

No Knowledge from Falsehood but from Reflective Knowledge

Frank Hofmann
Université de Lusembourg

Yannick Kohl
University of Luxembourg

Abstract: Some philosophers have claimed that there is knowledge from falsehood (KFF, in short), i.e., inferential knowledge that involves a relevant false premise. The main thesis of this paper is that there are *no standard cases of KFF*. By ‘standard cases’ we mean cases in which the subject employs a measurement procedure in order to determine the value of some quantity, such as the time or the number of people present in a room. If knowledge is attained at all, it is attained by inference not from a false premise but, instead, from two pieces of reflective knowledge. A brief look at putative non-standard cases of KFF reveals that they can be explained away in a different way, albeit a way that shares a common pattern with standard cases.

Key words: knowledge, inference, reflection, margin of error, measurement

1. INTRODUCTION

Knowledge from falsehood (KFF) is defined as knowledge with two features: it is purely inferential, and it involves a relevant false premise. This is how Ted Warfield, who sparked off the recent discussion about KFF, understands KFF, and we will follow him on this point.¹ *Standard cases* of KFF, as we will call them, are cases of KFF in which the subject uses some *measurement procedure* when forming the relevant premise belief, such as counting the number of people in a room or looking at one’s watch in order to determine the time. (Non-standard cases of KFF would be cases that do not involve any such measurement procedure.) Of course, the measurement procedure need not be an elaborate or scientific one. Ordinary life involves lots of measurement procedures. People just employ certain ways of determining the value of a quantity, be it number or time or what have you. They count, they consult clocks, they use tape measures to determine the size of their living room, and so on. It is fair to say that in the discussion about KFF the standard cases of KFF are the ones that are usually mentioned first and that these are the intuitively most compelling candidate cases. The falsity of the false premise belief is simply due to the fact that the measurement yields the wrong result in this instance. For example, on the basis of having counted the people in the room, the subject comes to believe that there are 53 people in the room, but in fact there are 52 people. She draws the conclusion that there are less than 100 people in the room. Proponents of (standard

Frank Hofmann, *Institute of Philosophy, Department of Humanities, University of Luxembourg, 11, Porte des Sciences, 4366 Esch-Belval, Luxembourg*; <https://orcid.org/0000-0001-9730-8930>, frank.hofmann@uni.lu.

Yannick Kohl, *Institute of Philosophy, Department of Humanities, University of Luxembourg, 11, Porte des Sciences, 4366 Esch-Belval, Luxembourg*; <https://orcid.org/0000-0002-4508-5525>, yannick.kohl@gmail.com.

cases of) KFF hold that the resulting conclusion—the belief that there are less than 100 people in the room—is an instance of knowledge from falsehood.²

Knowledge from falsehood account:

(KF1) There are 53 people in the room. (False premise)

(KFC) The number of people in this room is less than 100. (Known conclusion)

2. EXPLAINING AWAY STANDARD CASES OF KFF

Since the beginning of the discussion about KFF, there have been attempts at explaining away the impression that there are (standard) cases of KFF. A prominent explaining-away strategy has been to *postulate approximate beliefs*—call this the ‘proxy-belief account.’³

Proxy-belief account:

(P1) There are approximately 53 people in the room. (Known premise)

(PBC) The number of people in this room is less than 100. (Known conclusion)

We will present an alternative account that explains away putative KFF in standard cases without postulating proxy beliefs. We can call it the ‘*reflective account*.’ According to this proposal, the terminal knowledge is attained in a *reflective* or *meta-cognitive way* (in case there really is any such terminal knowledge).

Here is how it works. The subject holds a *true* belief about the result of her measurement, and together with some reflective knowledge of the margins of error of her measurement procedure she inferentially arrives at a final piece of worldly knowledge (that there are less than 100 persons, e.g.) *without relying on any false premise*. Obviously, ‘measurement outcome’ here does not mean the real value of the relevant quantity, since the subject does not have to believe (and perhaps does not believe) that the real value is the same as the measured value. Rather, the true premise belief is the belief that the measurement assigns a certain value to the measured quantity. Thus, two pieces of reflective knowledge are put together and yield further knowledge of the world. We can represent the inference in the following way:

The reflective account:

(1) The result of my counting procedure as to the number of people in this room is 53. (Known premise)

(2) The margin of error for my counting procedure in this context is approximately plus/minus five. (Known premise)

Thus,

(3) The number of people in this room is less than 100. (Known conclusion)

