

Saving the Appearances: Distinguishing Evidence from Knowledge

[Please note: this is a very rough draft that I wrote in the summer of 2009. Any published version of this paper will differ significantly from the present draft.]

1. Against $E = K$

1.1. Limited Extrapolation

The following appears to be a very plausible thesis:

LIMITED EXTRAPOLATION: If we possess the evidence that it has snowed in Moscow every year for at least one thousand consecutive years, and if we possess no contrary evidence suggesting a change in the Russian climate, and if, as a matter of fact, it snows in Moscow the next year not belonging to this set, then we are in a position to know that it snows in Moscow the next year. That is, if, on the basis of our evidence that it snows every year in this set, we form the belief that it snows the following year, then this belief will constitute knowledge.

By contrast, the following thesis appears to be very implausible:

UNLIMITED EXTRAPOLATION: If we possess the evidence that it has snowed in Moscow every year for at least one thousand consecutive years, and if we possess no contrary evidence suggesting a change in the Russian climate, and if, as a matter of fact, it will continue to snow in Moscow every year forever, then for any value of n , we are in a position to know that it snows in Moscow n years in the future.

UNLIMITED EXTRAPOLATION has the implausible implication that, if we have observed that it has snowed in Moscow every year for the past thousand years, and if, as a matter of fact, it will always snow in Moscow, then we are in a position to know that it will snow in Moscow a trillion trillion years in the future. Hence, while we should accept LIMITED EXTRAPOLATION, we should reject UNLIMITED EXTRAPOLATION. But if we accept $E = K$, then we cannot do this.

For suppose we have observed that it has snowed in Moscow every year from 1009 AD to 2009 AD, and hence we possess the evidence that it has snowed each of these years. And suppose that, as a matter of fact, it will forever continue to snow in Moscow every year. It follows from LIMITED EXTRAPOLATION that we are in a position to know that it snows in Moscow in 2010. Hence, assuming we form the belief that it snows in Moscow in 2010 on the basis of our evidence that it snowed each of the preceding thousand years, we will know that it snows in Moscow in 2010. And so, by $E = K$, we will possess the evidence that it snows in Moscow in 2010. Hence, we possess the evidence that it snows in Moscow every winter from 1009 to 2010. Thus, it follows from LIMITED EXTRAPOLATION that we are in a position to know that it snows in Moscow in 2011. And this

interates. Therefore, assuming LIMITED EXTRAPOLATION and $E = K$, for any value of n , we will be in a position to know that it snows n years in the future. For we can acquire this knowledge by repeatedly extrapolating from our evidence set and thereby extending our evidence incrementally into the future.

in Moscow in the winter of year 1,000,000,010 AD. And yet this contradicts LIMITED EXTRAPOLATION. Since the latter thesis, combined with $E = K$, gives rise to a contradiction, the two theses are incompatible. And since LIMITED EXTRAPOLATION is very plausible, we should reject $E = K$.

1.2. The Evidence-For Relation

If a proposition belongs to our body of evidence, then it can support, or serve as evidence for, other propositions. Thus, suppose that, every year for the past two hundred years, we have measured the temperature in Death Valley in July, and found it to be hot. And so our total evidence includes the following proposition:

PAST it has been hot in Death Valley every July for the past two hundred years.

This proposition will constitute evidence for various other propositions, such as the following

NEXT it will be hot in Death Valley next July

EVERY it is hot in Death Valley every July

Plausibly, p is evidence for q and r because p belongs to our evidence, and because p raises the probability of these other propositions. Williamson proposes just such an account of the evidence-for relation:

EV e is evidence for b for S if and only if S 's evidence includes e and $P(b|e) > P(b)$

where P represents a prior probability function that does not include e (Williamson 2000, p. 187).

This appears to be very plausible account of the evidence-for relation. However, it has very implausible consequences when combined with the thesis that $E = K$. Assume that we possess the evidence that PAST is true, and that this evidence suffices for us to know that NEXT is true. Assume, therefore, that since we believe that NEXT is true on the basis of the evidence that PAST is true, we know that q is NEXT. Thus, assuming $E = K$, NEXT must belong to our evidence. It follows, from EV, that NEXT will constitute evidence for EVERY so long as $P(\text{EVERY} | \text{NEXT}) > P(\text{EVERY})$, where P is some reasonable prior probability function that does not include the evidence that NEXT. And it seems that this inequality should be true. For, since P does not include the evidence that NEXT, $P(\text{NEXT})$ must be less than one. Hence, $P(\text{EVERY})$ is the weighted average of $P(\text{EVERY} | \text{NEXT})$ and $P(\text{EVERY} | \neg \text{NEXT})$. And clearly, $P(\text{EVERY} | \text{NEXT}) > P(\text{EVERY} | \neg \text{NEXT})$. And so it follows that $P(\text{EVERY} | \text{NEXT}) > P(\text{EVERY})$. Thus, assuming EV, and assuming $E = K$, so that NEXT belongs to our evidence, it follows that NEXT is evidence for EVERY.

Similarly, it follows that NEXT is evidence for PAST. For assuming EV, and assuming NEXT belongs to our evidence, it will be evidence for PAST so long as $P(\text{PAST} | \text{NEXT}) > P(\text{PAST})$, where P is some appropriate probability function that does not include the evidence that NEXT. But if

possessing the evidence that PAST is sufficient ground for knowing that NEXT, and if $E = K$, then if we lack the evidence that NEXT we will also lack the evidence that PAST. Thus, P must not include the evidence that NEXT, P must likewise not include the evidence that PAST. But, in the absence of the evidence PAST, we should to PAST be more probable conditional on NEXT than unconditionally. That is, we should expect that $P(\text{PAST} \mid \text{NEXT}) > P(\text{PAST})$. And so it follows that NEXT is evidence for PAST.

But these conclusions are highly counterintuitive, to say the least. Recall that in the case under consideration, we have observed that it has been hot in Death Valley for the past two hundred years, but we have not yet observed that it will be hot next July. Hence it seems that we have evidence for the claim that it will be hot next July, and for the claim that it will be hot every July, consisting in the fact that it has been hot for the past two hundred Julies. But surely we do not have evidence for the claim that it will be hot every July, consisting in the fact that it will be hot *next* July. And, equally surely, we do not have evidence for the claim that it has been hot for the past two hundred years, consisting in the fact that it will be hot next July.

Perhaps the problem lies, not with $E = K$, but rather with EV. Perhaps, in order for a proposition, p, to be evidence for another proposition, q, it does not suffice for p to belong to our evidence and for p to increase the probability of q. Perhaps we should say that, in addition, p must be known by observation. This would explain why the proposition that it will be hot in Death Valley next July can't constitute evidence for the other propositions considered above. But assuming something can be known without being known by observation, it follows that if we say that in order for a proposition to count as evidence *for* something, it must be known by observation, and if we also accept $E = K$, then we must conclude that there are propositions that belong to our evidence but that cannot be evidence *for* anything. And this is hardly a happy result.

