

Margins and Errors

Brian Weatherson

Recently, Timothy Williamson (2013) has argued that considerations about margins of errors can generate a new class of cases where agents have justified true beliefs without knowledge. I think this is a great argument, and it has a number of interesting philosophical conclusions. In this note I'm going to go over the assumptions of Williamson's argument. I'm going to argue that the assumptions which generate the justification without knowledge are true. I'm then going to go over some of the recent arguments in epistemology that are refuted by Williamson's work. And I'm going to end with an admittedly inconclusive discussion of what we can know when using an imperfect measuring device.

1 Measurement, Justification and Knowledge

Williamson's core example involves detecting the angle of a pointer on a wheel by eyesight. For various reasons, I find it easier to think about a slightly different example: measuring a quantity using a digital measurement device. This change has some costs relative to Williamson's version – for one thing, if we are measuring a quantity it might seem that the margin of error is related to the quantity measured. If I eyeball how many stories tall a building is, my margin of error is 0 if the building is 1-2 stories tall, and over 10 if the building is as tall as the World Trade Center. But this problem is not as pressing for digital devices, which are often very *unreliable* for small quantities. And, at least relative to my preferences, the familiarity of quantities makes up for the loss of symmetry properties involved in angular measurement.

To make things explicit, I'll imagine the agent S is using a digital scale. The scale has a **margin of error** m . That means that if the reading, i.e., the **apparent mass** is a , then the agent is justified in believing that the mass is in $[a - m, a + m]$. We will assume that a and m are luminous; i.e., the agent knows their values, and knows she knows them, and so on. This is a relatively harmless idealisation for a ; it is pretty clear what a digital scale reads.¹ It is a somewhat less plausible assumption for m . But we'll assume that S has been very diligent about calibrating her scale, and that the calibration has been recently and skillfully carried out, so in practice m can be assessed very accurately.

We'll make three further assumptions about m that strike me as plausible, but which may I guess be challenged. I need to be a bit careful with terminology to set out the first one. I'll use V and v as variables that both pick out the **true value** of the mass. The difference is that v picks it out rigidly, while V picks out the value of

[†] Unpublished. These are some reflections on a paper Timothy Williamson gave at the 2012 CSMN/Arché epistemology conference. Thanks to Derek Ball, Herman Cappelen, Ishani Maitra, Sarah Moss and Robert Weatherson for helpful discussions, as well as audiences at Arché and Edinburgh.

¹This isn't always true. If a scale flickers between reading 832g and 833g, it takes a bit of skill to determine what *the reading* is. But we'll assume it is clear in this case. On an analogue scale, the luminosity assumption is rather implausible, since it is possible to eyeball with less than perfect accuracy how far between one marker and the next the pointer is.

the mass in any world under consideration. Think of V as shorthand for *the mass of the object* and v as shorthand for *the actual mass of the object*. (More carefully, V is a random variable, while v is a standard, rigid, variable.) Our first assumption then is that m is also related to what the agent can know. In particular, we'll assume that if the reading a equals v , then the agent can know that $V \in [a - m, a + m]$, and can't know anything stronger than that. That is, the margin of error for justification equals, in the best case, the margin of error for knowledge. The second is that the scale has a readout that is finer than m . This is usually the case; the last digit on a digital scale is often not significant. The final assumption is that it is metaphysically possible that the scale has an error on an occasion that is greater than m . This is a kind of fallibilism assumption – saying that the margin of error is m does not mean there is anything incoherent about talking about cases where the error on an occasion is greater than m .

This error term will do a lot of work in what follows, so I'll use e to be the **error** of the measurement, i.e., $|a - v|$. For ease of exposition, I'll assume that $a \geq v$, i.e., that any error is on the high side. But this is entirely dispensible, and just lets me drop some disjunctions later on.

Now we are in a position to state Williamson's argument. Assume that on a particular occasion, $0 < e < m$. Perhaps $v = 830$, $m = 10$ and $a = 832$, so $e = 2$. Williamson appears to make the following two assumptions.²

1. The agent is justified in believing what they would know if appearances matched reality, i.e., if V equalled a .
2. The agent cannot come to know something about V on the basis of a suboptimal measurement that they could not also know on the basis of an optimal measurement.

