

COUNTER-CLOSURE

Federico Luzzi

5

The focus of this paper is the *prima facie* plausible view, expressed by the principle of Counter-Closure, that knowledge-yielding competent deductive inference must issue from known premises. I construct a case that arguably falsifies this principle and consider five available lines of response that might help retain Counter-Closure. I argue that three are problematic. Of the two remaining lines of response, the first relies on non-universal intuitions and forces one to view the case I construct as exhibiting a justified, true belief to which none of the usual diagnoses of knowledge failure in Gettier cases apply. The second line involves claiming that Fake Barns and its ilk are misdiagnosed by epistemological orthodoxy as Gettier cases. We are thus confronted by a trilemma: either the case I discuss undermines the first-blush plausible principle of Counter-Closure; or the case I discuss instantiates a novel kind of Gettier case; or a popular conception of a key range of alleged Gettier cases must be rejected. No matter which horn we choose, the case points to a philosophically curious conclusion.

10

15

20

1. Closure and Counter-Closure

It is a widely held view that knowledge must admit of extension via competent deductive inference. For a subject S and propositions p and q:

25

Closure: Necessarily, if (i) S knows p, (ii) S knows that p entails q, and (iii) S comes to believe q solely on the basis of competently deducing it from p, then S knows q.

30

The view Closure expresses is well-discussed.¹

My focus here is on a sister principle to Closure, one that expresses a related and *prima facie* very plausible view: that knowledge obtained on the basis of competent deductive inference alone must issue from known premises.

35

40

¹For prominent endorsements of this view see, e.g., DeRose [1995], Williamson [2000], and Hawthorne [2004] and [2005]. For prominent opposition see, e.g., Dretske [1970] and [2005] and Nozick [1981]. Like many others who discuss closure principles for knowledge, I do not put forward my formulation of Closure with the presumption that it is definitive: it may very well admit further refinement.

45 *Counter-Closure*:² Necessarily, if (i) S knows that p entails q and (ii) S comes to believe q solely on the basis of competently deducing it from p, and (iii) S knows q, then S knows p.

50 Counter-Closure sounds plausible, at least at first blush.³ It follows from the palatable view that in a knowledge-yielding *modus ponens* from p to q, if q is believed solely on the basis of the *modus ponens* itself, then the epistemic pedigree of q can be no better than the epistemic pedigree of p. Also, Counter-Closure connects to the plausible idea that the basing of one's belief plays an important role in determining whether that belief is known.
55 Counter-Closure dictates that when the basis for beliefs held via competent deduction is at issue, this *desideratum* is best captured by requiring that the basing premise be known.⁴

60 Despite these motivations and its intuitive appeal, a strong case can be made against Counter-Closure. Consider the following scenario.

60 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
65 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
70 it on. (B) When this happens, Humphrey is almost always in the lounge.

75 ²The name 'Counter-Closure' might suggest logical incompatibility with Closure. I do *not* intend the principle to be read this way—in fact, Counter-Closure is compatible with Closure. Rather, as will become apparent, the prefix 'Counter-' is best read as metaphorically indicating movement in the opposite direction, much like 'counter-clockwise' does with respect to 'clockwise'.

80 ³It may sound obviously true to some. This would partly explain why the view it expresses is sometimes assumed to be true without argument (see, e.g., Armstrong [1973: 198–9]; Nozick [1981: 231]; Stanley [2005: 89–90]). As with Closure, this particular formulation of Counter-Closure may admit of further refinement. However, my argument will depend in no way on such possible refinements.

85 ⁴As an anonymous referee has pointed out to me, the view expressed by Counter-Closure has come under recent fire: both Warfield [2005] and Klein [2008] construct and defend cases in which an agent comes to know the conclusion of a deductive inference that contains a false and therefore unknown premise. The challenge I raise to Counter-Closure below importantly differs from these attacks because in the example I will discuss the relevant knowledge-yielding deductive inference contains only true premises.

