Everettian Confirmation and Sleeping Beauty

Wilson, Alastair

DOI:

[10.1093/bjps/axt018](https://doi.org/10.1093/bjps/axt018)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Wilson, A 2014, 'Everettian Confirmation and Sleeping Beauty', *The British Journal for the Philosophy of Science*, vol. 65, no. 3, pp. 573-598. <https://doi.org/10.1093/bjps/axt018>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is a pre-copyedited, author-produced PDF of an article accepted for publication in British Journal for the Philosophy of Science following peer review. The version of record [Alastair Wilson. Everettian Confirmation and Sleeping Beauty. Br J Philos Sci (2014) 65 (3): 573-598 first published online August 30, 2013 doi:10.1093/bjps/axt018] is available online at: <http://bjps.oxfordjournals.org/content/65/3/573.full.pdf+html>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Everettian Confirmation and Sleeping Beauty

27 August 2012

ABSTRACT

Darren Bradley has recently appealed to observation selection effects to argue that conditionalization presents no special problem for Everettian quantum mechanics, and to defend the ‘halfer’ answer to the puzzle of Sleeping Beauty. I assess Bradley’s arguments and conclude that while he is right about confirmation in Everettian quantum mechanics, he is wrong about Sleeping Beauty. This result is doubly good news for Everettians: they can endorse Bayesian confirmation theory without qualification, but they are not thereby compelled to adopt the unpopular ‘halfer’ answer in Sleeping Beauty. These considerations suggest that objective chance is playing an important and under-appreciated role in Sleeping Beauty.

1. *Introduction*
2. *Confirmation in EQM*
3. *Sleeping Beauty*
4. *The selection model*
5. *Bradley’s argument*
6. *The right route to 1/3*
7. *The breakdown of the analogy*
8. *Alternative diagnoses*
9. *God’s Gambling Game*
10. *Non-chancy Sleeping Beauty cases*
11. *Conclusion*

1. Introduction

In a recent BJPS paper (Bradley [2011]), Darren Bradley fans the flames of the debate over confirmation theory in Everettian quantum mechanics (EQM). Bradley argues that, by taking the centred nature of our evidence and the corresponding observation selection effects into account in the right way, Everettians can tell a plausible story about how EQM gets confirmed. The spectre of automatic confirmation of many-worlds theories by any evidence whatsoever, which has worried authors like Barrett, Myrvold and Greaves, is then dispelled without fuss; the ‘evidential problem’ for EQM is shown not to need a solution of the sort proposed by Greaves [2007a] and by Greaves and

Myrvold [2010], which involves a framework of ‘quasi-credences’ and an associated update rule called ‘quasi-conditionalization’. Even if quasi-credence and quasi-conditionalization might be required for us to make sense of the pre-measurement credential state of a subject who countenances EQM¹, such notions are unnecessary for modelling the experimental support that past measurement results provide for the theory. Or so Bradley argues.

This discussion of centred evidence and observation selection effects in the Everettian case, if correct, is significant enough by itself. But Bradley follows Peter Lewis² in drawing a close parallel between confirmation in EQM and the puzzle of Sleeping Beauty (SB). According to Bradley, Lewis is right to think that Everettians have to be ‘halfers’ when it comes to SB, but he is wrong to think that this presents a problem for Everettians. Bradley in fact believes he can demonstrate that halving is the right response to SB. These are highly controversial conclusions; how does Bradley reach them? In brief, his claim is that in each case the strongest new evidence available to the relevant subject has the form *I learned that X by a random method*. Given this conception of the evidence, Bradley argues that EQM is not automatically confirmed and that halving is the correct response to SB. The key thesis here is that the procedure of observation selection involved is in each problem a *random* one, in a sense to be explained below. In defence of this thesis, Bradley offers an argument which turns on the notion of a *biased* observation selection procedure.

In what follows I will argue that Bradley is right about EQM but wrong about SB. Proper attention to observation selection effects does rule out any automatic confirmation of EQM, but it does not generate the halfer position in SB. The cases are disanalogous because Everettian confirmation scenarios involve no equivalent of the chancy coin toss that governs Sleeping Beauty’s awakenings. This is doubly good news for Everettians; their theory can be confirmed by evidence in the usual way, but they need not be saddled with the unpopular halfer conclusion.

In §2 and §3 I set out the details of two cases: an Everettian confirmation scenario that I call ‘Quantum Wombat’, and the slightly modified Sleeping Beauty scenario that Bradley calls ‘Technicolor Beauty’. §4 explains how Bradley models these cases as selections from a population, and §5 reconstructs his argument, which uses the selection model to support halfer-friendly and Everettian-friendly conclusions. In §6, I explain how I think thirders should

¹ Wilson [forthcoming b] argues that they are not required even for this purpose.

² See Lewis [2007] for details. Papineau and Durà-Vilà [2009] reply to Lewis; further epicycles are Lewis [2009] and Papineau and Durà-Vilà [2009]. I discuss this exchange in Section 8; my own account of this disanalogy between Everettian confirmation and Sleeping Beauty is given in Section 7.

respond. §7 is the core of the paper: in it I argue that the analogy between SB and confirmation in EQM breaks down in a crucial way. In §8 I compare my account of the disanalogy to extant accounts, and in §9 I describe a bizarre variably-many-worlds theory which genuinely is analogous to SB. §10 applies my analysis to a variant on SB which lacks the chancy element of the original case; §11 is a conclusion.

2. Confirmation in EQM

Confirmation theory in EQM is perplexing because the theories to be compared – EQM and some candidate one-world stochastic theory ST – have systematically different consequences for the number of observers in existence. Consider the following case:

Quantum Wombat: Wombat is unsure whether EQM or ST is correct. He has just performed a spin measurement with possible outcomes Up and Down, but he has not yet looked at the result. According to EQM, after the measurement there are two observers, located on branches of equal weight³, one of whom observes Up and the other of whom observes Down. According to ST, after the measurement there is one observer, who observes either Up or Down, with equal probability. Wombat is not sure whether *i)* EQM is true and he is one of the two observers, or *ii)* ST is true and he is the only observer.

Pictorially:

³ The metaphysics and epistemology of branch weight are controversial: see part 3 of Saunders et al. [2011] and Wilson [forthcoming b] for discussion. However, for our purposes only the relatively straightforward equally-weighted case will be needed.

Suppose Wombat observes Up. Should he take this observation to support EQM over ST? Plausibly not: the two theories are usually supposed to be (at least approximately) empirically equivalent. However, the probability that *Up is observed by somebody* is 1 according to EQM but is less than 1 according to ST. Accordingly, the credential update rule which Greaves [2007a] and Bradley [2011] dub ‘naïve conditionalization’ appears to break down when one of the options on the table is EQM. The following line of thought is tempting: if the many-worlds theory predicts all possible outcomes, then no possible observation can disconfirm EQM; and since no one-world theory likewise predicts all possible outcomes⁴, every observation confirms EQM over ST. In Bradley’s words:

The Ancients could have worked out that they have overwhelming evidence for MWI merely by realizing it was a logical possibility and observing the weather.

Bradley [2011], p.336

Something has gone wrong. And notice that it wasn’t any assumption that EQM was correct that landed us in this mess; it was just the assumption that EQM might be correct. Compensating by setting prior credences in EQM lower is an unattractively brute-force response, since branching is so ubiquitous⁵. We can only safely ignore the apparent problem that EQM generates for confirmation theory if we are willing to set our prior credence in EQM arbitrarily low, effectively ruling out EQM *a priori*. This seems an uncomfortably dogmatic position for philosophers to adopt, given how seriously physicists take EQM.