I'm assuming here that the optimal measurement displays the correct mass. I don't assume the actual measurement is *wrong*. That would require saying something implausible about the semantic content of the display. It's not obvious that the display has a content that could be true or false, and if it does have such a content it might be true. (For instance, the content might be that the object on the scale has a mass near to a , or that with a high probability it has a mass near to a , and both of those things are true.) But the optimal measurement would be to have $a = v$, and in this sense the measurement is suboptimal.

The argument then is pretty quick. From the first assumption, we get that the agent is justified in believing that $V \in [a - m, a + m]$. Assume then that the agent forms this justified belief. This belief is incompatible with $V \in [v - m, a - m]$. But if a equalled v , then the agent wouldn't be in a position to rule out that $V \in [v - m, a - m]$. So by premise 2 she can't knowledgeably rule it out on the basis of a mismeasurement. So her belief that $V \geq a - m$ cannot be knowledge. So this justified true belief is not knowledge.

²I'm not actually sure whether Williamson *makes* the first, or thinks it is the kind of thing anyone who thinks justification is prior to knowledge should make.

If you prefer doing this with numbers, here's the way the example works using the numbers above. The mass of the object is 830. So if the reading was correct, the agent would know just that the mass is between 820 and 840. The reading is 832. So she's justified in believing, and we'll assume she does believe, that the mass is between 822 and 842. That belief is incompatible with the mass being 821. But by premise 2 she can't know the mass is greater than 821. So the belief doesn't amount to knowledge, despite being justified and, crucially, true. After all, 830 is between 822 and 842, so her belief that the mass is in this range is true. So simple reflections on the workings on measuring devices let us generate cases of justified true beliefs that are not knowledge.

I'll end this section with a couple of objections and replies.

Objection: The argument that the agent can't know that $V \in [a - m, a + m]$ is also an argument that the argument can't justifiably believe that $V \in [a - m, a + m]$. After all, why should it be possible to get justification from a suboptimal measurement when it isn't possible to get the same justification from an optimal measurement?

Reply: It is possible to have justification to believe an outright falsehood. It is widely believed that you can have justification even when none of your evidential sources are even approximately accurate (Cohen, 1984). And even most reliabilists will say that you can have false justified beliefs if you use a belief forming method that is normally reliable, but which badly misfires on this occasion. In such cases we clearly get justification to believe something from a mismeasurement that we wouldn't get from a correct measurement. So the objection is based on a mistaken view of justification.

Objection: Premise 2 fails in cases using random sampling. Here's an illustration. An experimenter wants to know what percentage of F s are G . She designs a survey to ask people whether they are G . The survey is well designed; everyone gives the correct answer about themselves. And she designs a process for randomly sampling the F s to get a good random selection of 500. It's an excellent process; every F had an equal chance of being selected, and the sample fairly represents the different demographically significant subgroups of the F s. But by the normal processes of random variation, her group contains slightly more G s than the average. In her survey, 28% of people said (truly!) that they were G , while only 26% of F s are G s. Assuming a margin of error in such a study of 4%, it seems plausible to say she knows that between 25 and 32% of F s are G s. But that's not something she could have known the survey had come back correctly reporting that 26% of F s are G s.

Reply: I think the core problem with this argument comes in the last sentence. A random survey isn't, in the first instance, a measurement of a population. It's a measurement of those surveyed, from which we draw extrapolations about the population. In that sense, the only *measurement* in the imagined example was as good as it could be; 28% of surveyed people are in fact G . So the survey was correct, and it is fine to conclude that we can in fact know that between 24 and 32 percent of F s are G s.

