Acta Analytica
https://doi.org/10.1007/s12136-022-00511-6

Safety, Evidence, and Epistemic Luck

Michael J. Shaffer[1]

Received: 18 December 2021 / Accepted: 18 December 2021

Abstract
This paper critically explores Timothy Williamson's view of evidence, and it does so in light of the problem of epistemic luck. Williamson's view of evidence is, of course, a crucially important aspect of his novel and influential "knowledge-first" epistemological project. Notoriously, one crucial thesis of this project is that one's evidence is equivalent to what one knows. This has come to be known as the E = K thesis. This paper specifically addresses Williamson's knowledge-first epistemology and the E = K thesis in the context of anti-luck epistemology (i.e., the view that knowledge is not compatible with certain forms of epistemic luck) and the idea that knowledge is factive (i.e., the view that knowledge implies truth). Williamson's views on these matters are worth investigating in some detail because he subscribes to a well-worked out anti-luck view of knowledge that incorporates what is perhaps the most common anti-luck condition (i.e., the safety condition). But this paper is also of more general importance because the critique of Williamson's views on these matters reveals some important things about the nature of evidence and evidence is one of the most fundamental concepts in epistemology.

1 Introduction

This paper critically explores Timothy Williamson's view of evidence, and it does so in light of the problem of epistemic luck. Williamson's view of evidence is, of course, a crucially important aspect of his novel and influential "knowledge-first" epistemological project. Notoriously, one crucial thesis of this project is that one's evidence is equivalent to what one knows. This has come to be known as the E = K thesis. This paper specifically addresses Williamson's knowledge-first epistemology and the E = K thesis in the context of anti-luck epistemology (i.e., the view that knowledge is not compatible with certain forms of epistemic luck) and the idea that knowledge is factive (i.e., the view that knowledge implies truth). Williamson's

✉ Michael J. Shaffer
shaffermj66@outlook.com

1 Department of Philosophy, Gustavus Adolphus College, OM 104, 800 West College Avenue, Saint Peter, MN 56082, USA

views on these matters are worth investigating in some detail because he subscribes to a well-worked out anti-luck view of knowledge that incorporates what is perhaps the most common anti-luck condition (i.e., the safety condition). But this paper is also of more general importance because the critique of Williamson's views on these matters reveals some important things about the nature of evidence and evidence is one of the most fundamental concepts in epistemology.

Ultimately, it will be shown that Williamson's views on these matters have quite implausible implications about what can and cannot count as evidence. This critique will also reveal some important things about how evidence, knowledge, and epistemic luck are conceptually related. It is of special importance to note here that if Williamson's view of evidence is correct, then *good* evidence is not merely true, as per the standard account of evidence. Good evidence, on Williamson's view, is true *and modally safe*. The key here is Williamson's subscription to the modal safety of evidence/knowledge, and it is worth recognizing that requiring evidence/knowledge to be modally safe does have some decided merits. For example, on this basis, one can explain knowledge of improbable propositions. One can know improbable propositions in cases where the belief in a proposition is such that it would not be believed in close worlds where it is false. But, if propositional evidence must be true and modally safe, be it highly probable or improbable, then there is far less evidence than we normally suppose. These observations, in turn, show that $E \neq K$.

2 Gettier, Luck, and JTB+

The standard analysis of knowledge is that knowledge is justified true belief. So, this account of knowledge can be succinctly presented as follows:

(JTB) S knows that p, if and only if,

(i) S believes that p,
(ii) S's belief that p is justified, and
(iii) p is true.

(i), (ii), and (iii) are supposed to be individually necessary and jointly sufficient for knowledge. The JTB analysis is a decompositional analysis of the nature of knowing that is supposed to equate knowledge with more basic components. But, in 1963 Edmund Gettier demonstrated that the JTB account of knowledge is incorrect.[1] In that paper Gettier presented two cases in which conditions (i)-(iii) were met, but where the correct intuition is supposed to be that the agent in question does not know.[2] In this way, Gettier challenged the sufficiency of the JTB account of knowledge. Let us then look at one of the cases. Consider the case of Smith.[3] We

[1] Gettier 1963.
[2] It is worth mentioning that Gettier's case for the rejection of the JTB account only follows as a deductive consequence given the assumptions of epistemic closure and the idea that one can be justified in holding a false belief.
[3] Gettier 1963, 122.