3. Sleeping Beauty

Confirmation in EQM is analogous in certain ways to the Sleeping Beauty problem, introduced to philosophers by Elga [2000]. I will describe the variant⁶ that Bradley calls ‘Technicolour Beauty’. Here is the setup.

⁴ I leave aside what Bostrom [2001] calls ‘big-world cosmologies’: single-world theories in the traditional sense, but where a spatial or temporal infinity ensures that some plenitude of qualitative possibilities is realized.

⁵ The issue of how, if at all, to quantify branching in EQM is a vexed one: see Wallace [2012] and Greaves [2007b]. However, given that decoherence takes hold on timescales in the order of 10^{-20} seconds (Zurek [2003] gives some accessible models), any reasonable criterion will require initial prior credence in EQM to be well below 10^{-20} in order to insulate us from the automatic confirmation effect.

⁶ During the interval between waking and seeing the paper, Technicolour Beauty reproduces the original Sleeping Beauty scenario. Accordingly, a full solution to Technicolour Beauty determines a solution to the original Sleeping Beauty scenario, while a solution to the original Sleeping Beauty scenario constrains (though perhaps not uniquely) a solution to Technicolour Beauty. From this point on I will concentrate on Technicolour Beauty, and will refer to it as SB for short. I hope that the consequences of my argument for the original Sleeping Beauty scenario are clear.

Technicolour Beauty: Beauty will be put to sleep on Sunday night and a fair coin tossed. If the coin comes up Heads, Beauty will be woken on Monday. If the coin comes up Tails, Beauty will be woken on Monday and on Tuesday. Beauty’s memory of her Monday experience will be erased on Monday night; so each waking is initially subjectively indistinguishable from every other. However, shortly after each waking Beauty will be shown either a Red or a Blue piece of paper. If Tails comes up, she will be shown Red on Monday and Blue on Tuesday; if Heads comes up, she will be shown either Red or Blue on Monday, depending on the toss of a further fair coin. Beauty knows all this. Beauty sleeps. Beauty wakes. Beauty is shown the paper. It is Red. What should her credence be that the coin landed Heads, before and after seeing that the paper is Red?

Confirmation scenarios in EQM have a structural similarity to Technicolour Beauty. Let a spin measurement be made, which according to EQM results in an Up branch and a Down branch and according to ST results in a stochastic transition to either Up or Down. Then EQM predicts two branches containing Up and Down, just as Tails predicts two days containing Red and Blue; ST predicts one branch containing Up or Down, just as Heads predicts one day containing Red or Blue. Bradley provides the following table illustrating the analogy (he uses the acronym MWI to refer to EQM):

	Quantum Mechanics			Technicolour Beauty	
	Branch 1	Branch 2		Day 1	Day 2
ST	Up or Down		Heads	Red or Blue	
MWI	Up	Down	Tails	Red	Blue

Bradley [2011] p. 333

In Section 7 I will argue that the analogy between SB and confirmation in EQM breaks down; but I will pretend that it holds for the time being, while I set out Bradley’s central argument.

4. The selection model

Bradley models the procedure of making an observation as the taking of a sample from a population of locations of observation. A *location of observation* is an agent’s having a certain perceptual experience at a particular spacetime location (and, if EQM is correct, on a particular branch).

If ST is correct, then when Wombat observes the result of the spin measurement the population to be sampled contains one location of observation, which is either an Up location or a Down location. If EQM is correct, then when

Wombat observes the result of the spin measurement the population to be sampled contains two locations of observation, one of which is an Up location and the other of which is a Down location. Wombat's discovery of the result is then represented as him sampling himself from a population of two wombat-stages, one in each branch.

If the coin lands Heads, then the population to be sampled when Beauty observes the colour of the paper contains one location of observation, which is either a Red location or a Blue location. If the coin lands Tails, the population to be sampled when Beauty observes the colour of the paper contains two locations of observation, one of which is a Red location and the other of which is a Blue location. Beauty's discovery of the colour is then represented as her sampling herself from a population of two person-stages, one on Monday and one on Tuesday.

Note that there is no dubious generalized indifference principle being used here: it is not assumed that each epistemic possibility must be assigned the same credence. The route to $1/3$ in Sleeping Beauty which starts from the observation that there are three epistemically possible experiences and argues they must be equiprobable is not endorsed by any thirders I know of; and if anything, this argument is even less convincing in the case of Everettian confirmation. The population from which the observations are sampled is, in both examples, a population of experiences which do in fact occur: in different branches in one case and on different days in the other.

With this caveat, Bradley's sampling model looks innocuous. In sampling a single member from a population, one knows that the sample will be a member of the population, but one doesn't know which member it will be. If one does not know which population the sample is being drawn from, this creates further uncertainty. Prior to making their observations, Beauty and Wombat do not know which population their observation is to be drawn from, and they do not know which member will be drawn if the population contains two members. This is all there is to the sampling model, and I can't see any grounds for disputing its applicability.

Why does the sampling model help? Bradley's answer is that it highlights the role of the *observation selection effect* produced by the nature of the *observation selection procedure* that selects a sample from the population. We often know not only what the sample is – what the content of our location of observation is – but also how we acquired it – how the location was picked out from the population. An *observation selection procedure*, or a *method* for short, is a particular way of drawing a sample from a population of locations of observation. Bradley argues that when we make an observation of X , the total evidence we acquire is not just X occurred, and not even just *I learned that X*

occurred, but *I learned that X occurred by method M*. This of course entails that X occurred, but it can entail more or less in addition. For example, if method M only selects very small objects, then learning that it has selected a very small object is less revealing about the population than it would be if method M could select an object of any size. Insofar as knowledge of the method used affects the inferences that can be drawn after making a particular observation, we say that that method produces an *observation selection effect*.

Let's apply this to our examples. On Bradley's view, the evidence that Wombat gains when he sees either Down or Up has the form *I learned that X either by ST being true and X being the outcome produced by a stochastic process or by EQM being true and X being an outcome seen on one of the branches*. The evidence that Beauty gains when she sees either red or blue paper has the form *I learned that X either by Heads landing and X being seen on the one awakening, or by Tails landing and X being seen on one of the two awakenings*. In both cases this evidence is stronger than just *X*; and as Bayesians, we know always to conditionalize on the strongest evidence available⁷.

Not every method of sampling from a population is a random method. Bradley cites Eddington's famous example of the net which can only catch fishes larger than a certain size, and adds the example of an urn of balls with an opening through which only some of the balls can fit. Bradley contends that in the cases - Tails and EQM - where the populations sampled contain more than one member, we should treat the samplings as random. Then Beauty should be equally confident that it is Monday as that it is Tuesday conditional on the coin landing Tails, and Wombat should be equally confident that he is on an Up branch as that he is on a Down branch conditional on EQM being correct.

What is the probability of the new evidence acquired via our agents' observations, given the various hypotheses being tested? Bradley argues that in the quantum case the probability of seeing Up was 1/2, whether ST or EQM is true (likewise for Down); and that in the SB case the probability of seeing Red was 1/2, whether the coin landed Heads or Tails (likewise for Blue). Then in both cases the new evidence has the same probability on either hypothesis, so neither hypothesis is confirmed over the other in either case. Observing the colour of the paper does not provide Beauty with support for Tails; and observing the result of a spin measurement does not provide Wombat with support for EQM.