I don't claim that this argument is entirely conclusive. It may be that some account can be given of the evidence-for relation that is compatible with the thesis that $E = K$, though this account would need to be very different from the one offered by Williamson, and from the other standard accounts of this relation found in the literature [add footnote—ask Brandon!]. More serious problems, however, lie in the way of reconciling $E = K$ with an adequate account of evidential probability.

1.3. The Truth Condition

What justifies our confidence in propositions is our evidence. And so it seems that how confident we should be in a given proposition is a function of our evidence. This is Williamson's view, for he understands the evidential probability of a hypothesis, *b*, as $P(b \mid e)$, where *e* is our total evidence and P measures the intrinsic plausibility of hypotheses prior to investigation (compare pp. 211 and 220). Hence, on this view, for any two agents who have the exact same body of evidence, they are justified in having the same degree of confidence in any given proposition.

According to Williamson's conception of evidential probability, for any given hypothesis and any given body of evidence, there is a unique (though possibly vague) level of confidence that an agent would be justified in having in this hypothesis on the basis of this body of evidence. One might reject this strict view, and instead hold a more permissive view according to which a given body of

evidence justifies not a single level of confidence in a given proposition, but rather any level of confidence within some range. But the strict view and the more permissive view both have the following implication:

WEAK EVIDENCE THESIS: If two agents differ in the levels of confidence they could justifiably have in a given proposition, this difference must be explainable in terms of differences in their bodies of evidence.

Since this thesis is very plausible, and since it follows from Williamson's stated view, I will assume over the next two sections.

This thesis creates problems for the claim that $E = K$. One of these problems derives from the fact that whatever is known must be true, and so knowledge is factive. It follows that if two agents both have systematically false beliefs, they may both know the same set of propositions, even though, intuitively, it would seem that they ought to have different levels of confidence in certain propositions. Suppose that Andrei and Nicolai are both brains in vats who have recently been invatted. Suppose it seems to Andrei that he is seeing snow falling in the Kremlin, and it seems to Nicolai that he is seeing rain falling on the Kremlin. Suppose that Andy and Handy have the very same a priori knowledge, e.g., they both know that $2 + 2 = 4$. And suppose that neither Andrei nor Nicolai reflects upon his own mental states, and so their only empirical beliefs concern the external world. None of these beliefs, however, constitute knowledge, since they are both systematically deceived about the external world. Thus, neither Andrei nor Nicolai has any a posteriori knowledge. And so the two agents have the very same body of knowledge. And so it follows from $E = K$ that the two agents have the same total evidence. Hence, it follows from the PERMISSIVE VIEW that the range of levels of confidence they could justifiably have in any given hypothesis will be the same. It seems, however, that Andrei could justifiably have a high level of confidence in the proposition that it is snowing in Moscow, while Nicolai could not. Therefore, assuming the WEAK EVIDENCE THESIS, we must reject $E = K$.

One might object to my claim that Andrei and Nicolai have the very same body of knowledge. After all, Andrei presumably knows that he is himself, and thus he knows that Andrei is Andrei. And since Nicolai has no knowledge of the external world, he does not know of Andrei's existence, and so he does not know that Andrei is Andrei. Similarly, Nicolai presumably knows that he is himself, and hence that Nicolai is Nicolai, while Andy does not know this. But even so, surely it is not Andrei's knowledge that Andrei is Andrei that justifies his confidence in the proposition that it is snowing in Moscow. Thus, even if we insist that there is a difference between Andrei's and Nicolai's bodies knowledge, there is no difference that could explain the difference between these justified levels of confidence. Hence, assuming the Weak Evidence Thesis, we must assume that while there is no relevant difference in their knowledge, there is a relevant difference in their evidence. And so we must reject $E = K$.

1.4. The Belief Condition

Knowledge is belief-dependent, in the sense that one can only know what one believes. Hence, if $E = K$, then one's evidence, and hence the evidential probabilities of hypotheses, will depend, in surprising ways, on what one actually believes.

Suppose Petra and Olga are not brains in vats but embodied human beings. Petra is watching the snow falling onto the Kremlin, while Olga is sunbathing on an island off the coast of Turkey, and has not been keeping track of the weather in Moscow. But Petra and Olga are both practicing skeptics. After years of self-discipline, they managed to suspend judgment on all questions concerning the external world, so that neither one has any beliefs about the world around her. And after more years of self-discipline, they managed to suspend judgment on all questions concerning their own mental states, so that neither one has any beliefs about how things appear to her. Thus, since knowledge requires belief, and neither one has any knowledge of the external world or of her own mental states. We may therefore assume that in the case of Petra and Olga, as in the case of Andrei and Nicolai, there is no difference between their bodies of knowledge, or at least no difference that could explain a difference in their justifiable levels of confidence in the proposition that it is snowing in Moscow. And yet it seems clear that Petra could justifiably have a high level of confidence in the proposition that it is snowing in Moscow, while Olga could not. Hence, assuming the Weak Evidence Thesis, we must assume that while there is no relevant difference in their knowledge, there is a relevant difference in their evidence. And so we must reject $E = K$.

We have considered a case in which there is a difference in justifiable levels of confidence but no relevant difference in knowledge. But assuming $E = K$, the belief condition for knowledge will also imply differences in relevant knowledge where there would appear to be no difference in justifiable levels of credence. Suppose there is a jar of marbles. Petra knows that the jar either contains ninety red marbles and ten white ones, or ten red marbles and ninety white ones. She then draws ten marbles from the jar, replacing each marble after drawing it. The first five marbles she draws are red, and the last five are white. Assume that Petra sees each of these marbles clearly. It would seem that in this case, given Petra ought to have equal confidence in the two alternative hypotheses concerning the composition of the jar. But if $E = K$, then Petra's evidence depends not only on what she sees, but also on what she believes. Hence, she can change the levels of confidence she would be justified in having in the two propositions at will, by selectively suspending judgement in the testimony of her senses. If she suspends judgment in the proposition that the first five marbles she drew were red, and that she seems to remember that the first five marbles were red, then she will retain only the knowledge that the last five were white. Thus, if we hold that $E = K$, we must hold that Petra's relevant evidence consists only in the fact that the last five marbles she drew were white. Hence, we must say that she would be justified in being very confident in the hypothesis that most of the marbles in the jar are white. But if she tires of being justifiably confident in this proposition, then she can switch to being justifiably confident in the alternative hypothesis, by ceasing to suspend judgment in the proposition that the first five marbles she drew were red, and instead suspending judgment in the proposition that the last five were white. Since this conclusion is unacceptable, we should reject the claim that $E = K$.

