Knowledge from Falsehood

Branden Fitelson

Philosophy Department
Rutgers University

&

Center for Advanced Studies
Ludwig-Maximilians-Universität München

branden@fitelson.org

- **The Naive View** (TNV) of Inferential Knowledge (slogan):

(TNV) Inferential knowledge requires *known relevant premises.*

- One key aspect of (TNV) is "counter-closure" [9, 10]:

  (CC) If *S* comes to believe *q solely* on the basis of competent deduction from *p* and *S* knows that *q*, then *S* knows that *p*.

- It is useful to note how (CC) differs from *closure*:

  (C) If *S* comes to believe *q solely* on the basis of competent deduction from *p* and *S* knows that ***p***, then *S* knows that ***q***.

- I won't be discussing (C) today, but here is a useful *contrast.*

  - Entailment *does* preserve **some good**-making features of premises. Most notably, entailment preserves *truth.*

  ☞ *Why should it be* that entailment preserves ***any*** bad-making features of premises? [*e.g.*, entailment doesn't preserve *falsity.*]

- There are other, more concrete reasons to worry about (CC).

- There are various *(prima facie) counterexamples* to (CC).

  - *E.g.*, Think about NASA's inferential use of Newton's theory.

- It seems Saunders & Champawat [12, *p*. 9] were the first to raise an example of "knowledge from non-knowledge" (KFNK). Their example is *like* the following one (my spin):

  > An urn contains 2 balls of unknown (to Sam) color distribution. Sam samples one ball (with replacement) from the urn many, many times. He is a very reliable counter and observer (and Sam knows all of the above facts). Sam then reasons as follows: "I have sampled a red ball from the urn *exactly* 10^9 times in a row. ∴ Both balls in the urn are red."

- As it happens, Sam has (slightly) miscounted the number of consecutive red ball observations he has made. Sam *actually* observed 10^9 *plus one* such consecutive outcomes.

- S & C do not analyze their example — they merely present it as a case which shows that Clark's [1] "no false lemmas" requirement [6] (in response to Gettier's [5]) is *too strong*.

- This seems to be *inductive* inferential knowledge involving a false relevant premise. My focus today will be on *deduction*.

- It seems Hilpinen [7, *pp.* 163-4] was the first to discuss the sorts of examples I'll be focusing on. His example has the same structure as Warfield's, which I'll be discussing below.

  > A mother suspects that her child has temperature, and when she measures the temperature and looks at the thermometer, she takes it to read 40.0° C. . . . If the thermometer is fairly accurate and the mother has reasonably good eyesight, we can say under these circumstances that she knows that the child has temperature [*viz.*, that $t > 37°$ C]. . . . But the mother need not have perfect eyesight and the thermometer need not be completely accurate . . . the actual thermometer reading might be 39.7°, and the actual temperature of the child might be 39.2°. . . . This example suggests that a person can know things not only on the basis of (valid) inference from what he or she knows, but in some cases even on the basis of inference from what is not known (or even true), provided that the latter (evidential) propositions are sufficiently close to the truth.

- Since this example is mainly a digression for Hilpinen, he does not analyze it further. Such analyses came later [8].

- Warfield [13] discusses several examples of (KFF), and he defends (KFF) against various forms of "resistance".

- I'll focus on the following example from [13], which has (more or less) the same *formal* structure as Hilpinen's:

    I have a 7pm meeting and extreme confidence in the (exact) accuracy of my fancy watch. Having lost track of the time and wanting to arrive on time for the meeting, I look carefully at my watch. I reason as follows: "It is exactly 2:59pm. ∴ I am not late for my 7pm meeting." As it happens, it's exactly 3pm, not 2:59pm. [We may suppose that my fancy watch is running perfectly, but that I (unwittingly) set it so that it reads one minute early.]

- The rest of the talk will focus on variants of this case.

- Here is a natural thought about such cases. While they do seem to be cases of (KFF), the following *also* seems right.

    ($\star$) If S's belief that p had *not* been false (*i.e.*, if S's belief that p had been *true*), S would (still) have been in a position to know that q on the basis of a competent deduction from p.

- But, examples of (KFF) that *violate* ($\star$) can also be described.

- In [3], I offer the following variant of Warfield's watch case:

  I have a 7pm meeting and extreme confidence in the (exact) accuracy of both my fancy watch and the Campanile clock. Having lost track of the time and wanting to arrive on time for the meeting, I look out of my office window (from which the Campanile clock is almost always visible). As luck would have it (owing, say, to the fluke occurrence of a delivery truck passing by my window), the Campanile clock is obscured from view at that instant (which *is* exactly 2:59pm). So, instead, one minute later (at 3), I look carefully at my watch, which (because it happens to be reading one minute slow) reads exactly 2:59pm. I reason: "It is exactly 2:59pm (*p*); therefore (*q*) I am not late for my 7pm meeting." Thus (supposing Warfield is right), I have inferential knowledge that *q*, based on a relevant premise *p*, which is a falsehood. Now, for the twist. If my belief that *p* had been *true*, then (we can plausibly suppose) it would have been based on my reading (at exactly 2:59pm) of the Campanile clock, which would have read exactly 2:59. Unbeknownst to me, however, the Campanile clock has been (and would have been) *stuck at 2:59* for some time.

- It seems to me that I do *not* obtain inferential knowledge of *q*, on the basis of *p*, in the counterfactual scenario. [See Luzzi's [10] for an insightful diagnosis/discussion.]