⁷ Failure to conditionalize on an agent's strongest evidence is a failure of rationality for Bayesians: it leads to fallacies of exclusion, as when someone argues from their unblemished safety record on the road to the conclusion that it would be safe to drive home, neglecting the eight pints of beer recently consumed.

The claim that the probability of whatever new evidence is acquired is $1/2$ on both of the hypotheses to be tested is the crux of Bradley’s argument. Once it is established, automatic confirmation of EQM drops out of the picture, and the argument that Bradley discusses for the third solution lapses. So we must examine the motivation for this key claim. I’ll initially consider only the case of Everettian confirmation; once I’ve ironed out some wrinkles in the argument, I’ll move to consider SB also.

5. Bradley’s argument

Bradley defines a few technical terms for use in his argument. An observation selection procedure is *biased towards X* iff: if X is in the population, then X is selected for the sample. (He notes that this might be more aptly called ‘maximal bias’.) An observation selection procedure is *random* if each item in the population has an equal chance of being selected for the sample.

Armed with this terminology, Bradley proceeds as follows.

1. If EQM is correct, the observation made by Wombat can be modelled as a sampling from a population of two locations of observation – the Up branch and the Down branch. (Premise.)
2. If EQM is correct, after the measurement Wombat is only located on one branch, so he only observes a single outcome; he observes either Up or Down but not both. (Premise.)
3. If the observation selection procedure which selects Wombat’s observation is biased towards both Up and Down, then both Up and Down are selected and Wombat observes both Up and Down. (From 1, definition of biased towards X.)
4. If EQM is correct, the observation selection procedure which selects Wombat’s observation cannot be biased towards both Up and Down. (From 2, 3.)
5. If a selection procedure is not biased towards both of its possible outcomes, then it is random. (?)
6. If EQM is correct, the observation selection procedure which selects Wombat’s observation is random. (From 4, 5.)
7. If ST is correct, the observation selection procedure which selects Wombat’s observation is random. (Premise.)
8. Wombat knows that either ST or EQM is correct. (Premise.)
9. The strongest new evidence acquired by an agent on making an observation is ‘I learned that X by method M’. (Premise.)
10. The strongest new evidence E_w acquired by Wombat on making his observation is ‘I learned that (Up or Down) by a random procedure.’ (From 6, 7, 8, 9.)
11. For a random procedure, the probability of any possible outcome is equal. (Definition of random procedure.)

12. The probability of Wombat's strongest new evidence given EQM is 1/2. (From 10, 11.)
13. The probability of Wombat's strongest new evidence given ST is 1/2. (From 10, 11.)
14. Observing the result of the measurement does not confirm EQM over ST. (From 12, 13.)

Wombat's observation – whatever the result – does not confirm EQM over ST, or *vice versa*. This is the intuitively correct result.

There is an obvious gap in the argument as stated, and it is rather surprising that Bradley fails to plug it. In an introductory section he explicitly notes that biased procedures and random procedures do not exhaust the possible types of procedure; yet he moves freely from the conclusion that the observation selection procedure which selects Wombat's observation is not biased in both directions, to the conclusion that it is random⁸:

Suppose instead that there is no bias. That is, we have a random procedure – Up branches have the same probability of being observed as Down branches, so you are just as likely to be in an Up branch as a Down branch.

Bradley [2011], p. 331

For all that has been established so far, the procedure which selects Wombat's observation could be biased towards Up and not towards Down, or be biased towards Down and not towards Up, or be biased towards neither outcome but be such that Up is selected with probability 0.99. (Recall that 'bias', for Bradley, means *maximal* bias: if X is in the population then X is definitely in the sample.) So there seems to be no good motivation for 5, the premise that if a selection procedure is not biased in both directions then it is random. Bradley moves from lack of maximal bias in both directions (a trivial claim) to equiprobability (a highly non-trivial claim).

Bradley apparently intends⁹ this gap in the argument to be filled by the restricted principle of indifference proposed by Elga [2000]. This principle, which Bradley devotes the final section of his paper to defending, is that two centred

⁸ Bradley [2007] uses 'random' to mean not-maximally-biased. The argument given above does establish that the procedure is random in this weaker sense. But it doesn't establish that the procedure is random in the sense specified in the quote in the main text; and it is this stronger conclusion which is needed to establish 12 and 13.

⁹ Bradley does say that he doesn't want to rely too much on this principle, because the thesis that the selection procedures in question are not maximally biased in either direction doesn't require it. However, the (much stronger) thesis that the selection procedures are random does appear to require the indifference principle.

worlds deserve equal credence when they correspond to the same uncentred world and the agents at the centres are ‘subjectively indistinguishable’. However, this particular indifference principle is a poisoned chalice for Everettians since it conflicts with the Born rule, which says that agents should apportion their credences to the branch weights. Probabilities must be given by the Born rule if the predictions of EQM are to agree with the predictions of conventional quantum theory; and so Everettians have made sustained attempts to argue for it (see e.g. Deutsch [1999], Saunders [2004], Wallace [2002, 2006, 2010, 2012], Wilson [forthcoming b]). Bradley stipulates that his arguments should only be taken as applying to the case where the two branches have equal weight, thereby avoiding clashing with the Born rule; but *prima facie* Elga’s restricted principle of indifference applies to cases of unequal weight as well as to cases of equal weight.

I think the best way to patch up Bradley’s argument in an Everettian-friendly fashion is to appeal to the symmetry of the setup. In the case of equal weight, the physical situation is exactly symmetric between the Up outcome and the Down outcome. There are two possible outcomes; exactly one of them will be observed by Wombat; *there is no fact (such as an unequal weighting of the branches) which could provide a reason why one outcome and not the other is observed by Wombat*; so the selection procedure by which Wombat’s observation is selected must be random, on pain of arbitrariness. Bradley does hint at this line of reasoning:

Up branches are just as hospitable to life, and just as likely to be observed as Down branches.

Bradley [2011], p. 331

Inserting this sort of appeal to symmetry into Bradley’s argument gives rise to the following modified argument, which can safely be endorsed by Everettians:

15. If EQM is correct, the observation made by Wombat can be modelled as a sampling from a population of two locations of observation – the Up branch and the Down branch. (Premise.)
16. If EQM is correct, after the measurement Wombat is only located on one branch, so he only observes a single outcome; he observes either Up or Down but not both. (Premise.)
17. If the observation selection procedure which selects Wombat’s observation is biased towards both Up and Down, then both Up and Down are selected and Wombat observes both Up and Down. (From 15, definition of biased towards X.)
18. If EQM is correct, the observation selection procedure which selects Wombat’s observation cannot be biased towards both Up and Down. (From 16, 17.)

19. If a selection procedure is not biased towards both of its possible outcomes, and it is symmetric with respect to its possible outcomes, then it is random. (Premise.)
20. If EQM is correct, the observation selection procedure which selects Wombat's observation is symmetric with respect to its possible outcomes. (Premise.)
21. If EQM is correct, the observation selection procedure which selects Wombat's observation is random. (From 19, 20.)
22. If ST is correct, the observation selection procedure which selects Wombat's observation is random. (Premise.)
23. Wombat knows that either ST or EQM is correct. (Premise.)
24. The strongest new evidence acquired by an agent on making an observation is 'I learned that X by method M'. (Premise.)
25. The strongest new evidence E_W acquired by Wombat on making his observation is 'I learned that (Up or Down) by a random procedure.' (From 21, 22, 23, 24.)
26. For a random procedure, the probability of any possible outcome is equal. (Definition of random procedure.)
27. The probability of Wombat's strongest new evidence given EQM is $1/2$. (From 25, 26.)
28. The probability of Wombat's strongest new evidence given ST is $1/2$. (From 25, 26.)
29. Observing the result of the measurement does not confirm EQM over ST. (From 27, 28.)