1.5. Conjunction

Consider the following principle:

CONJUNCTION: If we know, of each proposition in some finite set of propositions, that it is true, then we could justifiably be fully confident that the conjunction of these propositions is true.

For example, there are many propositions that I believe because I seem to remember that they are true, where my evidence for their truth is sufficient that if they were true, I would know count as knowing they are true. Thus, for example, I seem to remember eating three large eggs for breakfast yesterday, returning *Protagoras* to the library last month, going to my highschool prom with Olivia Zaple, etc. And since my memory for such propositions is fairly reliable, it would seem that, in the absence of any defeating conditions, I will know that each of these propositions is true, assuming it is true. Let S be the set of all the propositions of this kind. Suppose it contains 100,000 propositions. Suppose my memory is 99.99% reliable, and that all but ten of the propositions in S is true. Let S' be this subset of S containing all and only the propositions. Since, ex hypothesi, every proposition in S is such that, if it were true, I would count as knowing that it is true, it follows that I know, of each proposition in S' , that it is true. And so it follows from CONJUNCTION that I could justifiably be fully confident that the conjunction of all the propositions in S' is true. But this seems false. After all, since S contains 100,000 propositions and ten false propositions, there will be x subsets of S containing 99,990 propositions, and only one of these subset will contain only true propositions. Moreover, these are precisely the sorts of numbers I should expect, assuming I have reason to believe that my memory is around 99.99% reliable. And so it would seem that since I should expect the vast majority of these subsets to contain some false propositions, and since I have no way to discern which of these subsets contains only true propositions, I should suspect, of each of these subsets, that it contains some false propositions. Hence, in particular, for the subset that actually contains only true propositions, I should not be fully confident that this is so, and thus I should not be fully confident that the conjunction of the propositions it contains is true. And so it seems we should reject CONJUNCTION.

Hawthorne and XXXX discuss a case which appears to be another counterexample to conjunction. If we know that a marble is about to be dropped onto a concrete floor, then it seems we are in a position to know that its descent will be stopped by the floor. And so if, on the basis of our knowledge that the marble will be dropped onto the concrete floor, we form the belief that its descent will be stopped by the floor, and if this belief is true, then it seems we will know that the marble's descent will be stopped by the floor. And it seems that we would know this even if we know that there is a minuscule chance (say, one in a billion trillion) that the marble will pass right through the floor by a process of quantum tunneling. Moreover, it seems that we would know, of any given marble, that its descent will be stopped by the concrete floor onto which it is dropped, regardless of how many marbles there are that we know will be dropped onto concrete floors. Suppose therefore, that we know that a trillion trillion marbles will be dropped sequentially, each onto a concrete floor. Suppose that, for each of the marbles in this sequence, I form the belief that its descent will be stopped by the floor. And suppose that each of these beliefs is true. It would seem that we will know that the first marble's descent will be stopped by the floor, and we will also know that the second marble's descent will be stopped by the floor, and so on for every marble in the sequence. And so it follows from CONJUNCTION that I would be justified in being fully confident in the conjunction of all these propositions. It seems, however, that since I know that each marble has a one in a billion trillion chance of tunneling through the floor, I should expect that some of the trillion trillion marbles in the sequence will tunnel through the floor, and so I should

not be fully confident in the proposition that the descent of every marble will be stopped by the floor. And so, once again, it appears that CONJUNCTION is false.

But if we accept $E = K$, and if, following Williamson, we hold that the evidential probability of a proposition, or the level of confidence we can justifiably have in it, is equal to its probability conditional on our total evidence, then we are forced to accept CONJUNCTION. For if we know, of each proposition in some set, that it is true, then it follows from $E = K$ that every proposition in this set belongs to our total evidence. Thus, if the evidential probability of a proposition is its probability conditional on our evidence, then every proposition in this set will have an evidential probability of one. And assuming evidential probability is coherent, and thus obeys the principle of finite additivity, it follows that if every proposition in a finite set has an evidential probability of one, then the conjunction of the propositions in this set will likewise have an evidential probability of one. And so it follows from $E = K$ that the conjunction of all the propositions in a finite set of known propositions will have an evidential probability of one. In other words, $E = K$ implies CONJUNCTION.

This problem is compounded by the fact that knowledge depends on truth. For consider two alternative scenarios. In the first case, a trillion trillion trillion trillion marbles will be dropped in sequence, each onto a concrete floor, and, by a strange coincidence, the descent of each will be stopped by the floor. In the second case, a trillion trillion trillion trillion marbles will be dropped in sequence, each onto a concrete floor, and a trillion trillion of them will pass through the floor by a process of quantum tunneling. In the first case, it seems that I am in a position to know, of each of the marbles in the sequence, that it will be stopped by the floor, and so it follows from CONJUNCTION that I am in a position to justifiably be fully confident that all the marbles in the sequence are stopped by the floor. But in the second case, a trillion trillion of these marbles will tunnel through the floor. And so only the non-tunneling marbles are such that I know, of each, that it will be stopped by the floor. Hence, by $E = K$, my total evidence will only include the stopping of those marbles that, as a matter of fact, will be stopped by the floor. And, given this evidence, and given what I know about the objective chance of quantum tunneling, it is overwhelmingly likely that some of the remaining trillion trillion marbles will tunnel through the floor. And so the evidential probability of the proposition that every marble in the sequence will be stopped by the floor is close to zero. Thus, in the first scenario, the evidential probability of this proposition is one, while in the second scenario, the evidential probability of this proposition is close to zero. And yet the two scenarios differ only in the outcome of a chance process that will occur in the future. Thus, the defender of $E = K$ appears to be committed to the view that the current evidential probability of a proposition, or level of confidence we would now be justified in having in it, can depend in dramatic ways in the outcome of a chance process that will occur in the future. Hence, she is committed to the view that whether our evidence now justifies a certain level of confidence in a proposition can depend on what happens in the future.

Note: if we reject Conjunction, we must reject Closure.

p. 118: we should in any case be very reluctant to reject closure, for it is intuitive. If we reject it, in what circumstances can we gain knowledge by deduction? [paraphrase first sentence, and quote second]

answer: when the premises belong to our evidence

1.6. Easy Evidence and Easy Knowledge

Stewart Cohen describes the problem of easy knowledge thus:

Suppose I have reliable color vision. Then [assuming reliabilism about knowledge] I can come to know, e.g., that the table is red, even though I do not know that my color vision is reliable. But then I can note that my belief that the table is red was produced by my color vision. Combining this knowledge with my knowledge that the table is red, I can infer that in this instance, my color vision worked correctly. By repeating this process enough times, I would seem to be able to amass considerable evidence that my color vision is reliable, enough for me to come to know that my color vision is reliable.