are to suppose that Smith has strong evidence for the claim that Jones owns a Ford. This evidence includes that Jones has always owned a Ford and that Jones has just offered a ride to Smith while driving a Ford. Suppose also that Smith has a friend Brown and that Smith does not know where Brown currently is. So, Smith formulates the following beliefs. Either Jones owns a Ford or Brown is in Boston. Either Jones owns a Ford or Brown is in Barcelona. Either Jones owns a Ford or Brown in Brest-Litovsk. All three are entailed by the claim that Jones owns a Ford. But, suppose that Jones does not in point of fact own a Ford, say he is presently driving a rental car. Moreover, by coincidence suppose that unknown to Smith Brown is actually in Barcelona. This means that Smith meets conditions (i)–(iii) of the JTB analysis, but intuitively we do not believe that Smith knows that either Jones owns a Ford or Brown is in Barcelona. What has happened is that Smith's belief has been caused in some inappropriate manner and the truth of his justified belief is, in some important sense of the term, *a matter of epistemic luck*. Gettier's cases have been widely taken to refute the JTB analysis of knowledge. Importantly, in light of this result, many practitioners of post-Gettier epistemology have then been concerned with the offering of an alternative analysis of knowledge, prominently including "fourth condition" analyses (JTB + analyses) that are intentionally designed rule out cases involving epistemic luck as bona fide cases of knowledge.[4]

3 E = K and Knowledge-First Epistemology

In this context, Timothy Williamson has proposed a re-orientation of epistemology in response to the post-Gettier dynamic just described. Central to this project is Williamson's rejection of decompositional analyses of knowledge. This is the idea that rather than approaching knowledge from the perspective of conceptual analysis where one aims to break the concept of knowledge down into more basic and familiar component concepts and then to test such compositional analyses against cases, he recommends that we approach epistemology by first identifying clear cases of knowledge. This is to be followed by the identification of epistemologically interesting properties that such states typically have. Given this very basic understanding of knowledge-first epistemology, it is then perfectly clear that in pursuit of this project, Williamson (2000) endorses a set of epistemic principles that jointly imply the view that all evidence must be true. In epistemological jargon, this is just the view that evidence is *factive*. Littlejohn (2013) refers to those who defend this position about the factivity of evidence as "truthers" and he usefully construes this view as follows:

(ET) p is evidence only if p is true.

"Falsies," on the other hand, endorse the denial of ET:

[4] See Unger 1968, Pappas and Swain 1978, Shope, and Neta 2009 for a survey of the variety of post-Gettier accounts of knowledge.

(EF) p can be evidence even if p is false.

More specifically, Williamson (2000) endorses the truther view of what constitutes *propositional* evidence and this version of ET can be succinctly stated as follows:

(PET) p is propositional evidence only if p is true.

Williamson is quite explicit in his commitment to this view and asserts that:

> One's total evidence in *w* can in turn be identified with the total content of what one knows in *w*.[5]

In contrast, the propositional version of the falsie view can be stated as follows:

(PEF) p can be propositional evidence even if p is false.

But, it is clear that Williamson rejects PEF and this can be seen quite easily by looking at some of Williamson's core epistemological commitments.

So, let's look at the relevant principles of Williamson's epistemology implicated in the problem to be raised here. As we have seen, Williamson defends ET in general and PET in particular. His defense of this stance is found primarily in Williamson, 2000. Let us quickly rehearse the simple derivation of Williamson's factive view of evidence. This is useful because looking carefully at this version of PET illustrates some of the main features of evidence that are at issue here that ultimately turn out to raise skeptical problems when conjoined with the safety condition. So, as part of his much discussed and radical knowledge-first epistemology, Williamson (2000) argues that one's evidence is equivalent to what one knows. Let "K_Sp" signify that S knows that p and "E_Sp" signify that p is evidence for S. This thesis can be captured simply and clearly as follows:

(W1) $E_s p \equiv K_s p$.

According to Williamson, knowledge is also the most general factive mental state operator. Thus, if a proposition is known, then it must be true. This is just the familiar and orthodox sort of factivity condition for knowledge. It can be simply stated as follows:

(W2) $K_s p \rightarrow p$.

What W1 and W2 entail is the following very interesting but perhaps implausible claim about evidence:

(W3) $E_s p \rightarrow p$.

[5] Williamson 2011, 150. See also Williamson 2000, 184-208.

W3 is then a *factivity condition for evidence* and so it is abundantly clear then that Williamson is committed to a form of ET, in fact a form of PET. He is a truther with respect to propositional evidence.