- If this is correct (and assuming that Warfield is correct about his original case), then we have a *stronger* KFF…

- ... we seem to have a case involving inferential knowledge of *q* on the basis of a false relevant premise *p*, *and such that*:
  - ☞ If *S*'s belief that *p* had *not* been false, then *S* would *not* have been in a position to know that *q* on the basis of a competent deduction from *p*.
- Now, *S*'s belief that *p* is *not merely* "*causally* essential" to the *production* of *S*'s knowledge that *q* [8]. The *falsity* of *S*'s belief that *p* seems explanatorily/epistemically relevant.
- There are several reasons why this is important:
  - Commentators (to date) have not focused on the precise role that the *falsity* of *S*'s belief that *p* can play.
  - Commentators (to date) seem to presuppose that it is *despite* the falsity of *S*'s basis belief that *S* knows *q*.
  - Some commentators presuppose that there must be a *specific* "*nearby* truth" that plays a certain epistemic role. This example (and variants) will call this into question.
- Next, I'll discuss one recent line of "resistance" to purported examples of (KFF)/(KFNK). This is E.J. Coffman's line [2].

- Coffman [2, *pp*. 190-1] conjectures that in *all* cases of (KFF)...

    ...we can identify a *true proposition* p' [I will refer to p' as p's "epistemic proxy"] with the following two features:

    - the subject is (at least) disposed to believe p',
    - if the subject's inferential belief (that q) had been based on a belief in p', the inferential belief would (still) have constituted knowledge.

- In the cases on which I am focusing, Coffman's proxy would be:

    (p') It is *approximately* 2:59pm (*e.g.*, 2:59pm ± 2 minutes).

- We can amend our last example, so as to *refute* (*this* version of) Coffman's conjecture. To wit, consider this amendment:

    I am confident that my fancy watch is *exactly* accurate, whereas I believe that the Campanile clock is only accurate to within (say) two minutes. And, as a result, I am disposed to come to believe "it is *approximately t*" when I look at the Campanile clock and it reads exactly t; whereas, I am disposed to come to believe "it is *exactly t*" when when I look at my fancy watch and it reads exactly t.

- Having said that, I think there is *something* right about this "approximate truth" idea (remember, Hilpinen thought so too).

- And, I think there are more promising "proxy" approaches.
- Recall, Coffman's choice of proxy for Warfield's case was:
  - (p') It is *approximately* 2:59pm (*e.g.*, 2:59pm ± 2 minutes).
- But, why not go for the following *alternative* proxy?
  - ($p^\star$) My *watch reads* 2:59pm, and it is "reasonably" accurate.
    After all:
    - (i) I am (plausibly) *disposed to believe* $p^\star$ in the example.
    - (ii) It could be argued (plausibly) that if my belief that q had been based on $p^\star$ (rather than p) then it would (still) have constituted inferential *knowledge* that q.
- I think this is the most promising "proxy" line. But — a dilemma:
  - (a) $p^\star$ *entails* q. Then, presumably, $p^\star$ must *also* entail *some* "approximate truth" claim *like p'*. And, then, it seems we've just slipped back into "Coffmanian resistance" territory.
  - (b) $p^\star$ does *not* entail q. Then, we would have a *non*-deductive inference "going proxy" for what seemed (*actually*) to be an instance of *deductive* inferential knowledge that q.

- This brings us (finally) to a more general problem with any of the "proxy"-style approaches to resisting (KFF)/(KFNK). I will call this problem *The Actual Reasons Problem*. To wit:

  (AR) The *actual* inferential route by which *S* arrives at a belief that *q* (and perhaps basing relations more generally) are always relevant to whether *S* (inferentially) knows that *q*.

- Like (CC), (AR) is something that many traditional accounts of inferential knowledge seem to take for granted.

- However, in light of (*prima facie*) examples of (KFF)/(KFNK), there seems to be a *tension* between (CC) and (AR).

  ☞ The most natural strategies for *resisting* (KFF)/(KFNK) — *viz.*, *defending* (CC) — seem to involve *repudiating* (AR).

- So, the main upshot of my talk today is that I'm afraid we're going to have to give up *at least one* of (CC) or (AR).

- Depending on one's theory of inferential knowledge (and one's other theoretical epistemological commitments), this may be a rather surprising and unwelcome result.

[1] M. Clark, *Knowledge and grounds: a comment on Mr. Gettier's paper*, *Analysis*, 1963.

[2] E.J. Coffman, *Warrant without truth?*, *Synthese*, 2008.

[3] B. Fitelson, *Strengthening the case for knowledge from falsehood*, *Analysis*, 2010.

[4] _____ , *Knowledge from non-knowledge*, unpublished manuscript, 2010.

[5] E. Gettier, *Is justified true belief knowledge?*, *Analysis*, 1963.

[6] G. Harman, *Thought*, Princeton University Press, 1973.

[7] R. Hilpinen, *Knowledge and conditionals*, *Philosophical Perspectives 2 (Epistemology)*, 1988.

[8] P. Klein, *Useful false beliefs*, *New Essays in Epistemology*, Q. Smith (*ed.*), 2008.

[9] F. Luzzi, *Counter-Closure*, *Australasian Journal of Philosophy*, September, 2010.

[10] _____ , *Knowledge from ignorance*, *Phd. Dissertation, U. Aberdeen*, December, 2010.

[11] D. Miller, *Out of Error: Further Essays on Critical Rationalism*, Ashgate, 2005.

[12] J. Saunders and N. Champawat, *Mr. Clark's definition of 'knowledge'*, *Analysis*, 1964.

[13] T. Warfield, *Knowledge from falsehood*, *Philosophical Persepctives 19 (Epistemology)*, 2005.