There are independent reasons for thinking this is the right way to talk about the case. If a genuine measuring device, like a scale, is off by a small amount, we regard that as a reason for tinkering with the device, and trying to make it more accurate. That's one respect in which the measurement is suboptimal, even if it is correct within the margin of error. This reason to tinker with the scale is a reason that often will be outweighed. Perhaps it is technologically infeasible to make the machine more accurate. More commonly, the only way to guarantee greater accuracy would be more cost and hassle than it is worth. But it remains a reason. The fact that this experiment came out with a deviation between the sample and the population is *not* a reason to think that it could have been run in a better way, or that there is some reason to improve the survey. That's just how random sampling goes. If it were a genuine measurement of the population, the deviation between the 'measurement' and what is being measured would be a reason to do things differently. There isn't any such reason, so the sample is not truly a measurement.

So I don't think this objection works, and I think the general principle that you can't get extra knowledge from a suboptimal measurement is right. But note also that we don't need this general principle to suggest that there will be cases of justified true belief without knowledge in the cases of measurement. Consider a special case where e is just less than m . For concreteness, say $a = v + 0.95m$, so $e = 0.95m$. Now assume that whatever is justifiedly truly believed in this case is known, so S knows that $V \in [a - m, a + m]$. That is, S knows that $V \in [v - 0.05m, a + m]$.

We don't need any principles about measurement to show this is false; safety considerations will suffice. Williamson (2000) says that a belief that p is safe only if p is true in all nearby worlds. But given how close v is to the edge of the range $[v - 0.05m, a + m]$. Rival conceptions of safety don't help much more than this. The most prominent of these, suggested by Sainsbury (1995), says that a belief is safe only if the method that produced it doesn't produce a false belief in any nearby world. But if the scale was off by $0.95m$, it could have been off by $1.05m$, so that condition fails too.

I don't want the last two paragraphs to leave too concessive an impression. I think the objection fails because it relies on a misconception of the notion of measurement. But I think that even if the objection works, we can get a safety based argument that some measurement cases will produce justified true beliefs without knowledge. And that will matter for the argument of the next two sections.

2 The Class of Gettier Cases is Disjunctive

There's an unfortunate terminological confusion surrounding gaps between knowledge and justification. Some philosophers use the phrase 'Gettier case' to describe any case of a justified true belief that isn't knowledge. Others use it to describe just cases that look like the cases in Gettier (1963), i.e., cases of true belief derived from justified false belief. I don't particularly have strong views on whether either of these uses is *better*, but I do think it is important to keep them apart.

I'll illustrate the importance of this by discussing a recent argument due to Jeremy Fantl and Matthew McGrath (Fantl and McGrath, 2009, Ch. 4). I've previously

discussed this argument (Weatherson, 2011), but I don't think I quite got to the heart of why I don't like the kind of reasoning they are using.

The argument concerns an agent, call her T , who has the following unfortunate combination of features. She is very confident that p . And with good reason; her evidence strongly supports p . For normal reasoning, she takes p for granted. That is, she doesn't distinguish between φ is best given p , and that φ is simply best. And that's right too, given the strong evidence that p . But she's not crazy. Were she to think that she was facing a bet on extreme odds concerning p , she would cease taking p for granted, and revert to trying to maximise expected value given the high probability that p . But she doesn't think any such bet is salient, so her disposition to retreat from p to *Probably* p has not been triggered. So far, all is going well. I'm inclined to say that this is enough to say that T justifiably believes that p . She believes that p in virtue of the fact that she takes p for granted in actual reasoning.³ She's disposed to stop doing so in some circumstances, but until that disposition is triggered, she has the belief. And this is the right way to act given her evidence, so her belief is justified. So far, so good.

Unfortunately, T really does face a bet on long odds about p . She knows she has to choose between φ and ψ . And she knows that φ will produce the better outcome iff p . But she thinks the amount she'll gain by choosing ψ if $\neg p$ is roughly the same as the amount she'll gain by choosing φ if p . That's wrong, and her evidence clearly shows it is wrong. If p is false, then φ will be *much* worse than ψ . In fact, the potential loss here is so great that ψ has the greater expected value given the correct evidential probability of p . I think that means she doesn't know that p . Someone who knows that p can ignore $\neg p$ possibilities in practical reasoning. And someone who could ignore $\neg p$ possibilities in practical reasoning would choose φ over ψ , since it is better if p . But T isn't in a position to make that choice, so she doesn't know that p .

