

# When Beauties Disagree: Why Halfers Should Affirm Robust Perspectivalism

John Pittard  
October, 2013

In this paper I present a variant of the “Sleeping Beauty” case that shows that the “halfer” approach to the original Sleeping Beauty problem is incompatible with an extremely plausible principle pertaining to cases of disagreement. This principle says that, in contexts where rationality is not “permissive,” the weight I give to my view on  $p$  and to my disputant’s view on  $p$  ought to be proportional to my estimation of the strength of our epistemic positions with respect to  $p$ . In requiring such proportionality, the principle denies the possibility of what I will call “robustly perspectival” contexts, contexts where two maximally rational disputants who are in perfect communication are rationally required to disagree despite knowing that their epistemic positions are equally strong. Given the plausibility and widespread acceptance of the proportionality principle, its incompatibility with the halfer approach to Sleeping Beauty gives us an apparently powerful new argument against the halfer position and for the alternative “thirder” view.

But I am a halfer, not a thirder. So I go on to argue that despite the principle’s intuitive plausibility, there are good reasons for thinking that the case I present here does involve a robustly perspectival context and that the principle should therefore be rejected. I suggest that the lesson that we should draw from this case is *not* that we should accept the thirder view, but rather that rationality can be perspectival in a robust way that many may find quite surprising.

The paper will proceed as follows. In section 1, I present the key example, a variant of the Sleeping Beauty case that involves multiple subjects. I then show that an intuitive way of reasoning about the case leads to the conclusion that it involves a robustly perspectival context. In section 2, I show that this conclusion is in tension with the proportionality principle mentioned above, a principle that is often taken for granted in discussions of the epistemic significance of disagreement. In section 3, I argue that halfers and thirders are committed to different views on the key case: halfers should endorse the intuitive line of reasoning outlined in section 1 that leads to the acknowledgment of robustly perspectival contexts and the rejection of the proportionality principle; thirders, on the other hand, should deny that the case involves a robustly perspectival context and can therefore maintain a commitment to the proportionality principle. The counterintuitive perspectivalism implied by the halfer position could be seen as a significant and as-yet unappreciated cost to the halfer view. But in section 4, I appeal to the rational significance of evidential selection procedures in order to defend the reasonability of the striking perspectivalism to which the halfer is committed. Finally, in section 5, I turn to another case that is arguably more difficult for the halfer to handle, a case where the halfer seems to be insufficiently deferential to an acknowledged expert. I point to two plausible lines of response available to the halfer.

It should be noted at the outset that I do not in this paper offer a positive argument for the halfer view. Rather, I show that the halfer view is committed to a counterintuitive form of perspectivalism and then attempt to defend the halfer view against the objection that its perspectivalist commitments are unreasonable.

## 1. DISAGREEING BEAUTIES

The original “Sleeping Beauty” case features Sleeping Beauty and some experimental philosophers who put her to sleep and then subject her to a certain number of awakenings and to memory tampering that prevents her from knowing which awakening she is presently experiencing (Elga 2000; more on this case later). The case that will be the primary focus of this chapter involves not one “Beauty” but four: Alvin, Brenda, Claire, and Dillon, who have each agreed to spend the night in the infamous Experimental Philosophy Laboratory.<sup>1</sup> Before they are put to sleep, the four subjects are brought together and informed of what will happen over the course of the evening. By means of random selection, one of the four subjects will be selected as the “victim.” The experimenters will know the identity of the victim, but will not disclose this information to any of the subjects. At three different times during the night, the experimenters will simultaneously awaken the victim and exactly one of the other subjects. Each non-victim subject will be awakened only once during the night. The victim and the other awakened subject will converse for a time and will be asked to discuss their views about the probable identity of the victim. After each of these conversations, the victim and the other awakened subject will be given a drug that will put them back to sleep and that will also erase their memory of the awakening. So while the victim will be awakened three times, there will be no way for the victim to determine whether a given awakening is the first, second, or third awakening. And while the non-victim subjects will be awakened only once, during their awakenings they will have no way of knowing whether they are in their first and only awakening or whether they are the victim and possibly in a second or third awakening. Steps are taken to make sure that there are no differences between the waking experiences of the victim and those of the non-victims that might serve as clues as to the identity of the victim. (So it is not the case, for example, that one of the awakened subjects will feel especially tired or “drugged up” in comparison to the other awakened subject.) When the morning arrives, all four of the subjects will be woken at the same time and sent on their way, none of them possessing any memories of their awakening (or awakenings) during the preceding night. After being made certain of all of these features of the experiment, the subjects are put to sleep in the lab and the night begins.

Let us suppose that Alvin, Brenda, Claire, and Dillon are each ideally rational agents, and that they know this about one another. It is clear that at the moment after Claire has learned the setup of the experiment but before the subjects have been put to sleep for the night, her credence that she will be the victim will be 0.25. And when Claire and the other subjects are leaving the lab the next morning, her credence that she was the victim will again be 0.25. But suppose that sometime during the night, the experimenters wake Claire, and when Claire looks across the room, she sees that Dillon has also just been awakened. At this time, what should Claire’s credence be for the proposition that she is the victim?

---

<sup>1</sup> So named by David Lewis (2001, 171).

A very intuitive line of reasoning suggests that Claire's credence that she is the victim should not change during this awakening, but should remain 0.25. Consider first the moment after Claire has been awakened but before she has looked across the room to see who else is awake. At this moment, it seems that Claire has not gained any new evidence that is relevant to the question of whether she is the victim. For before she went to sleep, Claire knew that she would be awakened sometime during the night, whether or not she turned out to be the victim, and that the features of a victim's awakening would not be discernibly different from the features of a non-victim's awakening. Moreover, it does not seem that Claire has *lost* any evidence that she possessed just before the experiment began. And the following **relevance principle** seems to be a very basic principle of rationality: if I start off with a credence for  $p$  that is maximally rational given my evidence, and if I neither gain nor lose any evidence that is relevant to the likelihood of  $p$ , then I should not change my credence for  $p$ .<sup>2</sup> This relevance principle, together with the claim that Claire has neither gained nor lost relevant evidence, implies that after waking (but before identifying her "co-waker") Claire's credence that she is the victim should still be 0.25.

Next, it seems that Claire's credence that she is the victim should not change upon scanning the room and seeing that Dillon is also awake.<sup>3</sup> For suppose that after seeing that Dillon is awake, the value of Claire's credence that she is the victim should be  $c$ . Since Claire's sharing an awakening with Dillon gives her no more or less reason to suspect that she is the victim than her sharing an awakening with Alvin or Brenda, Claire's credence should also be  $c$  if she were to look across the room and see Alvin or Brenda awake. And since before looking across the room, Claire knows that she will see one of these three subjects awake, she already knows that after seeing her co-waker her credence will be  $c$ . But if Claire knows that her credence will be  $c$  once she looks across the room, then she should already have a credence of  $c$ . Thus, after seeing that Dillon is awake, there should be no change in the value of Claire's credence. In looking across the room and identifying her co-waker, Claire does not learn anything that ought to raise or lower her credence for the proposition that she is the victim.

Taken together, the previous two paragraphs give us an argument for the conclusion that an ideally rational agent like Claire will, during her awakening or awakenings, continue to have a credence of 0.25 for the proposition that she is the victim. I will explore an objection to the first step of this argument later in the paper. But for the moment, let us assume that the reasoning just presented is correct and explore its implications. If Claire's credence that she is the victim ought to remain 0.25, then upon seeing that Dillon is awake, Claire's credence for the proposition that Dillon is the victim ought to go up from 0.25 to 0.75. For since Claire knows that in any given awakening, one of the two awakened subjects must be the victim, she now knows that the victim is either Dillon or herself. Given that she assigns a 0.25 credence to the proposition that she is the victim, probabilistic coherence requires that her credence that Dillon is the victim be 0.75. This increase in her credence for Dillon's being the victim does *not* seem to be in tension with

---

<sup>2</sup> In his defense of the halfer position on Sleeping Beauty, Hawley (2013) endorses this sort of relevance principle under the label INERTIA. Note that this relevance principle is very different from the "Relevance-Limiting Thesis" (so named in Titelbaum 2008, 556) that is endorsed by some halfers, a thesis that holds that learning "self-locating" information should never lead to changes in credences for non-self-locating propositions.

<sup>3</sup> My statement of the argument in this paragraph benefitted from the comments of an anonymous referee.

the relevance principle. For while it seems that Claire has not learned anything that ought to increase her suspicion of herself, in seeing that Dillon is her co-waker it is clear that she *has* learned something that ought to increase her suspicion of Dillon.

Since Dillon is, like Claire, a perfectly rational agent, he will reason in an analogous way. This means that, upon being awakened and seeing that Claire is being awakened at the same time, Dillon will maintain his 0.25 credence for the proposition that he is the victim and raise his credence for the proposition that Claire is the victim to 0.75.

We have, then, a case of disagreement. And it is a most perplexing one. For starters, the disagreement is not the result of one party possessing better evidence than the other, or of one party processing their evidence more rationally. For the case is perfectly symmetrical: Claire and Dillon know that the other possesses equally strong evidence and that they are both perfectly rational in how they respond to that evidence. Second, it seems to make no difference if Dillon and Claire go to the effort of fully and perfectly disclosing all of their reasons for their view. For since Dillon and Claire share in common all of the knowledge and information that led them to their views, and since they both know that they are equally informed and equally rational, discussing their reasons with one another will not reveal anything about the other's reasoning and views beyond what could already be anticipated. Claire will already know that Dillon will employ reasoning that is exactly analogous to hers, and Dillon will likewise know that Claire will engage in reasoning just like his. So discussion and debate will not bring about any changes in their respective levels of confidence.

It would be natural to describe the disagreement as follows: Claire and Dillon are mutually acknowledged "epistemic peers" who have the exact same evidence and who are perfectly confident that the other person is reasoning impeccably, and yet they are nonetheless rational in remaining steadfast and dismissing the view of their disputant.<sup>4</sup> But one might reasonably question whether we should describe Claire and Dillon as having the same evidence. For while Claire and Dillon can affirm the same "third person" facts about the way the world is, they obviously cannot make the same "first-person" affirmations about their own "location" within this world. Claire, for example, can assert that she opened her eyes and saw Dillon, while Dillon cannot assert this (though he can of course agree that *Claire* opened her eyes and saw Dillon). And if we suppose that Claire and Dillon will in fact disagree about the likelihood that Claire is the victim, then this seems to be a case where the credence that Claire should have for a hypothesis expressible in *third-person* terms is in part determined by Claire's *first-person* information about who she happens to be and thus what her vantage point on the world is. Since first-person information might possess such rational significance in this case, I do not want to simply assume that Claire's and Dillon's total relevant evidence is exhausted by their third-person information. Rather, I want to allow that the relevant evidence in this case could include information that is irreducibly "first-person." Following others, I will use '**uncentered information**' to refer to purely "third-person" information about the world that in no way implies anything about my "location" in that world (e.g., who I am and where I am in space and time). **Centered information**, on the other hand, has at least some bearing on my location in the world. The extent to which two subjects can share their centered information is limited. For example, though Claire and Dillon can both affirm the centered proposition, "I am in the lab," Claire can truly affirm, "I am Claire," while Dillon cannot.