90 An anonymous referee has also observed that the Gettierized Reliable Reporter case [Coffman 2008: 191] seems to undermine Counter-Closure and displays a deductive inference containing only true premises. While there are structural similarities between Gettierized Reliable Reporter and the example I will discuss below, it is worthwhile noting that, while he alludes to this possibility, Coffman does not provide an extended discussion of whether and how his case may undermine the view expressed by Counter-Closure; nor does he consider, as I will below, the interesting philosophical ramifications of insisting that Counter-Closure is true in the light of this kind of challenge. For these reasons, I take the detailed discussion of these issues in what follows to significantly advance the debate over whether we should endorse or reject Counter-Closure.

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

As a preliminary, note that argument (1)–(3) is valid and sound, and that the inductive argument in support of (1) contains only true premises. Let's examine the epistemic credentials of each proposition of the deductive argument.

Proposition (1) is true, and Ingrid's belief in (1) is justified by her inductive argument from (A) and (B). But intuitively it is not known, for her belief in (1) seems to be subject to too much luck to count as known. One way of providing support for this intuition is by pointing out that her belief in (1) seems 'Gettiered'—it fails several constraints on knowledge that theorists commonly appeal to in order to explain knowledge failure in a range of Gettier cases, i.e., a range of cases exhibiting a justified, true, yet unknown belief.⁵

For example, Ingrid's belief fails sensitivity. S's belief in p sustained via method M is sensitive just in case in the closest not-p world, S would not believe that p via M.⁶ To see why Ingrid's belief in (1) is not sensitive, consider that in the closest world where Humphrey is not in the lounge, Ingrid would not refrain from believing it via the same method deployed in the actual world; for presumably this world is one where Humphrey carries out his prank with his usual focus, in which case she would be fooled and believe (1) via the very same method.

Ingrid's belief also fails safety. S's belief that p sustained via M is safe just in case in all close possible worlds, if S believes that p via M then p is true.⁷ To see why Ingrid's belief in (1) is not safe, consider that Ingrid would believe (1) via the same method in the close world in which Humphrey carries out his prank with his usual focus, and that in such a world (1) is false.

A further reason to think that Ingrid does not know (1) is that the justification for her belief in (1) is undermined by an undefeated defeater:

⁵See Gettier [1963] for the first examples of Gettier cases.

⁶See, e.g., Nozick [1981] and Dretske [1971] for endorsements of similar constraints. Differences in their formulations of sensitivity have no significant impact on what follows.

⁷See, e.g., Sosa [1999], Williamson [2000], and Pritchard [2005] for conceptions of safety. Again, the slight variations in their formulations of safety have no repercussions on my argument. In particular, Ingrid's belief in (1) is not safe even on a formulation of safety that merely requires that in *most* (as opposed to all) close possible worlds, if S believes that p via M then p is true.

namely, the true proposition that Humphrey intends to mislead her as to his whereabouts.^{8,9}

I should flag that this may not be a universally shared view. Consider momentarily Alvin Goldman's well-known Fake Barn case (Goldman [1976: 773]).¹⁰ While driving down the highway Henry believes on the basis of eyesight that there is a barn in the field. Unbeknownst to him, farmers in the county have erected many fake barns that are invariably taken to be genuine by those driving through. The intuition is supposed to be that Henry does not know that there is a barn, even though he is actually looking at the only genuine barn in the county, and his justified belief is hence true. Henry's standing with respect to his belief seems equivalent to Ingrid's standing with respect to (1).¹¹ But of the Fake Barn case, Hetherington [1998] argues that we should not leap to the conclusion that it is 'Gettiered' and therefore unknown. In recent work, some experimental philosophers (Nichols, Stich, and Weinberg [2002]) have tried to erode epistemologists' traditionally stark confidence in the intuition that agents in similar scenarios lack knowledge. Are analogous responses appropriate in the case of Ingrid's belief in (1)?

I will leave it an open question whether—notwithstanding the intuition shared by many that Henry's belief that there is a barn is not known,¹² and which extends to Ingrid's relevantly similar belief in (1)—these lines of response are ultimately correct. The puzzle I sketch in the following pages can be evaded by those taking these views, and their ability to circumvent it may ultimately be taken by its endorsers as grist for their mill. For the moment, however, let us set them aside. I will return to the view that Ingrid's belief in (1) is not 'Gettiered' in the final section.

Proposition (2) is known solely by reflection on the features of the house.

