RANDOM AND SYSTEMATIC ERROR IN THE PUZZLE OF THE UNMARKED CLOCK

Abstract. A puzzle of an unmarked clock, used by Timothy Williamson to question the KK principle, was separately adapted by David Christensen and Adam Elga to critique a principle of Rational Reflection. Both authors, we argue, flout the received relationship between ideal agency and the classical distinction between systematic and random error, namely that ideal agents are subject only to the latter. As a result, these criticisms miss their mark.

1. The KK Principle

Our targets here are two papers (David Christensen 2010 and Adam Elga 2013) on "Rational Reflection". Our journey begins, however, with an author in whom one finds similarly eccentric (albeit more defensible) sinuosities–Timothy Williamson.

Williamson (2011, 2014) sought to develop formal epistemic models undermining the so-called *KK principle.* (If you know P, you know that you know P.) He wrote "As is well known, the non-transitivity of the epistemic accessibility relation R is necessary and sufficient for a frame to contain a counterexample to the KK principle." He then claims that the following scenario may be aptly modelled by a non-transitive epistemic accessibility relation:

> Imagine a plain, unmarked circular dial with a single pointer [that] can point at any one of n equally spaced unmarked positions on the perimeter of the dial (...) Number the positions $0, \ldots, n-1$ clockwise from the top. For simplicity, we individuate 'worlds'...just by the position of the hand. (...) Now imagine that you are looking at the dial from a fixed point of view (...) the difference between neighbouring positions is well below your threshold of discrimination. Consequently, if you are in fact in world w, the worlds that for all you know you are in (the epistemically possible worlds) are those at most h steps from w, for some natural number h; h is greater than 0, otherwise your discrimination would be perfect. We can regard h as the width of the margin for error you require in order to know something in the model (...) More formally, let R be the epistemic accessibility relation; then for all worlds w and x, Rwx if and only if the distance between w and x is at most h. (...) For any world w, $R(w)$ (the set of worlds epistemically possible in w) (...) is an interval with w as its midpoint and $2h + 1$ members. (...) $R(w)$ is known only at w.

Williamson says that you know $R(w)$, where w is the actual world, but don't know that you know $R(w)$. (And, if h is large, deem it *unlikely* that you so know $R(w)$; hence *Improbable Knowing.*)

This argument requires non-standard assumptions about what otherwise looks like a standard case of observational error. For concreteness, let $n = 60$ (minute hand of an unmarked clock) and consider $w = w_{23}$ (23 minutes past the hour) with $h = 2$. On the standard view (cf. Mott 1998), when you "guess" the position of the hand, e.g. determine a midpoint of a range of possible values, you are, in essence, conducting a "measurement". If Williamson's "margin for error" were observational error then the measurement you conduct could potentially return the value "21", in which case what you know (idealistically assuming the margin of error to be strict of course, as knowledge is factive) given that $h = 2$ would be $\{w_{19}, w_{20}, w_{21}, w_{22}, w_{23}\}$. So the epistemically possible worlds (assuming the measurement you conduct returns an integer value) would be those where the dial is within h steps not of the actual *position*, but of the actual *measurement*.

So, what Williamson means by "margin for error" must be something else. Indeed, he implicitly assumes that no observational error occurs at all. For suppose you are told (either before or after looking at the clock) that the actual position of the clock is either 20 or 23. If your measurement returns the value 21 but w_{23} is actual, you can't very well know the event $E = \{w_{21}, w_{22}, w_{23}, w_{24}, w_{25}\}$. You don't even believe it–your credence in w_{20} is at least 50%. So Williamson would not be able to say "if you are in world 23 then the epistemically possible worlds are those at most 2 steps from 23" if the measurement might return the value 21 (or, *mutatis mutandis*, any value other than 23). So the word choice "h is greater than 0, otherwise your discrimination would be perfect" is misleading; more consistent would be "h is greater than 0, though your discrimination is perfect". Since both Christensen and Elga capitalize on this very feature of Williamson's model we want to make careful note of it here. It's a model that can only apply to agents who are in one sense perfect, but do not realize that they are so.[1]

2. Rational Reflection, part one: Christensen

Williamson's framework was seized upon by David Christensen (2010), then Adam Elga (2013), each of whom sought implications of the unmarked clock for *Rational Reflection*. Letting Cr denote an agent's credences and Pr the credences "that would be maximally rational for someone in that agent's epistemic situation", Christensen formalizes this principle as follows:

RatRef: $Cr(A|Pr(A) = \alpha) = \alpha$

Christensen writes:

[1]Williamson (2014) gives an embellished frame, again non-transitive, allowing for the possibility of imperfect measurement. He also proves that non-transitive frames involve violations of a Reflection principle based on sufficiently regular "evidential" probabilities (these are obtained from a common prior at world w by conditionalization on $R(w)$). Since it is not possible in his model for agents to always assign $R(w)$ full credence, however, there is no certain relationship between his Reflection principle and those considered in this paper. Since (also) our primary targets drew no inspiration from the later model, there is no reason to consider it here.

> Suppose Chloe is looking at an unmarked clock a few feet away with just a minute hand. (...) Chloe is considering various propositions about the hand's position (...) if she had to bet on one position, she'd bet on P21. (...) So it seems that her confidence should be distributed over the possible hand positions in a tight, roughly bell-shaped curve, with the peak at 21. (...) ...trouble ensues, as follows: Chloe is .3 confident that P21, and .7 confident that it's somewhere else. But she also thinks that it's rational for her to have .3 confidence that P21 only if she's actually in the evidential situation that would be produced by the hand being at 21. And she believes that the rational credence for her to have that P21 is less than .3 if she's in any of the other evidential situations she might be in. Since she thinks she's probably in one of those other situations, it seems that she should think that .3 is probably too high a credence for her to have that P21, and certainly not too low. But that seems to suggest that her credence that P21 should be lower than .3.

Christensen is more explicit than Williamson about the fact that his agent is not subject to observational error. Indeed, the language "...it's rational for her to have .3 confidence that P21 only if she's actually in the evidential situation that would be produced by the hand being at 21" makes clear that he considers Chloe's evidential situation (her credences, in particular) to be a function of hand position; all of Chloe's evidential counterparts inhabit P21 scenarios. Christensen brings this point into relief by making a comparison to a more typical case:

> Consider...a doctor who has (and should have) .8 credence that her patient has hepatitis, based on her list of the patients symptoms and on statistics she knows which say that 80% of patients with these symptoms have hepatitis. In this sort of case, the fact that the doctor should not be maximally confident that her patient has hepatitis seems to derive from the fact that in evidential situations exactly like the present one, the patient has hepatitis only 80% of the time. The doctor's rational credence is limited, as are most credences most people have most of the time, by the fact that her evidence doesn't discriminate perfectly between cases in which the relevant proposition is true, and ones where it's false.

What's different about the doctor's situation? Recall the distinction between *random error* and *systematic error*. Suppose that one is attempting to estimate a random variable Y with an estimator X. If, conditional on $X = X_0$, Y is not constant almost surely, one says there is "random error" in the estimation. If, conditional on $X = X_0$, $E(Y) \neq X_0$ (here the expectation is taken with respect to ideal rational credence), one says there is "systematic error". (In the examples we'll look at Y will be an indicator function, but the idea is more general.)

To return to the example, "in evidential situations exactly like the present one, the patient has hepatitis only 80% of the time" tells us that ideal rational credence in *hepatitis* is .8. Letting $Y = 1$ when the patient has hepatitis and $Y = 0$ when the patient does not, the doctor "estimates" Y by setting $X = X_0 = .8$; hence there is no systematic error in the doctor's estimation. On the other hand there is random error, for "her evidence doesn't discriminate perfectly between cases in which the relevant proposition is true, and ones where it's false". That is to say, "in evidential situations exactly like the present one", Y isn't constant almost surely (it's sometimes 0 and sometimes 1).

Chloe, on the other hand, is prone to systematic, but not to random, error. Let $Y = 1$ when the absolute difference of the hand position and Chloe's estimate is positive; $Y = 0$ otherwise. Since Chloe thinks that her accuracy is independent of the perceived position and that there is a 70% chance that $Y = 1$, she will estimate Y by $X = X_0 = .7$. But $E(Y) = 0$, because in every evidential situation exactly like Chloe's, $Y = 0$. (Christensen is committed to this, for he says that it's rational for her to have .3 confidence that P21 only if she's actually in the evidential situation that would be produced by the hand being at position 21; she has .3 confidence that P21, and she's supposed to be rational.) So $E(Y) = 0 \neq .7 = X_0$; this is systematic error.