But, as has been established elsewhere, there are some compelling reasons to believe that at least some false propositions can constitute bona fide evidence.[6] One such compelling and important reason to believe that false propositions can (at least sometimes) constitute evidence is that false but approximately true measurement reports can serve as bona fide evidence.[7] This point is the crux of the argument from approximation introduced in Shaffer, 2012a, b and Shaffer, 2013 and developed more fully in Shaffer, 2015. The idea that this argument is based on is simply that reports of the results of measurements are, by their very nature, false but approximately true because of the presence of measurement errors.[8] If the view of evidence that this line of argument indicates is correct, then evidence is only *quasi-factive* in this important and specific sense and ET and EPT are false. If this line of thinking is correct, then Williamson is then simply wrong about the strict factivity of evidence. Some form of EF is true and the argument from approximation supports EF against ET (and PEF against PET). However, there are additional reasons to challenge Williamson's view of evidence that do not involve the issue of the factivity of evidence and the remainder of this paper is dedicated to the introduction of a new argument that specifically targets Williamson's views of evidence in light of his commitment to the safety condition on knowledge. So, the new argument presented here is intended to strengthen the case against Williamson's view of evidence. This new argument ultimately shows that Williamson's various epistemic views imply a strong and wildly implausible form of skepticism about measurement reports that is at odds with the common-place epistemic use of such evidence in the sciences and elsewhere.[9] Let us then turn our attention to the epistemic principle that Williamson endorses that is relevant to this argument (i.e., the safety condition).

4 Anti-luck Epistemology and the Safety and Factivity Conditions

Williamson believes that there are other important properties that states of knowledge exhibit in addition to factivity. The most important such property for the purposes of this paper is safety. The safety condition on knowledge is a necessary condition for knowing that, recently, has been most systematically defended by Williamson, Sosa, and Pritchard.[10] But it came into prominence in virtue of Nozick's post-Gettier analysis of knowledge, which was itself a reaction to earlier reliabilist accounts of knowledge and justification.[11] The safety condition is supposed to reflect the basic idea of the sort of reliability associated with bona fide knowledge that distinguishes knowledge from mere belief and lucky true belief. The gist of the idea

[6] See Shaffer 2012, Shaffer 2013 and Shaffer 2015.
[7] See Shaffer 2019.
[8] See Shaffer 2019.
[9] So, the argument also extends the result found in Shaffer 2019.
[10] See Williamson 2000, Sosa 1999, Pritchard 2005, Pritchard 2007, Pritchard 2008, and Pritchard 2009.
[11] See Nozick 1981.

is that knowledge is incompatible with a form of epistemic luck that Pritchard and Engel call *veritic luck* (i.e., knowledge defeating luck).[12] The safety condition is a modal condition, and it can be understood simply and informally as follows:

> If S knows that p, then S could not easily have falsely believed that p.

Williamson is then clear that there are essentially two different renderings of the safety condition that correspond to the ideas that knowledge is not compatible with *any* risk (i.e., the "no risk" conception) and that knowledge is compatible with some *small amount* of risk (i.e., the "small risk" conception). He explains that:

> The two conceptions disagree on whether a low but non-zero level of risk excludes or implies safety. Each conception of safety combines with a general conception of knowledge as safety from error to yield a more specific conception of knowledge. The safety conception and a 'no risk' conception of safety jointly imply a 'no risk of error' conception of knowledge. The safety conception of knowledge and a 'small risk' conception of safety jointly yield a 'small risk of error' conception of knowledge.[13]

Williamson favors the no risk of error conception of safety because the small risk of error conception of safety and knowledge is incompatible with factivity (which is a non-negotiable component of knowledge) and he believes that we ordinarily think if risk in terms of the no risk conception.[14] This yields up the idea that safety is a sort of necessity and this relatively non-technical gloss on safety and it can be made more precise as follows:[15]

$$\text{(Safety) } (w_i \vDash K_S p) \rightarrow \neg\, [\langle w_i \rangle \vDash (B_S p \;\&\; \neg p)].$$

Here '$\langle w_i \rangle$' is the set of worlds sufficiently close to w_i, '$K_S p$' represents S's knowing that p, and '$B_S p$' represents S's believing that p. So understood, the safety condition is the claim that if S knows that p at w_i, then S does not believe that p when p is false in worlds sufficiently similar to w_i. Ultimately, this regimentation captures the core idea of the safety condition well and it ties the safety condition to the anti-luck insight. Bona fide safe knowledge cannot be such that the proposition is believed and true as a matter of veritic luck.