---

<sup>4</sup> The term 'epistemic peer,' popular in the disagreement literature, was first introduced by Gutting (1982).

While Claire and Dillon cannot make all of the same first-person affirmations and thus do not share all of their centered information, this in no way should prevent them from acknowledging one another as epistemic peers, where (for present purposes) two people are *epistemic peers* with respect to  $p$  just in case they both possess the same quality and quantity of evidence bearing on  $p$  and are equally rational and reasonable in their assessment of that evidence.<sup>5</sup> For clearly neither Claire nor Dillon have any reason to think that they possess some sort of epistemic advantage merely in virtue of being who they are. It would be ludicrous for Claire to think that, simply due to the fact she happens to be Claire and not Dillon, she is more likely than Dillon to have accurate views on the identity of the victim. Surely *being Claire* can be considered epistemically advantageous only if we think there is some *other* epistemically relevant factor (whether access to more evidence, greater intelligence, possession of a particular insight, or whatever) that distinguishes her and Dillon. But we are supposing that they know there to be perfect parity with respect to such matters. So neither Claire nor Dillon can take their centered information to confer some sort of epistemic advantage. We can say, then, that Claire and Dillon are mutually acknowledged epistemic peers who perfectly share their *uncentered* evidence and yet rationally disagree.

The mere fact of rational disagreement between mutually acknowledged epistemic peers would not be surprising if we suppose that there are “permissive contexts,” contexts where, for one or more disagreeing subjects, the evidence and other rational factors bearing on  $p$  do not determine one credal attitude that is maximally rational for that subject to have towards  $p$ .<sup>6</sup> But it seems fairly clear that this is a *non*-permissive context where there is one maximally rational credence for Claire and one maximally rational credence for Dillon (even if there is, as we shall see, some debate as to what their credence assignments should be and whether or not they will differ). Even given that some context is non-permissive, it is clear that I can rationally affirm  $p$  while acknowledging that some epistemic peer denies  $p$ , as long as I am not in communication with this epistemic peer. For example, suppose I am one of ten subjects who each know that there is a 0.5 probability that an urn contains nine red balls and one black ball, and a 0.5 probability that it contains nine black balls and one red ball.<sup>7</sup> All of us are each allowed to draw one ball from the urn (without replacement), and initially we are not allowed to communicate with one another about the color of the ball we drew. If I draw a red ball, I am justified in being fairly confident that the urn originally contained nine red balls even though I know that there is at least one person who drew a black ball and who, on equally strong evidence, is justified in concluding that the urn probably contained only one red ball. The mere fact of peer

---

<sup>5</sup> I am using ‘peer’ in a more fine-grained and demanding sense than is sometimes customary. It is important to keep this demanding sense of ‘peer’ in mind in order to appreciate the apparent difficulty of maintaining confident and reasonable belief in a disagreement with a disputant who is an acknowledged peer.

<sup>6</sup> Note that given the way I am using ‘permissive,’ denying that there are permissive contexts does not amount to accepting the “Uniqueness Thesis” endorsed by Feldman (2007) and others. To deny that there are permissive contexts is equivalent to affirming that, for any given person and proposition, there is exactly one maximally rational credal attitude for that person to have towards the proposition. But the Uniqueness Thesis goes further than this, holding that for any two subjects with the same evidence, the *very same* credal attitudes are maximally rational. So using the terminology I will employ here, we can say that the Uniqueness Thesis denies both the possibility of “permissive” contexts *and* the possibility of “perspectival” contexts.

<sup>7</sup> This example is adapted from Michael Titelbaum’s “Mystery Bag” example (2012, 235–6), though Titelbaum’s example involves additional components and is used to make a different point.

disagreement need not pose an epistemic threat. But matters change once I am in *communication* with a disagreeing peer, so that that peer becomes my *disputant*, where by “disputant” I mean someone who not only disagrees with me, but who is also in communication with me about this disagreement. Suppose that on the next day I happen to cross paths with another one of the subjects who tells me that he drew a black ball. Earlier, my drawing a red ball gave me a reason for thinking that a subject who drew a black ball was not representative of the subjects as a whole. But now, my reason for thinking such a subject unrepresentative has been defeated by the fact that the one subject I happened to cross paths with drew a black ball. At this point, it seems as though I cannot rationally disagree with this subject on the likely number of red balls while acknowledging that he is my epistemic peer (in the very demanding sense of ‘peer’ stipulated above). Maintaining a disagreement requires thinking that I have some evidential advantage or that I am processing my evidence more rationally. So what *would* be surprising is if in a non-permissive context there could be a rational disagreement between two disputants who are mutually acknowledged epistemic peers. But if I am right that Claire and Dillon will rationally disagree about the likelihood that Claire is the victim, then the “multiple Beauties” example is precisely such a case.

Let us say that a context involving a dispute over some uncentered proposition  $p$  is **robustly perspectival** if and only if it is a non-permissive context where (i) there is full communication between two disputants and perfect sharing of their uncentered evidence, (ii) each disputant knows with full confidence that the other disputant possesses the same uncentered evidence and processes that evidence perfectly rationally, and (iii) the final credence for  $p$  that is rationally required of one disputant differs from the final credence for  $p$  that is rationally required of the other disputant. And let us use the label **robust perspectivalism** to refer to the view that there are robustly perspectival contexts. If it is true that Claire and Dillon will disagree upon awakening, then there are robustly perspectival contexts and robust perspectivalism is correct. As I will argue in the next section, robust perspectivalism is incompatible with an intuitively plausible principle that is often taken for granted in discussions of the epistemic significance of disagreement.

## 2. THE EPISTEMOLOGY OF DISAGREEMENT AND THE RULE OF PROPORTIONALITY

In this section, I characterize a highly intuitive epistemic principle that bears on disagreement and show that this principle is incompatible with the view that Claire ought to maintain her credence of 0.25 for the proposition that she is victim.

If I am highly confident that  $p$  and my disputant is highly confident that  $\neg p$ , consistency requires that I also believe that I am right about  $p$  and that my disputant is wrong about  $p$ . This is trivial and obvious. But can I rationally affirm that I am right and my disputant wrong while also affirming that my disputant is in at least as strong an epistemic position with respect to  $p$  as I am? Arguably not, at least not if we understand the strength of one’s epistemic position as taking into account *all* the dimensions of epistemic evaluation that bear on the likelihood of one’s arriving at a reasonable position on  $p$  in the present circumstances, so that the strength of one’s epistemic position takes into account such factors as the general reliability of one’s cognitive faculties, the current level of functioning of those faculties, the presence or absence of any errors in reasoning, the

presence or absence of bias, the quality and quantity of one's evidence, and the adequacy of one's overarching "epistemic framework." To be sure, there is nothing contradictory in affirming both that my disputant is in at least as strong an epistemic position with respect to  $p$  and that I am right and my disputant is wrong. For sometimes the epistemically disadvantaged can get lucky. An expert in probability theory may think it highly likely that in a particular series of coin tosses at least one coin has landed heads, and a toddler may firmly believe, without any evidence, that all of them have landed tails. It may turn out that, as luck would have it, the toddler is right. Nonetheless, the probability expert was clearly in a stronger epistemic position with respect to the question. But even if it is not contradictory to affirm both that I am right and that my disputant is in at least as strong an epistemic position, such an affirmation nonetheless seems epistemically problematic. For if I think that the superiority of my view on  $p$  is not due to any epistemic advantage I have over my disputant (such as greater insight or superior evidence), then I must think that my having the superior view is a matter of epistemic luck. And arguably, I cannot justifiably believe that I am lucky in this way without good evidence that I am lucky. Suppose I have such evidence. If my disputant does *not* have this evidence, then it would seem that this gives me reason for thinking that, in the present circumstances, my epistemic position with respect to  $p$  is superior. And if my disputant *does* have the evidence that I am the lucky one but does not give it proper weight, then this *also* gives me reason for thinking that my epistemic position is, on balance, superior (since I take my disputant to have improperly assessed a piece of evidence that ought to significantly shape her credence for  $p$ ). So it seems that I cannot reasonably be confident that  $p$  while also thinking that my disputant's epistemic position with respect to  $p$  is at least as strong as my own.

The above line of reasoning suggests that there is some sort of **rule of proportionality** governing my confidence in  $p$  in the face of disagreement over the plausibility of  $p$ : the weight I give to my initial opinion regarding  $p$  and to my disputant's initial opinion regarding  $p$  ought to be proportional (in some sense) to my assessment of the strength of our epistemic positions with respect to  $p$ . To the extent that I think my epistemic position is stronger, I should weight my opinion regarding  $p$  more heavily, and to the extent that I think my disputant's epistemic position is stronger, I ought to weight her opinion regarding  $p$  more heavily.

It is important to see that the rule of proportionality does *not* require that I assess my and my disputant's epistemic credentials in a way that is independent of my views on the disputed matter. *That* requirement, often called the "independence" requirement, is quite controversial and leads fairly directly to a "conciliatory" position according to which disagreement has significant skeptical force (Christensen 2009, 758). The rule of proportionality is completely open to the possibility, defended by opponents of conciliatory views, that the reasoning behind my belief that  $p$  may serve to ground my judgment that those who dispute  $p$  suffer from some sort of epistemic disadvantage (e.g., that they are prone to a certain error in reasoning, have some sort of cognitive defect, lack some piece of evidence, etc.) and are for that reason in a weak epistemic position with respect to  $p$ . Advocates of conciliatory views as well as their opponents can both affirm that my being steadfast in the face of disagreement is reasonable only if I think that those on my side of the dispute are, on the whole, in a stronger epistemic position.