Proposition (3) is true, and is justifiedly believed by Ingrid on the basis of competent deduction from (1). The intuition that Ingrid knows (3) finds strong support in its excellent epistemic standing: over and above being true and justified, it meets the constraints that rule out (1) as a case of knowledge. In particular, Ingrid believes (3) sensitively, since in the closest world where Humphrey is not in the house, this is presumably because of some exceptional circumstance such as a raging fire, in which case she would refrain from believing that he is in the house via any method; and her belief is also safe, since there is no close possible world in which he is not in the house (and *a fortiori* no close possible world where he is not in the house and

⁸See, e.g., Lehrer and Paxson [1969] and Klein [1971].

⁹While the three constraints just discussed have the virtue of explaining knowledge failure in a range of Gettier cases, I do not mean to suggest that this is their sole or even their main purpose. Their endorsers also believe them to be plausible for independent reasons.

¹⁰Goldman attributes this case to Carl Ginet.

¹¹Strictly speaking, to make the analogy tight we would need to posit that Henry, either implicitly or explicitly, relies on an inductive argument along the following lines to justify his belief that there is a barn. '(A*) That object looks like a barn. (B*) When things look like barns, they almost always are barns. So (C*) that object is a barn'. I think it's safe to say that reliance on this argument does not make any difference to our inclination or disinclination to ascribe knowledge to Henry.

¹²Apart from Goldman [1976], see—to name but a few—Plantinga [2000: 359–60], Lackey [1999: 487 and 2008: 68–9], Pritchard [2005: 161–2], Frances [2005: 114 (especially footnote 41)] and Gendler and Hawthorne [2005]. Gendler and Hawthorne report the view that 'the salient proximity of an indistinguishable facsimile is sufficient to indict the casual observer's knowledge' as widespread amongst philosophers [335]. Along similar lines, Fumerton [2006: 26] says of this case that 'many epistemologists don't want to allow that [Henry's] apparently justified true belief constitutes knowledge'.

she believes that he is via the method she deploys in the actual world). In terms of defeaters, while Humphrey's intention to mislead Ingrid acts as a defeater for her justification for (1), when it comes to (3) this defeater is itself defeated by the proposition that Humphrey is agoraphobic. So the reasons for deeming (1) not known lapse in the case of (3). 190

Moreover, (3) meets a host of further popular constraints on knowledge. For example, the deductive inference that sustains Ingrid's belief in (3) contains no false premises; there is an appropriate causal link connecting the fact that Humphrey is in the house with Ingrid's belief that Humphrey is in the house; in believing (3), Ingrid does not seem to act in an epistemically irresponsible manner. All things considered, (3) has an excellent epistemic pedigree. 195 200

By these remarks on (3) I do not wish to unduly ignore the opposite intuition that some might feel, namely that Ingrid does *not* know (3). It is rash to have one's argument depend on intuitions, and I will not do that here. However, I will postpone discussion of the intuition that Ingrid does not know (3) until §2.3, where it will be given due consideration. My immediate aim is to show there is no straightforward way of making Agoraphobia and Counter-Closure compatible other than by appealing to the intuition that (3) is not known or the view—momentarily set aside—that (1) is known. 205 210

2. Ways Out

Here, for convenience, is the relevant instance of Counter-Closure that Agoraphobia purports to falsify. 215

Necessarily, if (i) Ingrid knows that if Humphrey is in the lounge, then he is in the house, (ii) Ingrid comes to believe that Humphrey is in the house solely on the basis of competently deducing it from his being in the lounge, and (iii) Ingrid knows that Humphrey is in the house, then Ingrid knows that Humphrey is in the lounge. 220

How might one salvage Counter-Closure in the light of Agoraphobia? I will discuss four available strategies.¹³ 225

2.1. Denying Clause (ii)

Firstly, one might deny (ii) that Ingrid comes to believe that Humphrey is in the house solely on the basis of inferring it from Humphrey's being in the lounge. Now, admittedly, it is not uncontroversial, and maybe even implausible, that the standard way of coming to know a proposition like (3) is via deductive inference from (1). One easily available alternative route to (3) proceeds from the grounds for (1) without going through (1) itself. To 230