(I've said here that T is wrong about the costs of choosing φ if p , and her evidence shows she is wrong. In fact I think she doesn't know p if either of those conditions obtain. But here I only want to use the weaker claim that she doesn't know p if both conditions obtain.)

Fantl and McGrath agree about the knowledge claim, but disagree about the justified belief claim. They argue as follows (this is my version of the 'Subtraction Argument' from page 97 of their book).

1. T is justified in choosing φ iff she knows that p .
2. Whether T 's belief that p is true is irrelevant to whether she is justified in choosing φ .
3. Whether T 's belief that p is 'Gettiered' is irrelevant to whether she is justified in choosing φ .
4. Knowledge is true, justified, UnGettiered belief.

³There are some circumlocutions here because I'm being careful to be sensitive to the points raised in Ross and Schroeder (2014) about the relationship between belief and reasoning. I think there's less distance between the view they put forward and the view I defended in Weatherson (2005) than they do, but this is a subtle matter, and for this paper's purposes I want to go along with Ross and Schroeder's picture of belief.

5. So T is justified in choosing φ iff she is justified in believing that p .
6. T is not justified in choosing φ .
7. So T is not justified in believing that p .

I think this argument is only plausible if we equivocate on what it is for a belief to be ‘Gettiered’.

Assume first that ‘Gettiered’ means ‘derived from a false intermediate step’. Then premise 4 is false, as Williamson’s example shows. S has a justified true belief that is neither knowledge nor derived from a false premise.

Assume then that ‘Gettiered’ simply means that the true belief is justified without being known. In that case we have no reason to accept premise 3. After all, the class of true justified beliefs that are not knowledge is pretty open ended. Before reading Williamson, we may not have thought that this class included the beliefs of agents using measuring devices that were functioning properly but imperfectly. But it does. Prior to the end of epistemology, we simply don’t know what other kind of beliefs might be in this class. There’s no way to survey all the ways for justification to be insufficient for knowledge, and see if all of them are irrelevant to the justification for action. I think one way a justified belief can fall short of knowledge is if it is tied up with false beliefs about the stakes of bets. It’s hard to say that that is irrelevant to the justification of action.

It is by now reasonably well known that logical subtraction is a very messy and complicated business. See, for instance, Humberstone (2000) for a clear discussion of the complications. In general, unless it is analytic that F s are G s and H s, for some antecedently understood G and H , there’s nothing interesting to say about the class of things that are G but not F . It will just be a disjunctive shambles. The same is true for knowledge and justification. The class of true beliefs that are justified but not known is messy and disjunctive. We shouldn’t expect to have any neat way of overviewing it. That in part means we can’t say much interesting about it as a class, contra premise 3 in the above argument. It also means the prospects for ‘solving the Gettier problem’ are weak. We’ll turn to that issue next.

3 There is No Solution to the Gettier Problem

The kind of example that Edmund Gettier (1963) gives to refute the justified true belief theory of knowledge has what Linda Zagzebski (2009, 117) aptly calls a “double luck” structure. In Gettier’s original cases, there’s some bad luck that leads to a justified belief being false. But then there’s some good luck that leads to an inference from that being true. As was quickly realised in the literature, the good and bad luck doesn’t need to apply to separate inferential steps. It might be that the one belief that would have been false due to bad luck also ends up being true due to good luck.

This has led to a little industry, especially in the virtue epistemology section of the market, of attempts to “solve the Gettier problem” by adding an anti-luck condition to justification, truth and belief and hoping that the result is something like an analysis of knowledge. As Zagzebski (1994) showed, this can’t be an *independent* condition on knowledge. If it doesn’t entail truth, then we will be able to recreate the

Gettier cases. But maybe a ‘fourth’ condition that entails truth (and perhaps belief) will suffice. Let’s quickly review some of these proposals.