The presence of this systematic error precludes Chloe's being ideally rational in the end. To see this, note first that (given the setup) the ideal credence function is the one that assigns the observed position full measure. In the following passage, Christensen appears to concede this point:

> One possible reaction to this point would be to eliminate the tension in favor of holding that Chloe should be absolutely certain that P21. One might defend this initially unintuitive move by saying that (Chloe's) credences represent ideally rational beliefs, and that while we would not expect an actual human to live up to these ideals, that shouldn't undermine their status as ideally rational.

Christensen believes, however, that Chloe's status as an ideally rational agent does not preclude her adoption of a non-ideal credence function. He writes

> But it seems to me that our initial intuitive rejection of the rationality of Chloe's being certain is in the end correct. To be rationally certain that P21, Chloe would have to be certain that she wasn't mistaking the P20 visual experience for the P21 experience. Of course, it may be claimed that a perfectly rational agent would be immune from such cognitive errors. But this claim, it seems to me, is beside the point. For even if it is a fact that Chloe is cognitively perfect, and never misinterprets her experiences, she has no reason to be certain of that fact. Even if, say, she would in fact always pick the correct hand position if she were forced to pick just one, she has no grounds for being certain that this is so. So it seems that she cannot absolutely dismiss the possibility that she's made the sort of mistake in question. And to the extent that she can't dismiss that possibility, she must countenance the possibility that P21 is false.

That strikes us as confused. If Chloe's evidence discriminates perfectly between cases in which she is currently mistaking the P20 visual experience for the P21 experience and cases in which she is interpreting the P21 visual experience as such then she *does* have grounds for being certain that she isn't mistaking the P20 visual experience for the P21 experience. Christensen is therefore committed to holding that Chloe's evidence fails to discriminate perfectly between cases in which she is currently mistaking the P20 visual experience for the P21 experience and cases in which she is correctly interpreting the P21 visual experience, contradicting his earlier commitment that in every evidential situation exactly like Chloe's, the difference between the actual and perceived positions is zero.

Indeed, Chloe *can't* (robustly) possess the property that she always picks the correct hand position if forced to choose, yet routinely assign positive probability to being wrong. Suppose Chloe adopts credence functions given by a tight, roughly bell-shaped curves, with the peak at the measured position–say $(.05, .1, .2, .3, .2, .1, 05)$ over $(P(x-3), P(x-2), P(x-1), P(x), P(x+1), P(x+2), P(x+3))$ when it looks to her that the hand is in position x. Imagine now a scenario in which Chloe's prior distribution over $(P20, P21, P22, P23, P24)$ is given by $(\frac{9}{95}, \frac{8}{95}, \frac{18}{95}, \frac{24}{95}, \frac{36}{95})$. (Perhaps Chloe knows that the hand position was determined by a roll of a fair 95 sided die.) Chloe now looks at the clock and it looks to her like the hand is in position 21 (which it is in fact in, by prior assumption). A simple calculation shows that her posterior will be $(.15, .2, .3, .2, .15)$. That is, her posterior distribution will be represented by a tight, roughly bell-shaped curve, with the peak at...22. So now, "if she had to bet on one position", surely it would be position 22. But on the view that the field of light (in this case) determines the "evidential situation", Chloe's evidental situation is not consistent with P22.

Christensen's error involves confusion between "cognitive perfection" (always picking the correct hand position) and "perfect rationality". For Chloe's predicament to bear at all on **RatRef**, she must be perfectly rational, but needn't be "cognitively perfect". In a case where she isn't (or isn't certain that she is), however, her (ideal) credences needn't be a function of the actual hand position.

3. Rational Reflection, part two: Elga

Elga (2013) formulates naive Rational Reflection as follows:

RATIONAL REFLECTION $P(H|P' \text{ is ideal}) = P'(H)$.