Most parties to the current debate over the epistemic significance of disagreement seem to take it for granted that the rule of proportionality is correct, or at least that it is correct in

non-permissive contexts. If I am in a *permissive* context, and if I think that both my and my disputant's levels of confidence are in the range that is rationally permissible for me, then arguably I can remain steadfast in my level of confidence without thinking that my epistemic position is superior in some way.<sup>8</sup> But even if there are permissive contexts (which is controversial<sup>9</sup>), it seems that opposing sides in the disagreement debate are for the most part united in their acceptance of what we can call **MODEST PROPORTIONALITY**, which is the view that in non-permissive contexts, the rule of proportionality applies.<sup>10</sup>

There are, to be sure, epistemological positions that have adherents and that arguably conflict with **MODEST PROPORTIONALITY**. For example, **MODEST PROPORTIONALITY** might be incompatible with certain strong forms of "epistemic conservatism" that maintain that the fact that I believe *p* is a very resilient reason in favor of my continuing to believe *p*.<sup>11</sup> And certain forms of epistemic relativism that hold that there are no objective epistemic standards according to which the correctness of different "epistemic systems" can be measured might be in tension with the underlying motivations for **MODEST PROPORTIONALITY**.<sup>12</sup> But most philosophers currently engaged in the disagreement debate, even those who are not friendly to conciliatory views, have not taken such approaches. They have typically eschewed views on rationality that are explicitly perspectivalist, preferring instead views that are least moderately "objectivist." And those coming from such a perspective will be inclined to accept **MODEST PROPORTIONALITY**. What is interesting about the multiple Beauties case is that it challenges **MODEST PROPORTIONALITY** without relying on any premise (such as epistemic conservatism or relativism) that explicitly conflicts with a more objectivist understanding of rational norms. I will now make that challenge explicit.

**MODEST PROPORTIONALITY** implies that there cannot be a non-permissive context where two disputants rationally maintain different credences for *p* despite knowing with certainty that their epistemic positions with respect to *p* are equally strong. While **MODEST PROPORTIONALITY** is perfectly compatible with a subject's responding to a disagreement over *p* by concluding that her disputant is not in fact equally qualified to assess *p*, on the assumptions that each subject continues to know that the other subject is equally qualified and that the context is non-permissive, **MODEST PROPORTIONALITY** requires that both of the subjects give equal weight to the other's perspective. And if they both give equal weight to the other's perspective, then they will have the same final credence for *p*.<sup>13</sup> So any initial

---

<sup>8</sup> Goldman (2010, 195–6) briefly argues for this possibility, as does Kelly (2010, 118–119).

<sup>9</sup> For arguments against permissive contexts, see White (2005).

<sup>10</sup> That proponents of conciliatory views affirm **MODEST PROPORTIONALITY** is rather obvious. See, e.g., Feldman (2006), Elga (2007), and Christensen (2011). Several opponents of conciliatory views also affirm something along the lines of **MODEST PROPORTIONALITY**. Lackey (2010, 277, 281–2), for instance, suggests that confidence in the face of disagreement requires that there be some "symmetry breaker," a reason for thinking that one's own epistemic situation is favorable to that of one's disputant. Along similar lines, Bergmann (2009, 342) thinks that reasonable steadfastness requires thinking that one is *probably* "internally" or "externally" more rational than one's disputant. Van Inwagen (2010) argues in favor of several principles that are consonant with **MODEST PROPORTIONALITY**, though he stops short of firmly endorsing an anti-conciliatory view. Also see Enoch (2010, 975) and Fumerton (2010, 99).

<sup>11</sup> For a critical discussion of different versions of epistemic conservatism, see Vahid (2004).

<sup>12</sup> This characterization of epistemic relativism is based on Boghossian's discussion in (2006, 73).

<sup>13</sup> This does not mean that their new credence for *p* will be halfway between their original credences. In some cases, both disputants might agree that when both perspectives are given equal weight, the correct compromise credence is closer to one of the original views than the other. For example, if I think Horse A



difference in confidence levels will disappear after full communication. MODEST PROPORTIONALITY thus implies that there are no robustly perspectival contexts. But, as already discussed, the view that Claire and Dillon will disagree about the likelihood that Claire is the victim implies that there are robustly perspectival contexts. So if we are to affirm that Claire and Dillon will disagree, then we must reject MODEST PROPORTIONALITY.

If I am correct in maintaining that MODEST PROPORTIONALITY is incompatible with the view that Claire and Dillon will rationally disagree, we have two options: deny that Claire and Dillon will disagree, or deny MODEST PROPORTIONALITY. I will now show that the first of these two options is available to those who advocate the “thirder” position on the original Sleeping Beauty case. Arguments for the thirder position, when adapted to the multiple Beauties case, support the conclusion that Claire and Dillon will agree on the likelihood that Claire is the victim. But those who advocate the alternative “halfer” position on Sleeping Beauty will be under rational pressure to conclude that Claire and Dillon will disagree, and thus that MODEST PROPORTIONALITY is false.

### 3. AGREEING BEAUTIES? WHY THIRDERS AND HALFERS SHOULD DISAGREE ON WHETHER CLAIRE AND DILLON WILL DISAGREE

There are two basic positions on the original Sleeping Beauty problem: the “thirder” position and the “halfer” position. I will now argue that while “halfers” ought to affirm that Claire and Dillon will disagree, “thirders” ought to say that Claire and Dillon will not disagree but will both assign a credence of 0.5 to the proposition that Claire is the victim. Thus, the multiple Beauties case presents no problem for thirders, who can in fact use the case as a new argument for the thirder position, an argument that gains its strength from the apparent plausibility of MODEST PROPORTIONALITY. To show why halfers should maintain that Claire and Dillon will disagree and thirders should maintain that they will agree, I will describe the original Sleeping Beauty case, present the principal argument for the halfer position and one of the principal arguments for the thirder position, and show that analogous arguments in the multiple Beauties case lead halfers and thirders to disagree about whether Claire and Dillon will agree or disagree about the likelihood that Claire is the victim.

The original Sleeping Beauty case goes as follows. Sleeping Beauty is a perfectly rational agent who is about to be put to sleep on Sunday night by some experimental philosophers. The experimenters inform her that during the experiment she will either be awakened once (on Monday morning) or twice (on Monday and Tuesday morning), depending on the outcome of a coin toss. If the coin lands heads, Beauty will be awakened only on Monday. If the coin lands tails, Beauty will be awakened on Monday and Tuesday. Sleeping Beauty does not know the result of the coin toss, and thus does not know the number of awakenings she will experience. Moreover, the experimenters inform Beauty

---

won the race by a nose, and my friend thinks Horse B edged out Horse A at the finish line, we may agree that our final credence for a Horse B victory should be closer to my friend’s pre-disagreement credence if we know that in close races between the two horses, Horse B wins the large majority of the time. In this case, the evidential significance of each judgment is canceled out by the other, leaving us back at our prior probability (which is closer to my friend’s pre-disagreement credence than to mine). For another sort of example where equally weighting one another’s views does not lead to middle ground credences across all of the disputed propositions, see (Thurrow 2012).

that she will have no way of distinguishing a Monday awakening from a Tuesday awakening. In order to make these waking experiences indiscernible, the experimenters will erase Beauty's memory of her Monday awakening before putting her back to sleep, so that if she awakens on Tuesday, she will not know whether this is her first or second awakening. When the experiment is over on Wednesday, the experimenters will awaken Beauty in a manner that ensures that Beauty will know it is Wednesday and not an awakening during the course of the experiment. (We can imagine that Beauty knows that on Wednesday she will be awakened by a song whose lyrics are, "It's Wednesday, the experiment is over!") The controversial question, which has given rise to a small cottage industry producing papers on the subject, is this: when Beauty is woken on Monday and experiences an awakening that is indistinguishable from a Tuesday awakening, what should her credence be for the proposition that the coin landed heads? "Thirders" argue that the rational credence is one third; "halfers" contend that it is one half.

Before presenting arguments for the thirder and halfer positions and showing how those arguments translate to the multiple Beauties case, an important clarifying note is in order. It is an unfortunate feature of the multiple Beauties case that arguments for the *thirder* position in the original Sleeping Beauty case support the conclusion that Claire's credence that she is the victim should, upon waking, be *one half*; and the principal argument for the *halfer* conclusion in the original Sleeping Beauty case supports the conclusion that Claire's credence in the relevant proposition should remain *one fourth*. So "halfers" support 1/2 in the original case and should (I will argue) support 1/4 in the multiple Beauties case, and "thirders" support 1/3 in the original case and should support 1/2 in the multiple Beauties case. Having warned of this potential source of confusion, I will now characterize the principal halfer and thirder arguments and show how they bear on the multiple Beauties case.

The principal consideration in favor of the halfer solution to the Sleeping Beauty Problem is that it seems that Sleeping Beauty has not learned anything of evidential significance when she is awakened on Monday. On Sunday, when Beauty's credence for the coin's landing heads was one half, Beauty *knew* that she would soon have at least one waking experience and that no qualitative features of a given awakening would provide clues as to whether it is a Monday awakening or a Tuesday awakening. So it seems that she does not acquire any new evidence relevant to the outcome of the coin toss when she learns that she is presently in the midst of a waking experience with such and such qualitative features. And if this is right, then application of the relevance principle leads us to conclude that her credence for heads ought to remain one half. Halfers endorse this "no new relevant evidence" reasoning and the application of the relevance principle, contending that Beauty's credence for the coin's landing heads should not change upon awakening.

The same "no new relevant evidence" reasoning that motivates the halfer view also supports the view that, *after* waking up but *before* identifying her co-waker, Claire's credence that she is the victim ought to still be 0.25. Once this has been granted, arriving at the conclusion that Claire and Dillon will disagree merely requires showing that Claire's credence that she is the victim should not change upon looking across the room and identifying her co-waker. The argument for this claim, rehearsed in section 1, is extremely strong. Since Claire knows that she will see a co-waker when she looks across the room, and since the particular identity of the co-waker doesn't make a difference to how much

Claire suspects herself, in looking across the room and seeing that Dillon is awake, Claire does not learn anything that should lead her to change her credence that she is the victim. So it seems that halfers are committed to holding that Claire and Dillon will disagree, and thus that robust perspectivalism is true and that MODEST PROPORTIONALITY is false.

In saying that halfers are “committed” to the view that Claire and Dillon will disagree, I mean that any viable position that supports a halfer position in the Sleeping Beauty case will also imply that Claire and Dillon will disagree. There are, however, two reasons one might want to resist this claim. First, someone might argue that while halfers are rightly moved by the relevance principle in the Sleeping Beauty case, the relevance principle is not absolute and is trumped in circumstances where it threatens to come into conflict with MODEST PROPORTIONALITY. I concede that such a position is at least superficially consistent. But it also strikes me as quite ad hoc. The relevance principle seems to be a more basic rational requirement than MODEST PROPORTIONALITY, so it is hard to see why the latter should override the former. And if the relevance principle does admit of exceptions so readily, then it is also questionable whether halfers are reasonable in thinking that it holds in the original Sleeping Beauty case, given the force of the thirder argument to be considered below.

