

Coherence

Ram Neta
University of North Carolina – Chapel Hill

Rationality requires coherence. But what is coherence? The past two decades have witnessed the growth of a large literature dedicated to answering this question. But, as I shall argue here, the reductive ambitions of the extant literature cannot be satisfied. A correct account of coherence – at least of the kind of coherence that's required for rationality – cannot avoid appeal to the notion of rationality itself. In this paper, I will appeal to some intuitive desiderata for an account of coherence in order to do two main things: first, to criticize reductive approaches to coherence, and second, to defend a new, non-reductive account of coherence. I will then defend my account of coherence against objections. Along the way, I will say what's wrong with the Sleeping Beauty argument against the conditionalization requirement on credence updating, and suggest a similar diagnosis for analogous arguments against other coherence requirements.

One important source of interest in the topic of coherence has been a philosophical picture of rationality, according to which, while our practical or theoretical success might be hostage to fortune, there is some purely "internal" facet of rationality that is free from the influence of fortune: that purely "internal" fact has often been thought to consist in coherence. What I want to do here is to show that, once we understand what coherence is, we can no longer think of it as a refuge from the vicissitudes of fortune.

So: what is coherence? The word "coherence" simply means *sticking together*, and clearly, there can be forms of coherence that have nothing to do with rationality, e.g., the coherence of the parts of a complex machine. But the question to be addressed in this paper is: what is the kind of coherence that is required for rationality? It will help us to address this question if we begin by considering a pair of cases in which the relevant kind of coherence is radically absent.

Dreaming Dante.

Dante dreams a lot – in fact, roughly $\frac{1}{2}$ of his apparent experience takes place while he's dreaming. Furthermore, his dreams have the same force, vivacity, and coherence as the conscious experiences he has while awake, even though, in his dreams, he is always a non-human creature. And worse still for Dante, his dreams always begin with him waking up in the dream, they then pick up where they left off the previous night, and they end with him going to sleep in the dream. Dante, sadly, can never tell whether he's awake or dreaming, and so cannot rely on his own perceptual experiences to disclose much, if anything, about the physical world around him.

Vacillating Vanya.

Whenever Vanya decides on a particular course of action, then the preferences that made it rational for him to choose that course of action change in just such a way as to make

that chosen course of action irrational for him. For instance, since he now prefers chocolate to vanilla, he orders chocolate. But as soon as he places his order, his preferences change so that now he prefers vanilla. And if he then orders vanilla, his preferences will change so that now he prefers chocolate again. Vanya is thereby condemned to live a life in which he will disprefer whatever choice he makes to one of its alternatives: Vanya can never be satisfied doing whatever he prefers to do.

These cases are examples of extreme incoherence. Thankfully, we are not in any of these dire predicaments. But we resemble Dante and Vanya in a limited way: at least sometimes, we are not in a position rationally to rely on our own experiences, preferences, thoughts, or feelings. We admit to this when we write our prefaces, tie ourselves to the mast, or refuse to put our money where our mouth is. In these cases, we lack coherence: our experiences, preferences, thoughts or feelings – more generally, the mental states our being in which constitutes our possession of reasons – are in conflict with each other. And because we lack coherence, we cannot be fully rational.

Of course, I needn't have reached for such radical examples to illustrate the kind of coherence that I want to discuss here. But I did so for reasons that will become clear only towards the end of the paper.

Now that we have seen a couple of cases that exemplify the kind of coherence (or incoherence) that is required for (or incompatible with) rationality, I turn next to enumerating the desiderata for an account of the relevant kind of coherence.

A. Desiderata:

(1) *Variety of objects*. Coherence can be a feature of experiences, as indicated by the fact that it is present in each of Dante's dreams, but missing from his overall series of experiences. It can be a feature of preferences over time, as indicated by the fact that it is missing from Vanya's preferences over time. And, of course, coherence can be a feature of theories, stories, plans, promises, belief-systems, speeches, narrative films, and even pictorial representations – as indicated by the fact that any of these kinds of things can be, and in some cases are, incoherent. (Not that all of these things are subject in the same way to rational constraints; more about this below.) In general, then, coherence is a feature of composite representations, i.e., representations that have other representations as proper parts. I seek an account of coherence in general – one that explains what coherence is for any composite representation, no matter what sorts of representational parts it has.

