptions that

$P_1(\text{heads}/H_1 \text{ v } T_1) = P_1(H_1/H_1 \text{ v } T_1) = P_1(H_1)/P_1(H_1) + P_1(T_1)$, and that
$P(\text{tails}) = P_1(T_1) + P_1(T_2)$

The shift from P_0 to P_1 does not violate Reflection, however a question still arises about how belief change is to be characterized. Both Lewis and Elga take it that Beauty obtains centered evidence and accordingly conditions on centered propositions. On Sunday, Beauty's probability that it is Monday is 0, but conditionalizing on any proposition assumes that its initial probability is nonzero.[1] When Beauty is first awakened her probability that it is Monday is nonzero, and so once she is told that it is Monday it is presumably legitimate for her to conditionalize.[2] Still, we cannot analyze Beauty's probability function as evolving over the entire course of the experiment simply by conditionalizing over her experiential inputs. It seems that, with Elga and Lewis, we may treat some segments of Beauty's evolving belief state as arising by conditionalization. In general we do treat the evolution of belief in such a fragmentary way, in order to handle situations in which probabilities move from the extreme points of zero and one. It is not the fragmentary treatment of belief, nor even that the information obtained is centered, but rather what beauty should count as evidence in her induced state of uncertainty that is at issue.

The crucial assumption in Lewis' analysis is that Beauty learns nothing relevant to heads upon awakening. It will be useful to see that this assumption amounts to treating Beauty's predicament as fundamentally like the following case of drawing coins, which lacks the induced uncertainty:

Drawing Coins -Box 1:

This box contains two drawers, one of which contains a single silver coin and the other contains two silver coins. Suppose that a drawer is chosen at random by the toss of a coin, and that a silver coin is drawn. What is the probability that it came from the drawer with just one coin? Answer 1/2.

Now suppose that the coins in this box have either an 'M' or 'T' on them, but that these letters are not observed at first. The lone coin in the first drawer is marked with 'M', as is one of the coins in the other drawer. The second coin is marked with 'T'. After observing that the coin drawn is marked with 'M', what is the probability that it came from the drawer with just one coin? Answer 2/3.

---

[1] Some Bayesians will object that contingent propositions should never be assigned probability one. What is important here is that Lewis and Elga, in supposing that Beauty updates by strict conditionalization when she is told that it is Monday, are assuming that the claim that it is Monday is assigned probability one.

[2] Elga and Lewis make much of the fact that when Beauty learns that it is Monday, she is given centered, rather than the usual uncentered, information. However, what seems crucial to the problem is not that she loses track of her place in the world, but rather that she enters into a state of uncertainty.

It is crucial to obtaining a probability of 2/3 for having drawn from the drawer with one coin after learning that the coin is marked with an 'M' that there was initially a 50-50 chance of drawing from the drawer with one coin. The claim that Beauty's probability should rise to 2/3 similarly depends on the assumption that Beauty's probability of being in a world in which the coin lands heads is at the time is awakened is 1/2. It seems that Beauty's predicament is like that of drawing coins from box 1. Just as it is certain that a silver coin will be drawn, and thus doing so provides no information that is relevant to which drawer was chosen, Beauty knows for certain that she will be awakened, so it seems that being awakened cannot count in favor of tails. If all of this is right, then Elga's assumption that $P_2(H_1) = 1/2$ looks unjustified, as Lewis charges. However, tossing the coin on Monday night calls into question the analogy with drawing coins from box #1, for it is unclear that prior to the toss that Beauty has an equal chance of being in a one awakening trial as being in a two awakening trial. Further, whatever her belief upon awakening, it seems that learning that it is Monday should reset Beauty's probability for heads to what it was on Sunday night, for at that point the chance of heads depends on the future flip of a fair coin. [3]

There are plausible intuitions behind both Lewis' claim that Beauty receives no new relevant evidence for heads and Elga's assumption that $P_2(H_1) = 1/2$. Both agree that $P_1(T_1) = P_1(T_2)$, on the basis of a restricted principle of indifference, which Elga relies upon in claiming that $P_1(T_2) = 1/3$, and Lewis uses to get that $P_2(H_1) = 2/3$. Regardless of whether we accept this common assumption, it is not essential to the conflict. Lewis and Elga's other common assumptions that

(1) $P_2(\text{heads}) = P_1(\text{heads}_1/H_1 \vee T_1)$

(2) $P_1(\text{heads}_1/\text{-}(H_1 \vee T_1) = 0$

(3) $0 < P_1(/H_1 \vee T_1) < 1$

together with

(4) $P_2(H_1) = 1/2$

(5) $P_1(H_1) = 1/2$

[3] Lewis argues that there is a proviso to the Principle Principle, that applies in this case.

lead to contradiction and so at least one of these plausible assumptions must go. Before considering the matter directly, it is worth pausing to see whether a recent appeal to Dutch Book arguments can bolster the case for jettisoning one of these assumptions.