Following Cohen, let us say that if a view implies that the process just described would yield *evidence* for the reliability of one's color vision, then this view suffers from the *problem of easy evidence*. And if a view implies that this process would yield *knowledge* that one's color vision is reliable, then this view suffers from the *problem of easy knowledge*. Cohen argues that in order to avoid these problems, we deny that a belief source can deliver knowledge prior to one's knowing that the source is reliable. On his view, unless we already know that our color vision is reliable, we cannot come to know the colors of objects using our color vision, and so we cannot bootstrap our way to knowledge of the reliability of our color vision.

Quite apart from the implausibility of Cohen's proposal, it fails to provide a general solution to the problems of easy knowledge and easy evidence, as Jonathan Weisberg has shown. Consider the following case:

Suppose I know that I have reliable color vision. I don't know, however, whether my color vision is color vision is super-reliable (roughly 100% reliable), or only highly reliable (roughly 99.9% reliable). I justifiably regard each of these possibilities as equally likely. I then can come to know that the table is red, and I note that my belief that the table is red was produced by my color vision. Combining this knowledge with my knowledge that the table is red, I can infer that in this instance, my color vision worked correctly. By repeating this process many thousands of times, I would seem to be able to amass considerable evidence that my color vision is super-reliable, enough for me to come to know that my color vision is super-reliable.

Note that in this case, I antecedently know that my color vision is reliable. And so we cannot solve the problems of easy evidence and easy knowledge that this case presents by denying that a belief source can deliver knowledge prior to one's knowing that the source is reliable. Another solution is necessary.

An obvious solution is to reject $E = K$. If we do so, then we can say that while I know, in the case described, that the table is red, the proposition that the table is red does not belong to my

evidence. Hence, when I observe the colors of the objects around us, although I may come to know that my color vision is producing reliable beliefs, I do not acquire the evidence that my color vision is producing reliable beliefs. Hence, I do not amass any evidence for the claim that my color vision is perfectly reliable. And I can come to be in a position to know a proposition only if I acquire evidence for this proposition. Thus, since I acquire no evidence for the proposition that my color vision is perfectly reliable, I don't come to be in a position to know this proposition.

Thus, it seems we can solve both the problem of easy evidence and the problem of easy knowledge by rejecting $E = K$, and by denying that, in coming to know that the table is red, I acquire the evidence that it is red, we can avoid both the problem of easy evidence and the problem of easy knowledge. And there is independent reason to deny the latter claim. For assuming evidential probability is probability conditional on my evidence, it follows that if a proposition belongs to my evidence, then its evidential probability, or the level of confidence I would be justified in having in it, must be one. But in the case described, it seems that I could not justifiably have full confidence in the claim that the table is red. For such a degree of confidence would be unstable. Recall that we are assuming that my confidence is divided equally between the hypothesis that my color vision is roughly 100% reliable, and the hypothesis that it is roughly 99.9% reliable. Thus, even if my color vision generates a state of full confidence in the proposition that the table is red, my level of confidence in the proposition that the table is red, conditional on my having become fully confident that it is red on the basis of color vision, should be $.5 * 1 + .5 * .999 = .9995$. Hence, when I learn that my full confidence in this proposition was produced by my color vision (as I must, if the problems of easy evidence and easy knowledge are to get started), then I should conditionalize on this information, and acquire a level of confidence of .9995 in the proposition that the table is red. Since full confidence is unstable, I cannot justifiably be fully confident in this proposition. Hence, its evidential probability must be less than one. And so this proposition must not belong to my total evidence.

If the proposition that the table is red does not belong to my total evidence, then what sort of proposition does my total evidence consist in? This is a question we take up later.

1.7. Corroboration

We have seen that $E = K$ implies that we acquire evidence for reliability in cases in which, intuitively, it seems we do not. But conversely, $E = K$ also implies that we do not acquire evidence for reliability in cases in which, intuitively, it seems we do. Suppose Macbeth has a job doing inventory in Duncan's armory. Duncan has a million numbered dagger holders, each containing a dagger. Macbeth knows that his vision is either super-reliable, or only highly reliable. He then looks at each of the dagger holders in turn, and in each case he sees that it contains a dagger. It would seem that he thereby comes to learn, of each dagger holder, that it contains a dagger. Nonetheless, he decides to make another round, this time touching the dagger in each holder with his hand. And again, for each holder, he finds that it contains a dagger. It would seem that every time Macbeth feels that a given holder contains a dagger, he corroborates the testimony of his vision, thereby confirming the reliability of his vision. And it would seem that as he feels more and more daggers in their holders, he acquires more and more evidence that his vision is not only highly reliable, but

super-reliable. But if we grant that he knows of each holder that it contains a dagger, and if we claim that $E = K$, then it's hard to see how his touching the daggers could provide any new evidence. For if he knows that a given holder contains a dagger, then presumably he knows that when he reaches out to touch the dagger in this holder, he will feel it. And so if $E = K$, then for each holder, upon seeing the dagger in the holder, he already possesses the evidence that when he reaches out to touch the dagger in the holder he will feel it. And so when he touches the dagger, this will provide him with no evidence he didn't already have. To allow that his touching the dagger provides him with evidence, we must deny that what he knows in advance belongs to his evidence in advance. And so we must deny $E = K$.

As another illustration, it would seem that if we know that a vast sequence of a trillion trillion marbles will be dropped, each onto a concrete floor, and if, as a matter of fact, the descent of each marble is stopped by the floor, then it seems that we are in a position to know in advance, of each marble in the sequence, that its descent is stopped by the concrete floor. Now suppose we know that one of two hypotheses is correct: H1, according to which the objective chance that a given marble will tunnel through a concrete floor is minute (one in a billion trillion), and H2, according to which this objective chance is zero. Suppose we watch the vast sequence of marbles being dropped, and we see the descent of each marble being stopped by the floor onto which it is dropped. It would seem that the more such events we observe, the more evidence we acquire for H2. But if we know in advance, of each marble, that its descent will be stopped by the floor, and if $E = K$, then, prior to observing any given marble being stopped by the floor, we already possess the evidence that the marble will be stopped by the floor. And so seeing the marble being stopped by the floor will provide us with no new evidence, and will thus not change the evidential probabilities of H1 and H2.

To solve this problem, while granting that we know in advance, of each marble, that its descent will be stopped by the floor, we must claim that this knowledge does not constitute evidence, and so we must deny $E = K$. And this seems like the natural thing to say. For, intuitively, it would seem that the only way to acquire the evidence that the descent of a given marble is stopped by the concrete floor onto which it is dropped is to observe this event, or to observe some consequence of this event. And yet, unless we are to embrace skepticism about the future, we should allow that we can know in advance that an event will occur in the future, and hence that we can know that an event occurs without observing the event or its consequences. Thus, in the case of the falling marbles, distinguishing between our prior knowledge and our prior evidence not only solves the problem of corroboration, it also has independent plausibility, as it accords with our ordinary ways of thinking about knowledge and evidence.