(2) *Comparative*. Some composite representations are more coherent than others, even when neither is completely coherent. Sometimes, this same point is put by saying that coherence "is a matter of degree". But that formulation is misleading if it suggests that coherence can be measured. If the relation "more coherent than" induced a total ordering on composite representations, then it could be measured, but there is no reason to think that this is the case. So, in contrast to most philosophers who work on coherence today, I will not seek a measure of coherence. But I will seek an account of coherence that

predicts that, and explains why, some composite representations are more coherent than others, even if neither is completely coherent.

(3) *Rational Acceptability*. You might have reasons to propound an incoherent theory: perhaps you seek to amuse or confound your audience. But, whatever the benefits might be of doing so, it cannot be rational for you to accept an incoherent theory as true. Of course, each proposition of that theory might be such that it is possible for you rationally to accept that proposition as true. But the theory itself cannot be rationally acceptable if it is incoherent. Again, you might have reasons to propose an incoherent plan of action. But, whatever the benefits might be of doing so, it cannot be rational for you to accept that plan as your own if it is incoherent. Of course, each proper part of that plan might be such that it is possible for you rationally to accept it as your own. But the whole plan cannot be rationally acceptable if it is incoherent. In general, composite representations that are incoherent are not rationally acceptable – that is, it is not rational to accept that those representations perform the function that it is in their nature to perform, given the kinds of representations that they are.

(4) *Synchronic or Diachronic*. A thinker could be incoherent at a moment, or over time. Like Dreaming Dante, the thinker could have an incoherent series of experiences (proper subseries of which were themselves coherent). Like Vacillating Vanya, the thinker's preferences could evolve incoherently (even if, at any particular time, the preferences that he has at that time are themselves coherent). What these examples indicate is that coherence is a property of durable processes or activities as much as it is a property of momentary states or events. Epistemologists should be familiar with this phenomenon from the study of diachronic coherence: a thinker's credal state could be probabilistically coherent at each moment, even if the thinker is not diachronically coherent. An epistemic analog of Vacillating Vanya would be the thinker who updates her credal states in an incoherent way. If the Bayesian principle of conditionalization is a principle of coherent updating, then a thinker who updates her credences in a way that violates conditionalization is an incoherent thinker. I will say more about this below.

(5) *Support is unobstructed*. A composite representation is incoherent in so far as a reason to accept one part of it provides a reason to reject another part of it. If what I say in my preface does not cohere with the conjunction of claims that I make in my book, then any reason I have for believing what I say in my preface can provide a reason to reject what I say in the rest of my book, and vice-versa. If my preference for chocolate over vanilla does not cohere with my other preferences, then any reason I have for strengthening my preference for chocolate over vanilla will be a reason to reject my other preferences, and vice-versa. I'll say that, in an incoherent representation, support for one part is *obstructed*, for it is either opposed by support for other parts, or it opposes them, or both. And then I'll say that, in a coherent representation, support for one part is unobstructed.

I want an account of coherence that satisfies all five of these desiderata. But the literature on coherence has grown vast over the past two decades. Has such an account been offered already?

B. Why Reductive Accounts of Coherence Do Not Meet These Desiderata

Recent accounts of coherence are all reductive, and they attempt to account for coherence as a measure of some quantity or other, of sets of propositions. But there are at least three problems with this approach:

- (i) Not all of the representational states that are assessable for coherence have propositional content. For instance, it's not at all clear how to specify the propositional content of plans (one typically plans *to...*, one doesn't plan *that...*) or preferences (one typically prefers one thing to another, one doesn't prefer that ...). Thus, previous accounts of coherence do not satisfy our first desideratum.
- (ii) A particular set of propositions might be such that you can coherently distribute your confidence over all of them, but you cannot coherently believe all of them. For instance, consider this familiar sort of set:

{There will be precisely 1 winning ticket in the n-ticket lottery, ticket #1 will not win, ticket #2 will not win, ticket #3 will not win, ... (and so on for each of the other n tickets)}

One cannot coherently believe every element of this set of propositions, because they are jointly inconsistent. But one can coherently distribute one's confidence across this set, for instance by assigning credence 1 to the first proposition and by assigning credence $(n-1)/n$ to each proposition after the first. So is the set of propositions itself coherent or incoherent? I don't see how this question can be answered, except by specifying a kind of relation that one is to bear to each proposition in the set. But that is to specify something that extant accounts of coherence do not specify. Thus, such accounts do not satisfy our third desideratum.