The same argument can be applied to SB.

30. If the coin lands Tails, the observation made by Beauty can be modelled as a sampling from a population of two locations of observation: the Red day and the Blue day. (Premise.)
31. If the coin lands Tails, on awakening Beauty is only located on one day, so she only observes a single outcome; she observes either Red or Blue but not both. (Premise.)
32. If the observation selection procedure which selects Beauty's observation is biased towards both Red and Blue, then both Red and Blue are selected and Beauty observes both Red and Blue. (From 30, definition of biased towards X.)
33. If the coin lands Tails, the observation selection procedure which selects Beauty's observation cannot be biased towards both Red and Blue. (From 31, 32.)
34. If a selection procedure is not biased towards both of its possible outcomes, and it is symmetric with respect to its possible outcomes, then it is random. (Premise.)

35. If the coin lands Tails, the observation selection procedure which selects Beauty's observation is symmetric with respect to its possible outcomes. (Premise)
36. If the coin lands Tails, the observation selection procedure which selects Beauty's observation is random. (From 34, 35.)
37. If the coin lands Heads, the observation selection procedure which selects Beauty's observation is random. (Premise)
38. Beauty knows that the coin lands either Heads or Tails. (Premise)
39. The strongest new evidence acquired by an agent on making an observation is 'I learned that X by method M'. (Premise)
40. The strongest new evidence E_B acquired by Beauty on making her observation is 'I learned that (Red or Blue) by a random procedure.' (From 36, 37, 38, 39.)
41. For a random procedure, the probability of any possible outcome is equal. (Definition of random procedure.)
42. The probability of Beauty's strongest new evidence given Heads is $1/2$. (From 40, 41.)
43. The probability of Beauty's strongest new evidence given Tails is $1/2$. (From 40, 41.)
44. Observing the colour of the paper does not confirm Tails over Heads. (From 42, 43.)

How should thirders respond to this argument? I think they should grant the conclusion, but deny that it undermines their position. Bradley's argument is effective against someone who maintains that Beauty should have an initial credence in Heads of $1/2$, but who also maintains that this credence should alter to $1/3$ as soon as she sees any evidence which she is not guaranteed to see (such as the red or blue paper, or cloud patterns outside her window). However, this would be an unattractive position for a thirder to adopt. Sensible thirders ought to say that Beauty's initial credence on waking, before she opens her eyes and sees anything, should already be $1/3$, and that subsequent evidence should not automatically alter that credence.

What Bradley's argument does achieve, in the case of SB, is to show that thirders cannot motivate their position by arguing that some new evidence that Beauty acquires after awakening supports Tails over Heads. So Bradley's argument does undermine (Bradley's reconstruction of) the view of Titelbaum [2008], which motivates the answer $1/3$ via the thought that Red is guaranteed to be seen if the coin lands Tails. But thirders can, and should, motivate their solution differently.

6. The right route to $1/3$

In this section I will say how I think the $1/3$ answer should be motivated for the case of Technicolour Beauty. This solution will also entail that the $1/3$

answer is correct in the original Sleeping Beauty scenario. The resulting position is essentially that of Elga [2000], although I put a slightly different gloss on it.

Thirder should agree that E_B is the strongest new evidence that Beauty acquires on making her observation of the coloured paper. And they should agree that the probability of E_B given Heads is the same as the probability of E_B given Tails, and hence that Beauty's observation of the colour of the paper does not provide her with support for Tails. But they should deny that the total evidence that Beauty has after observing the colour of the paper is her evidence on Sunday night plus E_B . For Beauty's evidential state has undergone an additional change overnight. She has lost centred evidence, evidence about which day it is. As a result of this evidence loss, she may no longer treat the result of the coin toss as *effectively chancy*, in a sense to be explained below¹⁰.

Both my claim that Beauty has lost evidence and my claim that this has consequences for her credences about the outcome of the coin toss are contentious. In this section I will sketch the basis for thinking that Beauty has lost evidence, and then explore the consequences of this claim for Beauty's credential state.

Take first the claim that Beauty has lost irreducibly centred evidence overnight on Monday. This evidence loss can be most straightforwardly expressed by saying that she has lost evidence about *which day it is now*; it cannot be expressed in any fragment of language which does not contain indexical terms. Before going to sleep Beauty knew what day it was then – it was Sunday. On awakening she does not know which day it is now – it could be Monday or Tuesday. Normally, as the days pass without our losing track of their passing, we lose some centred evidence but gain corresponding new centred evidence. In cases like Sleeping Beauty, or like the case of O'Leary described by Elga [2004], there is a loss of irreducibly centred evidence which is not adequately compensated for. See Dieks [2007] for further discussion of this point.

Of course, to point to a way in which Beauty's evidential state changes overnight is not yet to show that she should have credence $1/3$ in Heads on waking. The change in evidential state might simply be irrelevant, or it might be relevant in a way which results in a credence other than $1/3$ in Heads. This brings us to the second claim; that the overnight change in evidential state means that Beauty should no longer regard the coin toss as effectively chancy. To explain what I mean by this, it will help to start by rehearsing one of the arguments given in Elga [2000] for the answer $1/3$.

¹⁰ See p.16 for the precise definition of this term.

At the heart of this argument are three assumptions¹¹. The first assumption is that chance is the norm of credence: that in a situation where an agent knows the chances and has no inadmissible information¹², the agent's credences should match the chances. Where a fair coin toss is in the future, an agent cannot have inadmissible information about it without the help of precognition or some other form of backwards causation.

The second assumption is that it doesn't matter when the coin is tossed, as long as it has been tossed by Tuesday morning. In the original version of SB, the coin is tossed on Sunday night. But since Beauty will be woken on Monday regardless, and the procedural need for the information about the result of the coin does not arise until the experimenters are deciding whether to wake Beauty on Tuesday morning, the coin could equally well be tossed on Monday night. Beauty knows this, so she ought to adopt the same credence on waking if she knows the coin is tossed on Monday as she adopts if she knows the coin is tossed on Sunday. (Indeed, it should make no difference to her credences if the coin is tossed even earlier, say on Saturday night.)

The third assumption, which Elga motivates via his restricted principle of indifference, is that the probability of it being Monday conditional on Tails is the same as the probability of it being Tuesday conditional on Tails. Whatever we may think about the restricted principle of indifference in general (and I have suggested above that it would be a mistake for Everettians to adopt it unrestrictedly), this particular application of it seems unproblematic. Beauty cannot distinguish between a Monday waking given Tails, and a Tuesday waking given Tails. She thus has no reason to assign them different credences, and so her credence that it is Monday conditional on Tails should be the same as her credence that it is Tuesday conditional on Tails.

Together, these assumptions entail that the thirder solution is correct. In the case in which the coin is tossed on Monday night, the chance on Monday of the coin landing Heads is $1/2$. Beauty knows this, so (by the first assumption) on awakening her credence in the coin landing Heads conditional on it being Monday is $1/2$. On awakening, Beauty does not know whether *a*) it is Monday (in which case the toss is in the future, and the chance of Heads and of Tails are both $1/2$) or *b*) it is Tuesday (in which case the toss has already occurred, and the chance of Heads is 0 and the chance of Tails is 1).