1.8. Margins of Error

Under ordinary circumstances, for any parameter that can vary in a continuous manner, our knowledge of this parameter will involve a margin of error. Thus, if a tower is exactly 65 feet tall, we will not know, by looking at it, that it is exactly 65 feet tall, and not 64.9 or 65.01 feet tall. Instead, we will know only that its height lies within some range (say, 65 to 70 feet) or on other words, that its height is 65 ± 5 feet. In this case, the margin of error for our knowledge of its height will be 5

feet. If we measure the tower, we will arrive at a more precise estimate of its height. Hence the range within which we know its height to lie will be smaller, and so, correspondingly, will be the margin of error for our best estimate of its height. The same goes for our knowledge of other parameters that vary continuously (mass, velocity, temperature, etc.). Indeed, even if the possible values of a parameter come in discrete units, our knowledge of this parameter will involve a margin of error, so long as the units in question are smaller than anything we can discern. Thus, under ordinary circumstances, our knowledge of an object's electric charge has a margin of error, even its electric charge is a function of the numbers of protons and electrons it contains, and thus does not vary continuously. In this section, I will argue that such margins of error create problems for the thesis that $E = K$.

Consider the following two cases:

Case 1: Tristan sees a tower that is 234 feet tall. Because his visual knowledge of height is not perfectly precise, what he knows about the tower's height on the basis of his visual perception is only that it is 234 ± 10 feet. But given the way the tower appears to him, his best estimate for the tower's height is 234 feet, and he thinks this height is more likely to be closer to 234 feet than further from 234 feet. Thus, he is more confident in the proposition that the tower is 234 ± 1 feet than in the proposition that it is 236 ± 1 feet tall. Tristan's friend Mary, who has measured the tower, knows that it is 234 ± 1 feet tall. But, rather than telling Tristan everything she knows, Mary simply tells him that the tower is 235 ± 2 feet tall. Since Tristan knows that Mary's testimony about the height of towers is reliable, he thereby comes to know that the tower is 235 ± 2 feet tall.

Case 2: Tristan sees a tower that is 236 feet tall. But what he knows about the height of the tower on the basis of his visual perception is only that it is 236 ± 10 feet. But given the way the tower appears to him, his best estimate for the tower's height is 236 feet, and so he is more confident in the proposition that the tower is 236 ± 1 feet than in the proposition that it is 234 ± 1 feet tall. Tristan's friend Mary knows that the tower is 236 ± 1 feet tall. But rather than telling Tristan everything she knows, she simply tells him that the tower is 235 ± 2 feet tall. Since Tristan knows that Mary's testimony is reliable, he thereby comes to know that the tower is 235 ± 2 feet tall.

In order to analyze these cases, let us define three propositions as follows:

p: the tower is 235 ± 2 feet tall

q: the tower is 234 ± 1 feet tall

r: the tower is 236 ± 1 feet tall

In Case 1, Tristan begins with more confidence in q than in r, while in Case 2, he begins with more confidence in r than in q. In both cases, he subsequently comes to know p. And so assuming that $E = K$, he acquires the evidence that p, and should thus conditionalize on this evidence. And since q and r each entail p, when he conditionalizes on p, this will not affect the ratio between his levels of confidence in q and r. And so, even after learning p, Tristan will continue to have more confidence in q than in r in Case 1, and he will continue to have more confidence in r than in q in Case 2. Thus,

there will remain a difference between his justified levels of confidence in the two cases. Assuming that differences in justified levels of confidence are to be explained in terms of differences in evidence, it follows that there must be a difference in Tristan's evidence between the two cases. And so if $E = K$, there must be a difference in his knowledge between the two cases. But what could this difference in knowledge be? For while there was initially such a difference in knowledge between the two cases (in Case 1, Tristan knew only that the tower was 234 ± 10 feet tall, while in Case 2, he knew only that it was 236 ± 10 feet tall), in both cases he subsequently learns p , which is stronger than what he initially knew about the height of the tower in either case. And so, after learning p , there will no longer be any difference in what he knows about the height of the tower between the two cases.

The defender of $E = K$ might reply as follows.

Since what Tristan knows on the basis of visual perception differs between the two cases, there must be a difference in the way the tower appears to Tristan between the two cases. And in both cases, Tristan will presumably know how the tower appears to him. Hence, his knowledge of the appearance of the tower will differ between the two cases. And this difference in his knowledge of appearance can explain why, even after learning that p , there is a difference in his justified levels of confidence between the two cases.

According to this response, since the difference between Tristan's justified levels of confidence is based in a difference in the appearance of the tower between the two cases, we can explain the former difference in terms of Tristan's knowledge of that latter difference. Unfortunately, this explanatory strategy will not always succeed. The problem is that our knowledge of the appearances of objects, like our knowledge of their objective properties, involves a margin of error, and this margin of error give rise to difficulties similar to the one just considered.

The reason that our knowledge of appearances involves margin of error is that appearances can vary in a continuous, or nearly continuous, manner. Suppose there is a 10 foot tall tower on the West side of a wall. At noon, the sun is directly overhead, and so the tower casts no shadow on the wall. But at 6pm, as the sun is setting, the tower casts a 10 foot tall shadow on the wall. Suppose that between noon and 6pm, Tristan is watching this shadow. In this case, it would seem that during this six hour interval, just as the height of the shadow will vary in a continuous manner, the appearance of this shadow to Tristan will likewise vary in a continuous, or nearly continuous, manner. And if this appearance varies in a continuous, or nearly continuous manner, then presumably so should Tristan's levels of confidence in propositions about the height of the shadow.

In general, a shadow will not appear to have any particular precise height. A shadow will not appear to be exactly six feet tall, but rather to be *around* six feet tall. Nonetheless, for some heights, the shadow will appear to be at least as tall as this height, and for some heights, the shadow will appear to be no taller than this height. Let us define the *apparent minimum height* of the shadow as the greatest height such that the shadow appears to be at least of this height. And let us define the *apparent maximum height* of the shadow as the smallest height such that the shadow appears to be no greater than this height. Let u be the centered proposition that the apparent minimum height of the shadow (for me now) is at least 9 feet. Suppose u becomes true for Tristan at exactly 5 pm. Will

Tristan know, at 5pm, that u is true? We can employ a Williamsonian argument to show that the answer is no.¹ For if Tristan knows that u is true at 5pm, then at 5pm he must have a very high level of confidence in u . Moreover, this confidence must be reliable, in the sense that he is not disposed to have a high level of confidence in u , on similar grounds, in a similar circumstance in which u is false. But if Tristan's confidence in u varies in a continuous manner, than at times before 5pm that are sufficiently close to 5pm, he will have a high level of confidence in u , on similar grounds, and yet, ex hypothesi, u is false. It follows that if Tristan's confidence in u at 5pm is unreliable, and hence that he does not know that u is true. Thus, Tristan will not know, at 5pm, that the apparent minimum height of the shadow is 5 feet. At most, he will know that the apparent minimal height of the shadow is within some range. And so his knowledge of the apparent minimum height of the shadow will have a margin of error. A similar argument shows that his knowledge of the apparent maximum height of the shadow likewise has a margin of error.