- (iii) As Moretti and Akiba 2007 prove, for all of the formal measures of coherence devised so far, the coherence of a particular set of propositions depends upon how the conjunction of those propositions is carved up into the elements of that set. But this is implausible: whether a composite representation is coherent cannot plausibly depend upon how we individuate the proper parts of that composite.

Consider Shogenji's measure of coherence (as presented in Shogenji 1999), according to which the measure of the coherence of a set of propositions is the ratio of the probability of their conjunction to the product of their individual probabilities. Now consider two sets of propositions:

$\{A_1, A_2, A_3, A_4, A_5\}$ and $\{A_1 \& A_2 \& A_3 \& A_4 \& A_5, A_2\}$

Although each set contains the same information, the Shogenji coherence of set 2 will be higher than the Shogenji coherence of set 1, because the denominator of the coherence measure will be higher for the second set than for the first. This is implausible.

Again, consider Olson's measure of coherence (as presented in Olson 2002), according to which the measure of coherence of a set of propositions is the ratio of the probability of their conjunction to the probability of their disjunction. Now consider these two sets of propositions:

$\{A_1, A_2, A_3, A_4, A_5\}$ and $\{A_1 \& A_2 \& A_3 \& A_4 \& A_5, A_1 \& A_2 \& A_3 \& A_4 \& A_5\}$

Again, although each set contains the same information, the Olson coherence of set 2 will be 1, no matter what the Olson coherence of set 1 is. That's because the probability of the conjunction of the two elements in set 2 is equal to the probability of their disjunction.

As Moretti and Akiba show, the very same point holds true also for the measures of coherence proposed in Fitelson 2003 and in Bovens and Hartmann 2003. (I omit those proofs here, since they are much more complicated.) In sum, then, if there is a formal measure of coherence that applies to a set of propositions and will generate the same result no matter how the total information in that set of propositions is carved up into elements of the set, it hasn't been discovered yet.

There are other problems of detail with each of the measures that has been devised so far, but those problems have been successfully identified already in the literature (see, e.g., Schippers 2014 for a survey of those problems), and I won't review them here. But, by way of summarizing my conclusion from those problems of detail, I will say this: If the post-Gettier literature provides the materials for a strong inductive argument against a reductive account of knowledge as JTB plus some fourth condition, then the literature of the last 15 years provides the materials for a strong inductive argument against a formal measure of coherence that applies to sets of propositions.

C. My non-reductive account of coherence:

I will first state my account of coherence briefly, then explain what it means, and finally show that it satisfies all five of the desiderata set out above.

So first, here's the account: Coherence is the most inclusive non-derivatively wide-scope constraint that rationality imposes on thinkers. A thinker is rational only in so far as she satisfies that condition.

Now I need to explain what this account means. In particular, I need to say something about what "constraints of rationality" are, about what it is for a constraint of rationality

to be “wide-scope”, about what it is for it to be “non-derivatively wide-scope”, and what it is for it to be “most inclusive”.

So first, constraints of rationality: we know that not just any old thing that we might do or think is rational. Rationality places constraints on what we may do or think. But how can we figure out what those constraints are?

To answer this question, let’s begin by noticing that rationality doesn’t guarantee success, nor does irrationality guarantee failure. But rationality is supposed to lead to success, and when it fails to do so that is a matter of the thinker’s bad luck. Similarly, irrationality can be expected to lead to failure, and when it doesn’t do so, that is a matter of the thinker’s good luck. Philosophers can exploit this last thought in arguing that something is a requirement of rationality: it is a necessary condition of a putative requirement’s being a requirement of rationality that violation of that requirement leaves it a matter of good luck that the thinker is not ruined. This is the basic idea of Dutch Book arguments in favor of principles of synchronic rationality (probabilism, Reflection) and principles of diachronic rationality (conditionalization). A similar idea guides other arguments in favor of other principles of rationality (e.g., the requirement that our preferences be transitive, or that our beliefs at a time be consistent). So, if we want to figure out what the constraints of rationality are, we should begin by asking: what are the constraints our violation of which would leave it a matter of good luck that we are not ruined?