But there is a more legitimate reason why one might want to resist the claim that halfers are committed to the view that Claire and Dillon will disagree. Some attempts to give a formalized generalization of the halfer approach (Halpern 2005; Meacham 2008; Briggs 2010) have associated the halfer approach with a policy that, when applied to the multiple Beauties case, does *not* support the conclusion that Claire and Dillon will disagree, but rather supports the conclusion that they will agree (a conclusion that I will associate with the thirder position). This policy, which Rachel Briggs calls the “Halfer Rule,” requires “conditionalizing the *uncentered portion* of one’s credence function on the *uncentered portion* of one’s total evidence, and then within each [uncentered] world, dividing one’s credence [for that world] among the doxastically possible centers [within that world]” (Briggs 2010, 9–10). Briggs thinks that halfers will be drawn to the view that only the *uncentered* component of my total evidence is of any relevance in determining the probability of some uncentered proposition, and the Halfer Rule requires that one form credences for uncentered propositions in a way that sets aside centered information as irrelevant. Since by hypothesis Claire and Dillon have the same uncentered evidence, clearly they will be in full agreement if they follow the Halfer Rule.

But even though the Halfer Rule does prescribe halfer credences in the original version of the Sleeping Beauty case, halfers have good reasons for rejecting the Halfer Rule. For in many cases, the Halfer Rule delivers prescriptions that are fundamentally at odds with the highly intuitive “no new relevant evidence” reasoning that drives people towards halfer conclusions in the first place (Titelbaum 2008, 591–7; Briggs 2010, 29). Consider, for instance, Titelbaum’s “Technicolor Beauty” case (2008, 591–7).<sup>14</sup> In the original Sleeping Beauty case, it was stipulated that Beauty’s Monday and Tuesday awakenings are qualitatively identical. The Technicolor Beauty example introduces qualitative differences that are evidentially irrelevant. In Technicolor Beauty, everything is just as it was in the original case except for the following addition. Beauty knows that at the same time the coin is to be tossed on Sunday night, a dice will also be rolled. If the dice roll comes out odd, a red piece of paper will be put in her room before she wakes up on Monday and then,

---

<sup>14</sup> A similar case can also be found in Meacham (2008, 263).

after she is put back to sleep, it will be replaced by a blue piece of paper that will be in the room throughout Tuesday. If the dice roll come out even, the colors will be reversed: a blue piece of paper will be put in her room on Monday and a red piece of paper will be put in her room on Tuesday. Beauty knows all of this. Quite clearly, the fact that Beauty sees a red piece of paper upon awakening should make no difference to her credence for the proposition that the coin landed heads. For Beauty knew she would see either a red or blue piece of paper, and a piece of red paper is no more or less likely to be present on a Monday (or on a Tuesday) than a piece of blue paper. So if Beauty's credence for heads in the original case should be 0.5, then it should also be 0.5 in Technicolor Beauty: the "no new relevant evidence" reasoning is just as strong in both cases. But while the Halfer Rule delivers halfer prescriptions in the original case, it delivers *thirder* prescriptions in Technicolor Beauty. This "instability" (Briggs 2010, 27) of the Halfer Rule's prescriptions surely shows it to be mistaken. Those who wish to stick by the halfer view on the original case therefore ought to reject the Halfer Rule and seek a way of generalizing the halfer approach that is more closely aligned with the kind of "no new relevant evidence" reasoning that makes the halfer view so intuitive in the first place.<sup>15</sup> So the Halfer Rule does not, in my judgment, supply a plausible way of being a halfer while affirming that Claire and Dillon will agree.

I will now show that thirderers ought to maintain that Claire and Dillon will find themselves in agreement during their awakening. While all of the arguments that I am aware of for the thirder position can be used to support the conclusion that Claire and Dillon will agree, I will focus here on an argument articulated in the piece by Adam Elga that introduced the Sleeping Beauty problem to the philosophical community. For purposes of presenting Elga's argument, I'll follow Elga in labeling the three possibilities open to Beauty upon awakening on Monday 'H1,' 'T1,' and 'T2,' with the following table showing which possibilities are designated by these labels:

	HEADS	TAILS
It's Monday	H1	T1
It's Tuesday	(Not possible)	T2

Let  $C$  be Beauty's (perfectly rational) credence function on Monday when she awakens. The thirder position follows from the following two premises:

- (1)  $C(T1) = C(T2)$
- (2)  $C(H1) = C(T1)$

Since H1, T1, and T2 are mutually exclusive and exhaust all possibilities, we know that  $C(H1) + C(T1) + C(T2) = 1$ . From this and from (1) and (2), it follows that  $C(H1) = C(T1) = C(T2) = 1/3$ . And since Beauty's credence for HEADS must be equal to her credence for H1, we can further conclude that her credence in HEADS upon awakening on Monday will be  $1/3$ .

Given the absence of any reason for thinking T1 more or less likely than T2, we have good reason for thinking that Beauty will assign the same credence to each of these two possibilities, and thus that (1) is correct. To be sure, some have contested the general

---

<sup>15</sup> For a defense of the halfer view that rejects the Halfer Rule, see the recent work of Darren Bradley (2011a; 2012; and especially 2011b).

“indifference principles” that have been offered as motivation for (1) (see, e.g., Weatherson 2005).<sup>16</sup> But even if such indifference principles are controversial, most everyone, halfers included, are inclined to accept that *in this particular case*, Beauty ought to assign equal credence to T1 and T2. The primary source of disagreement between halfers and thirders is (2).<sup>17</sup>

Elga’s argument for (2) may be summarized as follows. First, since Beauty will be awakened on Monday irrespective of whether the coin comes up heads or tails, we can imagine that the experimenters do not toss the coin until *after* Beauty is awakened on Monday. And whether they do in fact toss the coin on Sunday night or after waking up Beauty on Monday should not, it seems, make any difference to the credences Beauty assigns to H1, T1, and T2. So let’s imagine that Beauty knows that the coin is tossed on Monday, just after she is put back to sleep. Now, suppose that sometime after being woken up on Monday, Beauty is told that it is Monday. Upon learning this, Beauty learns that T2 is false, so that her credences for TAILS and HEADS will be identical to her updated credences for T1 and H1 (respectively). And what should her updated credences for H1 and T1 be? Well, her credence for H1 must be identical to her credence that a fair coin, *yet to be tossed*, will land heads, and her credence for T1 must be identical to her credence that a fair coin, *yet to be tossed*, will land tails. And surely, Elga contends, one’s credence that a *future* toss of a fair coin will land heads should be 0.5. So upon learning that it is Monday, Beauty’s credence for H1 ought to be 0.5.

From this, one more step is needed to arrive at the conclusion that (2) is correct. And it is a step that is frequently contested. Elga asserts the following:

- (3) The credence that Beauty has for H1 *after* learning that it is Monday (and thus that H1 or T1) should be equal to the conditional credence  $C(H1 \mid H1 \text{ or } T1)$  that she had *before* learning that it is Monday.<sup>18</sup>

Given (3) and the fact that Beauty’s credence for H1 after learning that it is Monday should be 0.5, it follows that *before* learning that it is Monday, it ought to be the case that  $C(H1 \mid H1 \text{ or } T1) = 0.5$ . And from this, it follows that, before learning that it is Monday, it

---

<sup>16</sup> Briggs (2010, 12) offers a weaker indifference principle that avoids Weatherson’s objections and motivates (1) on the condition that Beauty assigns both T1 and T2 some precise non-zero credence.

<sup>17</sup> One exception to this is (Hawley 2013). Hawley is a halfer who accepts (2) but denies (1), arguing that it ought to be the case that  $C(T1) = 0.5$  and  $C(T2) = 0$ .

<sup>18</sup> Since (3) simply asserts that Beauty ought to conditionalize on the information that it is Monday, it may seem that (3) should be accepted by anyone who is committed to Bayesian confirmation theory. But it is widely acknowledged that Bayesian conditionalization is not applicable to certain types of cases where what is learned is self-locating information. (Why? For starters, Bayesian conditionalization requires that certainties be preserved. But it is perfectly rational for an agent to go from being certain that it is morning to being certain at a later time that it is not morning. For discussion, see, e.g., (Meacham 2010; Bradley 2011b).) Whether or not conditionalization is appropriate when Beauty learns that it is Monday is contentious among halfers. In the original defense of the halfer view, Lewis (2001) affirmed that Beauty ought to conditionalize on this information, a view still held by some halfers (e.g., Jenkins 2005; Bradley 2011b). But most halfers today are “double halfers” who deny (3) and hold that Beauty’s credence for HEADS should remain  $\frac{1}{2}$  even after learning that it is Monday (see, e.g., Bostrom 2007; Meacham 2008; Pust 2012).

ought to be the case that  $C(H1) = C(T1)$ , thus completing Elga's argument for (2) and delivering us the thirder result.<sup>19</sup>

We are now in a position to see how an argument that is analogous to the one Elga gives in support of the thirder position can be given in support of the conclusion that, upon waking up together, Claire and Dillon will agree that the probability that Claire is the victim is 0.5. When Claire is awakened and sees Dillon being awakened at the same time, she knows that one of six mutually exclusive possibilities obtain, which I'll label C1–C3 and D1–D3, in accordance with the following chart:

	Claire is the victim.	Dillon is the victim.
It's the victim's 1 <sup>st</sup> awakening.	C1	D1
It's the victim's 2 <sup>nd</sup> awakening.	C2	D2
It's the victim's 3 <sup>rd</sup> awakening.	C3	D3

Let 'AGREEMENT' stand for the proposition that, upon being awakened during the night and seeing that Dillon has been awakened at the same time, Claire's credence for the proposition that she is the victim will be 0.5. (I'm calling this proposition 'AGREEMENT' because the exact same argument would also show that Dillon would assign a credence of 0.5 to the proposition that Claire is the victim, thus resulting in his agreeing with Claire.) And let  $C$  be Claire's credence function after waking and seeing that Dillon is her co-waker. AGREEMENT follows from the following two premises, which are analogous to (1) and (2) above:

- (4)  $C(C1)=C(C2)=C(C3)$  and  $C(D1)=C(D2)=C(D3)$
- (5)  $C(C1) = C(D1)$

Since C1–C3 and D1–D3 are mutually exclusive and collectively exhaustive possibilities, we know that  $C(C1)+C(C2)+C(C3)+C(D1)+C(D2)+C(D3) = 1$ . From this and from (4) and (5), it follows that  $C(C1)=C(C2)=C(C3)=C(D1)=C(D2)=C(D3)=1/6$ . And since Claire's credence that she is the victim must be equal to  $C(C1)+C(C2)+C(C3)$ , it further follows that upon awakening, Claire's credence that she is the victim will be 1/2, giving us our conclusion that AGREEMENT is correct.