So Zagzebski (1996) suggested that the condition is that the belief be true *because* justified. John Greco (2010) says that the extra condition is that the beliefs be “intellectually creditable”. That is, the primary that the subject ended up with a true belief is that it was the result of her reliable cognitive faculties. Ernest Sosa (2007) said that knowledge is belief that is true because it manifests intellectual competence. John Turri (2011) says that knowledge is belief the truth of which is a manifestation of the agent’s intellectual competence.

It should be pretty clear that no such proposal can work if what I’ve said in earlier sections is remotely right. Assume again that $v = 830$, $a = 832$ and $m = 10$. The agent believes that $V \in [822, 842]$. This belief is, we’ve said, justified and true. Does it satisfy these extra conditions?

My short answer is that it does. My longer answer is that it does if any belief derived from the use of a measuring device does, and since some beliefs derived from the use of measuring devices amount to knowledge, the epistemologists are committed to the belief satisfying the extra condition. Let’s go through those arguments in turn.

In our story, S demonstrates a range of intellectual competencies. She uses a well-functioning measuring device. It is the right kind of device for the purpose she is using. By hypothesis, she has had the machine carefully checked, and knows exactly the accuracy of the machine. She doesn’t form any belief that is too precise to be justified by the machine. And she ends up with a true belief precisely because she has so many competencies.

Note that if we change the story so a is closer to $v + m$, the case that the belief is true in virtue of S being so competent becomes even stronger. Change the case so that $a = 839$, and she forms the true belief that $V \in [829, 849]$. Now if S had not been so competent, she may have formed a belief with a tighter range, since she could easily have guessed that the margin of error of the machine is smaller. So in this case the truth of the belief is very clearly due to her competence. But as we noted at the end of section 1, in the cases where a is near $v + m$, the argument that we have justified true belief without knowledge is particularly strong. Just when the gap between justification and knowledge gets most pronounced, the competence based approach to knowledge starts to issue the strongest verdicts *in favour* of knowledge.

But maybe this is all a mistake. After all, the object doesn’t have the mass it has because of S ’s intellectual competence. The truth of any claim about its mass is not because of S ’s competence, or a manifestation of that competence. So maybe these epistemologists get the correct verdict that S does not know that $V \in [a - m, a + m]$?

Not so quick. Even had a equalled v , all these claims would have been true. And in that case, S would have known that V was within m of the measurement. What is needed for these epistemological theories to be right is that there can be a sense that a belief that p can be true in virtue of some cause C without C being a cause of p . I’m inclined to agree with the virtue epistemologists that such a sense can be given. (I think it helps to give up on content essentialism for this project, as suggested by

David (2002) and endorsed in Weatherson (2004).) But I don't think it will help. There's no real way in which a belief is true because of competencies, or in which the truth of a belief manifests competence, in the good case where $a = v$, but not in the bad cases, where a is in $(0, m)$. These proposals might help with 'double luck' cases, but there is more to the space between justification and knowledge than those cases. Of course, I think the space in question includes some cases involving false beliefs about the practical significance of p , but I don't expect everyone to agree with that. Happily, the Williamsonian cases should be less controversial.

4 What Can We Learn from Fallible Machines?

My presentation of Williamson's argument in section 1 abstracted away from several features of his presentation. In particular, I didn't make any positive assumption about what the agent can know when they find out that the machine reads a . Williamson makes a suggestion, though he offers it more as the most internalist friendly suggestion than the most likely correct hypothesis.

The suggestion, which I'll call the **Circular Reading Centred** hypothesis, is that the most the agent can know is that $V \in [a-(e+m), a+(e+m)]$. That is, the agent can know that V is in a region centred on a , the 'radius' of which is the margin of error m , plus the error on this occasion e . This is actually a quite attractive suggestion, though not the only suggestion we could make. Let's look through some other options and see how well they work.