By following this line of thought, we can identify several constraints of rationality. Now what is it for a constraint of rationality to be “wide-scope”? Constraints of rationality are “wide-scope” when they constrain which *combinations* of representational states we have. They are “narrow-scope” otherwise. (And it is a substantive question whether there are any narrow-scope constraints: a question that John Broome would answer in the negative and I would answer in the affirmative.) To illustrate the distinction, compare the following two constraints:

For any proposition p , one ought not believe $(p \ \& \ \text{not-}p)$.

For any proposition p , one ought not simultaneously believe p and believe $\text{not-}p$.

The first is narrow-scope, constraining us not to have a particular belief, while the second is wide-scope, constraining us not to have a particular combination of beliefs at the same time.

Now, what is it for a constraint of rationality to be “non-derivatively” wide-scope?

Consider the narrow-scope constraint:

(*) For any proposition p , one ought not believe $(p \ \& \ \text{not-}p)$

Since “S ought not F” implies “S ought not both F and G”, the constraint above implies the following:

(%) For any proposition p , one ought not believe $(p \ \& \ \text{not-}p)$ and also at the same time

believe that tomorrow is Tuesday.

But notice that (%) is wide-scope: it prohibits a certain combination of representations. But what generates (%) is simply that it follows from (*), which is narrow-scope. I will say that wide-scope constraints like (%), although wide-scope, are “derivatively” wide-scope: they can be derived from narrow-scope constraints in the same way that (%) can be derived from (*), by means of the general theorem that, if S ought not F, then S ought not F&G.

So much for “non-derivatively wide-scope”. Now what is it for a constraint of rationality to be the “*most inclusive*, non-derivatively wide-scope” constraint? Consider various non-derivatively wide-scope constraints of rationality, each of which is identifiable as such by means of the sort of consideration mentioned above (our violation of the constraint would leave it a matter of good luck if we are not ruined):

For any proposition p, one ought not simultaneously believe p and disbelieve other propositions that one knows to follow from p. (closure)

One ought not simultaneously desire that p is the case and disbelieve that it would be good for p to be the case. (desire as belief)

For any set of propositions P, one ought to distribute one’s confidence over the elements of P in a way that satisfies the Kolmogorov axioms of the probability calculus. (probabilism)

For any set of propositions P, and any increment of evidence E, when one gains E, one ought to update one’s distribution of confidence over the elements of P by setting one’s new level of confidence in each element p_i to one’s previous level of conditional confidence (p_i|E). (conditionalization)

To illustrate the sort of consideration that would show the principles above to be constraints of rationality, let me focus on the Bayesian Principle of Conditionalization:

$$C_{\text{post-E}}(\text{H}) = C_{\text{pre-E}}(\text{H}|\text{E})$$

The argument for this principle is that violations of the principle make you susceptible to a Dutch Book, i.e., a combination of bets that can result in a loss, but cannot result in a net gain come what may.

Example:

Suppose your confidence that the coin will land heads on the first flip is .5, and your confidence that the coin will land heads on the second flip if it lands tails on the first flip is also .5. Since your confidence in p fixes what you regard as a fair price for a bet on the truth of p, you will accept the following offer as fair: you pay me \$.50, and in return I give you two bets. On Bet A, you win \$1 if the coin lands heads twice in a row, and you win \$0 otherwise. On Bet B, you win \$.50 if the coin lands tails on the first toss, and you

win \$0 otherwise.

Now suppose the coin is flipped once and lands heads. What is your new level of confidence that the coin will land heads on the next toss? If you follow the principle of conditionalization, your new level of confidence is .5. But suppose you don't comply with the principle of conditionalization, and instead you allow your new level of confidence to be .4. In that case, you are committed to accepting the following offer as fair: I pay you \$.40, and in return you give me a third bet (Bet C), in which I win \$1 if the coin lands heads on the next toss and I win \$0 otherwise.