The same "indifference" reasoning offered in support of (1) above supports premise (4). No matter who the victim is, Claire has no reason for thinking that it is more likely that this is the first (or second or third) awakening for the victim as opposed to either of the other two possibilities. It seems that rationality requires her to assign each of these possibilities equal credence, as (4) requires.

Premise (5), too, can be supported with an argument along the lines of the one given in support of (2). We can imagine that the experimenters select the victim in the following way: first, they randomly select two of the four subjects; then they wake both of these

<sup>19</sup> We can spell out this last step more explicitly. By Bayes's Theorem,  $C(H1 | H1 \text{ or } T1) = C(H1 \text{ or } T1 | H1) \cdot C(H1) / C(H1 \text{ or } T1)$ . So on the assumption that Elga has shown that  $C(H1 | H1 \text{ or } T1) = 0.5$ , we know that  $C(H1 \text{ or } T1 | H1) \cdot C(H1) / C(H1 \text{ or } T1) = 0.5$ . Obviously,  $C(H1 \text{ or } T1 | H1)$  is 1; so simplifying we get  $C(H1) / C(H1 \text{ or } T1) = 0.5$ . Since H1 and T1 are mutually exclusive,  $C(H1 \text{ or } T1)$  must be equal to  $C(H1) + C(T1)$ . Thus,  $C(H1) / [C(H1) + C(T1)] = 0.5$ . From this, only algebra is needed to show that  $C(H1)=C(T1)$ .

subjects during the first awakening of the night; and finally, *after* the two subjects have debated the probable identity of the victim and have been put back to sleep, they flip a fair coin in order to determine which of the two subjects just put back to sleep will be the victim. It seems that whether the experimenters use this method to select the victim or select the victim ahead of time should make no difference to Claire's credences for C1–C3 and D1–D3. So let's suppose that the experimenters use the two-stage approach to victim selection just described, and that Claire knows this. Now, suppose that sometime after being awakened at the same time as Dillon, Claire is told that this is the first awakening, and that the identity of the victim will be chosen via a coin toss after Claire and Dillon go back to sleep, with Claire being selected if the coin lands heads, and Dillon being selected if it lands tails. Upon learning this, Claire learns that either C1 or D1 is true, so that her credence for her being the victim and her credence for Dillon being the victim will be identical to her new credences for C1 and D1 (respectively). And since her new credence for C1 must be identical to her credence that a fair coin, *yet to be tossed*, will land heads, and her credence for D1 must be identical to her credence that a fair coin, *yet to be tossed*, will land tails, it seems that upon being told that this is the first awakening and the victim has yet to be selected, Claire's new credence for C1 (and for D1) ought to be 0.5.

Again, we need one more premise in order to conclude that (5) is correct:

- (6) The credence that Claire has for C1 *after* learning that it is the first awakening (and thus that C1 or D1) should be equal to the conditional credence  $C(C1 \mid C1 \text{ or } D1)$  that she had *before* learning this information.

Given (6) and the reasoning just rehearsed, it follows that *before* learning that it is the first awakening of the night, it ought to be the case that  $C(C1 \mid C1 \text{ or } D1) = 0.5$ . And from this, it follows that, before learning that it is the first awakening, it ought to be the case that  $C(C1) = C(D1)$ , completing the argument for (5) and thus for AGREEMENT.

Given that thirders will think that Claire and Dillon will be in perfect agreement, thirders can continue to affirm MODEST PROPORTIONALITY. And given the intuitive appeal of MODEST PROPORTIONALITY, the fact that thirders can readily affirm it and halfers cannot (at least not without abandoning the principal motivation for their position) constitutes a new and not insignificant reason in favor of the thirder position. But however plausible MODEST PROPORTIONALITY may at first appear, I will argue that there are plausible reasons for thinking that the multiple Beauties case does involve a robustly perspectival context, and that MODEST PROPORTIONALITY is therefore false. If I am right, halfers should not be worried by the fact that the halfer approach can conflict with MODEST PROPORTIONALITY.

#### 4. EVIDENTIAL SELECTION PROCEDURES AND PERSPECTIVAL RATIONALITY

We can summarize the challenge to the halfer position in the following way. If we consider the perspective of Claire, there is an intuitive line of reasoning that leads to the conclusion that there is a 0.75 chance that *Dillon* is the victim. And if we consider the perspective of Dillon, an exactly analogous line of reasoning supports the conclusion that there is a 0.75 chance that *Claire* is the victim. But clearly we, as third party observers of Claire's and Dillon's situation, have no reason for privileging either Claire's or Dillon's perspective

over the other's. To prefer one side or the other would be completely arbitrary. We ought therefore to assign an equal probability to Claire's being the victim as to Dillon's being the victim. But (and here is the critical though misguided move) Claire has no more reason than we do for privileging her own perspective over Dillon's. Or at least she has no more *epistemic* reason than we do for privileging her perspective. For the mere fact that a perspective happens to be hers rather than Dillon's is no reason for thinking that that perspective will better serve the aim of true belief and accurate credences. Given that Claire knows that neither she nor Dillon possesses any epistemic advantage over the other, privileging her own perspective would amount to arbitrarily selecting one out of two perspectives that, from a disinterested point of view, are equally likely to be the more reliable guide to the identity of the victim. Rationality, it seems, would require that such arbitrary selection be avoided and that perspectives with equal epistemic standing be given equal weight. Thus, the halfer view on the multiple Beauties case, and the "no new relevant evidence" reasoning that motivates that view, ought to be rejected as fallacious.

Against the above line of reasoning, I will claim that Claire's privileging her own perspective is not epistemically arbitrary, but is rationally required given the appropriate background assumptions about the process by which Claire has come to acquire the evidence she receives during her awakening with Dillon. I am not the first to argue that the rationally-required credences can vary for two people even in contexts where the uncentered information that grounds those credences is shared by both parties. Arnold Zuboff (2000) and Nick Bostrom (2000) have both offered examples in support of this perspectivalist claim. But the robust perspectivalism that I will argue for using the multiple Beauties case goes beyond the merely moderate perspectivalism implied by Zuboff's and Bostrom's examples. For as I argue in the footnotes, in the examples of Zuboff and Bostrom, the perspectivalist results depend on the inability of the multiple parties to share certain bits of uncentered information that, while not themselves evidentially relevant, cannot be shared without generating *new* evidence that *is* evidentially relevant and that would lead both parties to converge on the same credences.<sup>20</sup> So for these examples to

---

<sup>20</sup> Zuboff explicitly acknowledges that in his example communication between the disagreeing subjects must be disallowed in order for the example to yield a (moderately) perspectivalist result. While Bostrom's example yields a perspectivalist result only if we assume that communication is not possible, Bostrom does not explicitly stipulate that there is no communication. But Bostrom does argue that the perspectivalist results do not support the possibility of mutually-agreeable bets between the disagreeing parties (105-6). And the same reasoning that shows why there are no mutually-agreeable bets also shows that the differences in credences will not persist through communication and the full sharing of uncentered information. To see why the perspectivalist result is undermined by the full sharing of uncentered information, consider the following case, which is structurally just like Bostrom's case. Suppose that in an experiment some scientists leave me on a desert island in the middle of an undisclosed ocean. I know the following: first, if a particular coin flip conducted by the scientists came up heads, then one subject has been left on a desert island in the Indian ocean and one subject has been left on a desert island in the Pacific ocean; if it came up tails, then one subject has been left on a desert island in the Indian ocean and *ten* subjects have been left on ten different desert islands in the Pacific ocean; second, each subject has an electronic device that at 12:00 a.m. on January 1<sup>st</sup> will display which ocean they are located in; third, no subjects will have any evidence beyond this that can help them determine which ocean they are in or how many subjects are involved in the experiment. Now suppose that at 12:00 a.m. on January 1<sup>st</sup>, my device informs me that I am in the Indian Ocean; at the same time, Fiona, another subject, gets her message indicating that she is in the Pacific ocean. According to Bostrom, I should at this point be more confident than Fiona that the coin landed heads. For the chance of my being the one person in the Indian Ocean is much more likely given heads than it is given tails. But Fiona knows that *someone* has just learned that he (or she) is in the Indian Ocean. But since *she* is not that person, it



yield results that could plausibly be called “perspectivalist,” it must be assumed that the parties with opposed views cannot communicate. The examples therefore do not establish the possibility of *robustly* perspectival contexts. The kind of perspectivalism implied by these examples is thus less surprising, and less significant for the epistemology of disagreement, than the robust perspectivalism for which I will now argue. Still, the diagnosis of the multiple Beauties case that follows can be understood as an application and elaboration of some of the insights of Zuboff and Bostrom.

To understand why the credences that are rational for Claire and Dillon are perspective-dependent, it will be helpful to consider cases where certain facts about the evidential selection procedure give a third party reason for preferring either Claire’s or Dillon’s perspective. Suppose that the experimental philosophers have concluded their experiment with Alvin, Brenda, Claire, and Dillon, and after their results are recorded into a database, the experimenters allow you to make certain queries of the database and to see the answers the database returns. You have no knowledge concerning the identity of the victim other than what you learn in response to your queries. Suppose, first, that you instruct the database to first display the name of one randomly selected subject and then to display the name of someone who shared an awakening with that subject. If the first subject displayed is *not* the victim, then the next subject displayed will be the only person that this person shared an awakening with, i.e., the victim; and if the first subject displayed *is* the victim, then the database will randomly select one of the other subjects (each of whom each shared an awakening with the victim). And suppose that in response to these instructions, the database program first displays the name “Claire” and then the name “Dillon.” Let’s call the procedure just described “Procedure 1.” It is uncontroversial that if Procedure 1 is the only basis for your knowledge that Claire shared an awakening with Dillon, then you should have a 0.75 credence for the proposition that Dillon is the victim. For in response to the first query that is part of the procedure, three times out of four the program will display the name of someone who was *not* the victim, which means that the next query will display the name of the victim (since every non-victim shares an awakening only with the victim). Matters would have been very different if instead you had employed Procedure 2, where Procedure 2 consists in your directing the database program to randomly select one of the awakenings and then display the names of the two subjects involved in that awakening. If you had learned by Procedure 2 that Claire and

---

should not have the same rational import (for her) as my learning that *I* am in the Indian Ocean. And of course I know that at least one person has just learned that his or her island is in the Pacific. But this knowledge does not have the same import (for me) as Fiona’s learning that *her* island is in the Pacific. So centered information is the difference-maker here. Nonetheless, the perspectivalist result depends on imperfect sharing of uncentered information. For suppose that somehow Fiona and I were able to share all of our uncentered information. In this case, I would know not only that some subject is on an island in the Pacific, but also that *Fiona* is on an island in the Pacific. And Fiona would know that I know this. But now, Fiona has a new piece of evidence that is relevant to heads or tails: namely, that while I know that *Fiona* is on a Pacific island, I do not know the names of any other subjects on Pacific islands. And from Fiona’s perspective, the chance of my knowing about *her* as opposed to some other subject is much more likely given heads (in which case, she is the only Pacific island subject I *can* know about) than it is given tails (in which case, there is only a one in ten chance that Fiona would be the only subject I know about). With this new evidence made possible by our sharing all of our uncentered information, Fiona’s credence for heads will converge with my credence for heads. Essentially, communication often enables interlocutors to gain evidence about how representative they and their perspective are, and this evidence undermines any pre-communication perspectivalism.