No false premise belief is relied on.⁴

One could object that the second premise does bring in a proxy belief after all. It is about an *approximate* margin of error. And so, its content qualifies it as a proxy belief. However, there is a crucial difference between this premise and the premise that the proxy-belief account postulates. The proxy account attributes the worldly belief that there are approximately 53 people in the room. This belief is about the value of an observable quantity (and it is true, we can suppose). The premise

belief (2), however, is about the margin of error of one's own measurement procedure (and it is also true, we can suppose). In fact, it qualifies as a piece of reflective knowledge. (In case the subject does not know (2), she would not attain knowledge of (3) either.) The subject knows something about the accuracy of her own act of measurement, where this knowledge might come from a variety of sources, such as previous experience, general background knowledge about human capacities, knowledge about the particular context, and so on. There is no unique way of arriving at such knowledge of margins of error of one's own methods of answering questions. Arguably, this piece of reflective knowledge is often not precise—though it might be on a particular occasion. So, the way in which approximation enters into the reasoning does not make the meta-cognitive account collapse into the proxy-belief account.

Is a first-person perspective needed to make the inference? It seems that the subject can also engage in a third-personal inference, without any epistemic loss, as follows:

The reflective account, third-personal formulation:

(1') The result of this counting procedure as to the number of people in this room is 53. (Known premise)

(2') The margin of error for this counting procedure in this context is approximately plus/minus five. (Known premise)

Thus,

(3) The number of people in this room is less than 100. (Known conclusion)

'My' has been replaced by 'this.' (Instead of 'this' one could also use a name to refer to the relevant measurement, or some other non-first-personal way of referring.) No *de-se* contents are involved in the premises anymore. However, we still have an interesting feature of reflectivity: the method of acquiring knowledge is (reflective in the sense of being) *self-referential*. The premises of the inference refer to some aspects of the very method that the subject is using. If, for example, the subject uses a counting machine, the subject knows something about its result and its margins of error, and thereby she can arrive at the terminal knowledge about the number of people. So, the method of acquiring knowledge is self-referential, according to the reflective account. (And this is not so according to the friends of knowledge from falsehood.)⁵

3. A DIAGNOSIS

Here is a diagnosis of what is going on. The inquiring subject is not so naïve as to take the result of her measurement at face value. The apparent result suggests perfect precision. But the subject is more reflective and does not succumb to this misleading impression, and rightly so. Through ordinary experience people are familiar with the fact that measurements are not perfectly precise, and they take account of the limited resolutions of ways of measuring things.⁶ This allows us to infer the conclusion from genuine knowledge. No false premise is needed. So, everything is kosher: knowledge arises from other bits of knowledge, not from falsehood (which would cry out for explanation). This favors the reflective account over the knowledge from falsehood account.

What if the subject is naïve and takes her measurement result at face value (in addition to believing the premises (1') and (2'))? In this case, we get knowledge *despite* falsehood but still no knowledge *from* falsehood. As long as the subject has and relies on the knowledge of (1') and (2') she can be said to attain knowledge of the conclusion (3). The additional false belief about the precise number of people might be there as well, but it drops out as irrelevant, arguably.⁷

On reflection, the idea of knowledge from falsehood would appear deeply puzzling. One big problem it leads to is that one must find a way of distinguishing between KFF and inferential Gettier cases where no knowledge is attained.⁸ In both cases, subjects essentially rely on a false

premise. In Gettier cases they fail to acquire knowledge; yet the proponent of KFF claims that in alleged KFF cases, the subject does acquire knowledge. How so? What is the crucial difference? A second problem is that in cases of KFF the subject would point at the content of a false belief as her motivating reason for believing the conclusion when asked about her reason, if we suppose that the reasoning is conscious, accessible to the subject herself—which is surely the case, at least typically. Arguably, however, a false content cannot be evidence. If this is true, the subject would think of herself as relying on a normative reason for belief even though in fact it is merely a motivating reason but no genuine evidence and, thus, no normative reason for belief. How could this lead to knowledge? It is preferable to think of these cases, if they are cases of inferential knowledge at all, as cases of knowledge from reflective knowledge or knowledge despite falsehood.

The issue is highly significant for a number of reasons. In our view, the most important reason is that the proposed meta-cognitive account of alleged standard cases of KFF shows the importance of *meta-knowledge* and *epistemic reflection*, i.e., the capacity of using reflective knowledge about one's own ways of forming beliefs for rationally improving one's view of the world. In many cases we can determine what the world is like by relying on knowledge about our own ways of forming beliefs. (This is in principle nothing puzzling. By epistemic reflection we revise our beliefs, and, in some good cases, we can thereby attain new knowledge about the world.) Thus, epistemic reflection is of central epistemic importance. Furthermore, because the employment of measurement procedures is ubiquitous in modern science, epistemic reflection can be seen to lie at the basis of modern science.⁹ And indeed, isn't this what we think about modern science? It proceeds empirically, to be sure, by gathering evidence through measurement, but it does so in a reflective way, i.e., by reflecting on the inherent limitations of measurement, instead of naïvely identifying measurement results and real facts.¹⁰ Some of our ordinary life knowledge, such as in cases of counting people, is similar in that it is equally guided by epistemic reflection. Ordinary life and science are continuous and share reliance on epistemic reflection.