In sum, you will have committed to accepting the following as fair:

Possible results of coin toss	Net gain for you
HH	-.60
HT	-.10
TH	0
TT	0

The result can be proven in general (cite Lewis, put general proof in footnote). In short, by failing to comply with the principle of conditionalization, your states of confidence have committed you to accepting a series of offers that can result in a net loss for you, but cannot result in a net gain no matter what happens. Your states of confidence are therefore not guiding your choices well, and this is something that reasoning alone can disclose to you. In short, reasoning reveals that you should not rely on your states of confidence to do what they're supposed to do, so long as you fail to comply with the principle of conditionalization. The principle of conditionalization is a wide-scope diachronic constraint of rationality.

A very similar argument would show that probabilism is a wide-scope synchronic constraint of rationality, and other analogous arguments involving acting on beliefs instead of credences would show that closure and consistency are also wide-scope synchronic constraints of rationality.

When I say that coherence is the most inclusive, non-derivatively wide-scope constraint of rationality, what I mean is that coherence includes all of these particular non-derivatively wide scope constraints of rationality. For a composite representation to be fully coherent is for it to comply with *all* of the non-derivatively wide-scope constraints of rationality that apply to the representations that are part of that composite. A composite representation is coherent to the extent that it complies with all such constraints.

This account of coherence is designed to apply, in the first instance, to those representational mental states that are assessable for rationality, viz., beliefs, preferences, credal states, etc. How does it apply to thinkers? And how does it apply to representational mental states that are not assessable for rationality, e.g., perceptual experiences? And how does it apply to representations that exist independently of any

particular person's thinking, e.g., stories, or narrative plays and films?

Let's call a thinker "coherent" to the extent that this thinker's mental states are coherent, according to the account above (i.e., they comply with all the non-derivatively wide-scope norms of rationality). Then, we'll say that perceptual experiences (or other non-rationality assessable mental states) are "coherent" just in so far as they are experiences that a coherent thinker could have – experiences that would allow the thinker to form coherent beliefs, preferences, and states of confidence about the world in which and on which she must act.

Finally, we'll say that a non-mental representation, like a story, is coherent just in so far as it is coherently credible. Thus, stories about time travel, or a political economy the participants of which are all beasts of different species, are not coherent. But, since rationality does not constrain stories in the same sort of way that it constrains thinkers, to say that a story is incoherent is not to say that it does not reward attention.

That is my account of coherence. Now, I need to show that it satisfies all five of the desiderata set out above.

(1) *Variety of objects*. Coherence, on the account above, can be a property of any composite representation, whether its parts are beliefs, states of confidence, preferences, stories, plans, or perceptual experiences. It thus satisfies this first desideratum.

(2) *Comparative*. Quite generally, satisfying constraints is something that can be done more or less well. (I have been pretty good at satisfying the prohibition against stealing, but why do I still have that pen labeled "Holiday Inn, Des Moines IA"?) *A fortiori*, satisfying the non-derivatively wide-scope constraints of rationality is something that composite representations can do more or less well. This is illustrated in the case of coherence norms by the fact that Dutch Books can be more or less dangerous for the person who is committed to their fairness.

(3) *Rational Acceptability*. A composite representation is rationally acceptable only in so far as it satisfies the constraints of rationality. *A fortiori*, it is rationally acceptable only in so far as it satisfies the non-derivatively wide-scope constraints of rationality.

(4) *Synchronic or Diachronic*. Coherence is a property of composite representations, whether they are representational states or events at a time (e.g., beliefs) or representational processes (e.g., belief updating).

(5) *Support is unobstructed*. Suppose a composite representation is incoherent according to the account above. In that case, the representation does not comply with some wide-scope constraint of rationality: if it is a composite of beliefs, they are either not consistent or not closed, if it is a distribution of confidence, it is not a probability, if it is a preference function, it is not transitive and irreflexive, and so on. But to say that the representation does not comply with wide-scope constraints of rationality is to say that rationality demands the representation to be otherwise than it is. And so if we have a

reason to accept one part of that composite representation, that reason will either be opposed by the reason that we have to reject the composite representation itself, or else it will oppose the reason that we have to reject the composite representation, or both. In other words, a representation that is incoherent on the present account will be a representation in which support is obstructed. A necessary condition to avoid such obstruction of support is for one's composite representation to comply with all of the non-derivatively wide-scope constraints of rationality. But that, on my account, is precisely for one's composite representation to be coherent.