We said above that the agent can't know more from a mismeasurement than they can know from an accurate measurement. And we said that given an accurate measurement, the most they can know is that $V \in [v-m, v+m]$. So here's one very restrictive suggestion: if $a \in [v-m, v+m]$, then the agent can know that $V \in [v-m, v+m]$. But we can easily rule that out on the basis of considerations about justification. The strongest proposition the agent is justified in believing is that $V \in [a-m, a+m]$. If the agent could know that $V \in [v-m, v+m]$, then she could know that $V \notin (v+m, a+m]$, even though she isn't justified in believing this. This is absurd, so that proposal is wrong.

We now have two principles on the table: S can't know anything by a mismeasurement that she knows on the basis of a correct measurement, and that she can only know things she's justified in believing. The first principle implies that for all $x \in [v-m, v+m]$, that $V = x$ is epistemically possible. The second implies that for all $x \in [a-m, a+m]$, that $V = x$ is epistemically possible. Our next proposal is that the epistemic possibilities, given a reading of a , are just that $V \in [v-m, v+m] \cup [a-m, a+m]$.

But this is fairly clearly absurd too. Assume that $a > v+2m$. This is unlikely, but as we said above not impossible. Now consider the hypothesis that $V \in (v+m, a-m)$. On the current hypothesis, this would be ruled out. That is, she would know it doesn't obtain. But this seems bizarre. There are epistemic possibilities all around it, but somehow she's ruled out this little gap, and done so on the basis of a horrifically bad measurement.

This suggests two other approaches that are consistent with the two principles,

and which do not have such an odd result. I'll list them alongside the proposal we mentioned earlier.

Circular Appearance Centred The strongest proposition the agent can know is that $V \in [a - (e + m), a + (e + m)]$.

Circular Reality Centred The strongest proposition the agent can know is that $V \in [v - (e + m), v + (e + m)]$.

Elliptical The strongest proposition the agent can know is that $V \in [v - m, a + m]$.

The last proposal is called **Elliptical** because it in effect says that there are two foci for the range of epistemic possibilities. The agent can't rule out anything within m of the true value, or anything within m of the apparent value, or anything between those.

Actually we can motivate the name even more by considering a slight generalisation of the puzzle that we started with. Assume that R is trying to determine the location of an object in a two-dimensional array. As before, she has a digital measuring device, perhaps a GPS locator trained on the object in question. And she knows that margin of error of the device is m . The object is actually located at $\langle x_v, y_v \rangle$, and the device says it is at $\langle x_a, y_a \rangle$. So the epistemic possibilities, by the reasoning given above, should include the circles with radius m centred on $\langle x_v, y_v \rangle$ and $\langle x_a, y_a \rangle$. Call these circles C_v and C_a . Unless $\langle x_v, y_v \rangle = \langle x_a, y_a \rangle$, the union of these circles will not be convex. If the distance between $\langle x_v, y_v \rangle$ and $\langle x_a, y_a \rangle$ is greater than $2m$, the union won't even be connected. So just as we 'filled in' the gap in the one-dimensional case, the natural thing to say is that any point in the convex hull of C_v and C_a is an epistemic possibility.

But now see what happens if we say those are all of the epistemic possibilities, i.e., that the agent knows that the true value lies in the convex hull of the two circles. Here's what it might look like.

Now consider the line from $\langle x_v, y_v \rangle$ to $\langle x_a, y_a \rangle$. No matter how bad the measurement is, the convex hull of the two circles C_v and C_a will include no points more than distance m from the line between $\langle x_v, y_v \rangle$ to $\langle x_a, y_a \rangle$. That is, the agent can know something surprisingly precise about how close V is to a particular line, even on the basis of a catastrophically bad measurement.

There are some circumstances where this wouldn't be counterintuitive. Assume that $x_v = x_a$, while y_v and y_a are very very different. And assume further that $\langle x_a, y_a \rangle$ is calculated by using two very different procedures for the x and y coordinates. (Much as sailors used to use very different procedures to calculate longitude and latitude.) Then the fact that one process failed badly doesn't, I think, show that we can't get fairly precise knowledge from the other process.

But that's not the general case. If the machine determines $\langle x_a, y_a \rangle$ by a more holistic process, then a failure on one dimension should imply that we get less knowledge on other dimensions, since it makes it considerably flukier that we got even one dimension right. So I think the space of epistemic possibilities, in a case involving this kind of errant measurement, must be greater than the convex hull of C_v and C_a .