¹¹ One further assumption, denied *inter alia* by Lewis [2010], is that on awakening Beauty's credence that *the coin lands Heads* cannot come apart from her credence that *this is a Heads-waking*. The motivation for this assumption ought to be clear: the setup of the case ensures that on awakening Beauty knows that these two propositions have the same truth-value, so she should assign them equal credence.

¹² See Lewis [1980] and Hoefer [2007] for classic discussions of the notion of inadmissible information.

Using a standard notion for subjective credence: $\text{Cr}(\text{Heads}|\text{Monday}) = 1/2$, by the Principal Principle. $\text{Cr}(\text{Heads}|\text{Tuesday}) = 0$, since Beauty is awakened on Tuesday only if the coin lands Tails. So if $\text{Cr}(\text{Tuesday}) > 0$, then $\text{Cr}(\text{Heads}) < 1/2$. This is already enough to establish the falsity of the halfer conclusion. By the third assumption and an application of Bayes' rule, the thirder conclusion follows.

The above argument applied to the case in which the coin is tossed on Monday night. But by the second assumption, the thirder conclusion is correct for the case where the coin is tossed on Monday night if and only if it is correct for the case where the coin is tossed on Sunday night. So the thirder conclusion follows for these cases also.¹³

On awakening, Beauty's uncentred evidence does not change, but her centred evidence does change: Elga's argument shows how this change results in a shift in credence in Heads from $1/2$ to $1/3$. The remaining difficulty is that of explaining why this evidential change should produce such a credential shift. After all, we are constantly undergoing changes in centred evidence which have no significant effect on our credences in uncentred propositions.

I suggest that the difference is that in SB the overnight loss of centred evidence is hooked up to a chancy memory-loss process. This means that different inferences can be drawn from the available self-locating evidence before and after going to sleep on Sunday. To explain this, I think we can usefully appeal to the notion of a proposition's being *effectively chancy*. On Sunday night, it is effectively chancy for Beauty whether Heads; on awakening, it is no longer effectively chancy whether Heads.

Effective chanciness isn't the same as chanciness in the sense of Lewis [1980], according to which all past propositions have chance zero and the chance of a proposition cannot vary from person to person. For a proposition to be effectively chancy is both a person-dependent and a time-dependent matter. The same proposition can be effectively chancy for one person and not for another (as when the latter but not the former has observed the result of a coin toss), and it can be effectively chancy for a given person at one time and another (before and

¹³ Arntzenius [2003] resists this argument on the grounds that Beauty knows she may suffer a 'cognitive mishap' on Monday night, and hence that she will not necessarily remain an ideal Bayesian agent. But it is unclear why this means that she should not do the best she can, by conditionalizing on her strongest evidence whenever possible. Arntzenius also objects that the argument hooks up Beauty's credence in Heads to the issue of whether she accepts causal or evidential decision theory. I do not find this connection either surprising or worrisome. (Nor does Briggs [2010].)

after observing the result of the coin toss, respectively). As I use the phrase, for a proposition P to be effectively chancy for agent A at time t is for P to be or have been chancy, and for A to possess no evidence at t which is or might be *inadmissible* with respect to A .

Since effective chanciness is an agent-dependent and time-dependent business, evidence that some proposition is effectively chancy is irreducibly-centred evidence. It is evidence that the proposition is effectively chancy *now, for me*. Such evidence will not be treatable in a framework where all evidence takes the form of untensed propositions with timeless truth-values. This is no great surprise: in response to SB, several authors have attempted to develop more general accounts of credential updating which can cope adequately with centred evidence. Such accounts can be found in Titelbaum [2008], Meacham [2008], Bradley [2011], Schulz [2010] and Schwarz [2012].

I assume that if a proposition P is effectively chancy for A at t , then at t A is bound *indirectly* by the Principal Principle with respect to T . Even though P may no longer be chancy at t , if P was chancy at some earlier time and if A both knows the chance that P had and has no evidence inadmissible with respect to P , then at t A is still bound to match her credence to the value that the chance of P is known to have had.

The notion of effective chanciness gives us the resources to say what sort of change in Beauty's evidence occurs overnight on Sunday. Where on Sunday night the result of the coin toss is effectively chancy for Beauty, when she awakes on Monday it is no longer effectively chancy. For all she knows on awakening, it could now be Tuesday, in which case the coin would have to have already been tossed and landed Tails. In that case, Beauty's new centered evidence on awaking is inadmissible with respect to the coin toss. So Beauty has lost evidence that the result of the coin toss is an effectively chancy matter, because she has lost evidence that she has no inadmissible evidence. As a result, she is no longer constrained by the Principal Principle to match her credences to the objective chances.

Since the only scenario in which Beauty possesses any inadmissible evidence is a scenario in which the coin landed Tails and it is Tuesday, learning that the coin toss is no longer effectively chancy should alter her credences in the direction of Tails. By the second assumption of the argument given above and an application of Bayes' rule, she should end up with credence $1/3$ in Heads.

Armed with this result, we can consider what would happen if Beauty learns that it is Monday. When she does, she learns that she has no evidence that is inadmissible with respect to the coin toss. She thereby re-acquires the evidence she lost overnight: that the coin toss is effectively chancy. Accordingly, she is once again bound by the Principal Principle with respect to the chance of

Heads, and must adopt credence $1/2$ in Heads. Although this line of thought has been developed in the context of Technicolour Beauty, it is essentially the same as Elga's solution to the original Sleeping Beauty case.

One way out for halvers is to maintain that it really matters whether the toss is past or future. Halvers are committed to the idea that if Beauty learns that it is Monday, she should increase her credence to $2/3$ that the coin will land Heads. Since it is crazy to think of a future fair coin toss that it is $2/3$ likely to land Heads¹⁴, halvers must say there is a difference between the toss-on-Sunday and toss-on-Monday cases. The resulting view has it that in a case in which Beauty knows the coin is tossed Sunday night, halving is the correct solution; but that in a case in which Beauty knows the coin is tossed Monday night, thirding is the correct solution. This requires Beauty to bet differently depending on whether a coin toss is future or past, even when she knows that the toss is fair¹⁵. I will not investigate here how halvers might seek to sweeten this pill.

In this section I have shown how thirders can motivate their position while simultaneously accepting Bradley's conclusions about the nature of the new evidence Beauty acquires on seeing the coloured paper, and about the bearing of this evidence on the result of the coin. The next section argues that the analogous move fails in the case of automatic confirmation in EQM. The analogy between Sleeping Beauty and Quantum Wombat breaks down, since there is no chanciness in the latter case to correspond to the chanciness of the coin toss.

7. The breakdown of the analogy

Beauty loses relevant centred evidence, and therefore should not continue to treat the result of the coin toss as effectively chancy. However, Wombat loses no similarly-relevant centred evidence. Wombat does not know, in advance of performing the spin measurement, that whether ST or EQM is true will be fixed by the result of a future fair coin. Indeed, Wombat knows that this is *not* the case – whichever theory is true is already true, and has chance 1. As a result, it is not the case that whether ST or EQM is true ceases to be an effectively chancy matter on performing the spin measurement, and Wombat's credences in ST and EQM should not change. EQM is not automatically confirmed.

¹⁴ Crazy or not, this was the view defended by Lewis [2001]. According to Lewis, on discovering that it is Monday Beauty acquires knowledge which is inadmissible with respect to a future coin toss. Given that SB involves no precognition or any other variety of backwards causation, this suggestion is implausible.