I conclude that my account of coherence satisfies all five of the desiderata mentioned. If there is another account that does so, I'm not aware of it.

In the next section, I address what seem to me to be the most tempting objections to the account.

D. Objections:

Objection 1. It is generally accepted that "ought" implies "can". But it is empirically obvious that actual humans cannot comply with all of the wide-scope constraints enumerated above (e.g., closure, probabilism, conditionalization) so how can it be true that we ought to do so?

Reply: "Ought" implies a certain kind of "can", but it does not imply every kind of "can". I ought to drive carefully, but I cannot do so when I'm asleep. The kind of "can" implied by "ought" is the "can" of competence (what Ernest Sosa calls the "seat" of our capacity), not the "can" of opportunity (what Sosa calls the "shape" or the "situation" of our capacity). While our grammatical capacity is infinite, in the sense that it operates according to principles that sort even infinitely long syntactic strings as well-formed or ill-formed, that capacity is never in a shape or a situation to be exercised properly on extremely long (let alone infinitely long) syntactic strings. And similarly, while our rational capacity is infinite, in the sense that it operates according to principles that sort infinitely complex composite representations as rational or irrational, that capacity is never in a shape or a situation to be exercised adroitly on extremely complex (let alone infinitely complex) composite representations.

Objection 2: If coherence involves compliance with all of the non-derivatively wide-scope constraints of rationality, and if consistency and probabilism are both non-derivatively wide-scope constraints of rationality, then, as Christensen 2004 has shown, complying with some demands of coherence can make it impossible to comply with others. And if this is right, then coherence is, by my own lights, incoherent.

Reply: Christensen is right that there are some conditions under which we cannot comply with all of the non-derivatively wide-scope constraints of rationality. To take the most familiar example, consider the set of propositions:

{There will be precisely 1 winning ticket in the n-ticket lottery ($n > 100$), ticket #1 will not win, ticket #2 will not win, ticket #3 will not win, ... (and so on for each of the other n tickets)}

One can coherently distribute one's confidence across this set, for instance by assigning credence 1 to the first proposition and by assigning credence $(n-1)/n$ to each proposition after the first. But one cannot coherently believe every element of this set of propositions, because they are jointly inconsistent. So, if one cannot coherently believe every element of the set above, then which elements should one believe? None but the first? That may seem to constitute an unreasonable skepticism, especially given the strength of evidence one has in favor of each of the other propositions. Some random selection of the other propositions, as long as it's not so large as to make its conjunction improbable? But that is to adopt an incoherent policy of doxastic response to evidence: sometimes believing and sometimes not believing, in response to what is, by one's own lights, relevantly similar evidence. So how can one comply with the wide-scope constraints of rationality in one's beliefs?

Christensen is right that one cannot do so in this case. But that doesn't imply that one can never comply with the wide-scope constraints of rationality in one's beliefs. It implies only that, if you want to be coherent, you should generally avoid forming any doxastic attitudes (positive or negative) about which lottery tickets are going to win. But this does not imply that you should avoid forming doxastic attitudes about other hypotheses for which you have at least equally strong statistical or probabilistic evidence.

In short, there are some sets of propositions, and some bodies of evidence, such that you cannot coherently adopt doxastic attitudes to all of those propositions given that body of evidence. But to say that there are some such pairs of propositions and evidence is not to say that coherence is impossible, but only to say that there are specific conditions of its possibility. Sadly for him, Dreaming Dante is not in such conditions, and neither are the players in a lottery.

Objection 3: Several of the principles of coherence that I've enumerated above – and that have been defended by appeal to the “Dutch Book” style of argument mentioned above -- have been refuted by counterexamples, and so the argumentative strategy that I've described for identifying constraints of rationality is unreliable. To take one famous example, Elga's famous Sleeping Beauty case refutes the principle of Conditionalization. Should we conclude that conditionalization and other apparent principles of coherence are not really principles of coherence? Or should we instead conclude that coherence is not a requirement of rationality for thinkers? Either way, my account of coherence will be forced to make an important concession in the face of cases like Sleeping Beauty.