Fortunately, there are a couple of natural generalisations of the elliptical proposal that avoid this complication. One of them says that the space of epistemic possibil-

ities forms an ellipse. In particular, it is the set of all points such that the sum of the distance from that point to $\langle x_v, y_v \rangle$ and the distance from that point to $\langle x_a, y_a \rangle$ is less than or equal to $2m + e$, where e again is the distance between the measured and actual value. As you can quickly verify, that includes all points on the line from $\langle x_v, y_v \rangle$ to $\langle x_a, y_a \rangle$, plus an extension of length m beyond in each direction. But it doesn't just contain the straight path between C_v and C_a ; it 'bulges' in the middle. And the considerations above suggest that is what should happen.

The other alternative is to drop the idea that the space of possibilities should be elliptical, and have another circular proposal. In particular, we say that the space of possibilities is the circle whose centre is halfway between $\langle x_v, y_v \rangle$ and $\langle x_a, y_a \rangle$, and whose radius is $m + e/2$. Again, that will include all points on the line from $\langle x_v, y_v \rangle$ to $\langle x_a, y_a \rangle$, plus an extension of length m beyond in each direction. But it will include a much larger space in the middle.

I think both of these are somewhat plausible proposals, though the second suffers from a slightly weaker version of the objection I'm about to mount to the Circular Reality Centred proposal. But they do share one weakness that I think counts somewhat against them. It's easy enough to see what the weakness is in the one-dimensional case, so let's return to it for the time being, and remember we're assuming that $a > v$.

Consider a case where e is rather large, much larger than m . This affects how far below v we have to go in order to reach possibilities that are ruled out by the measurement. But it doesn't affect how far above v we have to go in order to reach such possibilities. Indeed, no matter how bad e is, we can be absolutely certain that we know $V < a + 2m$, or that we know that $V > a - 2m$. That seems a little odd; if the measurement is so badly mistaken, it seems wrong that it can give us such a fine verdict, at least in one direction.

I don't think that's a conclusive objection. Well, I don't think many of the considerations I've listed here are *conclusive*, but this seems even weaker. But it is a reason to look away from the elliptical proposal and back towards the circular proposals that we started with.

If we just look at first order knowledge claims, it is hard to feel much of an intuitive pull towards one or other of the alternatives. Perhaps safety based considerations favour the Reality Centred over the Appearance Centred version, but I don't think the salient safety consideration is that strong.

If we look at iterated knowledge claims, however, there is a big problem with the Reality Centred approach. The intuition here is clearer if we use numerical examples, so I'll work through a case with numbers first, then do the general version next.

Assume, as above, that $v = 830$, $a = 834$ and $m = 10$. So we have a pretty decent measurement here. On the Reality Centred proposal, the strongest thing that S can know is that $V \in [816, 844]$. So it is an epistemic possibility that $V = 816$. Assume that that's the actual possibility. Then the measurement is rather bad; the new value for e is 18. Were V to equal 816, while a equalled 834, then on the Reality Centred approach, the epistemic possibilities would be a circle of radius $e + m$, i.e., 28, around the actual value, i.e., 816. So the strongest thing the agent could know

is that $V \in [788, 844]$. On the other hand, if V were 844, the strongest thing the agent could know is that $V \in [824, 864]$. Putting those together, the strongest thing the agent can know that she knows is that $V \in [788, 864]$. That's a very large range already. Similar calculations show that the strongest thing the agent can know that she knows that she knows is that $V \in [732, 904]$.

Now I'll grant that intuitions about second and third order knowledge are not always maximally sharp. But I think it is very implausible that a relatively accurate measurement like this could lead to such radical ignorance in the second and third orders of knowledge. So I think the Reality Centred approach can't be right.