Dutch Book Arguments

Hitchcock (Hitchcock forthcoming) argues that it is possible to construct a Dutch Book against an agent who commits in advance to either answer. However, he claims that the circumstances in which the Dutch Books are made impugns the cognitive rationality of Beauty if her answer is 1/2, but not if it is 1/3. He further argues that only the assignment of 1/3 can avoid vulnerability to the compelling kind of Dutch Book. Hitchcock anticipates that there will be some skepticism about his Dutch Book argument, at least in part because of the long history of attack on such arguments. I do find the argument less than fully convincing, but not on the grounds that such arguments are generally without force. In my view, Dutch Book vulnerability does involve the failure to satisfy some epistemic or rational ideal. Rather, my concern is with the sort of force that Hitchcock attributes to the argument.

Consider first how a Dutch Book can be made against Beauty if she plans to change her belief in heads upon awakening to 1/3. Here bet #1 is made Sunday night and bet #2 is made after Beauty is awakened.

Bets	Payoff	Cost	Heads	Tails
bet #1	$30 if heads	$15	$15	–$15
bet #2	$30 if tails	$20	–$20	$10
net	$30	$35	–$5	–$5

Given that the bookie knows in advance that Beauty will assign probability 1/2 to heads each time that she is awakened assures the bookie on Sunday of a net profit. Beauty intends to shift her beliefs in a way that violates Reflection, and so the Dutch Book strategy here is just a special case of the general Dutch Book strategy against agents who violate Reflection. See (van Fraassen 1984).

Interestingly, there is also a Dutch Strategy against Beauty if she sticks with the probability 1/2. The bets are given below, where bet #1 is placed before she goes to sleep, bet 2 is placed upon awakening and bet 3 only gets placed if she is awakened a second time.

Bets	Payoff	Cost	Heads	Tails
bet #1	$30 if tails	$15	–$15	$15
bet #2	$20 if heads	$10	$10	–$10
bet #3	$20 if heads	$10 if tails	0	–$10
net	$20 if heads	$25 if heads	–$5	
	$30 if tails	$35 if tails		–$5

What are we to make of these two strategies? It is generally thought that insofar as Dutch Books and Dutch Strategies reveal defects in an agent's belief system it is not because having such beliefs involves a pragmatic defect, in that they can lead to a loss of utility, but rather because beliefs that are correlated with Dutch Books exhibit some form of inconsistency.[4] Not all such forms of inconsistency are indicative of epistemic irrationality. For example if a bookie knows that the proposition S is true, where the agent has some confidence in S that is less than one, he can assure himself a profit by getting the agent to bet against S. In this case, the agent's beliefs need not be internally

[4] See Vineberg, S. (2001). "The Notion of Consistency for Partial Belief." Philosophical Studies **102**: 281-296.

inconsistent, but are merely inconsistent with the facts, and the latter sort of inconsistency is not epistemically irrational. Such vulnerability indicates an epistemic defect, but not necessarily a cognitive one. Hitchcock notes that "if the bookie can achieve his certain gain only by exploiting information that is unavailable to the agent, then the Dutch Book reflects an evaluation of the system of bets that is not the agent's own."

Hitchcock argues that the Dutch Strategy that can be devised if Beauty plans to change her belief in heads to 1/3 upon awakening cannot be carried out without exploiting information that is unavailable to Beauty, and hence does not show that her beliefs are epistemically irrational. In order for the bookie to guarantee himself a profit, he must show up on Monday to place the second bet, but then his ability to make the book involves information unavailable to Beauty. What about the second strategy that applies if Beauty sticks with 1/2? Hitchcock observes that here the bookie can make these bets without having knowledge that is unavailable to Beauty, by undergoing the same protocol as she does. He will wake up when Beauty does and in effect place bet # 2 whenever he wakes up. Observe that this strategy will work to ensure his profit in the case where Beauty sticks with the probability 1/2, but not in the case above where she switches to 1/3. Since the bookie needs no special knowledge unavailable to Beauty to implement the strategy in the former case, Hitchcock takes this Dutch Strategy to tell against the answer 1/2. This together with a further argument that a proper or telling Dutch Strategy cannot be make against Beauty if she switches her belief in heads to 1/3, leads Hitchcock to favor the latter solution to the Sleeping Beauty Problem.