Reply: No, my account of coherence will not be forced to make either of these concessions. The principle of conditionalization is a principle of coherence, and coherence is rationally mandatory for thinkers. But, sometimes – as I pointed out in my reply to Objection 2 above – one is in a situation in which it is impossible to satisfy all

rational mandates. Elga's Sleeping Beauty case, I claim, is just such a case. Recall the case:

It is Sunday. When Beauty falls asleep tonight, a fair coin will be flipped. If the coin lands heads, then Beauty will be woken up on Monday morning, and everything will go on as normal. If the coin lands tails, then Beauty will be woken up on Monday morning, go to sleep again Monday evening, and then given an amnesiac so that when she wakes up on Tuesday morning she will have forgotten all about Monday. Beauty knows all this, and will continue to know it after she is awoken.

Beauty goes to sleep, the coin is flipped. Now Beauty wakes up, and cannot recall waking up previously (except prior to the coin flip). She cannot see a calendar, so she doesn't know if today is Monday or Tuesday. How confident should she be that the coin landed heads?

Notice that the coin lands heads in precisely $1/3$ of the situations in which she has this very experience – and Beauty knows that. If her knowledge that this is so forms part of her evidence, then compliance with the principle of total evidence should lead Beauty to confidence $1/3$. But the principle of conditionalization says that her confidence that the coin landed heads should be $1/2$: she knew it to be a fair coin in advance, so her earlier confidence that it would land heads given that she was woken up was $1/2$, and the only new evidence that she's gained is that she was woken up. Thus, it seems, given everything she knows about her situation, Beauty should either violate the principle of conditionalization and assign credence $1/3$, or else violate the requirement of total evidence and assign credence $1/2$. Doesn't this show either that conditionalization is not a principle of coherence, or else that coherence is not rationally mandatory?

No, it shows neither. The case of Sleeping Beauty, I claim, is like the cases of Vacillating Vanya or Dreaming Dante in this respect: there is no way for the protagonist of the case to satisfy all of the demands of rationality simultaneously, given the situation she's in. In particular, in the Sleeping Beauty case, the narrow-scope demand for her to conform to her total evidence at the moment that she wakes up, and thus assign confidence $1/3$ to the proposition that the coin landed heads, conflicts with the demand for her to be diachronically coherent, and thus conditionalize on all and only her new increment of evidence and assign confidence $1/2$ to that proposition.

What holds true of Sleeping Beauty, I claim, holds true of all of the other protagonists of the standard counterexamples to coherence principles like conditionalization, Reflection, etc. In each case, we find that, under certain circumstances, a thinker has reason to violate a principle of coherence. But this doesn't show that the principle in question is not really a principle of coherence, nor does it show that coherence is not rationally mandatory. What it shows rather, is that in some possible situations, the requirements of rationality are not jointly satisfiable.

Rationality may be a good and unequivocal guide to the situations in which normal humans normally find themselves. But it is not a good guide to every possible situation

that a philosopher or motivated Dutch Bookie can dream up. And why should we expect it to be otherwise?

E. Conclusion: Lucky to be Rational

Amia Srinivasan has recently argued that the failure of our reasons to be luminous to us means that we are sometimes in a position in which we cannot knowingly comply with the requirements of rationality. In this paper, I've reached a similar conclusion but without reliance on the claim – the argument for which I have elsewhere challenged (Neta and Rohrbaugh 2004) – that our reasons are not luminous to us. Even if our reasons are luminous to us, they might still leave us in a position in which we cannot comply with all the requirements of rationality. When our reasons, however luminous they may be, leave us in a position in which it is possible to comply with the coherence norms (and indeed, with all the norms) governing the exercise of reason, then we are fortunate, though this sort of fortune is routine enough that we take it for granted.

Some philosophers are concerned to see our compliance with the demands of rationality as protected from the vicissitudes of fortune. Coherence, I conclude, cannot be any aid or comfort to such philosophers.