4. A BRIEF LOOK AT NON-STANDARD CASES

How about other, non-standard cases of KFF? It is hard to tell since no general taxonomy of all the cases of KFF has been found yet, at least to our knowledge. Some cases sound very much like standard cases, though it's not entirely clear. However, Peter Klein's Santa Clause Case is *not* like that.

The Santa Clause Case: Mom and Dad tell young Virginia that Santa will put some presents under the tree on Christmas Eve. Believing what her parents told her, she infers that there will be presents under the tree on Christmas morning. She knows that. (Klein 2008: 37)

Intuitively, in this case the premise is unnecessarily rich (implying the existence and action of Santa Claus), relative to the terminal conclusion. One might want to argue that it is a case of knowledge *despite* falsehood, and not of knowledge *from* falsehood.¹¹ According to such an interpretation, arguably, a corresponding weakened premise (leaving Santa Claus out of the set of events) is also believed by the child, and this *true* premise properly epistemizes the conclusion that there will be presents under the Christmas tree again this year.

Now, importantly, a common pattern becomes visible. Both in standard cases and in Klein's case the contents of the premise beliefs are too strong to be true. The subject naïvely endorses her measurement result that there are 53 people in the room; the child believes not only that there were presents under the Christmas tree in previous years but also that they were brought by Santa Claus. We can provide an alternative interpretation, however, according to which the relevant belief contents are weaker, true, and known. (They are weaker in the sense of not entailing any precise number of people.) In standard cases, the subject merely believes and knows something about her measurement procedure. In Klein's case, the child believes and knows that there were presents under the Christmas tree in previous years, and that is a good inductive basis for coming to know

that it will happen again this year. The falsehoods have dropped out of the picture, have become irrelevant. Charitably interpreted, the subjects has these weaker bits of knowledge. So, there is at most knowledge despite falsehood, but no knowledge from falsehood. To boldly generalize, we can say that whenever such a content weakening re-interpretation is possible, we can move from putative knowledge from falsehood to true knowledge despite falsehood. What is special about the standard cases is that reflective knowledge replaces the false belief, and rightly so, since the proper use of measurement procedures should involve reflective awareness of margins of error.¹²

NOTES

We would like to thank Thomas Grundmann, Susanne Mantel, Thomas Raleigh, and Kieran Salt for helpful discussions and constructive criticism. Special thanks to an anonymous referee for helpful comments. Financial support for this work was provided by the project “RREXAM—Rational Reflection Examined” funded by the FNR Luxembourg (O18/12880605/RREXAM).

Competing interests declaration: The authors declare no competing interests.

1. Cf. Warfield 2005. ‘Relevant’ means that the premise could not simply be left out while still attaining knowledge (as in cases of overdetermination). Cf. Warfield 2005: 405. Sven Bernecker speaks of the false premise being ‘essential’ (cf. Bernecker 2022: 1625). We take it that this is supposed to express the same idea of epistemic dependence of the attained knowledge on the ‘relevant’ premise.
2. This is a slightly simplified example from Warfield 2005 that has become the most popular example of KFF in the meantime.
3. Proponents of the proxy-belief account are Coffman (2008), Ball and Blome-Tillmann (2014), Schnee (2015), and Lee (2021). A problem for the proxy-belief account is that it is rather unclear whether the postulated beliefs really exist. On an interpretationist view (like Davidson’s) it might be the case, but on a more robust mental representation view, e.g., it might not be the case. A further problem is that it is far from clear that these beliefs, even if they exist, are relied upon by the subject.
4. One might think that a further, third premise is needed. Perhaps, the subject has to use some premise like the following: If the result of my measurement is 53 and my margin of error for counting people in this context is approximately plus/minus five, then the real value of the measured quantity is less than 100. We might include this further premise. This would not affect the main points of the argument. (Perhaps, such a third premise is sometimes present, and sometimes not. We should keep in mind that good inference need not be deductively valid.)
5. One might wonder whether the demonstrative concept expressed by ‘this’ introduces *de-se* content by the backdoor. An affirmative answer would surely require an extensive, non-trivial argument. For the present purposes we can leave this issue open, since an affirmative answer is not necessary for the reflective account to work. We are very grateful to an anonymous referee for pointing out that our original treatment here was flawed.
6. It is surely not entirely clear what exactly ‘perfectly precise’ means in this context. Values from the set of real numbers might be the paradigm, but we need not assume that this is the only possibility. In any case, fortunately, we need not settle this issue here in order to spell out the view.
7. What if the measurement result is off by much more than the margin of error but the terminal belief is very very weak and true? To us it is not clear that such a case is really possible. It depends on how exactly ‘margins of error’ are understood. If a margin of error guarantees (in the safety sense) that the real value is within the interval that it yields from the result of the measurement, it is impossible to have such a case. (Other construals of ‘margin of error’ might allow for the case.) We are inclined to classify it as a *non*-standard case in our sense. Thanks to an anonymous referee for bringing up this challenging case.
8. That this problem is much bigger than expected has been shown by Sven Bernecker recently. Cf. Bernecker 2022.
9. Recently, the importance of massive attempts at gathering evidence through experiments in modern science has been pointed out in Strevens 2020.
10. A similar point could be made about the use of *models*. Models are often said to be false, or even bound to be false (since idealizations). But we seem to be able to attain knowledge from models. So we get another case of KFF? No. The proper use of models has to invoke knowledge about how to interpret the model in order not to