The general form the case is as follows. The strongest thing the agent can know is that $V \in [v-(e+m), a+m]$. The strongest thing she can know that she knows is that $V \in [v-3(e+m), a+3m]$. And the strongest thing she can know that she knows that she knows is that $V \in [v-7(e+m), a+7m]$. In general, we have *exponential* growth of the possibilities as we add one extra order of knowledge. That seems absurd to me, so the Reality Centred approach is wrong.

Note that this isn't a problem with the Appearance Centred approach. The first-order epistemic possibilities are that $V \in [a-(e+m), a+e+m]$. If V is at the extremes of this range, then e will be rather large. For example, if V were equal to $a+e+m$, then the new error would be $e+m$, since the measured value is still a . So the range of possibilities would be that $V \in [a-((e+m)+m), a+((e+m)+m)]$. Somewhat surprisingly, those would also be the possibilities if V were equal to $a-(e+m)$, since the only feature of V that affects the epistemic possibilities for V is its distance from a . So for all S knows that she knows, V could be anything in $[a-(e+2m), a+(e+2m)]$. Similar reasoning shows that for all V knows that she knows that she knows, V could be anything in $[a-(e+3m), a+(e+3m)]$. In general, V has n 'th order knowledge that V is in $[a-(e+nm), a+(e+nm)]$. This linear growth in the size of the range of epistemic possibilities is more plausible than the exponential growth on the Reality Centred approach.

So all things considered, I think the Circular Appearance Centred approach is the right one, as Williamson suggests. Any simple alternative seems to have rather counterintuitive consequences.

References

- Cohen, Stewart, (1984). "Justification and Truth." *Philosophical Studies* 46: 279-295, doi:10.1007/bf00372907. (3)
- David, Marian, (2002). "Content Essentialism." *Acta Analytica* 17: 103-114, doi:10.1007/bf03177510. (8)
- Fantl, Jeremy and McGrath, Matthew, (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press. (4)
- Gettier, Edmund L., (1963). "Is Justified True Belief Knowledge?" *Analysis* 23: 121-123, doi:10.2307/3326922. (4, 6)
- Greco, John, (2010). *Achieving Knowledge*. Cambridge: Cambridge University Press. (7)

- Humberstone, Lloyd, (2000). "Parts and Partitions." *Theoria* 66: 41-82, doi:10.1111/j.1755-2567.2000.tb01144.x. (6)
- Ross, Jacob and Schroeder, Mark, (2014). "Belief, Credence, and Pragmatic Encroachment." *Philosophy and Phenomenological Research* 88: 259-288, doi:10.1111/j.1933-1592.2011.00552.x. (5)
- Sainsbury, Mark, (1995). "Vagueness, Ignorance and Margin for Error." *British Journal for the Philosophy of Science* 46: 589-601, doi:10.1093/bjps/46.4.589. (4)
- Sosa, Ernest, (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford: Oxford University Press. (7)
- Turri, John, (2011). "Manifest Failure: The Gettier Problem Solved." *Philosophers' Imprint* 11: 1-11, Available from: <http://hdl.handle.net/2027/spo.3521354.0011.008>. (7)
- Weatherson, Brian, (2004). "Luminous Margins." *Australasian Journal of Philosophy* 82: 373-383, doi:10.1080/713659874. (8)
- , (2005). "Can We Do Without Pragmatic Encroachment?" *Philosophical Perspectives* 19: 417-443, doi:10.1111/j.1520-8583.2005.00068.x. (5)
- , (2011). "Knowledge, Bets and Interests." In Jessica Brown and Mikkel Gerken (eds.), *forthcoming volume on knowledge ascriptions*, 75-103. Oxford: Oxford University Press. (5)
- Williamson, Timothy, (2000). *Knowledge and its Limits*. Oxford University Press. (4)
- , (2013). "Gettier Cases in Epistemic Logic." *Inquiry* 56: 1-14, doi:10.1080/0020174X.2013.775010. (1)
- Zagzebski, Linda, (1994). "The Inescapability of Gettier Problems." *Philosophical Quarterly* 44: 65-73, doi:10.2307/2220147. (6)
- , (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press. (7)
- , (2009). *On Epistemology*. Belmont, CA.: Wadsworth. (6)