Let us return to considering pairs of cases in which Tristan is looking at a very tall tower. For simplicity, let us assume that the apparent height of the tower, for Tristan, can be specified in terms of the tower's apparent minimum and maximum heights (even if other parameters were involved, Tristan's knowledge of these parameters would likewise involve a margin of error, and so including them would not affect the argument in any essential way). Let us define the *apparent medial height* of the tower as the average of its apparent minimal height and its apparent maximal height. Again for simplicity, let us assume that Tristan knows that the tower's apparent maximal height is ten feet greater than it's apparent minimal height, and hence that its apparent medial height is five feet greater than its apparent minimal height. Thus, we can specify the tower's apparent height by giving its apparent medial height, and we can specify Tristan's knowledge of the tower's apparent height by giving the range within which he knows the apparent medial height to lie. For example, if the apparent medial height of the tower is 235 feet, we can infer that it's apparent minimum height, for Tristan, is 230 feet, and that its apparent maximum height is 240 feet. And if Tristan knows that the apparent medial height of the tower is 235 ± 3 feet, then he knows that the apparent minimum height is 230 ± 3 feet, and that its apparent maximum height is 240 ± 3 feet.

Now let us define four propositions, as follows:

- a: the apparent medial height of the tower is 235 ± 2 feet
- b: the actual height of the tower is 235 ± 2 feet
- c: the actual height of the tower is 234 ± 1 inch
- d: the actual height of the tower is 236 ± 1 inch.

Now consider the following pair of cases.

Case 3: Tristan sees a tower that is 234 feet tall. The apparent medial height of this tower, for Tristan, is likewise 234 feet. But his visual knowledge of heights is not perfectly precise, and so what he knows about the actual height of the tower is only that it is 234 ± 10 feet. Similarly, his introspective knowledge of apparent heights is not perfectly precise, and so what he knows about the apparent medial height of tower is only that it is 234 ± 3 feet. Given the way the tower appears to him, Tristan's best estimate for the tower's actual

height is 234 feet, and so he thinks this height is more likely to be closer to 234 feet than further from 234 feet. Hence, he is more confident in the conjunction (a & c) than in the conjunction (a & d). Tristan's friend Mary is a super-scientist, who has not only measured the height of the tower, but has also scanned Tristan's brain. Thus, Mary has acquired more precise knowledge than Tristan both of the tower's actual height and of how the tower appears to Tristan. And so Mary knows that the tower's actual height, and the tower's apparent medial height for Tristan, are both 234 ± 1 feet. But rather than telling Tristan everything she knows, Mary simply tells him that the actual and apparent height of the tower are both 235 ± 2 feet. That is, she tells him the conjunction (a & b). Since Tristan knows that Mary's testimony is reliable, he thereby comes to know this conjunction.

Case 4: Tristan sees a tower that is 236 feet tall. And the apparent medial height of this tower is likewise 236 feet. But given the imprecision of Tristan's visual and introspective knowledge, he knows only that the tower's actual height is 236 ± 10 feet, and that its apparent medial height is 236 ± 3 feet. Given the way the tower appears to him, Tristan's best estimate for the tower's actual height is 236 feet, and so he thinks this height is more likely to be closer to 236 feet than further from 236 feet. Hence, he is more confident in the conjunction (a & d) than in the conjunction (a & c). However, Mary the super-scientist knows that the tower's actual height, and its apparent medial height for Tristan, are both 236 ± 1 . But rather than telling Tristan everything she knows, she simply tells him the conjunction (a & b). Since Tristan knows that Mary's testimony is reliable, she thereby comes to know this conjunction.

In Case 3, Tristan begins with more confidence in (a & c) than in (a & d), while in Case 2, he begins with more confidence in (a & d) than in (a & c). In both cases, he subsequently learns (a & b). Since (a & c) and (a & d) each entail (a & b), when he conditionalizes on (a & b), this will not affect the ratio between his levels of confidence in (a & c) and (a & d). And so, even after learning (a & b), Tristan will continue to have more confidence in (a & c) than in (a & d) in Case 1, and he will continue to have more confidence in (a & d) than in (a & c) in case 2. Assuming that differences in justified levels of confidence are to be explained in terms of differences in evidence, and assuming $E = K$, there must be a difference in his knowledge between the two cases. But what could this difference in knowledge be? For what he learns in both cases, namely (a & b), is stronger than what he initially knew in either case. And so, upon learning (a & b), there will be no difference between the two cases either in his knowledge of the actual height of the tower, or in her knowledge of the apparent height of the tower.

The defender of $E = K$ might appeal to additional resources in order to explain the difference in Tristan's justified levels of confidence in the above cases. Williamson writes:

If perceptual evidence consists in propositions, which propositions are they? Consider an example. I am trying to identify a mountain by its shape. . . No description of the mountain in words seems to capture the richness of my perceptual experience of its irregular shape. But it does not follow that my evidence is non-propositional. If I want to convey my evidence, I might point and say "it is that shape" (pp. 197-198).

And so perhaps, if Tristan wanted to convey his evidence in Case 3, he could point to the top of the 234 inch tall tower and say “the tower is *that* tall.” Similarly, if he wanted to convey his evidence in Case 4, he could point to the top of the 236 inch tall tower and say “the tower is *that* tall.” Since the tokens of the word “that” would refer to different heights in the two cases, Tristan’s utterance would express different propositions in the two cases. And so if his utterances accurately conveyed his evidence, it seems he would have different evidence in the two cases. And one might suppose that this difference in evidence could explain the difference in justified levels of confidence in the two cases. Alternately, the defender of $E = K$ might claim that Tristan has indexical knowledge of how the tower appears to him that he can express by saying “the tower appears to me in *this* manner.” And since this utterance will express a different proposition in the two cases, this will create a difference in Tristan’s knowledge that can explain the difference in his justified levels of confidence.

Now on Williamson’s view, this sort of de re knowledge, like our ordinary descriptive knowledge, has a margin of error. He says

[I]n ordinary circumstances I can know that the mountain is that shape ... when ‘that shape’ does not refer to an absolutely specific shape. Of course, I cannot see *exactly* what shape the mountain is ... That shape must be unspecific enough to give my knowledge that the mountain is that shape an adequate margin for error...