Observe that Hitchcock has not given an argument that Beauty should attach a probability of 1/3, rather than 1/2, to heads upon awakening, but instead provided an argument for changing from a probability on Sunday of 1/2 for heads to 1/3 upon awakening, over staying with 1/2. There would be no guaranteed loss in the latter case, without the bets made on Sunday night. To have a Dutch Strategy, there must be a time at which a betting strategy can be put into place that will guarantee a net profit through a series of bets placed over time. Typically, such strategies have been invoked to show that there is something wrong with the agent's initial beliefs, in that her subsequent beliefs will be, in some significant way, inconsistent with the agent's initial set of beliefs.[5] What is peculiar here is that a Dutch Strategy is invoked against an agent who does not change her beliefs at all. Dogmatism may be an epistemic defect, but it is surely not that of internal inconsistency!

According to Hitchcock, there is a particular form of inconsistency involved in Dutch Book vulnerability that comes into play in Beauty's case.

> Susceptibility to Dutch Book is symptomatic of an underlying evaluative inconsistency. On the one hand, the agent views a book of bets as fair — she judges each individual bet as yielding no expected loss or gain for either side. On the other hand, she views the book of bets as unfair — she can determine that a loss is inevitable using purely deductive reasoning, which does not presuppose probabilistic coherence.

By these standards, Beauty would apparently exhibit an evaluative inconsistency if she sticks with probability 1/2. Each bet will look fair, according to her beliefs, yet she can see, as well as the bookie, that there is a Dutch Strategy against her. It is important here that Beauty's degrees of belief for heads alone do not permit the construction of the Dutch Strategy, which relies on the fact that she will only be offered a third bet under the condition that the coin has already come up tails. However, Beauty has knowledge of the experimental setup and so is in a position to deduce from these facts and her beliefs that she will suffer a sure loss.

Although the bookie does not have information unavailable to Beauty at the time he places the bets, the occurrence of a second awakening and subsequent third bet are linked to the outcome of that bet. It is this fact of the experimental setup that the bookie exploits, rather than an internal inconsistency in the agent's degrees of belief. What matters is not what the bookie knows at the time that he makes the bets, but whether he exploits something beyond the agent's own beliefs. The example shows that it is possible to do this even when the bookie knows no more than the bettor at the time that the bets are placed. Most importantly here, the Dutch Book vulnerability of someone in the experiment, whose confidence remains at 1/2 upon awakening, isn't produced just by

---

[5] See Vineberg, S. (1997). "Dutch Books, Dutch Strategies and What They Show About Rationality." Philosophical Studies **86**: 185-201.

her degrees of belief, but involves vulnerability with those degrees of belief in the experimental situation. As such the vulnerability looks more like that in which a bettor is vulnerable to a bookie is superior information. This is not to say that sticking with 1/2 would not constitute a defect under the circumstances, but that it is not entirely an internal one.

Now, consider the Dutch Book Strategy that can be made against Beauty if she changes from 1/2 to 1/3. As a special case in which Reflection is violated, it is just her beliefs that the Bookie exploits. Since the loss is guaranteed after Beauty's first awakening, it is unclear why it should matter what the bookie does afterwards. The bookie doesn't need special information to guarantee himself a profit, only to keep it, and it would seem to be former that is relevant to questions about the fitness of Beauty's epistemic state.

The Dutch Book argument to show that Beauty should change her belief in heads from 1/2 to 1/3 upon awakening is less than fully convincing insofar as it is supposed to show this is a requirement of internal consistency. It may seem to suggest that regardless of what beliefs Beauty should have, she should change her betting quotient for heads to 1/3. If indeed Lewis were correct that her degree of belief should remain at 1/2, then we would have a new example of a case where degrees of belief and betting quotients should diverge.[6] However, a bit of reflection shows that Beauty's Dutch Book vulnerability, should she retain a belief of 1/2, does not even produce an argument that this is an inferior betting quotient for heads. It is correct that unless Beauty changes her betting quotient to 1/3, she will be vulnerable to a bookie that has no more knowledge than she does. But, when the issue is prudential, rather than epistemic, it doesn't matter whether the bookie possesses more knowledge than Beauty. Indeed, if she were to plan to change her betting quotient to 1/3, she would be open to exploitation, because of her violation of Reflection, to a bookie that is not put to sleep. How Beauty should set and adjust her betting quotients will depend on the bets that she can expect to be offered. Assuming that she is forced to post odds, she is in a predicament in which she is