One main issue involved in the debate about safety is determining what worlds count as close worlds and there is considerable controversy both about how to parse closeness and whether particular accounts of the factors involved in judging closeness are intuitively supported. For the purposes of this paper, this does not, however, matter. Whatever turn out to be the correct factors involved in judgments of closeness it should be clear that any such account of closeness must be reflexive, that is to say $w_i \in \langle w_i \rangle$. This is because, whatever the details involved, closeness is a similarity relation, and every world is *maximally* similar to itself.

12 See Engel 2011, Pritchard 2005 and Pritchard 2007. Williamson is keenly aware of the incompatibility of knowledge and epistemic luck, as his extensive discussion in his 2009 makes clear.

13 Williamson 2009, 10.

14 See Williamson 2009, 11 and 13.

15 This is the formalization of Willaimson's "no close risk" conception of Safety. See Williamson 2009, 10-19. See also Rabinowitz 2019.

In any case, according to those who defend this condition on knowledge, safety is supposed to have independent merit as an intuitively plausible condition on knowledge and it is supposed to provide the resources to handle Gettier cases and the like. But, it is advantageous to have a substantial argument in favor of this condition rather than having to depend on such weak and merely intuitive support for the principle and/or in light of conflicting and accounts of the closeness relation. Such an argument also makes the problems raised here for Williamson's views more acute, as rebutting them cannot be accomplished by simply disavowing commitment to the safety condition. Fortunately, there is such an argument, and it is based on Kripke's recognition that safety and factivity are intimately related. Kripke made the relevant observation that is crucial to this argument in a 1986 talk in reference to Nozick's account of knowledge. In short, the argument presented here in support of safety involves the Kripke-inspired recognition that denying safety entails denying the factivity (or veridicality) condition of knowledge. It proceeds then by showing that since we should not deny factivity, we should endorse safety.

Nozick introduced the following addition to the JTB account of knowledge as a particular form of epistemological reliabilism:

(iv) If p weren't true, S wouldn't believe that p.[16]

(iv) is, of course, Nozick's version of the safety condition. But, Kripke has pointed out that (ii) and (iii) jointly entail (i), in addition to pointing out a variety of other problems plaguing Nozick's analysis.[17] This point about the relationship between (i), (ii), and (iii) is particularly interesting because Kripke's observation can be leveraged into a substantive argument for the safety condition on knowledge that anticipates Williamson subscription to the no risk conception of safety we have been exploring.[18] This can be accomplished chiefly by considering what the denial of safety involves. What does denying safety entail? Denying safety entails this:

$$\text{(Unsafe Knowledge) } (w_i \vDash K_S p) \,\&\, [\langle w_i \rangle \vDash (B_S p \,\&\, \neg p)].$$

According to unsafe knowledge, knowing p at a given world is compatible with falsely believing p in worlds close to that given world. What then is the problem with respect to factivity? In order to see the problem, we must have a clearer understanding of factivity in hand. The factivity condition on knowledge can be simply and informally understood as follows:

If S knows that p, then p is true.[19]

As it is typically understood in epistemic logic, the factivity condition can then be parsed quasi-formally as follows:

$$\text{(Factivity) } (w_i \vDash K_S p) \rightarrow [(w_i \vDash p) \,\&\, (w_j \vDash p, \text{ for all } w_j \text{ that are accessible from } w_i)].$$

[16] Nozick 1981.
[17] Kripke 2011.
[18] See Shaffer 2017 on this argument.
[19] See Shaffer 2021 on the issue of the factivity of knowledge.

To see the important implications of factivity, consider the following basic model theory for standard epistemic logic. Let W be a set of worlds such that each $w_i \in$ W, and R be the relation of epistemic possibility relating worlds. $< W, R >$ is then a frame in the usual sense and propositions will be subsets of W such that p is true in w_i if and only if w_i p. Let $R(w_i)$ be defined as follows: $R(w_i) = \{x \in W: R\ w_i\ x\}$. p is known at w_i then if and only if p follows from $R(w_i)$. In other words, p is known at w_i if and only if p is true in all worlds that are epistemically accessible from, or are epistemic alternatives to, w_i. A world w_i is an *epistemic alternative* to world w_j for A just in case the accessibility relation holds between w_i and w_j. A bit more formally, factivity is the following condition on knowledge:

$$\text{(Factivity)}\left(w_i \vDash K_S p\right) \rightarrow R\left(w_i\right) \subseteq p.$$

Factivity holds in all frames in which the accessibility relation is reflexive. That is to say, factivity is an axiom of epistemic logic just in case w_i is accessible from itself. This is the case for all systems of epistemic logic at least as strong as the system KTD. So, it is clear then that knowledge is factive.