¹⁵ What if the time of the toss itself is unknown: for example, if it will occur during the first 10 minutes of the Monday waking? What if it is unknown in a self-centred way: if it will occur 10 minutes after the Monday waking, and there is no clock in Beauty's room? I leave these awkward questions to halvers.

The disanalogy between SB and QW is, on reflection, a straightforward one. Whether EQM or ST is true does not depend on any chance process, and Wombat knows that. In contrast, whether the coin lands Heads or Tails does depend on a chance process, and Beauty knows that. Consequently Beauty loses relevant evidence when she is put to sleep, but Wombat loses no relevant evidence on performing the measurement. The effect which generates the answer $1/3$ in the case of Sleeping Beauty is absent in the case of Everettian confirmation scenarios.

Wombat is correct to think that the probability of whichever experimental result he sees given ST is the same as the probability of his new centred evidence given EQM. And Beauty is correct to think that the probability of whichever colour of paper she sees is the same given Heads as it is given Tails. But Wombat can, after the measurement, still take into account all the evidence he had before the measurement that was relevant to EQM vs. ST. Beauty cannot, after awakening, still take into account all the evidence she had before going to sleep that was relevant to Heads vs. Tails. For some of that was irreducibly centred evidence: that the result of the coin toss was effectively chancy.

Bradley at one point implies that Wombat does lose information, endorsing Lewis' evocative phrase 'gets lost in the branches' (Bradley [2011] p. 333). The idea seems to be that before the spin measurement, Wombat knows exactly where he is – on *the* initial branch, about to press *the* button, while afterwards he could be in either of two places:

The problem with this response is that the information loss in question is irrelevant. It doesn't provide evidence which bears on whether a particular proposition is currently effectively chancy, so it does not give rise to any third-style shift in credence.

The irrelevance of the self-locating uncertainty which Lewis and Bradley point to is underlined by considering the *diverging* interpretation of EQM, defended by Saunders [2010] and by Wilson [forthcoming a]. According to overlapping EQM, different branches have earlier segments in common: the ‘splitting worlds’ metaphor is apt. According to diverging EQM, different branches are mereologically isolated: the ‘parallel worlds’ metaphor is apt. Diverging worlds may match one another up to a time, but they have no segments in common. Wombat, pre-measurement, is on one branch and one branch only, although he has no idea whether it is a branch on which Up will be measured or a branch on which Down will be measured. He is just as lost in the branches as he is after the measurement. Pictorially:

The evidence which defenders of the analogy must suppose that Wombat loses is evidence which, on the diverging picture, Wombat never even had. But the choice between branching and diverging versions of EQM, even if it is relevant to pre-measurement uncertainty¹⁶, should not be relevant to the bearing of past observations on the likelihood of a many-worlds theory. It is a problem with the analogy that it says that the choice between branching and divergence is relevant in this way: the analogy entangles issues of metaphysics and epistemology which are better kept apart.

I have argued that the source of the disanalogy between confirmation of EQM and Technicolour Beauty is that whether Heads or Tails is true is determined by a chance process, while whether EQM or ST is true is not so determined. In the next section I compare this suggestion to some alternative accounts of the disanalogy.

¹⁶ Wilson [forthcoming a] argues that it is.

8. Alternative diagnoses

Other authors to diagnose disanalogies between Everettian confirmation and SB are Papineau & Durà-Vilà [2009a, 2009b] and Peterson [2009]. As far as I can see, these authors give competing accounts of the disanalogy, both of which differ from my own account. The issue is complicated because these authors primarily discuss the ‘simplified Sleeping Beauty’ case presented in Lewis [2007], which involves no coin toss but rather has Beauty awaken on both Monday and Tuesday, with her memory erased in between. The simplified Sleeping Beauty case is not analogous to Everettian confirmation scenarios, but is supposed to be analogous to branching events for agents who are certain of the truth of EQM.

Papineau & Durà-Vilà take the difference between the simplified SB case and the case of Everettian branching to be that there are ‘two branches of reality after the spin measurement’, while there is ‘only one branch of reality in the simplified Sleeping Beauty case’. They claim that this allows Everettians to rationally assign credences of less than 1 to each outcome of the spin measurement, while Beauty in the simplified scenario must assign credence 1 to waking on both Monday and Tuesday. However, this claim appears to rest on a conflation of the probability *that someone wakes on Monday* with the probability *that it is now Monday* for a just-awakened agent. Papineau and Durà-Vilà correctly claim that the probabilities *that someone wakes on Monday* and *that someone wakes on Tuesday* are both 1 in the simplified Sleeping Beauty case. But likewise, the probabilities *that someone sees Up* and *that someone sees Down* are both 1 in the case of Everettian branching. Papineau & Durà-Vilà have not explained why the difference between days and branches leads to a difference between Everettian branching and the simplified Sleeping Beauty scenario.

Once we resist the conflation of the hypotheses *that it is now Monday* and *that someone is awake on Monday*, then (as Lewis [2009] replies) it is altogether unclear why the metaphysical difference between days and branches to which Papineau and Durà-Vilà point is a relevant difference. After the measurement but before the results have been examined, Wombat knows that – if EQM is correct – he is one of two observers, but does not know which. Likewise, after awakening Beauty knows that – if the coin landed Tails – she is on one of two days, but does not know which.

I suspect that the reason that Papineau and Durà-Vilà locate the disanalogy where they do is that they have been distracted by the pre-measurement credential state of the Everettian subject. On the account of Everettian probability defended by Papineau (see Papineau [1996], [2010]), the pre-measurement subject should set her credences according to the branch weights even though she knows, with certainty, that there will be an Up branch and that

there will be a Down branch. But the analogy between SB and confirmation in EQM can be made out entirely in terms of the post-measurement state of the observer, so pre-measurement credences are beside the point.

Peterson [2011] offers an alternative diagnosis of the disanalogy. Although his discussion is rather hedged, the suggestion appears to amount to this. On awakening, Beauty is uncertain about whether she has been previously awakened; however, on making a measurement, an agent certain that EQM is correct is sure that there is another branch on which another agent observes another result. Beauty's two wakings are connected in some way - Peterson suggests either by personal identity or by causal or counterfactual dependency - in which two agents in different branches are not. Why is this connection epistemically relevant? According to Peterson, it is because this connection provides something additional for Beauty to be uncertain about - whether or not she has been previously awakened. Unfortunately, Peterson doesn't say why this additional uncertainty should affect Beauty's credences about the result of the coin. So as things stand, his account of the source of the disanalogy is no advance on that of Papineau and Durà-Vilà.

The diagnosis of the disanalogy that Peterson offers does have a distinctive consequence, which points to a potential way of motivating this differential treatment of times and worlds. Following through on his account of the disanalogy, Peterson suggests that in SB variant cases where the agents who awaken after a toss of tails are distinct subjects, rather than being distinct time-slices of the same subject, the motivation for being a thirder disappears:

...in the Beauty case, the thirder's position seems correct; however, were the situation to involve not one person waking up twice but two successors being awoken once each... the halfer solution would be the correct one.

Peterson [2009]

A similar view is defended by Schwarz [MS], who argues that the correct answer to SB depends on whether it is construed as an episode of fission or not - that is, on whether the Tuesday waker should be regarded as a 'successor' of the Monday waker. Schwarz motivates the view from general considerations about diachronic belief updating¹⁷; and Peterson could help himself to this sort of motivation.

However, the resulting picture is highly problematic. Personal identity over time is a vague and arguably context-dependent business, and we can run a

¹⁷ Related arguments are given by Meacham [2010]. Meacham has confirmed (p.c.) that he accepts Schwarz' conclusions about rational belief in fission cases, though he'd resist Schwarz' further claims about chance and admissibility.

