King's Research Portal

DOI:

[10.1111/phpr.12311](https://doi.org/10.1111/phpr.12311)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Bird, A. (2018). Evidence and Inference. *PHILOSOPHY AND PHENOMENOLOGICAL RESEARCH*, 96(2), 299-317. <https://doi.org/10.1111/phpr.12311>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

EVIDENCE AND INFERENCE

Alexander Bird
University of Bristol

Abstract

I articulate a functional characterisation of the concept of evidence, according to which evidence is that which allows us to make inferences that extend our knowledge. This entails Williamson's equation of knowledge with evidence.

1 Introduction

What is evidence? As is well known, Timothy Williamson (1997, 2000) argues that S's evidence is precisely what S knows. While we may regard this as a substantive and necessarily true identity claim, akin to 'water is H₂O', we should note that Williamson's argument is an *a priori* one. If it is *a priori* that evidence and knowledge are one and the same, why do we have a concept of evidence at all?

I propose that we answer our initial question by posing another. What is the purpose of evidence? Whereas knowledge may be intrinsically valuable, we do seek evidence not for its own sake. Rather we want evidence because evidence will allow us to draw conclusions in the course of enquiry. It is evidence that allows us to infer that a theory is true or to know that a defendant is guilty. So evidence is sought and is used for the purpose of making inferences. Indeed, most of the philosophical discussion of evidence concerns this feature: above all we ask when will evidence support an inference to a conclusion; the term 'evidence' is most naturally used in the context of a proposition's being 'evidence *for*' some proposition (or sometimes 'evidence against' and other cognates). The significance of evidence for the task of inference suggests an answer to our question, why do we have a distinct concept of evidence: we have a concept of evidence in order to characterise certain propositions *functionally*. Put in general terms, the concept of evidence serves to characterize propositions in terms of their role in inference.

This paper looks at the relationship between evidence and inference in detail in order to better understand what evidence is. I have two objectives:

- (a) To articulate a functional characterisation of the concept of evidence: evidence is that which allows us to make inferences that extend our knowledge.
- (b) To provide an indirect argument for my characterisation by showing that it entails Williamson's claim that all and only knowledge is evidence, for which he provides independent argument.

2 The concept of evidence characterised

The characterisation of the concept of evidence I propose is this:

(EC) p is in S 's evidence if and only S can gain knowledge by inference from p .

(EC) captures the intuition that evidence is what we infer from in order to get knowledge; it acknowledges an asymmetry between evidence and what evidence is evidence for. Yet the asymmetry afforded by (EC) is local not global. Given some particular inference, it is the 'what is inferred from' that is the evidence. (EC) is consistent with its being the case that the knowledge inferred from evidence itself thereby becomes evidence, evidence that can be used as the premise of some further inference.

(EC) is the conjunction of the two implications:

(EC $\rightarrow$) if p is in S 's evidence then S can gain knowledge by inference from p ;

and:

(EC $\leftarrow$) if S can gain knowledge by inference from p then p is in S 's evidence.

I shall go on to give reasons why (EC $\rightarrow$) and (EC $\leftarrow$) provide accurate characterisations of the concept of evidence. But first some clarifications.

In the above p must play a *non-redundant* role in the relevant inferences. Take any knowledge-extending deductive inference: we can include p as a premise in that argument, and indeed we can construct an argument in which p might appear to play a role (e.g. using and-introduction followed by and-elimination). In such cases p would be redundant, and (EC $\leftarrow$) would not license inferring that p is in S 's evidence. The concept of redundancy also extends to the components of p if p is complex, e.g. a conjunction: the proposition $(r \wedge s)$ would not be part of one's evidence just because one can gain knowledge by inference from r .¹

We may consider a simple token inference to be composed of three elements: (i) a set of premises, (ii) an inference procedure, and (iii) a conclusion. Extended inferences are concatenations of simple inferences. What we may call the *epistemic quality* of a subject's belief in a proposition p concerns whether subject knows p , or has a justified belief in p that falls short of knowledge, or unjustifiably believes p . Let us take the case where a subject's belief in the conclusion comes about solely as a result of the inference. In that case the epistemic quality of the belief will depend on epistemic characteristics of the premises and of the inference procedure. The relationship between the epistemic characteristics of the premises and the epistemic quality of the conclusion will be the focus of this essay and will be discussed at length. Here I shall consider the epistemic characteristics of the inference procedure so that we can set it aside.

The epistemic characteristics of the inference procedure will depend on a number of factors. Consider, for example, a subject carrying out a deductive inference. The quality of the procedure will depend on formal features of the deduction—e.g. whether it is valid. It will also depend on factors concerning the subject, e.g. whether she is intellectually capable of entertaining the premises and conclusion, and that

¹I do not think that the notion of redundancy can be formalised. For example one might think that p is a redundant premise in the set $\{\Gamma \cup p\}$ in a derivation of q if q can be derived from Γ alone (and $\Gamma \not\vdash p$). But it might be that any derivation from Γ alone is very different from the former (all such proofs might be much longer). In such a case p would not be redundant in the sense relevant to the above principles, and should be included amongst S 's evidence for q .

she does not make errors in implementing the inference rule. If the inference procedure is non-deductive, its quality will also depend on relevant features of the world. Let us call an inference procedure *faultless* iff it suffers from no shortcomings in its epistemic characteristics. A simple deductive procedure carried out carefully by a competent logician will typically be faultless. Likewise, an electrician who infers from the movement of the needle in a properly functioning ammeter that the battery to which it is connected is not dead will typically be employing a faultless inference procedure. A faultless inference procedure will be one that leads from true premises to true conclusions. If one uses the procedure and reaches a false conclusion from true premises because of the procedure, then the procedure suffers from an epistemic shortcoming. This means that the properties of being faulty and of being faultless are strongly externalist epistemic properties: a procedure might be faulty simply in virtue of the way the world is: another procedure that is identical as regards the subject may be faultless. While I do not claim to have characterised 'faultless' with perfect precision, the notion does capture the intuitive idea that when an inference leads us to a false conclusion, we may sometimes regard the inference procedure, construed broadly so as to include external features of the world, to be at fault. In this essay, we shall ignore such possibilities of epistemic fault in the inference procedure, in order to focus attention on the epistemic quality of the premises in an inference.

While the epistemic characteristics of the premises and of the inference procedure will be important in determining whether an inference gives the subject new knowledge, they are not the only relevant factors. S might already know the conclusion to be true, so S cannot gain any knowledge by the inference. Or S might have strong evidence against q , so that it would be irrational to believe q despite this new inference. We may regard a subject as *receptive* if she has no thoughts relevant to q other than those employed in this inference.