¹³The strategy that denies (i) that Ingrid knows that if Humphrey is in the lounge, then he is in the house is too implausible to deserve discussion. Furthermore, as stated above, I will leave discussion of the view that affirms the consequent of this instance of Counter-Closure principle to the final section. 235

clarify: instead of reasoning as she does in Agoraphobia, Ingrid might simply reason ‘The TV is on in the lounge; when this TV is on, Humphrey is usually in the lounge/house. So Humphrey is in the house’ without explicitly believing that Humphrey is in the lounge on this occasion. Or perhaps, her belief that Humphrey is in the lounge might be formed spontaneously and instantly upon her hearing the TV, without the aid of any explicit inference.

While all this is true, it is hard to see how this impinges on the mere possibility that Ingrid reason as she does in Agoraphobia and arrive at her belief in (3) solely through deductive inference from (1). In so far as any such piece of reasoning is possible, then this strategy cannot constitute a satisfactory response to Agoraphobia.

2.2 Denying Clause (iii) for Theoretical Reasons (Take 1): the No-Analysis Strategy

A second strategy is to deny (iii) that Ingrid knows that Humphrey is in the house, and consider it independently of any intuition one might have that this proposition is not known. We can imagine endorsers of any (combination of the) constraint(s) met by Ingrid’s belief in (3) providing theory-based reasons for this strategy. Such theorists might argue along these lines (where ‘X’ stands for any subset of the set of constraints that Ingrid’s belief in (3) meets, e.g., sensitivity, safety, epistemic responsibility, etc.): ‘We never intended X to provide the missing tile in an analysis of knowledge in terms of individually necessary *and jointly sufficient* conditions: we do not hold that justified, true beliefs that also satisfy X must amount to knowledge. So it is overly quick to move from the fact that Ingrid’s belief in (3) is justified, true, and meets X to its being known. And in so far as it is not known, then Agoraphobia is no counterexample to Counter-Closure.’

We should immediately concede to this line of argument that Agoraphobia provides no conclusive *proof* that Ingrid knows (3). What it does provide, however, is good reason to believe that (3) is known, given the excellent epistemic standing of (3) reflected in its satisfying a range of proposed constraints on knowledge. Its great pedigree, I submit, places on defenders of Counter-Closure the burden of proof. While of course it is a possibility that (3) is not known, it is difficult to see in what respect Ingrid’s belief is epistemically deficient; what plausible constraint on knowledge does Ingrid’s belief in (3) not satisfy?

Spurred by this challenge, some partial to modal constraints might turn to a recently developed version of safety custom-built to explain knowledge-failure in the following case.

MATHS: Buster’s mathematical skills are appalling. Whenever he forms a belief in a mathematical proposition, he takes a wild guess. Today, he guesses that the cube root of 729 is 9.

Intuitively, Buster does not know that the cube root of 729 is 9, even though his belief is true. However, both sensitivity and safety have difficulty in yielding a verdict of ignorance. It is unclear what Buster would believe in

the closest possible world where 9 is not the cube root of 729, since there is no such world. Hence it is unclear that Buster's belief is insensitive. Furthermore, in all close possible worlds where Buster believes that the cube root of 729 is 9 via the method deployed in the actual world, his belief is still true (since his belief is true in *all* close possible worlds). His belief is therefore clearly safe.

Weatherson's 'belief-safety' allows for the content of one's belief to vary across possible worlds when evaluating the safety of that belief (Weatherson [2004]), and thereby motivates a clear verdict of ignorance concerning Buster's belief. Roughly, the idea is that there are close possible worlds in which Buster believes that the cube root of 729 is, say, 8, and that these make for a failure of belief-safety for his belief that the cube root of 729 is 9. If belief-safety is required for knowledge, then we can explain Buster's knowledge failure in Maths.

It's not implausible to think that, in some sense, Ingrid's belief in (3) is relevantly similar to Buster's belief that the cube root of 729 is 9: both beliefs have contents that are *modally robust*, i.e., they could not easily have been false. And you might plausibly think that it is in virtue of this fact that these beliefs meet ordinary safety. Perhaps, then, Weatherson's tailored belief-safety can be brought to bear on Ingrid's belief that Humphrey is in the house and rule it out as a case of ignorance, just as it rules out Buster's belief.