Here P is the credence function of a possible rational subject and "ideal" means "perfectly rational"). Elga also assumes that "every situation determines a unique ideally rational probability function". He buys into Christensen's critique of naive Rational Reflection and proposes the following variant:

NEW RATIONAL REFLECTION $P(H|P' \text{ is ideal}) = P'(H|P' \text{ is ideal})$

On the (standard, again) view we've been espousing, Elga's "new" Rational Reflection principle is equivalent to the original, as $P'(P' \text{ is ideal}) = 1$ whenever P' is ideal. To see this, recall that a credence function is ideal if and only if it is calibrated to one's evidential position; it's ideally rational to presently assign A a credence of .37 (say) if and only if in evidential situations exactly like the present one, "A" comes out true 37% of the time. So if, for each counterpart x, you consider the function T_x taking propositions to $\{0,1\}$ (true propositions get mapped to 1, false to zero), the credence function defined by $P'(q) = \int T_x(q)\, d\mu(x)$ is ideal. Here μ is ideal credence, realized as a density over counterparts–something like "the measure employed by nature to finger said counterparts for instantiation".[2]

Elga however insists that **NEW RATIONAL REFLECTION** really is "new". In support of this he lays out the following familiar thought experiment purporting to establish that there are situations where uncertainty regarding what attitudes are perfectly rational is the perfectly rational attitude to have.

> **HYPOXIA** Bill the perfectly rational airline pilot gets a credible warning from ground control: "Bill, there's an 99% chance that in a minute your air will have reduced oxygen levels. If it does, you will suffer from hypoxia (oxygen deprivation), which causes hard-to-detect minor cognitive impairment. In particular, your degrees of belief will be slightly irrational. But watch out–if this happens, everything will still seem fine. In fact, pilots suffering from hypoxia often insist that their reasoning is perfect–partly due to impairment caused by hypoxia!" A few minutes later, ground control notices that Bill got lucky–his air stayed normal. They call Bill to tell him. Right before Bill receives the call, should he be uncertain whether his degrees of belief are perfectly rational?

Elga says "no", which leads him to acceptance of the following principle:

MODESTY In some possible situations, it is rational to be uncertain about what degrees of belief it is rational for one to have. Furthermore, it can be rational to have positive degree of belief that one is in such a situation.

To repeat, it strikes us as a truism that all counterparts of an ideally rational agent are ideally rational; they share the agent's evidential situation–in particular her degrees of belief–and what it means to be ideally rational is just that one's credences are ideally calibrated to one's evidential situation. Since it is insisted upon in this thought experiment that Bill, if hypoxic, will have "slightly irrational" degrees of beliefs, it follows that if Bill is ideally rational then he has no hypoxic counterparts and will therefore assign credence 1 to the proposition *I am not in the hypoxic state*. Elga glosses this view, but disappointingly rejects it as "desperate".

It isn't desperate. Elga wants Bill to be ideally rational and yet have hypoxic counterparts, and that just isn't possible if the hypoxic versions of Bill have irrational credences.[3] It is possible (perhaps this is what Elga's dissenting intuitions

[2] What μ is precisely may be a difficult question in general, but it doesn't matter in this case because P' is ideal for you just in case it is ideal for your evidential counterparts. (They share both your evidential situation and your credences.) So if P' is ideal for you (and hence for your counterparts), $\int T_x(P' \text{ is ideal})\, d\mu(x)$ is going to come out 1, regardless of what μ is.

[3] Contra Elga's implicit suggestion it doesn't suffice that Bill cannot identify himself as hypoxic in those cases where he is. By that reasoning, since some have the thought "I am not

are hearkening to) for Bill to be ideally rational and have counterparts (the hypoxic ones, if you like) who are not "cognitively perfect", however. On this view it is only hypoxic Bill's status as an *instrument* that suffers impairment in the hypoxic state. Hypoxic Bill is prone, that is, to heightened levels of random error.

To see how, suppose for concreteness that the hypoxia only affects Bill in his judgments about positions on an unmarked clock. Hypoxic Bill looks at the clock and records the position he thinks is most likely. In fact, the distribution of his choice is $(.15, .2, .3, .2, .15)$ over $(P(x-2), P(x-1), P(x), P(x+1), P(x+2))$, where x is the actual position. Non-hypoxic Bill's distribution, meanwhile, is $(.3, .4, .3)$ over $(P(x-1), P(x), P(x+1))$. Assume that Bill knows all of this.

Suppose further that all of Bill's information about his state (hypoxic or not) comes from an internal, $\{0, 1\}$ detector that returns the correct reading 99% of the time (conditional on either being in the hypoxic state or not). A quick calculation tells us that if the reading comes back "non-hypoxic", Bill will assign credence $\frac{1}{2}$ to his being in the hypoxic state and will, accordingly, adopt credence function $(.075, .25, .35, .25, .075)$ for the actual hand position (centered around the observed position). If on the other hand, the reading comes back "hypoxic" then Bill will assign credence $\frac{9801}{9802}$ to being in the hypoxic state and adopt a credence function only modestly more concentrated than $(.15, .2, .3, .2, .15)$.