We can now explicate the proposition (C) 'S can gain knowledge by inference from p ' as it is used in (EC), (EC $\rightarrow$), and (EC $\leftarrow$). (C) is true precisely when it is possible for a receptive subject, S, as a consequence of some faultless inference procedure from p , a non-redundant premise, to come to know the proposition that is the conclusion of the inference.

3 All evidence supports knowledge-producing inferences

In this section I shall provide the argument for:

(EC $\rightarrow$) if p is in S's evidence then S can gain knowledge by inference from p .

The idea behind (EC $\rightarrow$) is that evidence has a sufficiently high-grade epistemic quality that a subject is capable of gaining knowledge by inference from it. (EC $\rightarrow$) contrasts with the claim that evidence might be epistemically low-grade and might be incapable of supporting knowledge-producing inferences. Hence the clarification of (EC $\rightarrow$) in Section 2 excluded purely accidental reasons why a subject might fail to produce knowledge from her evidence. A counterexample to (EC $\rightarrow$) must be one where a failure to produce knowledge must be attributable to the epistemic weakness of the evidence.

Thus (EC $\rightarrow$) would be false if either of the following were true:

(A) A receptive subject S infers q from evidence p by a faultless procedure, but S 's resulting belief, q , is not justified.

(B) A receptive subject S infers q from evidence p by a faultless procedure, but S 's resulting belief, q , is not true.

That (A) is false is clear. It is at the very least the function of evidence to justify our hypotheses. If a subject uses a faultless inference procedure to gain a new belief from her evidence, then the new belief is justified.

The falsity of (B), although less immediately apparent, follows straightforwardly if all evidence is true, since all propositions inferred by a faultless procedure from true premises will be true, as was discussed above. One might wonder, however, whether this claim about faultlessness is too strong. Might it not be sufficient for faultlessness that it justify beliefs, without the additional requirement that the procedure is perfectly reliable? For example, the inference from 'Jane has only one ticket in a ten million ticket fair lottery' to 'Jane will not win the lottery' might be thought to employ a procedure that is not perfectly reliable but nonetheless justifies its conclusion. Should not such a procedure count as faultless in our sense? In the case where Jane does win the lottery, may not the falsity of the conclusion be regarded as epistemic bad luck external to the inference as opposed to a result of an epistemic fault of inference? This is an issue that has been much discussed and cannot be settled here. It is my view (Bird 2004) as well as that of Nelkin (2000), Sutton (2007), et al. that in lottery cases, a belief (on this evidence) that Jane will not win is not justified. Such an inference is faulty because its procedure is imperfectly reliable.

So if the evidence propositions are true, a faultless inference procedure will lead to a true conclusion. But are all evidence propositions true? A central function of evidence is the testing of hypotheses. In particular it is a requirement on any hypothesis that to be credible it must be consistent with the evidence. Dick Heuer (1999: 47), an intelligence analyst, states 'A hypothesis may be disproved ... by citing a single item of evidence that is incompatible with it.' A century and a half earlier, the legal theorist Thomas Starkie (1834: 444) tells us 'It is essential *that all the facts should be consistent with the hypothesis*. For as all things which have happened were necessarily congruous and consistent, it follows, that if any one established fact be wholly irreconcilable with the hypothesis, the latter cannot be true. Such an incongruity and inconsistency is sufficient to negative the hypothesis, even although it coincide and agree with all the other facts and circumstances of the case to the minutest extent' [emphasis as in the original]. Although Starkie uses the term 'facts', it is clear from the context that he is talking about evidence; indeed this passage comes from a section entitled 'Circumstantial evidence' in his book *A Practical Treatise on the Law of Evidence*.

That all evidence is true is the best explanation of the requirement on any hypothesis that to be credible it must be consistent with the evidence. If evidence must be true, then this requirement excludes a set of propositions all of which must be false. The requirement that hypotheses be consistent with the evidence would lose its force if false propositions were among our evidence, for then perfectly true propositions would be excluded as being inconsistent with our evidence.

A response to this argument suggests that the injunction that a hypothesis must be consistent with the evidence is derived from the prohibition on believing contradictions. If one permitted hypotheses inconsistent with the evidence one would be permitting contradictions. But this response is inadequate, since the requirement of consistency with evidence is stronger than the prohibition on contradictions. If

one is inclined to believe both of two propositions that contradict one another, the prohibition on contradictions means that one cannot fully believe both. In order to fully believe either of the propositions, one will have to reject the other. The prohibition on contradictions does not say which to reject. It may be that further evidence is required in order to decide between them. The requirement of consistency with the evidence is stronger since it does tell us which of the two contradictory propositions to reject. If a hypothesis and the evidence contradict one another, then the hypothesis is to be rejected. No further evidence is required. Of course, one may not always know what one's evidence is. In which case one might be in the position of doubting that one's evidence is indeed evidence when it conflicts with a favourite, well-confirmed hypothesis. But it remains that case that if one did know in such a case that one's supposed evidence is indeed one's evidence, then one would be able to reject the hypothesis as false without further ado. This makes sense only if one's evidence is true. If one's evidence is true, then what is reliably inferred from it is also true. Hence (B) is false.

(That evidence is true makes it apparent that evidence must be conceived of in an externalist fashion. I shall return to this in more detail in Section 6. For the moment I shall reiterate that a strong conception of evidence is forced on us by the principle 'reject a hypothesis inconsistent with the evidence'. For example, we do not accept the principle 'reject a hypothesis inconsistent with what you believe'—for it may be appropriate to drop the belief instead. Hence evidence must be more than mere belief.)

Since (A) and (B) are false, we may conclude that evidence must be capable of supporting inferences that yield beliefs that are both true and justified. We know that not all true justified beliefs are knowledge. But it would be very odd if evidence could support inferences to true justified beliefs but not to knowledge. Furthermore, Gettier cases typically arise because of reliance on justified but false lemmas. But the falsity of (B) means that we can rule out false lemmas. So there is no reason to suppose that evidence is capable of supporting inferences producing justified true beliefs without being able to support knowledge-producing inferences. We may conclude that evidence is indeed capable of supporting knowledge-producing inferences.

4 Propositions supporting knowledge-producing inferences are evidence

In this section I shall discuss:

(EC $\leftarrow$) if S can gain knowledge by inference from p then p is in S's evidence.