Unfortunately for this response, belief-safety cannot help to motivate the view that Ingrid doesn't know (3) in Agoraphobia. This is because it is not clear that Ingrid would hold different beliefs that are false in close possible worlds. For what else would Ingrid believe via the method deployed in the actual world, if she didn't believe (3)? Possibly some false things (e.g., that Humphrey is sitting in his favourite armchair, that Humphrey is watching TV) and possibly some true things (e.g., that Humphrey is nearby, that Humphrey is not in the attic). But unlike Buster's scenario, there is no definite understanding that Ingrid's believing something true is modally unstable. In fact, we could always make further stipulations on Agoraphobia, such that Ingrid would only believe things that happen to be true on the basis of the method that sustains her belief in (3). This would clearly neutralize an appeal to belief-safety as an explanation of knowledge-failure.

In sum, my objection to the reply to Agoraphobia under consideration highlights an inability to motivate the verdict of ignorance concerning (3) that the response at issue deems correct. This inability stems from the fact that Ingrid's belief in (3) satisfies a wide range of popular constraints on knowledge, and places the burden of proof squarely on defenders of Counter-Closure.

2.3. Denying Clause (iii) on the Basis of Intuition

Having considered these possible replies to Agoraphobia, it is now time to return to the response mentioned at the end of section 1, according to which

philosophical intuition alone adequately motivates the view that Ingrid does not know (3). This is not an intuition that all can be expected to share. Nonetheless, those who feel its pull will be reluctant to abandon Counter-Closure in the light of Agoraphobia. Can anything be said against an intuition-based argument that is not a mere statement of the putative incorrectness of the intuition? Can we avoid a philosophically unsatisfactory ‘clash of intuitions’?

I don’t think this line of response can be undermined. But we can point out that, if this response were ultimately correct, then although Agoraphobia would not constitute a counterexample to Counter-Closure, it would nonetheless be a philosophically interesting case. This is because Agoraphobia would have to be viewed as a Gettier case: according to this line of response, in fact, Ingrid’s belief in (3) is true, justified, but not known. Curiously, none of the familiar diagnoses of why knowledge fails in Gettier cases would work for Agoraphobia. Thus, for example, because Agoraphobia displays a sound inference, it is immune to the no-false-premises response. It would also not be susceptible to a diagnosis in terms of sensitivity or safety. Furthermore, as mentioned above, knowledge failure in Agoraphobia could not be diagnosed in terms of the presence of an undefeated defeater or in terms of the absence of an appropriate causal link between Ingrid’s belief and the relevant fact. Thus, if we take that line that Ingrid does not know (3), we are forced to view Agoraphobia as exemplifying a Gettier case to which none of the familiar diagnoses of knowledge failure applies; that is, we would have to understand Agoraphobia as exemplifying a *novel* kind of Gettier case.

2.4. Denying Clause (iii) for Theoretical Reasons (Take 2): the Luck-Based Strategy

Before closing, I would like to mention a fourth possible line of resistance, one that denies Ingrid knowledge of (3) by exploiting the extremely plausible thought that, in some sense, knowledge is incompatible with luck. Now, the truth of Ingrid’s belief in premise (1) seems hostage to too much luck for the belief to be knowledge. Her belief is apparently ‘Gettiered’, and a popular general explanation of what is wrong with Gettier cases is precisely that although the agent’s belief is true, its truth is down to too much luck. The fact that Ingrid’s belief fails to meet constraints like safety and sensitivity, one of whose alleged virtues is their ability to explain knowledge failure in Gettier cases, speaks in favour of this view.

The following principle also appears to be plausible, at least at first blush:

Luck Transmission: If the truth of S’s belief that p is hostage to too much luck to count as known and S believes q solely via competent deductive inference from p, then the truth of S’s belief that q is hostage to too much luck to count as known.

The line of resistance under consideration appeals to the luck involved in Ingrid’s belief in (1) coupled with Luck Transmission to explain why Ingrid

does not know the conclusion of her deductive inference: she doesn't know (3) because its truth is also hostage to too much luck. There are three points to be made against this strategy.