[6] Arntzenius argues that it may be perfectly reasonable for Beauty's degrees of belief and betting quotients to diverge, though he ultimately concludes that the most reasonable response to her predicament may be to have completely vague degrees of belief Arntzenius, F. (2002). “Reflections on Sleeping Beauty.” Analysis **62**(1): 53-62..

vulnerable to some sort of bookie if she announces a betting quotient for when she awakes that differs from 1/2, but also if she doesn't change it to 1/3. However, which book she is vulnerable to will depend on the bookies that she is most likely to confront. Under these circumstances, her betting quotients can reasonably come apart from her degrees of belief, and so do not determine what she should believe.

What Should Beauty Believe?

It has been suggested that the Dutch Book argument for switching from 1/3 to 1/2 is not entirely convincing. Moreover, even it were we would need to explain the force behind Lewis' assumption. As reinforcement for switching, consider the following: Beauty's predicament arises because she is put into a situation in which she becomes uncertain of her place in the world.[7] This suggests that we might apply the basic idea in treating uncertainty, which is to weigh how we would believe or act under the different possibilities by our estimate of their probability. This idea can be applied both to cases in which there is an initial uncertainty, or to fix probabilities in cases where one moves from certainty to uncertainty. When Beauty is awakened she knows that it is either Monday or Tuesday. If she was not in a state of uncertainty, and knew that it was Monday, she would take the probability of heads to be 1/2. If she knew that it was Tuesday, she would know that she had been awakened for the second time, and so would attach a probability of 0 to heads. In her state of uncertainty, Beauty takes it as possible that it is Tuesday when she awakes. Accordingly, she attaches some positive probability to the claim that it is now Tuesday, and thereby will take the probability of heads to be less than 1/2. Symmetry considerations can be invoked to argue that her probability should be 1/3.

Lewis would agree that in a situation with no uncertainty about the day, Beauty should attach probability 1/2 to heads on Monday. However, he must dispute that this fixes the probability of heads conditional on its being Monday at 1/2 in Beauty's actual predicament, and accordingly should reject the above reasoning. In order to maintain the answer of 1/3, a case must be made that being awakened really is evidentially relevant to heads, despite the intuition to the contrary. Notice that it is important not only to rejecting Lewis' argument, but also the position that Beauty should not change her belief in heads either upon awakening or upon being told that it is Monday.[8]

As previously noted Beauty knew before being put to sleep that her experience would be one of three indistinguishable awakening experiences. This suggests that finding herself awake within the experiment cannot serve to rule out or make less probable the proposition that the coin landed heads. A plausible assumption about evidence, which goes along with the Bayesian and various other accounts, maintains that evidence E for hypothesis H must rule out, or at least make less probable, at least some way in which H could be false. Thus it appears that that although Beauty learns that she is awake this does not count as evidence for or against heads. To see how awakening, contrary to appearances, could produce the right sort of evidence, consider a second coin drawing case:

Box 2: This is a box with two drawers, where each drawer contains two coins. Suppose that one drawer contains two silver colored coins and the other contains a silver coin and a gold coin. A drawer is selected at random, say by the toss of a fair coin, and that a silver coin is selected. What is the probability that it came from the drawer with 1 silver coin? Answer 1/3.

This example involves having an experience that rules out one of the initial alternatives. It seems Beauty's situation is better modeled by that of box 2, since drawing a gold coin, which would have to model it's being Tuesday with the coin coming up heads, is not one of the experiential possibilities. However, as the experiment begins that is a possibility for the future. We can get a somewhat better model of Beauty's situation by imagining that instead of a gold coin, there is an invisible coin in the drawer that the experimenter is

---

It has been suggested by Monton that the chief difficulty in the problem is that Beauty is subject to forgetting and that such changes cannot be described by conditionalization, but the problem is not one of forgetting per se, but of uncertainty.

[8] An alternative to the position of Elga and Lewis is to maintain that being awakened offers no genuine evidence about the toss of the coin, but learning that it is Monday acts to reset Beauty's degrees of belief in a way that violates conditionalization.

unable to detect.[9] In such a case, drawing a silver coin counts as evidence that one has drawn from the drawer with two silver coins because the state of having selected the invisible coin has been ruled out.