Dillon shared an awakening, you would know that one of these two was the victim, but would have no basis for thinking either one of them is more likely to be the victim than the other.

Interestingly, Procedure 1 and Procedure 2 are both random procedures that are equally likely to result in your learning that Claire and Dillon shared an awakening; for Procedure 1 is equally likely to turn up information about the subjects in the first awakening as it is the second or third awakening. Nevertheless, if Procedure 1 is your method for arriving at the information that Claire and Dillon shared an awakening, it would be irrational for you to respond to this information as though all you had learned was that a randomly selected awakening involved Claire and Dillon. For you have learned something else that is evidentially relevant, namely that Dillon was a randomly selected “co-waker” of a randomly selected subject. This additional knowledge changes the rationally required response. The key point, one recently defended by Darren Bradley (2012), is that the process by which the evidence was selected is often itself a critical piece of evidence. And if this process is not known, background views and assumptions about the *likely* process will often play a critical role in determining the rational response to a piece of evidence.

Claire’s epistemic situation upon awaking with Dillon is, I will argue, relevantly like that of someone who, employing Procedure 1, has randomly selected Claire from among the four subjects and then learned that Dillon is a co-waker of this randomly selected subject. So upon waking up and seeing Dillon, Claire is justified in adopting the same credences as someone who has employed Procedure 1. Of course it is true that Claire did not randomly select herself from among four subjects. Nor is it possible, given the constraints of the experiment, for Claire to randomly choose one of the subjects in order to learn the identity of one of that subject’s co-wakers. Since Claire is not privy to information about any awakenings that do not involve herself, she cannot expect to learn the identity of any particular subject’s co-wakers except for herself. So Claire’s epistemic situation is relevantly like someone who has employed Procedure 1 only if she is justified in thinking of herself as a “randomly” selected subject. Claire can, I suggest, legitimately think of herself this way. Even though Claire is constrained in which of the subjects she is able to learn about, and even though this constraint biases Claire toward learning one of her own co-wakers rather than another subject’s co-wakers, this biasing constraint does not undermine the analogy between Claire’s situation and Procedure 1. For the fact that constrains Claire to “select” herself (namely, the fact that subjects cannot gain information about any awakenings in which they are not involved) is probabilistically independent of whether or not Claire is the victim. And such probabilistic independence is all that the “randomness” of Procedure 1 was meant to achieve. Given the lack of probabilistic correlation between Claire’s reason for selecting herself and the identity of the victim, Claire may legitimately think of herself as a “randomly” selected subject.

To help illustrate why such probabilistic independence is sufficient, imagine that instead of using the researchers’ database to carry out Procedure 1, you are going to randomly select one of the four subjects and hypnotize that subject in order to retrieve the memory of one (randomly selected) forgotten awakening. To your dismay, it turns out that Alvin, Brenda, and Dillon are not susceptible to hypnosis. Claire, however, is able to be hypnotized. As long as you know that whether a subject is susceptible to hypnosis is probabilistically independent of whether that subject was the victim, then there is no problem in your thinking of Claire as a “randomly selected” subject. Upon hypnotizing her

and learning the identity of a randomly selected co-waker, the rational implications will be the same as Procedure 1 as originally described. Similarly, Claire's reason for "selecting" Claire is a result of the epistemic constraints imposed by the experiment and the centered fact about her identity—facts that have no probabilistic correlation with Claire's being the victim. Thus, it seems that Procedure 1 is a fully adequate model for Claire's situation and that her credences should be identical to someone who has employed Procedure 1, randomly selected Claire, and then learned that Dillon is a co-waker of Claire's.

Of course Dillon is also justified in treating himself as a randomly selected subject and in adopting the credences of someone who, performing Procedure 1, selected Dillon randomly and then learned that Claire is a co-waker of Dillon's. And since Claire knows this, one might think that Claire's epistemic situation is best modeled by someone who has performed Procedure 1 twice, the first time randomly selecting Claire and then learning that Dillon is a randomly selected co-waker of this randomly selected subject, and the next time randomly selecting Dillon and learning that Claire is a randomly selected co-waker of this randomly selected subject. If this were the best model of Claire's situation, then it would indeed be the case that Claire ought to put equal credence in her being the victim and in Dillon's being the victim. But it would be a mistake for Claire to think that her situation is analogous to the situation of someone who has performed Procedure 1 twice. As we have seen, Procedure 1 is an adequate model only if the basis for the selection of the first subject is probabilistically independent of the identity of the victim. The fact that accounts for Claire's "selection" of Dillon from among the other subjects is the fact that she is currently sharing an awakening with Dillon. And this fact *is* probabilistically correlated with the identity of the victim, since having an awakening with Dillon is three times more likely if Dillon is the victim than if he is not the victim.

We are now in a position to appreciate why the rational credences are perspectival in this case, despite the perfect sharing of all uncentered evidence. The evidential significance of the information that Claire and Dillon share during their awakening depends on the process by which that information has been acquired. If the process yields information about a randomly selected co-waker of *Claire's* for reasons that are probabilistically independent of whether or not Claire is the victim, then the evidential significance of the information will be different than if the process yields information about a co-waker of *Dillon's* for reasons that are probabilistically independent of whether or not Dillon is the victim. Learning that Claire and Dillon share an awakening by the first kind of process (as in an instance of Procedure 1 where Claire is the randomly selected subject) can have no bearing on the likelihood of Claire's being the victim, and learning this information by the second kind of process (as in an instance of Procedure 1 where Dillon is the randomly selected subject) can have no bearing on the likelihood of Dillon's being the victim. But whether or not one has learned this evidence by a process of the first type or of the second type (or by some other type of process, like Procedure 2) depends on features of one's causal history that can vary from one subject to another. Thus, the rational significance of the evidence can depend on one's observational standpoint. Such seems to be the case in the multiple Beauties example. Claire and Dillon can share all of their uncentered information, but they cannot share their causal histories and thus cannot share the same observational standpoint. As a result, the evidential significance of their shared information differs for each of them, and different credences are called for.

One might worry that the robust perspectivalism I defend here stands in tension with Robert Aumann's agreement theorem, which says that two people who have the same prior credences will also have the same posterior credences for an event  $A$  if their posteriors are "common knowledge" (where persons 1 and 2 have common knowledge of event  $E$  if and only if "both know it, 1 knows that 2 knows it, 2 knows that 1 knows it, 1 knows that 2 knows that 1 knows it, and so on") (1976, 1236).<sup>21</sup> Essentially, Aumann shows that two ideal Bayesians with common priors will converge on the same posteriors if they fully disclose their credences to one another. On the supposition that ideally rational agents have the same priors, the agreement theorem would rule out robust perspectivalism, since according to the latter view, ideally rational agents can be required to disagree even in a context of perfect communication. But the view on the multiple Beauties case I have been defending here gives us reason to reject the supposition that maximally rational agents must have the same priors. I have argued that centered information like "I am Claire" can be evidentially relevant to an assessment of an uncentered proposition, since knowledge of who you are (and what your causal history is) may be needed to determine what evidential selection procedure you have used to acquire your evidence.<sup>22</sup> Thus, the prior credences that collectively determine Claire's views on the identity of the victim will include credences for certain centered propositions. Since Claire's prior for "I am Claire" is 1, and Dillon's prior for "I am Claire" is 0, their relevant prior credences are not the same and Aumann's agreement theorem does not apply. And in this case the divergent priors are clearly not due to any rational shortcoming.

## 5. A HARDER CASE?

In the previous section, I gave what I take to be a satisfactory justification for the striking perspectivalism that halfers seem committed to in the multiple Beauties case. But I would now like to look briefly at another variant of the Sleeping Beauty case that is arguably more difficult for halfers to accommodate. Like the multiple Beauties case, this case also shows that halfers are committed to a surprising disconnect between how we assess someone's epistemic credentials and how much weight we give their views.

Let us suppose that alongside Sleeping Beauty during her time in the Experimental Philosophy Laboratory is a second subject named Informed Beauty. Informed Beauty will be put to sleep at the same time as Sleeping Beauty, and he will be awakened on Monday at the same time as Sleeping Beauty and then put to sleep at the same time after this awakening. But no matter what the outcome of the coin toss, Informed Beauty will be awakened on Tuesday, either along with Sleeping Beauty if her coin landed tails, or by himself if her coin landed heads. Another critical difference between Informed Beauty's situation and Sleeping Beauty's situation is this: Informed Beauty will not be subject to any memory tampering, and both he and Sleeping Beauty know this. So when Informed Beauty is awakened on Monday, he will know what day it is, and likewise when he is

---

<sup>21</sup> Thanks to an anonymous referee for suggesting that I show how robust perspectivalism can be squared with Aumann's theorem.

<sup>22</sup> The claim that one's evidence for an uncentered proposition is not always exhausted by one's uncentered information can be supported without committing to any particular stance on robust perspectivalism. For example, see the "Mystery Bag" example in (Titelbaum 2012, 235–6).

awakened on Tuesday. Unfortunately, the two Beauties are not allowed to communicate, so Sleeping Beauty cannot benefit from Informed Beauty's information.

Like Sleeping Beauty, Informed Beauty is known to be a paragon of rationality. Given that this is the case, it seems that when Sleeping Beauty awakens along with Informed Beauty on Monday morning, she ought to regard Informed Beauty as an "expert" (relative to her) with respect to the probability of HEADS, where someone is an **expert** on  $p$  relative to  $S$  just in case his epistemic position with respect to  $p$  is at least as good as  $S$ 's in every respect (including possession of evidence, rationality of judgment, and functioning of cognitive faculties) and is superior to  $S$ 's in at least one respect. Sleeping Beauty ought to regard Informed Beauty as an expert because he has all of the evidence that Sleeping Beauty has, is just as rational as Sleeping Beauty, and possesses one piece of pertinent knowledge that Sleeping Beauty does *not* have: namely, knowledge of what day it is. Since knowledge of what day it is is relevant to assessing the likelihood of HEADS (since if one knows that it is Tuesday and whether or not Sleeping Beauty is awake, one can confirm whether HEADS or TAILS is true), it seems that Sleeping Beauty ought to prefer Informed Beauty's credence to her own, and thus ought to mirror Informed Beauty's credence as best as she can. For it seems that rationality requires that our current credences "reflect" the credences of acknowledged experts. Specifically, the following seems to be a rational constraint:

**EXPERT REFLECTION:** If I know that  $S$  is an expert on  $p$  relative to myself, then my credence for  $p$  conditional on  $S$ 's credence for  $p$  being  $x$  should also be  $x$ . Formally, if  $C$  is my credence function and  $C_S$  is  $S$ 's credence function, then it ought to be the case that  $C(p \mid C_S(p)=x) = x$ .