This is because if one simultaneously accepts factivity and unsafe knowledge then one is committed to contradiction. This will be the case if there is at least one world where p is false that is close to a given world where p is known that is also an epistemic alternative to that world, and there is always at least one such world.[20] Consider a given proposition p known at w_1 and the definition of unsafe knowledge. Since the notion of closeness involved in the safety condition is reflexive, if p is known at w_1, then it can be the case that p is false at w_1. Why? This is simply because unsafe knowledge permits an agent to have knowledge of a proposition in a given world w_1 even when the agent falsely believes the proposition in worlds that are close to w_1. But, since closeness is reflexive, w_1 is itself one of those close worlds. So, unsafe knowledge permits an agent to know in w_1 even when the agent falsely believes the proposition in question in w_1. However, by factivity and the reflexivity of the epistemically access relation, if p is known at w_1, it also follows that p is true at w_1, since w_1 is a member of the set of worlds that are epistemically accessible from w_1. So, jointly endorsing unsafe knowledge and factivity leads to contradictions and one must go. But, since factivity is such a deeply entrenched and orthodox condition on knowledge and its denial invites all sorts of Morrean-like worries about false knowledge claims of the form "I know that p, but ¬p", we should simply treat Kripke's observation about conditions (i), (ii) and (iii) as a reductio of the denial of safety and thereby as a substantive argument in favor of safety. In other words, since such Moorean "knowledge" claims clearly involve contradictions and are infellicitous, we should maintain factivity and reject the denial of safety.[21] What Kripke's recognition allows us to see then is that arguments that support factivity are, ipso facto, arguments that support safety. Williamson is also keenly aware that the "small risk of error" conception of safety is incompatible with factivity and so he ultimately commits to the "no close risk of error" version of the "no risk of error"

20 There will actually be many such worlds.
21 See Shaffer 2017.

conception of safety.[22] So, in light of the unavoidability of safety, the retention of (strict) factivity, and the concomitant commitment to the "no close risk" rendering of safety, what do Williamson's various claims about knowledge, factivity, and evidence amount to?

5 Epistemic Luck and Factive Safe Evidence

We have seen then that Williamson is committed to E = K, the factivity of knowledge, and the safety of knowledge (understood in terms of the "no close risk" model). This set of commitments jointly implies that anything that is bona fide evidence is not only true, but also safe and known. This means then that a believed proposition can be bona fide evidence only if it is true in the actual world and all worlds where it is also believed that are close to the actual world. Thus, Williamson is committed to the following view of evidence:

> (PETS) with respect to $B_s p$, p is propositional evidence only if p is true and $B_s p$ is safe.

PETS is simply a consequence of E = K and Williamson's conception of knowledge. However, it does not take much to see that this view of evidence is implausibly strong. First, it is important to acknowledge that Williamson holds that epistemic probabilities are evidential probabilities. As he explains,

> The evidential probability of a proposition in a world *w* is identified with its probability conditional on one's total evidence in *w*. One's total evidence in *w* can in turn be identified with the total content of what one knows in *w*.[23]

Similarly, he tells us that,

> I have argued that the sort of probability most relevant to the epistemology of science is probability on the evidence, and that the evidence is simply what is known.[24]

Moreover, Williamson notes that the following epistemic principle is a plausible component of standard theories of evidence:

> (HC) One knows something only if it has a high chance.[25]

Here chances are supposed to be objective probabilities and not subjective probabilities or evidential probabilities.[26] But, Williamson argues that HC is ultimately

22 Williamson 2009, 11.
23 Williamson 2011, 150. See also Williamson 2000, 184-208.
24 Williamson, 2009, 2
25 Williamson 2009, 6.
26 Willaimson 2009, 4-5.

incompatible with anti-skepticism and his favored epistemic closure principle (MPC).[27] Ultimately, he rejects HC and simultaneously endorses the "no close risk" conception of safety and factivity.

So, Williamson recognizes that knowledge, evidence, and probability are intimately related in terms of the common view of knowledge. But, this relationship is tricky and complex. Importantly, he appears to be committed to the views that all evidence is known, that known propositions must be safe, and that known propositions have high probabilities conditional on one's total evidence (i.e., high conditional evidential probabilities). As we shall see, this combination of views raises certain problems for Williamson (some of which he explicitly acknowledges).[28] At this juncture, the most important things to note is that he is committed to the view that knowledge does not require high chance, but it does require safety of the no risk sort.