If this is right, then presumably we don’t know, in looking at a tower, exactly how tall a tower is, and so we can know that a tower is *that* tall only if “that tall” does not refer to an absolutely specific height. Similarly, given the imprecision of our introspective knowledge, we presumably do not know exactly how a tower appears to us, and so we will know that a tower appears to us in *this manner* only if ‘this manner’ does not refer to an absolutely specific manner of appearing. But if our de re knowledge of heights and appearances involves a margin of error, then it will not solve our problem. For, as I will illustrate shortly, it will still be possible for people to begin with different indexical knowledge, and then both learn something that is stronger than what either of them initially knew, thereby eliminating the difference in their knowledge, without eliminating the differences between their justified levels of confidence.

But the defender of $E = K$ might reject the view that such de re knowledge involves a margin of error. For the reasons for thinking that there is a margin of error to our ordinary descriptive knowledge of continuously variable parameters do not carry over, in any obvious way, to our de re knowledge of such parameters. If I look at a tower and, on the basis of how it appears, say “this tower is exactly 234 feet tall,” then even if what I say is true, it will not constitute knowledge, for I could easily have been wrong. If I am disposed to have this belief in my actual circumstances, then I could easily have this belief in a similar circumstance in which it is false, because the tower is only 233.99 feet tall. Hence, absolutely precise descriptive claims about the height of a tower will not constitute knowledge. But suppose I look at a tower and say “the tower is exactly *that* tall,” referring to the exact height of the tower. And suppose I am fully confident in the proposition I thereby express. In this case, it seems my belief could not easily have been wrong. For the proposition in question is not one that I would believe in a circumstance in which it is false. For suppose that

when I confidently said “the tower is exactly *that* tall,” the tower to which I was pointing were only 233.99 feet tall. In this case, I would be referring to a different height, and thus expressing a different proposition. And the proposition I would be true, and so my belief would be veridical.

It seems, therefore, that one might plausibly claim that our de re knowledge of heights and of appearances has a margin of error. And so let us assume, for the sake of argument, that this is so. Even granting this assumption, such de re knowledge will not help the defender of $E = K$, for it cannot adequately explain differences in our justified levels of confidence in ordinary descriptive propositions. To show this, let us assume, for the sake of argument, that we have an extremely strong form of de re knowledge. Let us define an agent’s epistemic situation as the conjunction of all the agent’s properties that are relevant to what she is justified in believing. If what one is justified in believing depends on one’s relations to external states of affairs, then one’s epistemic situation will include these relations. And let us assume, for the sake of argument, that if an agent, *s*, is in epistemic situation, *E*, then *s* has the de re knowledge, of *E*, that she is in it. She can convey this knowledge by saying “I am in this epistemic situation,” where “this epistemic situation” refers to an absolutely specific epistemic situation. One might suppose that if we grant this very strong assumption about our de re knowledge, then our de re knowledge could play any explanatory that could be played by anything else. Suppose, for example, than one could explain an agent’s justified levels of confidence by appealing to her experiential states. Let *E* denote this agent’s epistemic situation. Since her experiential state will be a component of her epistemic situation, only someone in her experiential state can be in her epistemic situation. And so if she knows that she is in *E*, then she knows something she could only know if she were in this experiential state. One might think, therefore, that her knowledge that she is in *E* could explain everything that could be explained by her experiential state. I will argue, however, that this is not so. For there are pairs of cases in which agents have different levels of justified confidence that are grounded in different experiential states, and where this different levels of justified confidence cannot be explained in terms of the de re knowledge of their epistemic situation, even granting that they have such knowledge.

Let us consider a case involving two identical twins, Beth and Liz. Assume that neither Beth nor Liz is aware of the existence of the other. Beth knows that “Beth” refers to her, but she doesn’t know whether “Liz” refers to her or to someone else. Similarly, Liz knows that “Liz” refers to her, but she doesn’t know whether “Beth” refers to her or to someone else. Let us define four centered propositions, as follows (these are the centered analogues of propositions a through d, defined above):

a’: the apparent medial height of the tower I am looking at is 235 ± 2 feet

b’: the actual height of the tower I am looking at is 235 ± 2 feet

c’: the actual height of the tower I am looking at is 234 ± 1 inch

d’: the actual height of the tower I am looking at is 236 ± 1 inch.

And let us define a further, uncentered proposition, as follows:

e: Beth and Liz are in the same epistemic situation.

Now consider the following case, Case 5, which involves three consecutive stages.

Stage 1: Beth and Liz are each looking at a tower, and neither one can see the other, or the tower the other one is looking at. Beth's tower is 234 feet tall, and its apparent medial height is likewise 234 feet. Because of the imprecision of her visual and introspective knowledge, her only descriptive knowledge of the tower's height is that it is 234 ± 10 feet, and her only descriptive knowledge of the tower's apparent medial height is that it is 234 ± 3 feet. Given the way the tower appears to her, Beth's best estimate for the tower's actual height is 234 feet, and so she thinks this height is more likely to be closer to 234 feet than further from 234 feet. Hence, she is justifiably more confident in the conjunction (a' & c') than in the conjunction (a' & d'). And this remains true even when she conditionalizes on the proposition that Beth and Liz are in the same epistemic situation. That is, Beth's justified conditional level of confidence in (a' & c') given e exceeds her justified conditional level of confidence in (a' & d') given e. Liz, on the other hand, is looking at a tower that is 236 feet tall, and its apparent medial height is likewise 236 feet. But her only descriptive knowledge of the tower's height is that it is 236 ± 10 feet, and her only descriptive knowledge of the tower's apparent medial height is that it is 236 ± 3 feet. Given the way the tower appears to her, Liz's best estimate for the tower's actual height is 236 feet, and so she is justifiably more confident in the conjunction (a' & d') than in the conjunction (a' & c'). And this remains true even when she conditionalizes on e. That is, her justified conditional level of confidence in (a' & d') given e exceeds her justified conditional level of confidence in (a' & c') given e. However, Beth and Liz both have precise de re knowledge of their epistemic situation. Thus, where EB and EL denote Beth's and Liz's epistemic situations, respectively, Beth knows, of EB, that she is in it, and Liz knows, of EL, that she is in it.

Stage 2: Mary tells Beth that the tower she is looking at is 235 ± 2 feet tall, and that its apparent medial height is likewise 235 ± 2 feet. And she also tells Liz that the tower she is looking at is 235 ± 2 feet tall, and that its apparent medial height is likewise 235 ± 2 feet. Beth and Liz both know that Mary is reliable, and so they both come to know the centered proposition (a' & b'). Since this conjunction is stronger than the relevant descriptive knowledge either of them initially possessed, there remains no relevant difference in their descriptive knowledge. But there remains a difference in their de re knowledge.