While (EC $\leftarrow$) is a natural counterpart of (EC $\rightarrow$) neither implies the other. (EC $\leftarrow$) can be true without requiring (EC $\rightarrow$). One could deny that evidence must be capable of supporting knowledge-producing inferences, yet agree that if some inference does produce knowledge, its premises must be among one's evidence.

(EC $\leftarrow$) asserts the epistemological significance of evidence. If (EC $\leftarrow$) were false, then there are propositions that are not evidence but from which one can nonetheless make knowledge-producing inferences. Thus to deny (EC $\leftarrow$) would be to deny the place given to evidence in epistemology. If (EC $\leftarrow$) were false, it would be possible to gain knowledge by inferring from propositions not among our evidence, in

which case there would be no reason for us to care about which propositions are among our evidence. A criticism of S's claim to know q may be made thus, "Proposition q is not supported by your evidence!". The falsity of $(EC \leftarrow)$ would permit S to reply, "True, q is not supported by my evidence. But I did infer it from non-evidence propositions that support knowledge-generating inferences."

One might reject $(EC \leftarrow)$ while still giving evidence a key role. This role would be foundational. An inferred proposition may be thought of as standing at the near end of a chain (or tree) of inferences. For the inferred proposition to be known the chain of inference must start from evidence propositions. But the intermediate propositions need not count as evidence also (on this view). An adjustment to $(EC \leftarrow)$ that encompasses this view would be:

$(EC \leftarrow)'$ if S can gain knowledge by inference from p and p is not itself known by inference then p is in S's evidence.

According to the view of evidence just considered, inferred propositions cannot count as evidence, whereas non-inferred propositions, such as propositions known by direct perception, may so count. It should be noted that on this view, what counts as evidence may be very vague. The theory-ladenness of perception means that it will be unclear whether or not a known proposition should count as inferred or not. For example, it seems as if the fact that a tomato is ripe can be known by direct perception. But for some people, e.g. children, the fact that the difference between ripe and unripe tomatoes is marked by a difference in colour is a fact that is learned, and so their knowledge that the tomato is ripe is inferred. As they become accustomed to making this inference, it passes from a conscious to an unconscious process, and until it would be inappropriate to call it an inference at all. So knowledge that the tomato is ripe will be something that is non-inferential at one stage but is inferential at a later stage. That some theory of a concept entails that it is vague is no objection to that theory on its own. But this vagueness does suggest that evidence could not play a central role in epistemology. Does it make a deep difference that in one case the proposition that the tomato is ripe is inferred from its perceived redness (and so is not evidence, according to $(EC \leftarrow)'$) and in another case its ripeness is directly perceived (in which case the proposition is evidence)?

More of a problem for $(EC \leftarrow)'$ is the fact that evidence propositions can be forgotten. Let q be inferred from a proposition p that is evidence by the lights of $(EC \leftarrow)'$. Then r is inferred from q . We may imagine that all the propositions are known. Let it be the case that between the first inference and the second inference the subject forgets the proposition p , so when inferring r she has no longer has any knowledge of p . According to $(EC \leftarrow)'$ she has no evidence for r : her evidence cannot be p , since she has forgotten p ; nor can it be q , since q is inferred and is thus not knowledge. That seems distinctly odd, since this subject certainly has a good reason for her belief in r (viz. her knowledge of q), and it again undermines the idea that evidence is epistemically significant. As I (Bird 2004) argue, this and other cases show that the restriction of evidence to non-inferential beliefs is implausible.²

To sum up the argument for $(EC \leftarrow)$: evidence is epistemologically significant; inferences depending on premises that are not part of our evidence are flawed; in particular they don't lead to knowledge. If that were wrong, i.e. $(EC \leftarrow)$ is false, then it becomes unclear why we care about evidence. If something else other than evidence, call it 'schmevidence', could lead to knowledge, then it might be perfectly

²Strictly, I argue against the idea that evidence is restricted to non-inferential *knowledge*. Those arguments may be extended to the proposal that evidence is a any subset of non-inferred propositions.

alright to ignore the evidence so long as one had enough schmevidence, since that could be sufficient for us to expand our knowledge. But we don't think that. The rhetorical force of terms such as 'Evidence-Based Medicine' and 'Evidence-Based Policy' is precisely that it is absurd to suggest that it is satisfactory to form beliefs on the basis of anything other than evidence.³ I considered the response that we might care about evidence nonetheless, if it is the *ultimate* source of premises for our inferences, i.e. if evidence is not itself inferred. This proposal ignores the fact that we generally do not remember the non-inferred propositions from which our inferred knowledge ultimately derives. Nor is that failure to remember to be regarded as a mere wrinkle, a departure from the ideal depicted by the theory. For our memories are designed *not* to store such information. Psychologists distinguish three types of memory: perceptual memory, short-term memory, and long-term memory. What we learn non-inferentially, e.g. in perception, will enter the first and often the second of these kinds of memory (whose duration is short) but only enter the third if particularly significant. Long-term memory will store information we are likely to use again, typically information we infer from what we perceive (or from other knowledge) rather than the perceptual experiences themselves. Given the way that memory works, evidence could not fulfil its epistemic role if it were restricted to non-inferential knowledge—we just do not have enough non-inferential knowledge to support, even in an ancestral way, the knowledge-generating inferences we make. As I shall explain in Section 7, scientists do not use 'evidence' to refer only to non-inferential beliefs or knowledge; for example, proponents of Evidence-Based Medicine are not asking us to base clinical decisions on what we know through perception but rather on the results of properly conducted randomized trials or meta-analyses of such studies—such evidence is the outcome of sophisticated statistical inference, doubly so in the cases of meta-analysis.⁴

5 Evidence and knowledge

The preceding sections articulated and defended the view, (EC), that p is in S 's evidence if and only S can gain knowledge by inference from p . That completes my first task, to provide a functional articulation of the concept of knowledge. In the light of that account, we may ask: what must something be like in order to fulfil that function? I shall defend the claim that all and only knowledge can fulfil this function. So we may conclude that all and only knowledge is evidence. That is a conclusion already reached by Williamson (1997). Since Williamson presented independent arguments for his conclusion, we may regard his arguments as providing indirect support for (EC).

I claim that if $(EC \rightarrow)$ is true, then evidence itself cannot be less than knowledge:

$(E \rightarrow K)$ if p is in S 's evidence then S knows p .