In any case, evidence is generally characterized in terms of the following fundamental probabilistic notions:

$$\text{(CON)p is confirming evidence for q, if and only if, } P(q|p) > P(q).$$

$$\text{(DCON)p is disconfirming evidence for q, if and only if, } P(q|p) < P(q).$$

These principles are simply consequences of the probability calculus and the evidentiary value of p with respect to q is given by Bayes' Theorem:

$$\text{(BT)}P(q|p) = P(q)P(p|q)/P(p).$$

Notice then that evidence p enters into BT as the denominator on the right side of the equation and BT requires that P(p) is non-zero (for obvious reasons). Moreover, CON and DCON jointly imply that p is evidence with respect to q only if $P(q \mid p) \neq P(q)$. Also, note that reasonable conditions for p's being good evidence for q are that p confers a higher probability on q and that p itself is highly probable. This latter condition is the case because evidence is knowledge (i.e., $E = K$) and a proposition must—in addition to meeting various other conditions—itself be highly probable conditional on one's total knowledge to itself be knowledge/evidence. As we have seen already, given PETS, Williamson is committed to the idea that p can be evidence only if p is safe and true. But, we can further characterize Williamson's view of evidence in terms of these probabilistic notions. For Williamson, p is evidence for q only if p is safe, p is true, $P(p) \neq 0$ and $P(q \mid p) \neq P(q)$. As we shall see subsequently, this makes for an excessively strong notion of what can and cannot be evidence. Crucially, it importantly implies that good evidence cannot be such that it is known in the actual world on the basis of sufficient evidence, but believed and false in close possible worlds. But, if this is the case, then most measurements cannot be evidence.

27 Williamson 2009.

28 Williamson 2009, 19.

6 The Measurement Problem for Safe and Factive Evidence

First, let us note that Williamson relativizes his notion of safety to the use of specific methods, but, for the most part, we will put that matter to the side in what follows for the sake of simplicity. But, whether a belief that p is safe or not has to do with what belief-forming mechanism gives rise to the belief that p. This is informative and it makes it much easier to discuss the sorts of cases to be raised here that are problematic for Williamson's views on evidence. Strictly speaking, Williamson holds that with respect to belief-forming mechanism M and where both wi $\vDash B_S p$ on the basis of M and <wi> $\vDash B_S p$ on the basis of M, $B_S p$ is safe at w_i iff $\neg$ [<w_i> $\vDash$ ($B_S p$ on the basis of M & $\neg$p)]. Let us then explore what can/cannot count as evidence for Williamson in light of his particular epistemological views, especially how this raises problems concerning measurements and evidence.

In the typical operations of the developed empirical sciences, we often deal with exceptionally accurate scientific measurements based on various mechanisms for measuring quantities of scientific interest. In light of various uncertainties that afflict measurement, there is a fixed degree of accuracy associated with all measurements and no measurement method is without some such degree of error. This aspect of measurement (i.e., error) is the focus of measurement theory.[29] However, there are, in fact, two sorts of measurement errors: systemic errors and random errors.[30] Random errors are the results of episodic variations in the conditions present during specific instances of measurement and they can be treated statistically so as to maximize reliability. Systemic errors, on the other hand, are pervasive errors that are not episodic results of variation in measurement conditions from measurement to measurement and they cannot be dealt with in terms of statistical methods. In fact, there is no standard way to deal with systemic measurement error other than to know what sort of physical influences affect a given measurement apparatus so that we can correct for them by controls, shielding, and data adjustments based on known systemic influences.[31] But, in real cases, we never are in a position to identify, let alone deal with, all (physically) possible systemic errors. There are simply too many such factors (e.g., gravitational effects, EMF interference, vibrations) that might influence singular or repeated measurements. Let us then examine a bit more closely how measurements and the associated errors are typically dealt with in measurement theory and why the possibility of unknown systemic errors turns out to be such a problem for Williamson's theory of evidence and knowledge.

The standard measurement theoretic approach to dealing with the sorts of errors associated with *singular* measurements is to repeat the measurement operation in question and to generate a statistical distribution of measurement results that has an associated margin of error. This is a complex procedure, and it involves a variety of important steps, but they are designed to increase the reliability of measurement evidence. Let us then consider a case where we are dealing with a situation where a

[29] See Taylor 1997.

[30] See Taylor 1997.