Stage 3: Mary gives Beth a microphone and a pair of headphones, and Mary gives the same to Liz. Mary explains to Beth that, through her headphones, she will hear the voice of Liz. And Mary explains to Liz that, through her headphones, she will hear the voice of Beth. However, since Beth doesn't know whether she is Liz, she doesn't know whether she will hear her own voice or someone else's, and since Liz doesn't know whether she is Beth, she likewise doesn't know whether she will hear her own voice or someone else's. They each know, however, that whoever they hear has precise de re knowledge of her own epistemic situation. Then, simultaneously, Beth and Liz utter the words "I am in this epistemic situation," each one referring to the epistemic situation she is in. And so as Beth is uttering these words, through her headphones she hears Liz uttering these words, and as Liz utters these words, through her headphones she hears Beth uttering these words. But since they utter these words simultaneously in the exactly the same manner, they both continue to be

uncertain as to whether they are hearing their own voice or someone else's. But when Beth hears Liz say "I am in this epistemic situation," Beth comes to have knowledge that she can convey by saying "Liz is in that epistemic situation," where "that epistemic situation" refers anaphorically to the epistemic situation referred to by Liz. Hence, Beth acquires the de re knowledge that Liz is in EL. Similarly, when Liz hears Beth say "I am in this epistemic situation," Liz comes to have knowledge that she can convey by saying "Beth is in that epistemic situation," where "that epistemic situation" refers anaphorically to the epistemic situation referred to by Beth. Hence, Liz acquires the de re knowledge that Beth is in EB.

Now if our de re knowledge of our epistemic situation were not perfectly precise, Beth could know that she is in EB, and Liz could know that she is in EL, only if EB and EL were not absolutely specific epistemic situations. And in this case, they could presumably both be in epistemic situations EB *and* EL, in spite of the slight difference in their epistemic situations. And so, with the help of Mary the super-scientist, each one could come to know that she is in epistemic situations EB and EL. In this way, the difference between their de re knowledge of their epistemic situations would be eliminated, without eliminating the difference between their justified levels of confidence in propositions concerning the height of the tower. And so we would have a counterexample to $E = K$ similar to the one presented by cases 3 and 4. However, we are now assuming, for the sake of argument, that our de re knowledge of our epistemic situation is absolutely precise. And, given this assumption, so long as the epistemic situations of Beth and Liz differ, there will be an ineliminable difference in their de re knowledge of their epistemic situations.

But there remains a problem. For at Stage 3, the only relevant knowledge Beth possess and Liz lacks is the centered, de re knowledge, of EB, that she is in it. Call the centered proposition that is the content of this knowledge iB . And the only relevant knowledge Liz possesses that Beth lacks is the centered, de re knowledge, of EL, that she is in it. Call the centered proposition that is the content of this knowledge iL . And so it follows from $E = K$ that, in Stage 3, the only relevant evidence Beth possesses that Liz lacks is iB , and the only relevant evidence Liz possesses that Beth lacks is iL . And while iL is not metaphysically compatible with the rest of what Beth knows (since it is metaphysically impossible that she is both EB and EL), iL is *epistemically* compatible with the rest of what she knows, in the sense that she is not in a position to rule it out. Because for all Beth knows, she is Liz. Hence, Beth is not in a position to rule out the possibility that she is Liz. And Beth knows, of EL, that Beth is in it. And so Beth is not in position to rule out the possibility that she is in EL. That is, she cannot rule out iL . And so she could coherently conditionalize on iL . Similarly, Liz is not in a position to rule out the possibility that she is in EB, and so she could coherently conditionalize on iB .

But if justified levels of confidence are a function of one's total evidence, then Beth's and Liz's justified levels of confidence in any centered or uncentered proposition should converge if they were to conditionalize on all the relevant evidence centered or uncentered evidence possessed by either. Thus, for any centered or uncentered proposition, p , if Beth were to conditionalize on every proposition that is relevant to p that belongs to Liz's evidence but does not belong to her own, and Liz were likewise to conditionalize on every relevant proposition that belongs to Beth's evidence but not to her own, then there would remain no difference in their justified levels of confidence in p .

Consequently, if Beth were to conditionalize on iL , and Liz were to conditionalize on iB , then there would remain no difference in their justified levels of confidence in $(a' \& c')$ or in $(a' \& d')$. And so Beth's justified levels of confidence in $(a' \& c')$ and in $(a' \& d')$ conditional on iL must be equal to Liz's justified levels of confidence in these centered propositions conditional on iB .

Moreover, Beth knows, of EL, that Liz is in it. And so Beth knows, of EL, that she is in it if and only if she is in the same epistemic situation as Liz. And since Beth knows, de dicto, that she is Beth, she knows, of EL, that she is in it if and only if Beth and Liz are in the same epistemic situation. In other words, she knows that iL is true if and only if e is true. Consequently, assuming $E = K$, for any centered or uncentered proposition, p , Beth's justified level of confidence in p conditional on e must be equal to her justified level of confidence in p conditional in iL . By similar reasoning, it follows that for any centered or uncentered proposition, p , Liz's justified level of confidence in p conditional on e must be equal to her justified level of confidence in p conditional on iB . And since we have seen that Beth's justified levels of confidence in $(a' \& c')$ and in $(a' \& d')$ conditional on iL must be equal to Liz's justified levels of confidence in these centered propositions conditional on iB , it follows that Beth's justified levels of confidence in these propositions conditional on e must be equal to Liz's.

But this conclusion is unacceptable. For, ex hypothesi, at Stage 1, Beth should be more confident, conditional on e , in $(a' \& c')$ than in $(a' \& d')$. And for Liz, the reverse should be true. But the only change that occurs between Stage 1 and Stage 2 is that Beth and Liz both learn a centered proposition, $(a' \& b')$, that is weaker than either $(a' \& c')$ or $(a' \& d')$. And so, for both Beth and Liz, this change should not affect the ratio between their justifiable levels of confidence in $(a' \& c')$ and $(a' \& d')$, either unconditionally or conditionally on e . Thus, at Stage 2, it should remain the case that Beth should be more confident, conditional on e , in $(a' \& c')$ than in $(a' \& d')$, and that the reverse should be true for Liz. And all that happens between Stage 2 and Stage 3 is that Beth hears Liz say "I am in this epistemic situation" and Liz hears Beth utter the same words. And it is hardly plausible that this change should affect either Beth's or Liz's conception of the height of the tower she is looking at. Nor is it plausible that this change should affect either Beth's or Liz's conception of the height of the tower she is looking at, conditional on the supposition that Beth and Liz are in the same epistemic situation. And so, at stage 3, it must remain the case that Beth should be more confident, conditional on e , in $(a' \& c')$ than in $(a' \& d')$, and that the reverse should be true for Liz. But on the basis of $E = K$, we derived the conclusion that, conditional on e , Beth's justified levels of confidence $(a' \& c')$ and in $(a' \& d')$ are the same as Liz's. And so we must reject $E = K$.

As this case illustrates, differences in experiential state can give rise to differences in justifiable levels of confidence that cannot be explained if we assume $E = K$, even if we grant the assumption that we have de re knowledge of our precise epistemic situation.

¹ See Williamson 2000, section 4.3.