The first is that, in relying on Luck Transmission, the endorser of this strategy does not flesh out the specifics of the notion of luck at issue here. A tendency in epistemology, in fact, has been to try to home in on the kind of luck at issue in Gettier cases and other luck-infected cases of knowledge failure by proposing constraints on knowledge that rule out luck. A prime example is Pritchard [2005], who defines safety precisely as an antidote to luck; but it goes for most proposals of constraints on knowledge that they can be viewed as having at least implicitly the goal of ruling out cases of luck-infected true belief. That no such constraints can rule out (3) is a significant point.

Of course, one could insist that the truth of (3) is down to too much luck, and that this explains why Ingrid doesn't know (3). The point I am trying to press, however, is that, so long as the notion of non-luckiness that knowledge requires is not further explicated, this can constitute nothing but a superficial, uninformative, and thereby unsatisfactory explanation. This dissatisfaction would be roughly on a par with that incurred were we to define knowledge—in the light of the incompatibility of luck and knowledge highlighted by Gettier cases—as non-lucky true belief, and then take ourselves to have met the challenge of defining knowledge.

The second point is connected to the first. Suppose for the sake of argument we agree with this fourth strategy that (3) is not known in virtue of Luck Transmission and the luck affecting Ingrid's belief in (1). Then still the point remains that (3) meets a host of hitherto proposed constraints on knowledge, and that it is a belief that is justified and true. So the point made regarding the previous intuition-based strategy is applicable: Agoraphobia in any case represents a novel and interesting kind of Gettier case, one which would pressure endorsers of this last strategy to come up with and explicitly endorse an informative constraint on knowledge capable of ruling out the brand of case Agoraphobia exemplifies.

But hopefully—and this is the third point—we will not need to go this far, since in any case there is good reason to think that Luck Transmission is false.¹⁴ Agoraphobia shows that the kind of luck that, for instance, safety rules out does not transmit across deductive inference, for Ingrid's belief in (1) is clearly unsafe, yet her belief in (3) is clearly safe. Further support for the falsehood of Luck Transmission may be garnered by looking to non-epistemic variants of luck: as Dretske [1970: 1008] points out, it may be lucky that my dart hit the bull's-eye; my dart's hitting the bull's-eye entails that it either hit the bull's-eye or the side of the barn; yet it may not be at all lucky that my dart either hit the bull's-eye or the side of the barn.¹⁵

For these reasons, the line of response that appeals to the luck in Ingrid's belief in (1) conjoined with Luck Transmission to rescue Counter-Closure is not viable.

¹⁴Thanks to an anonymous referee for raising this third point.

¹⁵Other epistemic kinds of luck may fail to transmit across deductive inference. See Pritchard [2005: 133–41] for descriptions of these varieties of luck.

3. Conclusion

On the one hand, the view that knowledge-yielding deductive inference must issue from known premises elicits strong *prima facie* approval. On the other, Agoraphobia indicates that problems can be raised for this principle.

There are three ways out of the challenge that Agoraphobia raises. The first is to concede its force and abandon Counter-Closure. The second is to insist that Ingrid's belief in (3) is not known, thus defusing the threat to Counter-Closure, at the cost of having to introduce Agoraphobia into one's epistemology as a novel kind of Gettier case. As mentioned in section 1, there is a third option. This involves insisting that Ingrid's belief in (1) is known. Running this line forces one to claim that the relevantly similar belief held by Henry in Fake Barns is, despite a widespread understanding to the contrary, not 'Gettiered', and hence a case of knowledge. Thus, no matter which horn of this trilemma we choose, the conclusion Agoraphobia leads to is of significant philosophical interest.¹⁶