The principle 'No False Lemmas' states that an inference that depends non-redundantly on a false premise cannot generate knowledge (Harman 1965: 92). The principle can be extended. If an inference depends non-redundantly on a premise

³Which is not to say that the claims of, for example, the Evidence-Based Medicine movement are trivial, because their claims also involve contentious assertions about what counts as good evidence in medicine.

⁴A meta-analysis provides a statistical analysis of a set of several first-order studies thereby aiming for a more robust result than any individual study provides.

that is unjustified, then the conclusion of the inference will fail to be knowledge. (This will be so whether justification is conceived of externally, for example in terms of reliability, or internally, for example in terms of rules of rational belief-formation.) When an inference depends on several premises non-redundantly, then a subject is justified in believing the conclusion of the inference only if she is justified in believing all the premises (assuming that the subject has no other reason for believing the conclusion). So in addition to the No False Lemmas principle we may state a No Unjustified Lemmas principle. Consequently, for an inference to generate knowledge, its essential premises must be both true and justified. If the old JTB account of knowledge were correct, then we could simply infer from the conjunction of the two principles (No False Lemmas and No Unjustified Lemmas), a third principle, No Unknown Lemmas:

(NUL) S can gain knowledge by non-redundant inference from p only if
S knows p .

As it happens a Gettier belief (one that is true and justified, but does not amount to knowledge) could satisfy No False Lemmas and No Unjustified Lemmas without satisfying No Unknown Lemmas. Nonetheless, the best explanation of why No False Lemmas and No Unjustified Lemmas are correct principles is that they follow from the more general principle, No Unknown Lemmas. There is nothing special about true justified beliefs. If p is a Gettier belief, then we should not expect an inference that depends on p to be knowledge generating. Typically p is a proposition that is accidentally true, e.g. $p = q \vee r$, where q is false but justified but r is true but unjustified. But if p is accidentally true its truth cannot make any difference to the epistemic quality of what is inferred from p . But since propositions inferred from false lemmas are not knowledge, neither then are propositions inferred from accidentally true propositions. In the kind of Gettier case just mentioned, an inference from p that yields a justified conclusion will depend essentially on the contribution of the disjunct q of p (since only the disjunct q is justified). But in that case the inference depends on a false lemma, and the conclusion is not knowledge. Hence an inference from a premise which is a Gettier case of a justified true belief that is not knowledge will not yield knowledge. Nothing short of a known premise will yield knowledge. (In Section 6 I consider an objection to this claim.)

Having argued for No False Lemmas, (NUL), we now note that this principle plus (EC $\rightarrow$) gives us that all evidence is knowledge, (E $\rightarrow$ K).

Hence, the left to right implication of (EC), (EC $\rightarrow$), entails that all evidence is knowledge. Correspondingly the right to left implication of (EC), (EC $\leftarrow$), entails that all knowledge is evidence, when conjoined with the highly plausible principle that inference from knowledge is knowledge generating:

(KI) if S knows p then S can gain knowledge by inference from p ;
(EC $\leftarrow$) if S can gain knowledge by inference from p then p is in S's evidence;
therefore
(E $\leftarrow$ K) if S knows p then p is in S's evidence.

(KI) is simply the idea that inference from what we already know has the power to extend our knowledge. The phrase 'can gain knowledge by inference from p then p ' is still interpreted as discussed in Section 2, i.e. where the subject is receptive to p and the inference procedure is faultless; so (KI) is not refuted by cases where the subject has other beliefs relevant to the conclusion of the inference or where

the inference is epistemically unsatisfactory. Under the intended interpretation, a counterexample to (KI), would be a case where S comes to believe q as a result of an inference from p where S knows p , yet S fails to come to know q because of some inadequacy in the epistemic quality of p (and not because of some other facts about S or facts about the inference procedure).

It is difficult to see how (KI) could be false: what more could be required of the premise of an inference that aims at extending knowledge than that the premise itself be known? However, the fact that (KI) looks sufficiently similar to the closure principle for knowledge might make one think that someone such as Nozick (1981) who rejects closure should reject (KI) as well. Alleged counterexamples to closure claim that (known) valid inferences from known premises can sometimes lead to conclusions that cannot be known—for example where the premise is a known proposition about everyday circumstances ('S is typing at a computer in a university office') and the conclusion is a consequences of that proposition that denies some relevant sceptical scenario ('S is not a brain-in-a-vat on a planet orbiting Gamma Cephei being deceived into believing that she is typing at a computer in a university office'). Are such cases, if we accept them, counterexamples to (KI)? It is difficult to say with confidence that they are, precisely because it is difficult to say what has gone wrong for the subject in such cases. For, *prima facie*, what could be epistemically faulty about a valid inference that is furthermore known to be valid? And the subject in such a case may unproblematically be assumed to be receptive. That said, I think we can be confident that the fault cannot be laid at the door of the *premise* of the subject's inference. The counterexample posits that S *knows* that she is typing as a computer in a university office. Could S's epistemic relation to the proposition that S is typing at a computer in a university office be of any higher quality, such that a higher quality epistemic attitude to that proposition *would* permit inferred knowledge that S is not a brain-in-a-vat on a planet orbiting Gamma Cephei etc.. Clearly not; this is not analogous to the case making an inference from a premise that is false or unjustified; it is not that the premise is epistemically not up to the task of supporting a knowledge-generating inference. For that reason, I believe that a denier of closure must say that in some cases a known valid inference can be epistemically faulty. The oddity in so saying is just why many find it difficult to deny closure; but it is at this locus, on the inference procedure, that the debate is centred. If so, then such counterexamples to closure, even if accepted, need not be taken as counterexamples to (KI). For example, one might argue, as some epistemic contextualists do, that closure is a correct principle when restricted to a given context; what goes wrong in the counterexamples is that there is a cross-context application of closure; such applications constitute faulty inference procedures. In any case, the worst damage that a denial of closure could do to my argument is to require a modification to (KI) and the other principles under discussion. For the closure denier does regard many inferences as potentially knowledge extending, and so needs to supply some replacement for closure. Whatever that replacement is, it will suggest an appropriate replacement for (KI). *Mutatis mutandis* the arguments of this paper will stand. For example, one modification would be to say that say that the 'can gain knowledge' in all the principles applies only to cases where the conclusion, when true, is knowable at all to S. The corresponding alteration to the central claim of this paper would be that the function of evidence is to provide the premises of inferences that can extend knowledge *where knowledge is possible at all*. Likewise, if one prefers a contextualist approach, then the changes will require restriction to a given context, and the corresponding conclusion will

give a contextualist account of evidence (which is just what one would expect a contextualist to endorse). (I note that the argument for $(E \leftrightarrow K)$ stands even if we have to have modify preceding arguments to accommodate the denial of closure. If we adjust the interpretation of ‘can gain knowledge’ that applies equally to (KI) and to $(EC \leftarrow)$, and so $(E \leftrightarrow K)$ still follows. If we take a contextualist approach, then the same contextualist restrictions will apply to what counts as evidence in $(E \leftrightarrow K)$ as applies to what counts as knowledge.)