[31] See Taylor 1997, chapter 4.

measurement is not repeatable, and we have only a single measurement result. Say, for example, that we are measuring the speed s of a car driving down the highway with a radar gun and we do not have multiple sources of measurement simultaneously applicable to that task. Also, suppose that the speed limit on this stretch of highway is set at 70 mph. In this case, supposed that, using the radar gun, we determine that the value of s, $v(s) = 70$ mph. In this case, the measured value of s is true at the real world and seems to be (good) evidence that the car is not violating the speed limit on that stretch of road. Now, in measurement circumstances such as this one, we would normally want to determine some margin of error associated with measurements of speed using the radar gun. Suppose then that we determine that the margin of error for measurements using the radar gun is ± 0.001 mph. The real measurement report of the value of s should include this information and should really be stated as 70 ± 0.001 mph. But, the only way to determine the margin of error here is by a priori estimation, by the collection of other measurements of independently known speeds of objects using the radar gun, or by comparing the measured results of the radar gun with another more accurate source of measurement of speed (i.e., by calibration).[32] This cannot always be done in cases of singular measurement. However, we might find that in previous measurements using the gun, or by comparison with another source that the associated errors were -0.0004, $+0.0003$, -0.0002, etc. Each of these measurements itself is subject to error and so there is always an associated and ineliminable margin of error in any measurement. All of this indicates that, where possible, it is preferrable to repeat measurements of any quantity using the same measurement apparatus in order to manage error possibilities.

This is because even if the radar gun determines the value of s to be 70 mph on the basis of a singular measurement and it in fact is the correct value, this result is unsafe. It is always possible that the agent would believe this to be the case in close possible worlds to the actual world where $v(s) \neq 70$ mph. Say in one of those non-actual worlds the real value of s is 69.99999 mph, but there is an unknown factor that makes the reading of the value of s on the radar gun 70 mph. There will then be any number of possible but non-actual scenarios where the measurer believes that $v(s) = 70$ mph, but where that is false due to the result of some such influence that has not been ruled out. So, the agent who believes truly that $v(s) = 70$ mph in the actual world does so only as a matter of veritic luck. Despite the intrinsic uncertainty of the measurement, the agent just happens to have formulated a true belief about the speed of the car. But, she could easily have had a false belief. So, again, we find that the measured value of the speed of the car in this case, though true, is not evidence. This is because that belief is not safe, and, as we have seen, according to Williamson evidence has to be safe as well as true. But, this seems to be an absurdly restrictive concept of evidence, for the agent's measurement of the speed of the car is clearly evidence that the car is not exceeding the speed limit.

Let us then consider a case where we are dealing with a situation where a measurement is repeatable, and we have an array of results from making such repeated measurements using the same mechanism. For example, consider measuring the voltage of a battery using a voltmeter. In such cases, we can minimize measurement error by making repeated measurements and applying statistical methods to generate a distribution of the

32 See Taylor 1997.

results. This also makes explicit the error associated with the particular measurement mechanism. The standard practice in such cases is to employ a statistical distribution of the measurement results in order to correct for non-systematic measurement errors afflicting individual cases of measurement.[33] Thus, when we are dealing with random errors in measurement and we can repeat measurements, there are methods to increase reliability of the measurement. One might think that this would secure the possibility of having safe beliefs about such statistically conditioned evidence by eliminating random errors in the individual measurements. But, this turns out not to be the case.

The real problem arises in all these cases when we begin to look at them and not that safety is intrinsically modal in nature. So, if a singular measurement or an ensemble of such measurements is to serve as evidence in terms of Williamson's views it must have a high evidential probability, it must be true at the actual world and safe at the actual world. As pointed out in Shaffer, 2012a, b, Shaffer, 2013 and Shaffer, 2015, it is not really clear at all that evidence must be true at the actual world and the truther view of evidence seems to be false. But let us put that aside and focus on the issue of safety here.[34] If a measurement or ensemble of measurements is safe, then it is true at the actual world and it must not be the case that there are close worlds where that evidence is believed and false. However, this condition will virtually never be met in real-world cases of measurement. There is always the possibility that there are unknown systemic errors afflicting any measurement or ensemble of measurements and while in some situations we can control for known influences, systemic or otherwise, it is not possible to control for *unknown* systemic errors even with statistical methods.[35] So, it appears to be the case that we cannot know the results of measurement given this account of evidence. If p is known at the actual world on the basis of M and p is evidence at the actual world, then there cannot be *any* close possible world where p is falsely believed on the basis of M. This is just a consequence of the no risk conception of safety.