Sorites series from cases in which the sleeper who awakes on one day is definitely continuous with the sleeper who awakes on the next day to cases in which they are definitely distinct. Peterson and Schwarz are committed to saying that somewhere along this series there is a sharp discontinuity, where the credence that the subjects should have alters from $1/3$ to $1/2$, even though the subjects may be unaware of where this discontinuity lies.

Moreover, this view does not altogether escape the pathological consequences associated with the halfer solution. For example, Schwarz admits that on his account one of two fission products, about to toss a fair coin to determine whether the other fission product will be destroyed before awakening, should have $2/3$ confidence that the fair coin will land Heads (Schwarz [MS] p.22). I take this to be a fatal difficulty for the approach; so I conclude that Peterson has not correctly located the disanalogy between Everettian confirmation and SB.

9. God's Gambling Game

In Section 7 I argued that the analogy between SB and confirmation in EQM breaks down because the result of the coin toss depends on a chancy process, while whether EQM or ST is correct does not depend on any chancy process. The analogy can be restored if we consider a modified branching theory which I will call *God's Gambling Game*.

God's Gambling Game. Whenever a spin measurement is made, God tosses a fair coin to select either Down or Up. He then creates a branch in which this selected outcome is observed. He then tosses a further fair coin. If it lands Heads he rests, his creative work done. If it lands Tails he creates a second branch, in which the outcome not selected by the first coin toss is observed.

This scenario is illustrated below:

Call the theory that God does carry out the procedure just outlined GGG. Since this situation genuinely is analogous to the Technicolour Beauty scenario, does an analogue of the thirder effect arise? The answer is yes. If an agent knows that GGG is true, then after making a measurement, then even before he observes the result, that agent should become uncertain as to whether he is on the first branch to be created (in which case the second coin might have landed Heads or Tails), or whether he is on the second branch (in which case the second coin must have landed Tails.) On making the measurement, then, the agent loses the evidence that whether there will be one branch or two is an effectively chancy matter.

Does this lead to a problematic automatic confirmation for GGG? The answer is no. It's true that, via the standard thirder reasoning run through above, evidence that experiments have taken place does raise the probability that the world has multiplied *conditional on GGG*. But this result is irrelevant to our credence in GGG itself. Reason to reassign credence amongst the various ways things could go conditional on a theory isn't reason to change unconditional credence in that theory. This applies equally to Sleeping Beauty: if Beauty isn't certain that the procedure will go as promised (if she thinks maybe the experimenters were lying about what they'd do), then awakening doesn't provide Beauty with evidence – for example – that the coin used was fair.

Automatic confirmation of the many-worlds hypothesis over the single-world hypothesis for agents who are certain of GGG is altogether unsurprising, and it need cause no alarm for Bayesian confirmation theorists. We have no good reason to assign any significant credence to GGG, so it matters not at all what our credences in the universe having branched conditional on GGG might be. If we were somehow to come to believe that God tosses a coin to decide, for every quantum interaction, whether the outcome was to be chancy or branchy, then we would indeed have reason to favour the hypothesis that there are many branches over the hypothesis that there is one branch. But this is all exactly as it should be.

10. Non-chancy Sleeping Beauty cases

The foregoing arguments have an interesting upshot for variants of Sleeping Beauty cases which lack the chancy element of the original case. In the *mathematical Sleeping Beauty case* (MSB), uncertainty about the result of a fair coin toss is replaced by uncertainty about the truth of a mathematical proposition. On Sunday night Beauty has credence $1/2$ that Fermat's Last Theorem is true. She will be awakened on Monday if the theorem is true, and on both Monday and Tuesday (again with her memories from Monday erased) if the theorem is false. Beauty knows all this. The puzzle is to say what credence Beauty should have on Monday in the proposition that Fermat's Last Theorem is true (call this proposition True.)

The setup of MSB ensures that Beauty is not a fully rational Bayesian agent, since such agents are required to be logically omniscient. Is there then any sense in asking what Beauty's credences *ought to be*? – after all, in one obvious sense, her credence in True ought to be 1. I think this question does make sense; it is common to distinguish diachronic and synchronic constraints on credences, and to evaluate an agent's performance with respect to the diachronic constraints independently of her performance with respect to certain synchronic constraints (such as logical omniscience). Epistemologists had certainly better hope that something like this is possible: we're not in fact perfect Bayesians, and agents with non-trivial degrees of belief in mathematical propositions need to be modelled by any plausible epistemology for mathematics.

The difference between SB and MSB may seem somewhat incidental. General considerations relating to topic-neutrality perhaps make it natural to assume that the two puzzles have the same solution. Regardless, discussions of SB have usually paid little attention to the role that chance plays in the story¹⁸. SB is usually described as a puzzle about self-location; and our two cases seem to involve structurally similar self-locating uncertainty. In both SB and MSB, when she awakens on Monday Beauty is unsure whether it is Monday or Tuesday.

In this section I will give some reasons for suspecting that the two puzzles may in fact have different solutions. Two of the most powerful arguments for the answer $1/3$ in SB (both of which appear in Elga's original article) are inapplicable to MSB. Moreover, the analogy between SB and confirmation in EQM gives us reason to prefer the answer $1/2$ to MSB.

Consider first the *long-run frequency argument* for the answer $1/3$. Beauty knows that if SB is repeated infinitely many times, the ratio of Tails-awakenings to Heads-awakenings will (with chance 1) tend to 2:1. Since any awakening is indistinguishable from any other, her credence that a randomly-chosen awakening is a Heads-awakening should be $1/3$. This argument lapses in MSB. The truth-value of mathematical propositions cannot vary between awakenings, so Beauty knows that in the long-run either all the awakenings will be True-awakenings or all the awakenings will be False-awakenings¹⁹.

¹⁸ For example, Christopher Meacham remarks that 'the chanciness of the coin toss only plays a superficial role in the argument.... the argument goes through just as well if heads and tails are replaced by two different hypotheses we have other reasons for having $1/2$ / $1/2$ credences in.' (Meacham [2008]).

¹⁹ What if we pose different mathematical questions on each run of the experiment? This variable-question mathematical case will probably pattern with SB rather than with MSB, since it seems inevitable that the procedure by which the questions are selected will involve a chance process at some stage or other. Be that as it may, in this paper I will only address MSB.

Consider now the *Principal Principle argument* for the answer $1/3$, set out in detail in Section 6. This argument too lapses in MSB. Since mathematical propositions have objective chances of either zero or one²⁰, the constraint that the Principal Principle places on credence in mathematical propositions amounts to the requirement of logical omniscience, which is explicitly suspended in MSB. Accordingly, nothing in the setup of MSB prohibits Beauty from adopting credence $1/3$ in True on being informed that it is Monday, and hence nothing prohibits her from adopting credence $1/2$ in True when she initially awakens on Monday.

I conclude that there are important disanalogies between MSB and SB. Two influential arguments for the thirder position in SB are inapplicable to the mathematical case. This provides reason to doubt whether our two cases have a uniform solution. Moreover, the analogy between MSB and confirmation in EQM provides an indirect argument for the answer $1/2$ in MSB.