In this section I have argued that the following are both correct:

$(E \rightarrow K)$ if p is in S’s evidence then S knows p .

$(E \leftarrow K)$ if S knows p then p is in S’s evidence.

which together entail:

$(E \leftrightarrow K)$ p is in S’s evidence if and only S knows p .

If we combine (KI) with (NUL) we get the following extended ‘knowledge by inference’ principle:

(KI+) S can gain knowledge by inference from p if and only if S knows p .

Recalling:

(EC) p is in S’s evidence if and only S can gain knowledge by inference from p .

(KI+) and (EC) directly give us:

$(E \leftrightarrow K)$ p is in S’s evidence if and only S knows p .

In summary (where ‘ $\Rightarrow$ ’ denotes entailment):

$$\begin{array}{ccccc}
 (\text{NUL}) & \wedge & (\text{EC} \rightarrow) & \Rightarrow & (E \rightarrow K) \\
 \wedge & & \wedge & & \wedge \\
 (\text{KI}) & \wedge & (\text{EC} \leftarrow) & \Rightarrow & (E \leftarrow K) \\
 \Downarrow & & \Downarrow & & \Downarrow \\
 (\text{KI}+) & \wedge & (\text{EC}) & \Rightarrow & (E \leftrightarrow K)
 \end{array}$$

The fact that (EC), along with the plausible (KI+), entails $(E \leftrightarrow K)$ is indirect evidence in favour of (EC). For Timothy Williamson (1997, 2000) has argued persuasively for $(E \leftrightarrow K)$ on independent grounds. The grounds are independent since Williamson’s arguments proceed principally on the basis of whether we would regard certain propositions as part of a subject’s evidence, whereas the arguments above concern principles regarding the function of evidence. In so doing those arguments illuminate the concept of evidence in a way that Williamson’s arguments do not (and are not intended to). There is no reason to suppose that ‘evidence’ and ‘knowledge’ are the really the same concept. If they were, one might ask, why do we have a concept of evidence? Why don’t we make do just with the concept of knowledge? We talk of ‘evidence for’ but not of ‘knowledge for’. We have a concept of evidence distinct from that of knowledge to mark a special role for evidence. Evidence is characterised functionally as that which supports knowledge-yielding inferences. (KI+) tells us that it is knowledge that supports knowledge-yielding inferences, which explains why Williamson’s equation is right.

6 Knowledge inferred from premises that are not known?

Above I articulated a principle, No Unknown Lemmas, (NUL), which on the face of it is highly plausible. It has, however, recently come under scrutiny with apparent counter-examples. (NUL) is close to what Federico Luzzi (2010) calls ‘Knowledge Counter-Closure’:

(KCC) if (i) S comes to believe q solely on the basis of competent deduction from p and (ii) S knows q , then S knows p .

(NUL) was an extension of the No False Lemmas principle. Corresponding to the latter is Truth Counter-Closure:

(TCC) if (i) S comes to believe q solely on the basis of competent deduction from p and (ii) S knows q , then p is true.

I motivated the extension of No False Lemmas by arguing that one cannot gain knowledge by non-redundant inference from a premise, if that premise is not justified. The latter principle, No Unjustified Lemmas, corresponds to Justified Belief Counter-Closure:

(JBCC) if (i) S comes to believe q solely on the basis of competent deduction from p and (ii) S knows q , then S is justified in believing p .

Luzzi (2010: 674–5) provides this case that claims to be a counterexample to (KCC), and which would also be a counterexample to (NUL):

Agoraphobia Unbeknownst to Ingrid, her new and only housemate Humphrey is something of an epistemic prankster. One evening, while Ingrid is in the kitchen cooking dinner, Humphrey mischievously decides to mislead her as to his whereabouts in the house. He therefore turns on the TV in the lounge so that she will believe, as she typically would, that he is in the lounge watching TV. However, also unbeknownst to her, Humphrey is agoraphobic, and hence would leave the house under very few circumstances; any circumstance in which he would leave the house (e.g., because of a raging fire) is undoubtedly one in which Ingrid would be aware that he is leaving the house. Suppose that Humphrey subsequently momentarily forgets about his ploy (something quite out of character for him), and accidentally wanders for a few seconds back into the lounge. During that interval, Ingrid forms the belief that Humphrey is in the lounge on the basis of hearing the TV on by relying (whether implicitly or explicitly) on this inductive argument: ‘(A) The TV is on and I didn’t turn it on. (B) When this happens, Humphrey is almost always in the lounge. So (1) Humphrey is in the lounge’. She then carries out the following valid and sound deduction:

1. Humphrey is in the lounge.
2. If Humphrey is in the lounge, then Humphrey is in the house.
- Therefore,*
3. Humphrey is in the house.

Suppose, too, that Ingrid believes (3) via no other epistemic route.

Luzzi claims that Ingrid knows that Humphrey is in the house. However, she does not know that he is in the lounge. So she has knowledge thanks to a deduction from a premise she does not know to be true. So (KCC) is false and so is (NUL).

Why is it that Ingrid knows that Humphrey is in the house? Luzzi argues primarily on the basis that we have a strong intuition here that Ingrid knows. But he also points out that none of the normal reasons that preclude beliefs from being knowledge apply here, including the reasons we find for articulating why Gettier cases are not cases of knowledge. In particular, Ingrid's belief satisfies sensitivity—the nearest possible world in which it is false that Humphrey is in the house (e.g. there is a fire) is one in which she does not believe that he is in the house. And it satisfies safety—her belief is true in all nearby possible worlds.

The fact that Humphrey is agoraphobic and that this is not known to Ingrid is important. If Ingrid knew that Humphrey were agoraphobic, that knowledge would provide an alternative route for Ingrid to reach the conclusion, and we could not be sure that reflecting on this does not influence the intuition that she knows that Humphrey is in the house. That Humphrey is agoraphobic is what makes Ingrid's belief that Humphrey is in the house satisfy the externalist considerations of sensitivity and safety. For if he were not agoraphobic then there would be nearby worlds in which he leaves the house, and in particular nearby worlds where he leaves the house and in which she continues to believe that he is in the house.