University of Aberdeen
University of St Andrews

Received: June 2009
Revised: August 2009

References

- Armstrong, David 1973. *Belief, Truth, and Knowledge*, London: Cambridge University Press.
- Clark, Michael 1963. Knowledge and Grounds: A Comment on Mr. Gettier's Paper, *Analysis* 24/2: 46–8.
- Coffman, E. J. 2008. Warrant without Truth?, *Synthese* 162/2: 173–94.
- DeRose, Keith 1995. Solving the Skeptical Problem, *The Philosophical Review* 104/1: 1–51.
- Dretske, Fred 1970. Epistemic Operators, *The Journal of Philosophy* 67/24: 1007–23.
- Dretske, Fred 1971. Conclusive Reasons, *Australasian Journal of Philosophy* 49/1: 1–22.
- Dretske, Fred 2005. Is Knowledge Closed under Known Entailment? The Case against Closure, in *Contemporary Debates in Epistemology*, ed. Matthias Steup and Ernest Sosa, Oxford: Blackwell: 13–26.
- Frances, Bryan 2005. *Scepticism Comes Alive*, Oxford: Clarendon Press.
- Fumerton, Richard 2006. *Epistemology*, Oxford and Malden, MA: Blackwell.
- Gendler, Tamar Szabó and Hawthorne, John 2005. The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies, *Philosophical Studies* 124/3: 331–52.
- Gettier, Edmund 1963. Is Justified True Belief Knowledge?, *Analysis* 23/6: 121–3.
- Goldman, Alvin 1976. Discrimination and Perceptual Knowledge, *The Journal of Philosophy* 73/20: 773–91.
- Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hawthorne, John 2005. Is Knowledge Closed under Known Entailment? The Case for Closure, in *Contemporary Debates in Epistemology*, ed. Matthias Steup and Ernest Sosa, Oxford: Blackwell: 26–46.
- Hetherington, Stephen 1998. Actually Knowing, *The Philosophical Quarterly* 48/193: 453–69.
- Klein, Peter 1971. A Proposed Definition of Propositional Knowledge, *Journal of Philosophy* 68/16: 471–82.
- Klein, Peter 2008. Useful False Beliefs, in *Epistemology: New Essays*, ed. Quentin Smith, Oxford: Oxford University Press: 25–62.
- Lackey, Jennifer 1999. Testimonial Knowledge and Transmission, *The Philosophical Quarterly* 49/197: 471–90.
- Lackey, Jennifer 2008. *Learning from Words*, Oxford: Oxford University Press.
- Lehrer, Keith and Thomas Paxson 1969. Knowledge: Undefeated Justified True Belief, *Journal of Philosophy* 66/8: 225–37.

¹⁶Versions of this paper were presented at the Australian National University Philosophy Society, the Australasian Association of Philosophy Conference at La Trobe University, the 5th Biennial Rochester Epistemology Graduate Conference, the 8th Shapiro Graduate Conference at Brown University, the Arché Basic Knowledge Seminar at the University of St Andrews, the XI Taller d'Investigació en Filosofia at the Universitat de València and the Yale/Uconn Graduate Philosophy Conference. I am grateful to commentators and audiences on these occasions. Thanks to David Chalmers, Paul Dimmock, Douglas Edwards, Stephen Hetherington, Mireia López Amo, Aidan McGlynn, Andreas Stokke, Elia Zardini and two anonymous referees for helpful comments. I am particularly grateful to Leon Leontyev for numerous discussions on the issues of this paper, and to Patrick Greenough and Crispin Wright for discussion, guidance and constant encouragement.

- Nichols, Shaun, Stephen Stich and Jonathan Weinberg, 2002. Meta-skepticism: Meditations in Ethno-epistemology, *The Skeptics: Contemporary Essays*, ed Steven Luper, Aldershot and Burlington, VT: Ashgate: 227–48.
- Nozick, Robert 1981. *Philosophical Explanations*, Oxford: Oxford University Press.
- Plantinga, Alvin 2000. *Warranted Christian Belief*, Oxford: Oxford University Press.
- Pritchard, Duncan 2005. *Epistemic Luck*, Oxford: Oxford University Press.
- Sosa, Ernest 1999. How to Defeat Opposition to Moore, *Philosophical Perspectives* 13: 141–53.
- Stanley, Jason 2005. *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Warfield, Ted 2005. Knowledge from Falsehood, *Philosophical Perspectives* 19/1: 405–16.
- Weatherston, Brian 2004. Luminous Margins, *Australasian Journal of Philosophy* 82/3: 373–83.
- Williamson, T. 2000. *Knowledge and Its Limits*, Oxford: Oxford University Press.

480

485

490

495

500

505

510

515

520