In either case Bill assigns positive probability to the proposition that he is in the hypoxic state. This does not preclude his being ideally rational, however. Equivalently, it doesn't preclude his being certain of that fact.

4. Appendix: Objections Considered

In this section we entertain various possible objections.[4]

1. "The arguments depend on a conception of ideal rationality that is not defended or argued for."

We think we are repeating ourselves, but an agent's credence function is ideal if it is calibrated to the agent's evidential position. Ideal credence in A may be realized as chance of A, conditional on the agent's evidence. Here we are speaking

dreaming" whilst dreaming, (even though I am ideally rational) many of my current counterparts may be asleep, and I cannot know that I am not dreaming. That's plainly wrong, provided (as we may assume) these dreams aren't as vivid (etc.) as my current experience. Similarly, if I fantasize that I am Feynman, that doesn't disqualify Feynman from knowing that he isn't me; my wistful imaginings aren't so Feynmanesque as to qualify me as his counterpart.

[4]These are actual objections of a single anonymous referee to a draft essentially identical to the current one. (Apart from this new appendix, we previously used "epistemic counterpart" rather than "evidential counterpart". This terminological change accentuates the fact that credences are degrees of belief, not degrees of knowing. There is, in particular, nothing "factive" about credence. This distinction may be relevant to Christensen/Elga's) misappropriation of Williamson's model.) We thank this referee and apologize for stealing (in minor paraphrase) these objections. (We weren't, unfortunately, able to determine how the referee imagines Christensen and Elga might respond to our primary charge of systematic error.)

as if local initial conditions, i.e. the initial conditions giving rise to everything the agent will ever see, are "chancy". Alternatively, order the agent's evidential counterparts by physical nearness (or by any other nearness relation having no reason to be correlated with truth of A). Now ideal credence in A is the almost sure asymptotic frequency of "A scenario inhabitors" in this sequence. (Almost surely, then, it is the actual asymptotic frequency, about which there is a "fact of the matter".) Ideal credences exist and are unique, on our view, even when there aren't obvious symmetries allowing non-ideal agents (like us) to compute them. Elga (2010) invites consideration of a "no obvious symmetries" case:

> A stranger approaches you on the street and starts pulling out objects from a bag. The first three objects he pulls out are a regular-sized tube of toothpaste, a live jellyfish, and a travel-sized tube of toothpaste. To what degree should you believe that the next object he pulls out will be another tube of toothpaste?

What one should shoot for here is the almost sure asymptotic frequency of "toothpaste scenarios" in situations that are as-described. (And in an infinite universe, of course there are such situations, rare as they may be...infinitely many of them.)

2. "Both Christensen and Elga would reject this conception of ideal rationality."

To reiterate, our reading is that Christensen's paper essentially defers to our conception of idealness, at least as this term is applied to credence functions. (See also point 4 below.) In particular, given the supposition that all of Chloe's evidential counterparts inhabit P21 scenarios, Christensen grants that her ideal credence function assigns this center hand position full credence. What is being more generally conceded here, we think, is that in a case where proportion r (in an asymptotic sense, almost surely) of one's evidential counterparts inhabit A scenarios, the ideal credence function assigns A probability r. If one doesn't make this concession, one is committing to attitudes such as that of holding as ideal a credence different than .54 in the proposition P that the next object pulled from the stranger's bag will be a tube of toothpaste, when in fact .54 is the precise objective chance of toothpaste conditional on the agent's evidence. In such a case .54 will almost surely be the asymptotic frequency of P scenarios in situations that are evidentially like that of the agent, and so .54 will maximize expected long run accuracy in the class of evidential counterparts. Elga (2010) explicitly rejects imprecise credences, and Elga (2013) grants that "every situation determines a unique ideally rational probability function". Deferring to these attitudes, the position here would then presumably be that the "ideal credence function" assigns P some different sharp value. But what different sharp value? Something less than .54? Greater than .54? This choice looks arbitrary.

3. "The paper argues that in the unmarked clock case, an ideally rational agent would be certain that the clock pointer is where it appears to be."