As I shall now show, we are able to construct numerous cases that involve a single premise deduction with premise p and conclusion q , where the subject believes q (only) because of the deduction from p , the belief in q satisfies safety and sensitivity, but the belief in p does not. Such cases ought to be good candidates as additional counterexamples to (KCC) if Luzzi is right in his diagnosis of *Agoraphobia*.

Since q is true in all the worlds where p is true, q will be at least as safe as p . Consequently, for an unsafe proposition one ought to be able to find a safe proposition q that is deducible from p . Here are one case:

Complex tautology p is an unsafe proposition. S works out by the truth table method that $(p \rightarrow q)$ is a tautology. In fact this is because q is itself a tautology, but this fact is not apparent to S. S believes q because q is deduced from p .

In this case p is unsafe and so unknown. But q is safe—it is a tautology. Here is a case with a contingent proposition in place of q :

Spurious precision Theory T predicts that the value of parameter ϵ is $8.1 \times 10^{-12} F/m$. S measures the value of ϵ using a very reliable piece of equipment, and gets the result $\epsilon = 8.85419 \times 10^{-12} F/m$. So S believes that $\epsilon = 8.85419 \times 10^{-12} F/m$ and deduces that T is false. It happens that by chance the belief is true—although the equipment is reliable it is not nearly as precise as the readout suggests; S is unaware that its limit of precision is five significant digits.

S's belief concerning the value of ϵ is unsafe. But the proposition deduced from this, that T is false, may be a safe proposition.

That there can be pairs of such propositions playing the role of p and q in the above seems clear. As a result it is very tempting to think that there ought to be cases where p is not known, just because it is unsafe, and where q , being safe, is known.

Now consider sensitivity. Where q is deducible from p , trivial cases excepted, q will tend to be more sensitive than p . As above q will be true in a superset of the possible worlds in which p is true, and so the closest possible world where p is false will be closer than the closest possible world where q is true. By taking us to a more distant possible world, the falsity of q is more likely to generate a world in which S's beliefs are different from those in the actual world. Consequently, it seems plausible that there are many pairs of propositions p and q such that q is deducible from p where p fails to be known on grounds of failing to satisfy the sensitivity requirement, but where q does satisfy the sensitivity requirement and so succeeds in being known. In the case of *Complex tautology*, belief in p might not satisfy sensitivity, whereas belief in q does so trivially. In *Spurious precision*, it is a fluke that the equipment gives the correct answer to this degree of precision; it would have given the same result had the value of ϵ been slightly different, e.g. $8.85418 \times 10^{-12} F/m$. So S's belief that $\epsilon = 8.85419 \times 10^{-12} F/m$ fails sensitivity. But the world might have been very different had T been false, and so S's belief in T would satisfy sensitivity.

Consequently, we can indeed construct cases where the belief in the premise of a deduction does not satisfy safety and sensitivity but belief in the conclusion does satisfy those conditions. It is tempting to think that such cases constitute counterexamples to (KCC) and (NUL). Nonetheless, I am not sure that we should yield to the temptation to accept the plausibility of the forgoing reasoning. While it is true for the above conditions that it is possible to find cases where p deductively entails q and p fails the conditions and q satisfies it, it does not follow from this that there are cases where p fails to be known for just that reason, and thus q succeeds in being known. That line of thought seems to rely upon the idea that knowledge is a matter of meeting certain conditions, and that once all the conditions are met, the proposition is known. But this is just the view that is denied by Williamson.

It should be noted how easy it is to satisfy the conditions of sensitivity and safety. The simpler versions of the safety principle make necessary truths trivially safe; they also trivially satisfy sensitivity (some versions of sensitivity just decline to cover truths that are in any way necessary). For the reasons given above, it is possible to find non-trivial (but almost trivial) examples with contingent propositions. Consider propositions such as "the Earth is more than 10,000 years old" or "atoms are composed of electrons, protons, and neutrons". Such propositions, being deeply tied to the history of the world or the laws of nature, look to be true in all the closest possible worlds. In particular, were those propositions false we human believers wouldn't exist. So there are no worlds in which I falsely believe either of those propositions.

Furthermore, although those conditions are used to avoid some Gettier cases, I don't think that they can avoid all. For if the propositions above satisfy the safety and sensitivity conditions then their disjunction with some false but justified proposition will also satisfy the conditions. Consider a late seventeenth century follower of Newton who believes 'F=ma or the Earth is more than 10,000 years old'. He believes this only because he believes the first disjunct, which is justified, but strictly false for relativistic reasons. He doesn't believe the second disjunct. His overall belief is a true justified belief, but which is not knowledge. So it is a Gettier case. But it also happens to satisfy sensitivity and safety.

Note that there are more subtle variants on this. Consider the following: 'F=ma or Isaac Newton senior is the father of Sir Isaac Newton.' Now imagine that this is believed by some late seventeenth century follower of Newton just because they believe the first disjunct. In fact because of the some spurious scandal they mistakenly

deny the second disjunct. Consider the nearest possible world in which the proposition is true. That is one in which Isaac Newton senior does not have a son who becomes the famous mathematician, Sir Isaac Newton. In that world Sir Isaac Newton does not exist. But we may suppose that in this world it was sometime after the end of the seventeenth century that anyone proposed that ‘ $F=ma$ ’. And so our subject, does not believe $F=ma$. Hence, the disjunctive proposition satisfies sensitivity, and also safety.

Now we can return to Luzzi’s case. Luzzi suggests that if we deny that Ingrid knows that Humphrey is in the house, then we are left with a new kind of Gettier example. It is a new kind of Gettier because none of the normal diagnoses apply, including sensitivity and safety. But as the above shows, some cases following an entirely familiar pattern of Gettier cases—ones formed by disjoining a false but justified belief with a true but unjustified belief—cannot be diagnosed as violating those requirements.