When Beauty is awakened, she not only learns that she is awake, but that it is not now true that it is Tuesday and that the coin came up heads. Beauty rules this possibility out upon being awakened, even though she could never find herself in the situation of heads on Tuesday. This is like the case where by drawing a silver coin, one rules out the possibility of picking the invisible coin.

Beauty's peculiar situation in which she is abruptly put to sleep and awakened encourages the idea that her awakening is not evidentially relevant to heads. If Beauty could reason in her sleep, while remaining uncertain about the day and time, she would view its being Tuesday under the condition that the coin came up heads as a possibility that would be ruled out if she were to be awakened. [10] The evidence available to Beauty does not depend on whether she reasons continuously, and so she should not treat the case where she is jerked back into consciousness differently from the one in which in her thinking is uninterrupted.

Beauty's Lessons

It has been argued here that Beauty's loss of consciousness and circumstances of awakening serve to obscure the relevant evidence. As such, the case serves as an example of how we can fail to see the importance of unrealized possibilities in assessing evidence.[11] Beyond this, the example involves an especially interesting violation of Reflection. Elga points to the fact that the violation arises because the agent receives centered information. However, the fact that Beauty's information is centered in the

[9] I am grateful to Sean Stidd for suggesting this variation.

[10] What her probabilities should be in this situation will depend on exactly how we are to understand the claim that she is unaware of the time. If she is aware of time passing, then the longer she goes without being awakened, the greater her probability for Tuesday and heads.

example is simply a means of inducing a kind of uncertainty. The anticipated uncertainty reveals that Reflection can be violated in cases where the agent's circumstances are to blame rather than his deficiency. Imagine a variation in which Beauty is not put to sleep, and simply awaits the entry into the room by the experimenter, either once or twice depending on the flip of a fair coin. We can imagine that the conditions in the room are such that she will be in doubt about the time and whether the experimenter has entered the room and that indeed her perception of the passing of time will be unreliable. After the experimenter enters the room, Beauty will be made to forget. Until then she does not forget anything, although she will be uncertain about whether or not she has forgotten anything and uncertain about whether the experimenter has entered the room. Before entering the room, Beauty contemplates the future uncertainty that she will experience, prior to the experimenter first entering the room. Neither initially nor prior to his entering is she cognitively deficient. At the later time, she will be in a situation in which her uncertainty stems from unfavorable epistemic conditions, but they are not ones in which she is hampered from gaining information about her situation. We can imagine as well that Beauty will not forget her initial beliefs about her situation. Violating Reflection seems both entirely rational in this case and different from cases in which it stems from pure forgetting, as in Talbott's spaghetti dinner example (Talbott 1991). It also places pressure on the disability defense,[12] for Beauty is not impaired in evaluating the evidence that is available to her. What is unusual is that her uncertainty is causally tied to a future loss of information. Such conditions define a class of cases where it may be rational to violate Reflection.

Acknowledgements: I am grateful to Sean Stidd for many helpful discussions on this problem.

References

[11] The Monty Hall problem is similar in that the evidential significance of the door that Monty reveals depends on understanding how Monty may be constrained, which is tied to the possible but unrealized possibilities.

[12] Van Fraassen, B. (1995). "Belief and the Problem of Ulysses and the Sirens." **77**: 7-37..

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

Elga, A. (2000). “Self-locating Belief and the Sleeping Beauty Problem.” Analysis **60**(2): 143-7.

Hitchcock, C. (forthcoming). “Beauty and the Bets.” Synthese.

Lewis, D. (2001). “Sleeping Beauty: reply to Elga.” Analysis **61**(3): 171-76.

Monton, B. (2002). “Sleeping Beauty and the Forgetful Bayesian.” Analysis **62**(1): 47-53.

Talbott, W. (1991). “Two Principles of Bayesian Epistemology.” Philosophical Studies **62**: 135-150.

Van Fraassen, B. (1995). “Belief and the Problem of Ulysses and the Sirens.” **77**: 7-37.

Van Fraassen, B. C. (1984). “Belief and the Will.” Journal of Philosophy **81**(May): 235-56.

Vineberg, S. (1997). “Dutch Books, Dutch Strategies and What They Show About Rationality.” Philosophical Studies **86**: 185-201.

Vineberg, S. (2001). “The Notion of Consistency for Partial Belief.” Philosophical Studies **102**: 281-296.