Since we cannot rule out the possibilities of unknown systemic errors that impact p in close possible worlds, the belief that p *in the actual world* cannot be safe and so cannot be known there. The only way this could be effectively accomplished would be by first by making repeated measurements and ruling out all random errors using statistical methods, and then by identifying all physically possible sources of systemic error and controlling for them. But, we are rarely, if ever, in a position to do these things and yet we often use all sorts of singular and ensembles of measurements as evidence. In fact, we often use singular measurements as evidence and often do not even attempt to use statistical methods to deal with random errors. More importantly, even if we can repeat measurements and use the relevant statistical methods, we cannot deal with the possibilities of systemic errors, especially unknown ones. We are typically not even certain of (all of) the relevant physical possibilities and/or physical laws involved in a given measurement setup. So, even if we grant the issue of the factivity of evidence and the ability to repeat measurements and use statistical methods to deal with random errors, Williamson's views on evidence imply that virtually all measurement reports still cannot be evidence. This is because they are unsafe *even if they are repeated, statistically conditioned, true, and highly probable on the evidence*.

[33] See Taylor 1997.
[34] See Shaffer 2012, Shaffer 2013 and Shaffer 2015.
[35] See Taylor 1997, chapter 4.

7 Conclusion

The upshot of all of this is that Williamson's view of evidence as both factive and safe (whatever one says about the high chance requirement) implies that virtually all measurement reports cannot be evidence. This is because such reports are often just approximately true (and so strictly false) and even if such reports are true, they are typically not safe. This is then a far too strict view of evidence and these faults also show also that the E = K thesis is false. While knowledge *does* imply truth and safety, it appears that this is not the case for evidence. So, E cannot be equivalent to K.

References

Gettier, E. (1963). Is justified true belief knowledge? *Analysis, 23*, 121–123.

Kripke, S. (2011). "Nozick on knowledge," in *Saul Kripke: Collected Papers vol. 1* (Oxford: Oxford University Press, 162–224.

Littlejohn, C. (2013). No evidence is false. *Acta Analytica, 28*, 145–159.

Mylan Engel, M. (2011). "Epistemic luck," *The Internet Encyclopedia of Philosophy*.

Neta, R. (2009). Defeating the dogma of defeasibility. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 161–182). Oxford University Press.

R. Nozick (1981). *Philosophical explanations* (Cambridge: Harvard University Press).

Pappas, G. and M. Swain (1978). *Essays on knowledge and justification* (Ithaca: Cornell University Press).

Pritchard, D. (2007). Anti-luck epistemology. *Synthese, 158*, 277–298.

Pritchard, D. (2008). Knowledge, luck, and lotteries. In V. Hendricks & D. Pritchard (Eds.), *New waves in epistemology* (pp. 28–51). Palgrave Macmillan.

Pritchard, D. (2009). Safety-based epistemology: Whither now? *Journal of Philosophical Research, 34*, 33–45.

Pritchard, D. (2005). *Epistemic luck* (Oxford: Oxford University Press).

Rabinowitz, D. (2019). "The safety condition for knowledge," *The Internet Encyclopedia of Philosophy*.

Shaffer, M. (2012a). Not-exact-truths, pragmatic encroachment and the epistemic norm of practical reasoning. *Logos & Episteme, 3*, 239–259.

Shaffer, M. (2012b). Moorean sentences and the norm of assertion. *Logos & Episteme, 3*, 653–658.

Shaffer, M. (2013). E does not equal K. *The Reasoner, 7*, 30–31.

Shaffer, M. (2015). Approximate truth, quasi-factivity and evidence. *Acta Analytica, 30*, 249–266.

Shaffer, M. (2017). An argument for the safety condition. *Logos & Episteme, 8*, 517–520.

Shaffer, M. (2019). Rescuing the assertability of measurement reports. *Acta Analytica, 34*, 39–51.

Shaffer, M. (2021). Can knowledge really be non-factive? *Logos & Episteme, 12*, 215–226.

Shope, R. (1983). *The analysis of knowing* (Princeton: Princeton University Press).

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives, 13*, 141–54.

Taylor, J. (1997). *An introduction to error analysis 2nd ed.* (Sausalito: University Science Books).

Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy, 65*, 157–170.

Williamson, T. (2000). *Knowledge and its limits* (Oxford: Oxford University Press).

Williamson, T. (2009). "Probability and danger," *Amherst Lecture in Philosophy*.

Williamson, T. (2011). "Improbable knowing," in T. Daugherty (ed.), *Evidentialism HYPERLINK "*https://philpapers.org/rec/DOUEAI*" and its Discontents* (Oxford, Oxford University Press), 147–164.