Luzzi does also mention that *Agoraphobia* does not suffer from the No False Lemmas diagnosis of the failure of Gettier cases to be knowledge (while my Newton cases do suffer from that analysis). Nonetheless, *Agoraphobia* is very close to be a case where the No False Lemmas principle does rule out knowledge. Imagine that we retell the story slightly differently, so that Ingrid’s initial belief is “Humphrey is watching the television” from which she infers that he is in the house. But, continuing the story just as Luzzi does, Humphrey wanders back in to the lounge but not to watch the television. In this case Ingrid’s inference will have started from a false belief. But it seems implausible that whether Ingrid knows should depend on which of these two propositions she believes at the beginning “Humphrey is in the lounge” or “Humphrey is watching the television”. Both bear the same cognitive relationship to Ingrid—she believes them for the same reason (the sound of the TV) and they can both play the same role in her inference to the conclusion that Humphrey is in the house. If she starts with ‘Humphrey is watching the television’ I think it is intuitively clear that she does not know the conclusion of her inference; hence she does not know if instead she starts with “Humphrey is in the lounge”. It seems a plausible principle that if a belief fails to be knowledge because it fails the No False Lemmas condition, then it cannot be transformed into knowledge by replacing the false lemma by one that is true but only accidentally so (i.e. a true lemma that bears a very similar cognitive relationship to the subject). As I have argued above, the most coherent way of articulating the thought that if one makes an inference with false premises or accidentally true ones, then the conclusion cannot be knowledge, is the fact that in each case the premise fails to be knowledge. It may be claimed that what I am doing here is wielding my principle (NUL) in order to re-interpret a potential counterexample to (NUL) and that I am therefore begging the question. That isn’t quite what I intend. Rather I suggest that insofar as Luzzi is appealing to intuitions to assert that Ingrid knows, we can also elicit intuitions going in the opposite direction, that deny that Ingrid knows. While these intuitions are best explained by (NUL), that fact is not supposed to be part of their probative force of the intuitions. Consequently my conclusion is that the intuitions in favour of Luzzi’s reading are not sufficiently robust to refute (NUL).

Some authors (Hilpinen 1988; Warfield 2005; Arnold 2011) claim to find cases where knowledge can be gained by inferences from *false* propositions, thereby appearing to refute No False Lemmas also. Consider a case like *Spurious precision* where the precise belief is false. Hilpinen argues that such a case is one where we can gain knowledge by inference from a falsehood but notes also that in such cases

the false proposition is close to the truth. Plausibly therefore there is a nearby true proposition that the subject also believes or is disposed to believe and this can be made to do the epistemic work of supporting a knowledge-producing inference (cf. Coffman 2008). Warfield criticises this strategy for defending the requirement of truth. The best versions of the strategy that Warfield considers are:

(S1) S has knowledge despite inferring from a falsehood if there is a justified and (at least) dispositionally believed truth entailed by the falsehood that serves as the premise in the inference.

(S2) S has knowledge despite inferring from a falsehood if there is a justified and (at least) dispositionally believed truth evidentially supported by the evidence for the involved falsehood.

Warfield points out that these strategies are defeated by the fact that they yield the wrong answer in standard Gettier cases. I seem to see that there is a dog in the yard (in fact it is only a toy) and infer that there is an animal in the yard. Unseen there is a squirrel in the yard. In both cases the proposition ‘there is a dog or a squirrel in the yard’ serves as the relevant justified and (at least) dispositionally believed truth, and so both yield the conclusion that S knows, whereas this is a clear case of Gettier non-knowledge.

The two strategies are bound to fail because they require the relevant proposition only to be true and justified, inviting us to chose a Gettier-style proposition; propositions inferred from such propositions will not themselves be known. A natural requirement on the ‘replacement’ for the false proposition is one that is itself known. The idea, a further strengthening of the approach of Coffman and Hilpinen, is that where an apparent inference uses a premise that is false, the inferred proposition may still be knowledge if in the neighbourhood of the false proposition is a true proposition that the subject knows, which the subject comes to know by a mechanism similar to that by which the false proposition is believed, and which can substitute for the false proposition in an inference very similar to the one in question. For short:

(SK) S has knowledge despite inferring from a falsehood if there is a cognitively similar process leading to the same conclusion starting from a proposition that S knows.

(SK) escapes Warfield’s Gettier response. It addresses cases like *Spurious precision*. When S looks at the equipment and sees the display say that ϵ has the value $8.85419 \times 10^{-12} F/m$, S believes, but does not know, that $\epsilon = 8.85419 \times 10^{-12} F/m$; S also believes, and does know, that ϵ is greater than $8.8 \times 10^{-12} F/m$. In the kind of case we are considering, S does not infer the second from the first; rather both beliefs comes about as a direct result of seeing the display. From both beliefs S can infer, in the same way, that the theory under test, T, is false.

One objection to this approach (and indeed to other approaches of the same kind) is that it doesn’t really explain what needs explaining. We were wondering how an inference from a false (or true but not known) proposition can lead to knowledge. How does it help us understand this to show that a the subject *could* have engaged in a *different* inference starting from a *different* premise? Furthermore, does not this strategy just amount to admitting that there are counterexamples to (KCC) and (NUL)?

This objection assumes a misleading conception of the cognitive psychology of inference that is a consequence of the way that philosophers represent inferences.

We represent inferences as sequences of precise propositions related by precise inference rules—which for many purposes is a useful and revealing thing to do. If that were also a correct representation of the cognitive psychology of inference, then it would be correct to say that S actually made this inference but didn't in fact make this very similar inference. Yet, the psychology of inference is a much coarser, less precise matter. I pick up a PhD thesis in order to judge how long it will take me to read it; I conclude that it will take me more than two days. Which propositions formed the basis of this inference? That the thesis clearly has many more than 50,000 words? That it probably has more than 70,000 words? That it looks longer than most theses I have read? It does not appear correct to say that it is this one of these propositions, but not the others, that I used in my inference. More plausibly I make all of these judgments and all play some role in influencing my conclusion. I conjecture that, recognising this, we understand that although some inferences may be *represented* in a precise way as starting from a proposition that is not known, this is misleading, especially in the case of inferences starting from beliefs acquired by perception. We understand in such cases that the psychological reality is that there are other propositions that we do know and that are also involved in the cognitive process of inference and that the presence and role of these can suffice to lead to knowledge of the inferred proposition (or propositions). Of course, such a picture means that there will be considerable vagueness about when an inference process really does involved known premises in the right way. But that seems to be no more vagueness than we actually encounter in such cases, and may explain why cases such as Luzzi's are contentious.