This isn't exactly an objection, but we fear it may be a misunderstanding. Given the supposition that Chloe has .3 confidence that P21 only if she's actually in the evidential situation that would be produced by the hand being at 21, etc., yes. This however is Christensen's supposition, not ours. We worry that Christensen may conflate (under the heading "cognitive perfection") being ideally rational with being a perfect instrument. We would regard as more instructive a case in which an ideal rational agent is assumed to be an *imperfect* instrument. In such a case, this ideally rational agent would be unsure of the pointer's position.

4. "The paper's conception of ideal rationality is a coherent one, but so apparently is that of Christensen and Elga, who both want to leave room for cases where an ideal agent's counterparts are non-ideal. The interesting philosophical issue at the heart of the debate concerns the decision between the two conceptions."

Coherence doesn't enable an inference from "want to" to "may". (The function assigning A credence equal to the probability of $\neg A$ conditional on the evidence is coherent, but not defensibly "ideal".) Our guess is that this objection is a rehash of the idea that an ideal agent in the P21 scenario might have (non-ideal) evidential counterparts in a different clock position scenario. We have, however, disposed of that idea in Section 2 (see paragraph beginning "That strikes us as confused".) In the absence of a more perspicuous formulation we see no reason to honor the claim that there is some "interesting philosophical issue" in play here.[5]

5. "The claim that Chloe can't possess the property that she always picks the correct hand position if forced to choose misinterprets Christensen, who presumably means that Chloe is always correct when she choices just based on her visual experience, with flat priors. So the argument here isn't a problem for Christensen."

To the contrary, this observation misinterprets *our* point. The threat Christensen is addressing is that if an agent is rational then her credence function ought to be given by her prior conditional on her evidence. Chloe fails that criterion; her evidence entails P21, yet she only assigns P21 probability .3. Let us reexamine a portion of Christensen's parry:

> To be rationally certain that P21, Chloe would have to be certain that she wasn't mistaking the P20 visual experience for the P21 experience. (...) (but) even if...Chloe is cognitively perfect, and never misinterprets her experiences, she has no reason to be certain of that fact. Even if, say, she would in fact always pick the correct hand position if she were forced to pick just one, she has no grounds for being certain that this is so.

[5]How could there be? All of the philosophical subtlety in Williamson's work revolves around the non-transitive epistemic accessibility relation R, and the evidential sameness relation at the root of the current discussion isn't plausibly intransitive. Though outside our scope, incorporating appearances into the model and relaxing the stipulation that Chloe isn't subject to observational error brings this point out nicely; cf. Williamson's (2014) discussion of the (transitive) relation R_B tracking "blameless belief" (p. 986). It is natural that the transitivity of these relations should stand or fall together, for on the view that credences are quantified beliefs, ideally rational credences are just maximally blameless quantified beliefs.

Christensen relies here on an implicit inference from "Chloe would always pick the correct hand position if she were forced to pick just one" to "Chloe never misinterprets her experiences". (Recall, it is essential that Chloe never misinterprets her experiences; otherwise, her case doesn't bear on the veracity of **RatRef**.) As the non-flat prior example shows, however, "Chloe would always pick the correct hand position if she were forced to pick just one" isn't true. What is true (by assumption) is that "Chloe would always pick the correct hand position if she were forced to pick just one in a case where she had flat priors," but there is no valid inference from that much weaker proposition to "Chloe never misinterprets her experiences". The latter is in fact false here, as in the non-flat scenario Chloe's experience entails P21, yet she would choose P22 if forced to pick one position.

These objections fail to clear room for an "alternative conception of ideal rationality" on which it would be possible for ideal agents to have non-ideal evidential counterparts. Until such time as more satisfying grounds for such views are offered, then, one ought to look upon conclusions derived from them with suspicion.

References

Christensen, D. 2010. Rational Reflection. *Philosophical Perspectives* 24:121-140.

Elga, A. 2013. The Puzzle of the Unmarked Clock and the New Rational Reflection Principle. *Philosophical Studies* 164(1):127-139.

Mott, Peter. 1998. Margins for Error and the Sorites Paradox. *Philosophical Quarterly* 48:494-504.

Williamson, T. 2011. Improbable Knowing. In *Evidentialism and its Discontents*, T. Dougherty (ed.). Oxford University Press, Oxford. 147-164.

Williamson, T. 2014. Very Improbable Knowing. *Erkenntnis* 79:971-999.