Like the rule of proportionality, EXPERT REFLECTION requires that my credences be calibrated with my views on the epistemic credentials of myself and others. An upshot of EXPERT REFLECTION is that if  $S$  is an acknowledged expert on  $p$ , and if I have precise credences for all the possibilities for  $S$ 's credence for  $p$ , then my current credence for  $p$  should be equal to the expected value of  $S$ 's credence for  $p$ .

EXPERT REFLECTION has been endorsed by some philosophers, and it initially seems quite plausible.<sup>23</sup> And the claim that Sleeping Beauty ought to regard Informed Beauty as an expert is also very plausible. But halfers must reject one of these claims or embrace a highly implausible version of the halfer view. To see why, suppose that upon awakening, Sleeping Beauty's credences for H1, T1, and T2 are those endorsed by most halfers: 0.5, 0.25, and 0.25, respectively. What will Sleeping Beauty's expected value be for Informed Beauty's credence for HEADS? Well, if it is Monday, then Informed Beauty will not know if the coin landed heads or tails and will surely have a credence for HEADS of 0.5. And if it is Tuesday, then Informed Beauty will know this (since he'll remember the Monday

---

<sup>23</sup> While I have defined an "expert" as someone whose epistemic position is at least as strong as mine in every respect (including both information and judgment) and superior in at least one respect, and then offered EXPERT REFLECTION as an intuitive proposal about how we ought to reflect the credences of such an expert, Elga (2007, 479–80), adapting the "expert" terminology from Gaifman (1988), *defines* an expert as someone whose credences we reflect, and suggests that someone ought to be treated as an expert if she has all of the information I have and more, and if her judgment is at least as good as mine. So Elga does seem to endorse EXPERT REFLECTION. Titelbaum (2012, 147–8) also endorses an "interpersonal Generalized Reflection principle" that has EXPERT REFLECTION as a consequence.

awakening the day before), and, upon seeing that Sleeping Beauty is awake, will assign HEADS a credence of 0. Since Sleeping Beauty believes with 0.75 confidence that it is Monday, her expected value for Informed Beauty's credence for HEADS is  $0.75 \cdot 0.5 + 0.25 \cdot 0 = 0.375$ . Since this value differs from Sleeping Beauty's credence for HEADS of 0.5, Sleeping Beauty will violate EXPERT REFLECTION if she takes Informed Beauty to be an expert relative to her with respect to HEADS.

On the assumption that Sleeping Beauty should have precise credences for H1, T1 and T2, the only positions that allow Sleeping Beauty to reflect Informed Beauty's credences are those that maintain that her credence for H1 should be equal to her credence for T1.<sup>24</sup> So thirders have no problem accommodating EXPERT REFLECTION in this case. An "optimistic halfer" position that maintains that, upon awakening on Monday, Beauty ought to be certain that it is Monday (with a credence of 0.5 for H1 *and* a credence of 0.5 for T1) also has no problem accommodating the view that Sleeping Beauty ought to reflect Informed Beauty's credences. But while this view has at least one defender (Hawley 2013), I think most would judge it exceedingly implausible. If we set optimistic halfism aside, the halfer is left with two options: deny EXPERT REFLECTION, or deny that Sleeping Beauty ought to count Informed Beauty as an expert.

The latter option is a promising response for those halfers who affirm that Sleeping Beauty ought to conditionalize upon learning that it is Monday (e.g., Lewis 2001; Jenkins 2005; Bradley 2011b). If these halfers are correct in holding that Sleeping Beauty's credence for HEADS should move to 2/3 after learning that it is Monday (even if the coin has yet to be tossed), then at this stage Sleeping Beauty must have some sort of evidence that is not possessed by Informed Beauty and that justifies a credence that departs from the objective probability of the coin toss. The fact that Sleeping Beauty alone possesses this evidence gives her reason for not taking Informed Beauty to be an expert. But what could this evidence be? Here again, the differences in evidence could be explained in terms of differences in evidential selection procedures. From Sleeping Beauty's perspective, the current awakening could have been a Monday or Tuesday awakening. She thus possesses the following piece of evidence: this awakening that had to be a Monday awakening conditional on HEADS and that was equally likely to be a Monday or Tuesday awakening conditional on TAILS turned out to be a Monday awakening. From Sleeping Beauty's perspective, a "randomly" selected awakening proved to be a Monday awakening.<sup>25</sup> This evidence, which serves to confirm HEADS, is perspectival evidence since it cannot be possessed by Informed Beauty. While Informed Beauty can affirm that the present awakening could have been a Monday or Tuesday awakening *from Sleeping Beauty's perspective*, from his perspective the awakening was guaranteed to be a Monday awakening irrespective of whether HEADS or TAILS obtains, so that the fact of its being a Monday awakening has no evidential significance. In short, on this view we have another robustly perspectival context: because the day of the present awakening was initially uncertain for Sleeping Beauty but not for Informed Beauty, the fact that it is Monday has evidential significance for Sleeping Beauty that it cannot (and should not) have for Informed Beauty. And because on Monday there is evidence that is accessible to Sleeping Beauty and inaccessible to Informed Beauty, Sleeping Beauty should not count Informed

<sup>24</sup> Let C be Sleeping Beauty's credence function; Sleeping Beauty will reflect Informed Beauty's credences if and only if  $C(H1) = (C(H1) + C(T1)) \cdot 0.5 + C(T2) \cdot 0$ . Solving, we get  $C(H1) = C(T1)$ .

<sup>25</sup> This response was proposed by Darren Bradley in personal correspondence. See also (Jenkins 2005).

Beauty as an expert, and her failing to reflect Informed Beauty's credences therefore does not violate EXPERT REFLECTION.

However plausible the above response may or may not be, it is not available to the majority of halfers who deny that Sleeping Beauty ought to conditionalize on the information that it is Monday. These "double halfers" (e.g., Bostrom 2007; Meacham 2008; Pust 2012) maintain that after learning that it is Monday, Sleeping Beauty's credence for HEADS ought to remain at  $1/2$  (rather than moving to  $2/3$ , which was her prior conditional credence for HEADS conditional on its being Monday). Since double halfers deny that Sleeping Beauty has any epistemic advantage relative to Informed Beauty once Sleeping Beauty has learned what day it is, the claim that she has some epistemic advantage lacked by Informed Beauty *before* learning that it is Monday is unmotivated. Given that Sleeping Beauty knows that she will converge on Informed Beauty's credences once she acquires knowledge of what day it is, so that Informed Beauty's credences constitute an epistemic target for her, surely Sleeping Beauty must acknowledge that Informed Beauty is an expert. So the double halfer is stuck with having to reject EXPERT REFLECTION. Perhaps this result does not significantly intensify the already substantial worries about the double halfer position stemming from its rejection of conditionalization in this case. But since conditionalization is a *diachronic* constraint and EXPERT REFLECTION is a *synchronic* constraint, there is no guarantee that an explanation for the inapplicability of conditionalization will also provide the material for an explanation of the inapplicability of EXPERT REFLECTION.<sup>26</sup> So the apparent incompatibility of the double halfer position and EXPERT REFLECTION does provide additional grounds for worry.

While EXPERT REFLECTION is extremely plausible upon first inspection, I shall argue that there are cases where one should not reflect the credences of known experts, and further suggest that Informed Beauty is such a case. To see why EXPERT REFLECTION is false, consider the following case.<sup>27</sup> On Friday morning, Natasha announces to her online network of hundreds of friends that on Friday night she is going to perform a very important coin toss. She also informs her friends of the following: If the coin lands heads, Natasha will tell no one of the result. If the coin lands tails, Natasha will immediately share

---

<sup>26</sup> Indeed, some ways that double halfers have attempted to argue for the inapplicability of conditionalization when Sleeping Beauty learns that it is Monday quite clearly do not help with explaining why EXPERT REFLECTION should also fail to apply. Consider, for example, Pust's recent (2012) attempt to defend the double halfer position. Pust first notes that conditionalizing on evidence  $E$  requires that one have a prior credence for  $E$ , and thus that  $E$  be in the domain of one's prior credence function. He then argues that when Sleeping Beauty acquires the temporally indexical evidence that it is Monday *now*, the proposition that expresses that evidence is one that, due to its temporally indexical nature, could not be grasped by Sleeping Beauty at any other time. The proposition is therefore not in the domain of any prior credence functions, rendering conditionalization completely inapplicable. Or, if we are working in a framework where propositions express only eternal truths, then according to Pust, whatever else we may substitute in the place of a proposition to serve as the object of Sleeping Beauty's knowledge that it is Monday *now* will also be outside of the domain of her prior credence function. I don't think this attempt to defend the double halfer position succeeds. But the important point here is that even if it *did* succeed, it would not explain why EXPERT REFLECTION should fail to apply to Sleeping Beauty. For unlike conditionalization (as applied to temporally indexical knowledge), EXPERT REFLECTION does not require that one respond to temporally indexical evidence in a way that is coordinated with one's credal attitude towards that very same (not yet known) evidence at some *prior* time. Rather, it merely requires that one's *current* credence for some proposition be coordinated with one's *current* views on what an expert *currently* thinks about that proposition.

<sup>27</sup> Weatherson (2009) offers a somewhat similar counterexample to EXPERT REFLECTION in an online weblog.

this result with exactly one person who has been randomly selected in advance from the large pool of Natasha's out-of-state friends. No one besides Natasha knows the identity of this person, and Natasha ensures that none of her friends will talk to others about whether or not Natasha has contacted them about the coin toss.

As it happens, I am an in-state friend of Natasha's, so there is no chance that I will learn the result of the coin toss. But on Saturday morning, I know that there is one privileged friend of Natasha's who, conditional on the coin coming up tails, knows the result. Call this person "Special Friend." Let's suppose that I know that all of Natasha's friends are supremely rational thinkers, and that all of the information I have that could possibly bear on the outcome of Natasha's coin toss is information that is also possessed by all of Natasha's friends. This enables me to conclude that, relative to me, Special Friend is an expert with respect to the outcome of Natasha's coin toss. For Special Friend has all of my evidence plus knowledge of whether he or she has been contacted by Natasha, knowledge that will either very slightly confirm heads if he or she has not been contacted (since an out-of-state friend's not being contacted is *guaranteed* conditional on heads and only highly likely conditional on tails) or, if he or she *has* been contacted, fully confirm tails. So whatever the outcome of the coin toss, I can expect that Special Friend's credences for heads and tails will be more accurate than my own.