go beyond what can be justifiably inferred about the target system. Importantly, properly interpreted models in general represent empirical observables only up to a certain resolution, and so they don't have to be false. For example, classical mechanical models in physics can accurately represent systems at smaller speed than the speed of light when the assigned degree of accuracy is not extremely high. The skill of properly interpreting models, including the assignment of an appropriate resolution to the contents represented by the relevant features of the model, is required for gaining new knowledge from the model. The model must be interpreted as placing the real value of an observable in a certain interval, and not as identifying it with a precise number (a number of the set of real numbers, say). The skill of the model user is to identify this interval with the help of antecedent background knowledge. Needless to say that these thoughts could profit from further elaboration.

11. Here we are in agreement with Schnee's diagnosis of Klein's case. Cf. Schnee 2015: 71–72.
12. What about cases in which the subject infers the conclusion from a false imprecise belief about an interval of an observable (e.g., that a tree is between 90 and 110 inches tall, when in fact it is 120 inches tall) on the basis of perception merely? No measurement procedure is used, just perception, and so it is not a standard case as we conceive of it here. (In contrast, finger-pointing does involve a measurement and not just perception.) Then, the subject could still come to know that the tree is less than 200 inches tall, say. This would be knowledge despite falsehood if the subject also holds some weakened (imprecise) belief that constitutes knowledge (e.g., that the tree is between 70 and 130 inches tall) and the stronger imprecise belief becomes irrelevant. Intuitively, if no such weakened (imprecise) belief is held, no terminal knowledge is attained (unless some further background knowledge, like the knowledge that trees of the kind at hand are always less than 200 inches tall, is relied on). So depending on how the case is spelled out we get no knowledge from falsehood, but merely knowledge despite falsehood, or no knowledge at all. We are grateful to an anonymous referee for bringing up this case.

REFERENCES

- Ball, B., and M. Blome-Tillmann. 2014. "Counter Closure and Knowledge Despite Falsehood," *Philosophical Quarterly* 64(257): 552–568. <https://doi.org/10.1093/pq/pqu033>
- Bernecker, S. 2022. "Knowledge from Falsehood and Truth-closeness," *Philosophia* 50: 1623–1638. <https://doi.org/10.1007/s11406-022-00479-y>
- Coffman, E. J. 2008. "Warrant Without Truth," *Synthese* 162(2): 173–194. <https://doi.org/10.1007/s11229-007-9178-5>
- Klein, P. 2008. "Useful False Beliefs," in *Epistemology: New Essays*, ed. Q. Smith, 25–61. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199264933.003.0003>
- Lackey, J. 1999. "Testimonial Knowledge and Transmission," *The Philosophical Quarterly* 49(197): 471–490. <https://doi.org/10.1111/1467-9213.00154>
- Lee, K. Y. 2021. "Reconsidering the Alleged Cases of Knowledge from Falsehood," *Philosophical Investigations* 44(2): 151–162. <https://doi.org/10.1111/phn.12296>
- Schnee, I. 2015. "There Is No Knowledge from Falsehood," *Episteme* 12(1): 53–74. <https://doi.org/10.1017/epi.2014.26>
- Strevens, M. 2020. *The Knowledge Machine*. New York: Liveright Publishing Corporation.
- Warfield, T. 2005. "Knowledge from Falsehood," *Philosophical Perspectives* 19: 405–416. <https://doi.org/10.1111/j.1520-8583.2005.00067.x>