7 Evidence and science

The conception of evidence advanced by (EC) tells us that to achieve the status of evidence a proposition must be capable of supporting inferences that deliver knowledge; furthermore any proposition that is so capable counts as among the subject's evidence. The latter tells us that evidential propositions are not restricted by content nor by the means by which they come to be believed. The proposal thus rejects a conception of evidence that might appeal to certain empiricists, whereby evidence propositions must concern what a subject perceives or observes. I mentioned and rejected another conception of evidence that would appeal to some empiricists (and perhaps some non-empiricist foundationalists), viz. that evidence is restricted to propositions that are not themselves inferred from other propositions. Instead *any* proposition, including propositions that may have been inferred from other propositions and including propositions that may be quite removed from perception, may serve as an evidence proposition so long as it is capable of supporting knowledge-generating inferences. In the light of (KI+), we see that the propositions that fulfil this role are precisely the propositions that constitute the subject's knowledge.

It is worth noting that the use of 'evidence' in science accords with the conception promoted here rather better than with the empiricist view(s) of evidence. Part of the evidence for Einstein's general theory of relativity is the anomalous precession of the perihelion of Mercury. The precession is *inferred* from many individual observations of Mercury and that inference involves considerable mathematical work as well as substantive auxiliary hypotheses. Secondly, the precession is *not perceptible*. So this standard use of 'evidence' in science does not accord with an empiricist view, but does with the inferentialist conception.

It is nonetheless true that the precession of Mercury is often referred to by scientists as an ‘observation’. This shows only that the scientists’ conception of observation is one that is only remotely linked to perception. Historically, observation has typically been heavily perceptual, but even then it has rarely been exclusively perceptual. An astronomical observation may involve looking at a planet through a telescope, but the key information in the recorded observation will include the elevation of the telescope, its geographical location, and the time at which the planet was seen, which are not perceived through the telescope, and in the case of time, not perceptible at all (even if one has a watch). Likewise observations by nineteenth century chemists and physicists recorded values of variables such as temperature and electric current, which cannot be perceived. Positivists may argue that the real observations are perceptions of the instruments. Even so, it remains true that the use of even perceptual reports needs to be given context, which is not obviously perceptual: is what did the observer do to avoid a parallax error something that is itself perceptible by her? The empiricist proposal to limit ‘observation’ to perceptual encounters is a revision to our use of the term ‘observation’. Observation in modern science is even less perceptual and more heavily theory-laden. The Higgs boson has been observed, but no one has perceived it; nor indeed has anyone seen its trail on a cloud chamber photograph. Rather a vast array of different kinds of detectors, all designed in a theory-intensive way, are connected to a computer which delivered its results in a statistical form. In less esoteric science, modern radio telescopes deliver nothing of significance that is perceptible, but rather the evidence is in the form of data collected in a computer, which, furthermore, has been processed by statistical software. The same goes for the evidence that is used in constructing weather models, which is collected and processed automatically.

Furthermore, if the evidence is strong enough for the physicists to conclude that the Higgs exists, then its existence becomes a fact that itself is evidence that can be used as required, for example in confirming the standard model of particle physics. This illustrates the import of $(E \leftarrow K)$ and of $(EC \leftarrow)$. What makes something evidence is not that it was produced in some specific way, but rather that it can be used for a certain purpose: making knowledge-producing inferences. Correspondingly what scientists care about when they consider whether a proposition (such as an assertion that the precession of the perihelion of Mercury is 5557 seconds of arc per century, or that the Higgs boson exists) counts as evidence is not what general *kind* of source (perception, observation, inference from theory) produced the proposition, but whether the particular process that produced it warrants that proposition sufficiently that it can be relied upon in making further inferences. I have argued that such warrant is sufficient if and only if it is enough to make the proposition in question known.

8 Conclusion

At the outset I gave myself two objectives: (a) to articulate a functional characterisation of the concept of evidence that relates to inference; (b) to relate my characterisation to Timothy Williamson’s claim that all and only knowledge is evidence.

The guiding idea behind (a) is that evidence is something we want for a purpose. We don’t collect evidence for its own sake, but in order to make inferences from it (e.g. concerning the truth of a hypothesis). Of course we can make inferences from any old propositions and beliefs. But not just any old proposition or belief is ev-

idence. We want evidence for epistemically worthwhile inferences, the ones that produce knowledge: evidence is that from which knowledge-producing inferences can be made.

This conception allows any proposition to be evidence, so long as one can infer knowledge from it. This stands in contrast to a conception of evidence as that from which inference ultimately starts: evidence must be non-inferential. The inadequacies of this view of evidence were exposed by considering the possibilities of forgetting one's evidence or having it undermined by further evidence.

The conception of evidence as that which supports knowledge-producing inferences entails that all and only knowledge is evidence. This equation of knowledge and evidence is the conclusion reached by Williamson. My account and his are mutually supporting: mine provides an account of the concept of evidence, his provides a substantive identity concerning evidence. The idea that belief aims at knowledge is implicit in Williamson's epistemology (but not explicitly argued for). The initial intuition I employed concerning evidence is that we want evidence for a purpose: to make inferences from that evidence, inferences whose conclusions we can believe. Therefore, if belief aims at knowledge, then what we want from evidence is that it will permit knowledge-producing inferences.

References

- Arnold, A. 2011. Some evidence is false. *Australasian Journal of Philosophy*.
- Bird, A. 2004. Is evidence non-inferential? *The Philosophical Quarterly* **54**: 252–65.
- Bird, A. 2007. Justified judging *Philosophy and Phenomenological Research* **74**: 81–110.
- Coffman, E. J. 2008. Warrant without truth? *Synthese* **162**: 173–94.
- Harman, G. H. 1965. The inference to the best explanation. *Philosophical Review* **74**: 88–95.
- Heuer, R. J. 1999. *Psychology of Intelligence Analysis*. Washington, D.C.: Center for the Study of Intelligence, Central Intelligence Agency.
- Hilpinen, R. 1988. Knowledge and conditionals. *Philosophical Perspectives* **2**: 157–82.
- Luzzi, F. 2010. Counter-closure. *Australasian Journal of Philosophy* **88**: 673–83.
- Nelkin, D. K. 2000. The lottery paradox, knowledge, and rationality. *Philosophical Review* **109**: 373–409.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Starkie, T. 1834. *A Practical Treatise on the Law of Evidence* (5th ed.). Philadelphia PA: Nicklin and Johnson.
- Sutton, J. 2007. *Without Justification*. Cambridge MA: MIT Press.
- Warfield, T. A. 2005. Knowledge from falsehood. *Philosophical Perspectives* **19**: 405–16.

Williamson, T. 1997. Knowledge as evidence. *Mind* **106**: 1–25.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.