While Special Friend is clearly an expert in the stipulated sense, this is a case where I should *not* conform to the dictates of EXPERT REFLECTION. Because Special Friend's "expertise" is strongly skewed towards a particular direction, so that Special Friend is highly privileged with respect to evidence for tails but only slightly privileged with respect to evidence for heads, my "reflecting" Special Friend's credence in accordance with EXPERT REFLECTION would inevitably skew my credence for heads downwards.<sup>28</sup> And once I lower my credence for heads in order to reflect Special Friend's credence, I will be even *more* confident that Special Friend has a credence of 0 for heads, and I will have to lower my credence again, prompting yet greater confidence that Special Friend's credence for heads is 0 and calling for yet another decrease in my credence for heads. This process will continue indefinitely: given the extreme way in which Special Friend's expertise is skewed, the only credence for heads that stably satisfies EXPERT REFLECTION in this case is 0.<sup>29</sup> But clearly, rationality does not in this case require me to be perfectly confident that Natasha's coin landed tails! Thus, EXPERT REFLECTION is false.

I suggest that there is a similar sort of skewing effect in the Informed Beauty case that gives us reason for thinking that the requirement posited by EXPERT REFLECTION does not apply. In this case, however, what is skewed is not Informed Beauty's expertise, but rather

---

<sup>28</sup> For Special Friend's expertise to be skewed in this way, Special Friend must be ignorant of the fact that he or she is Special Friend (unless he receives the call from Natasha). Otherwise, the absence of definitive evidence for tails would itself be definitive evidence for heads. But this ignorance does not disqualify Special Friend from being an expert relative to myself on the outcome of the coin toss. For I also am ignorant of the identity of Special Friend, and there is no doubt that Special Friend's epistemic position is superior to my own.

<sup>29</sup> To see this, let  $c$  be my credence that Natasha's coin landed heads and  $r$  be the value that I expect Special Friend's credence for heads to be if the coin in fact landed heads. In this case, EXPERT REFLECTION requires that  $c = c \cdot r + (1 - c) \cdot 0$ . Solving, we get  $c = c \cdot r$ . Since I know that  $r$  is not 1 (recall: given heads, Special Friend will not know that he or she is Special Friend, and will only know that he or she has not been contacted—evidence that will justify an  $r$  slightly above 0.5), the only way that this constraint can be satisfied is if  $c = 0$ .



his expertise *conditional on Sleeping Beauty being awake and able to reflect his credences*. For while Informed Beauty can confirm either HEADS or TAILS on Tuesday, when Sleeping Beauty is awake he can only confirm TAILS. Thus, if Sleeping Beauty is awake and is considering the expected value of Informed Beauty's credence for HEADS *right now*, the fact that he might at some point be certain that HEADS can make no difference (since he is *not* certain of HEADS right now given that Sleeping Beauty is awake), while the fact that he might currently be certain of TAILS *will* make a difference to the expected value. Evaluated from Sleeping Beauty's vantage point, Informed Beauty is effectively a biased expert (in the manner of Special Friend), even though Informed Beauty is not biased when evaluated from the perspective of someone who is awake on both Monday and Tuesday. This gives us grounds for doubting whether EXPERT REFLECTION expresses a genuine rational requirement in this case.

Clearly, more must be said in order to articulate a corrected "expert reflection" principle and to determine whether this principle vindicates the halfer or thirder position. I think, though, that enough has been said to significantly blunt any worries that may have resulted from the realization that the majority halfer position is in tension with EXPERT REFLECTION.

## 6. CONCLUSION

I have argued that the highly intuitive reasoning behind the halfer solution to the Sleeping Beauty problem also leads to counterintuitive perspectivalist results. Halfers who continue to stand by this reasoning must affirm that the multiple Beauties case involves a failure of MODEST PROPORTIONALITY, and it seems that double halfers are committed to the view that the Informed Beauty case involves a failure of EXPERT REFLECTION. While these counterintuitive results may appear to constitute a significant challenge to the halfer position, I have attempted to show that there are plausible reasons for thinking that MODEST PROPORTIONALITY and EXPERT REFLECTION do in fact fail in the cases described. But whether or not one finds my diagnoses of the cases convincing, I hope to have at least demonstrated that there is a rather surprising connection between the debate concerning Sleeping Beauty and the apparently orthogonal debate concerning the epistemic significance of disagreement: halfers are committed to robust perspectivalism and therefore must deny MODEST PROPORTIONALITY, a principle that plays an important role in both conciliatory and non-conciliatory approaches to disagreement.<sup>30</sup>

## REFERENCES

- Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* 4 (6) (November): 1236–1239.
- Bergmann, Michael. 2009. "Rational Disagreement after Full Disclosure." *Episteme* 6 (3): 336–353.

---

<sup>30</sup> Thanks to Darren Bradley, Lara Buchak, David Christensen, Keith DeRose, Blake Roeber, Alex Worsnip, for discussions of the issues raised here and (in some cases) for comments on earlier versions of this paper. Two anonymous referees also provided very helpful comments.

- Boghossian, Paul Artin. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. New York: Oxford University Press.
- Bostrom, Nick. 2000. "Observer-Relative Chances In Anthropic Reasoning?" *Erkenntnis* 52 (1): 93–108.
- . 2007. "Sleeping Beauty and Self-location: A Hybrid Model." *Synthese* 157 (1): 59–78.
- Bradley, Darren. 2011a. "Confirmation in a Branching World: The Everett Interpretation and Sleeping Beauty." *The British Journal for the Philosophy of Science* 62 (2): 323–342.
- . 2011b. "Self-location Is No Problem for Conditionalization." *Synthese* 182 (3): 393–411.
- . 2012. "Four Problems About Self-Locating Belief." *Philosophical Review* 121 (2): 149–177.
- Briggs, Rachael. 2010. "Putting a Value on Beauty." *Oxford Studies in Epistemology* 3: 3–34.
- Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5): 756–767.
- . 2011. "Disagreement, Question-Begging and Epistemic Self-Criticism." *Philosopher's Imprint* 11 (6) (March): 1–22.
- Elga, Adam. 2000. "Self-locating Belief and the Sleeping Beauty Problem." *Analysis* 60 (266): 143–147.
- . 2007. "Reflection and Disagreement." *Noûs* 41 (3): 478–502.
- Enoch, David. 2010. "Not Just a Truthometer: Taking Oneself Seriously (but Not Too Seriously) in Cases of Peer Disagreement." *Mind* 119 (476): 953–997.
- Feldman, Richard. 2006. "Epistemological Puzzles About Disagreement." In *Epistemology Futures*, edited by Stephen Hetherington, 216–236. New York: Oxford University Press.
- . 2007. "Reasonable Religious Disagreements." In *Philosophers Without Gods: Meditations on Atheism and the Secular Life*, edited by Louise M. Antony, 194–214. New York: Oxford University Press.
- Fumerton, Richard. 2010. "You Can't Trust a Philosopher." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 91–110. New York: Oxford University Press.
- Gaifman, Haim. 1988. "A Theory of Higher Order Probabilities." In *Causation, Chance and Credence*, edited by Brian Skyrms and William L. Harper, 191–219. Dordrecht: Kluwer Academic.
- Goldman, Alvin. 2010. "Epistemic Relativism and Reasonable Disagreement." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 187–215. New York: Oxford University Press.
- Gutting, Gary. 1982. *Religious Belief and Religious Skepticism*. Notre Dame, IN: University of Notre Dame Press.
- Hájek, Alan. 2003. "What Conditional Probability Could Not Be." *Synthese* 137 (3): 273–323.
- Halpern, Joseph. 2005. "Sleeping Beauty Reconsidered: Conditioning and Reflection in Asynchronous Systems." *Oxford Studies in Epistemology* 1: 111–142.
- Hawley, Patrick. 2013. "Inertia, Optimism and Beauty." *Noûs* 47 (1): 85–103.

- Jenkins, C. S. 2005. "Sleeping Beauty: A Wake-Up Call." *Philosophia Mathematica* 13 (2) (June 1): 194–201. doi:10.1093/philmat/nki015.
- Kelly, Thomas. 2010. "Peer Disagreement and Higher Order Evidence." In *Disagreement*, edited by Richard Feldman and Ted A. Warfield, 111–174. New York: Oxford University Press.
- Lackey, Jennifer. 2010. "What Should We Do When We Disagree?" *Oxford Studies in Epistemology* 2: 274–93.
- Lewis, David. 2001. "Sleeping Beauty: Reply to Elga." *Analysis* 61 (271): 171–76.
- Meacham, Christopher J. G. 2008. "Sleeping Beauty and the Dynamics of de Se Beliefs." *Philosophical Studies* 138 (2): 245–269.
- . 2010. "Unravelling the Tangled Web: Continuity, Internalism, Uniqueness and Self-Locating Belief." *Oxford Studies in Epistemology* 3: 86.
- Pust, Joel. 2012. "Conditionalization and Essentially Indexical Credence." *The Journal of Philosophy* 109 (4): 295–315.
- Thurrow, Joshua C. 2012. "Does Religious Disagreement Actually Aid the Case for Theism?" In *Probability in the Philosophy of Religion*, 209–224. Oxford: Oxford University Press.
- Titelbaum, Michael G. 2008. "The Relevance of Self-Locating Beliefs." *Philosophical Review* 117 (4): 555–606.
- . 2012. *Quitting Certainties: A Bayesian Framework Modeling Degrees of Belief*. Oxford University Press.
- Vahid, Hamid. 2004. "Varieties of Epistemic Conservatism." *Synthese* 141 (1): 97–122.
- Van Inwagen, Peter. 2010. "We're Right. They're Wrong." In *Disagreement*, edited by Richard Feldman and Ted Warfield, 10–28. New York: Oxford University Press.
- Weatherson, Brian. 2005. "Should We Respond to Evil with Indifference?" *Philosophy and Phenomenological Research* 70 (3): 613–635.
- . 2009. "Trusting Experts." *Thoughts Arguments and Rants*.  
<http://tar.weatherson.org/2009/05/06/trusting-experts/>.
- White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19 (1): 445–459.
- Zuboff, Arnold. 2000. "The Perspectival Nature of Probability and Inference." *Inquiry* 43 (3): 353–358.