11. Conclusions

The analogy between SB and Everettian confirmation breaks down in a crucial way. Whether Heads or Tails is true depends on the outcome of a chance process, while whether EQM or ST is true does not depend on the outcome of any chance process. Consequently, SB involves the loss of relevant evidence – evidence that the result of the coin is effectively chancy – while Everettian confirmation scenarios involve no such loss. Although Bradley’s argument fails to establish the halfer conclusion in the case of SB, it does succeed in dispelling the spectre of automatic confirmation of EQM. Everettians need not adopt any novel confirmation theory; given a correct conception of the evidence, EQM is confirmed just in the same way as ordinary stochastic theories.

This result is doubly favourable to Everettians. If I am right, then Bayesian confirmation theory can be combined with EQM in a straightforward and non-pathological way, and Everettians are not obliged to adopt the unpopular halfer position in SB. Opponents of EQM would do better to target its metaphysics than to target its epistemology.

²⁰ This is so even according to ‘compatibilist’ conceptions of chance such as those of Eagle [2011], Glynn [2010] and Handfield & Wilson [2014].

References

- Arntzenius, Frank [2003]. 'Reflections on Sleeping Beauty'. *Analysis* 62(1): 53-62.
- Bostrom, Nick [2001]. *Anthropic Bias: Observation Selection Effects in Science and Philosophy*. London: Routledge
- Bradley, Darren [2007]. 'Four Problems about Self-Locating Evidence'. *Synthese* 182(3): 393-411.
- Bradley, Darren [2011]. 'Self-location is no problem for conditionalization'. *Synthese* 182(3): 393-411.
- Bradley, Darren [2011]. 'Confirmation in a Branching World: The Everett Interpretation and Sleeping Beauty'. *British Journal for the Philosophy of Science* 62: 323-342
- Briggs, Rachael [2010]. 'Putting a Value on Beauty', in T.S Gendler & J. Hawthorne (eds.), *Oxford Studies in Epistemology, Volume 3*. Oxford: Oxford University Press.
- Deutsch, David [1999]. 'Quantum Theory of Probability and Decisions'. *Proc. R. Soc. Lond. A*, 455(1988): 3129-3137
- Dieks, Dennis [2007]. 'Reasoning about the future: Doom and Beauty'. *Synthese* 156: 427-39.
- Eagle, Antony [2010]. 'Deterministic Chance', *Noûs* 45(2): 269-299.
- Elga, Adam [2000]. 'Self-locating belief and the Sleeping Beauty problem'. *Analysis* 60: 143-47.
- Elga, Adam [2004]. 'Defeating Dr. Evil with Self-Locating Belief'. *Philosophy and Phenomenological Research* 69:383-396.
- Glynn, Luke [2010]. 'Deterministic Chance'. *British Journal for the Philosophy of Science* 62(2): 343-392.
- Greaves, Hilary [2007a]. 'On the Everettian Epistemic Problem'. *Studies in History and Philosophy of Modern Physics* 38 (1):120-152.
- Greaves, Hilary [2007b]. 'Probability in the Everett Interpretation'. *Philosophy Compass*, 2: 109-128.
- Greaves, Hilary and Myrvold, Wayne [2010]. 'Everett and Evidence', in eds. Saunders *et al.*, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.
- Handfield, Toby & Wilson, Alastair [2014]. 'Everett and Evidence', in A. Wilson (ed.), *Asymmetries of Chance and Time*. Oxford: Oxford University Press.

- Hofer, Carl [2007]. 'The Third Way on Objective Probability: A Skeptic's Guide to Objective Chance'. *Mind* 116:549-596.
- Lewis, David [1980]. 'A Subjectivist's Guide to Objective Chance', in ed. Jeffrey, *Studies in Inductive Logic and Probability*. Berkeley: University of California Press.
- Lewis, David [2001]. 'Sleeping Beauty: reply to Elga'. *Analysis* 61: 171-76.
- Lewis, Peter [2007]. 'Quantum Sleeping Beauty'. *Analysis* 67: 59-65.
- Lewis, Peter [2009]. 'Reply to Papineau and Durà-Vilà', *Analysis* 66: 114-119.
- Lewis, Peter [2010]. 'Credence and Self-Location', *Synthese* 175: 369-382.
- Meacham, Christopher [2008]. 'Sleeping Beauty and the Dynamics of de Se Beliefs', *Philosophical Studies* 138(2): 245-269.
- Meacham, Christopher [2010]. 'Unravelling the Tangled Web: Continuity, Internalism, Non-Uniqueness and Self-Locating Beliefs', in Tamar Szabo Gendler & John Hawthorne (eds.), *Oxford Studies in Epistemology, Volume 3*. Oxford: Oxford University Press.
- Olum, Ken [2002]. 'The doomsday argument and the number of possible observers', *Philosophical Quarterly* 52(207): 164-184.
- Papineau, David [1996]. 'Many Minds are No Worse Than One', *British Journal for the Philosophy of Science* 47(3): 233-241.
- Papineau, David [2010]. 'A Fair Deal for Everettians', in eds. Saunders *et al.*, *Many Worlds? Everett, Quantum Theory, and Reality*.
- Papineau, David and Durà-Vilà, Victor [2009a]. 'A Thirder and a Everettian: A Reply to Lewis' 'Quantum Sleeping Beauty'', *Analysis* 69(1):78-85.
- Papineau, David and Durà-Vilà, Victor [2009b]. 'Reply to Lewis: Metaphysics Versus Epistemology', *Analysis* 69(1): 89-90.
- Peterson, Daniel [2011]. 'Qeauty and the Books: A Response to Lewis' Quantum Sleeping Beauty Problem', *Synthese* 181(3): 367-374.
- Saunders, Simon [2004]. 'Derivation of the Born Rule from Operational Assumptions', *Proceedings of the Royal Society of London A*, 460, 1771-88.
- Saunders, Simon [2010]. 'Chance in the Everett Interpretation', in eds. Saunders *et al.*, *Many Worlds? Everett, Quantum Theory, and Reality*.
- Saunders, Simon; Barrett, Jeffrey; Kent, Adrian; Wallace, David, eds. [2010]. *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford University Press.
- Schulz, Moritz [2010]. 'The Dynamics of Indexical Belief', *Erkenntnis* 72(3):337-351.

- Schwarz, Wolfgang [2012]. 'Changing Minds in a Changing World', *Philosophical Studies* 159(2): 219-239.
- Schwarz, Wolfgang [MS]. 'Belief Dynamics across Fission'. Available online at <http://www.umsu.de/papers/>
- Titelbaum, Michael [2008]. 'The Relevance of Self-Locating Beliefs', *Philosophical Review* 117(4):555-606.
- Wallace, David [2002]. 'Everettian Rationality: defending Deutsch's approach to probability in the Everett interpretation'. *Studies in the History and Philosophy of Modern Physics* 34: 415-438.
- Wallace, David [2006]. 'Epistemology Quantised: circumstances in which we should come to believe in the Everett interpretation', *British Journal for the Philosophy of Science* 57: 655-689.
- Wallace, David [2010]. 'How to Prove the Born Rule', in eds. Saunders *et al.*, *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.
- Wallace, David [2012]. *The Emergent Multiverse*. Oxford: Oxford University Press.
- Wilson, Alastair [forthcoming a]. 'Everettian Quantum Mechanics Without Branching Time', forthcoming in *Synthese*. Available at <http://dx.doi.org/10.1007/s11229-011-0048-9>
- Wilson, Alastair [forthcoming b]. 'Objective Probability in Everettian Quantum Mechanics', to appear in *British Journal for the Philosophy of Science*.
- Zurek, Wojciech [2003]. 'Decoherence and the Transition from Quantum to Classical: REVISITED'. Available at <http://arxiv.org/abs/quant-ph/0306072>