

CHAPTER I

*Knowledge from Knowledge***1.1 The Default View**

Suppose a subject competently infers a conclusion from a set of premises, each of which is essential to their inference. It is received wisdom in epistemology that in order for that subject to thereby acquire knowledge of the conclusion, they must know the premises they proceed from. This results from the seemingly compelling thought that no matter how competently the subject performs an inference, their belief in the conclusion can be no better, epistemically speaking, than the beliefs in the premises it is drawn from; consequently, failure to know one of the essential premises will result in failure to know the conclusion drawn from those premises.

This thought is particularly attractive when we focus on inferences proceeding from a single premise. For the time being, let us concentrate on this kind of inference and call ‘knowledge from knowledge’, or KFK, the view that requires the premise of a single-premise inference that yields knowledge of its conclusion to be knowledge.

There are several facts indicating that this first-blush highly intuitive view enjoys default status within contemporary epistemology. Firstly, until relatively recently KFK was neither the target of critical discussion nor the subject of defence. On the rare occasions KFK has been mentioned, philosophers have usually been happy to rely on it without providing an argument or acknowledging the need for one, as a stepping stone in their arguments towards further conclusions. For example, Timothy Williamson uncritically relies on KFK in his margin-for-error principle (1994: 222). In other work, while discussing whether imaginative exercises can be recast as pieces of reasoning, he writes in passing:

Normally, someone who believes a conclusion on the sole basis of inference from some premises knows the conclusion only if they know the premises ... The principle applies only to essential premises, those that

figure in all the inferences on which the relevant belief in the conclusion is based. (2007: 145)

Secondly, and relatedly, the view is deemed uncontroversial enough to be presented to epistemology students as received wisdom. In his epistemology textbook, for example, Robert Audi writes:

[w]e can extend our justification and knowledge by inference, but it appears that if we have none to start with, inference ... can give us none. (2010: 184)

and

[o]ne kind [of condition on inferential justification and knowledge] concerns the premise(s) of the inference – its foundations, so to speak ... First there are source conditions ... one needs justification or knowledge in the first place. (2010: 185)

Audi gives no explicit defence of KFK – only some examples that illustrate this view at work.

Thirdly, some prominent epistemologists have casually attached conditions that entail KFK to their theories of knowledge without substantial argument or independent motivation, as if it were a natural *desideratum* of one's theory that KFK follows logically from it. For example, Robert Nozick begins his search for an analysis of inferential knowledge by laying down two conditions:

S knows via inference (from p) that q if and only if:

- (1) S knows that p
- (2) q is true, and S infers q from p (1981: 231)

Nozick give no serious argument in favour of condition (1) and limits himself to swiftly remarking that it must be true '[o]therwise, there is no knowledge to transmit' (1981: 239).¹

Together, these three considerations suggest that philosophers do not typically consider KFK to be in need of substantial defence, and that they assume their audience shares this perception. In other words, KFK is perceived to be not only plausible but also *uncontroversial*. It is testament to the power of this perception that even after the publication of Warfield's (2005) first influential attack on this view and the literature that has come in its wake, some philosophers have advanced theories of knowledge that

¹ See also Stanley (2005: 90) for a move of this sort.

1.1 The Default View

3

stand in tension with Warfield's cases of knowledge from falsehood without acknowledging the need to address these tensions.²

Yet when we look for KFK's credentials, it is not clear what they are. There is the thought, previously articulated, that the epistemic pedigree of a conclusion can be no better than that of the premises it is inferred from. But this seems too close to a restatement of KFK, and it seems plainly false on certain understandings of 'epistemic pedigree'. For example, in Bayesian settings, when p entails q , the credence commanded by q is *no lesser* than the credence commanded by p . Alternatively, some might think that inference is a way of merely 'teasing out' knowledge that is in some sense already 'contained' in the premises; consequently, this thought runs, if there is no knowledge of the premises there can be no knowledge of the conclusion to be 'teased out'. But this consideration seems too metaphorical to be of dialectical use. Finally, the Nozickean thought that inferential knowledge requires knowledge of the premise because otherwise 'there is no knowledge to transmit' depends the idea that inference is essentially a matter of transmitting epistemic properties enjoyed by the premises to the conclusion – which, while *prima facie* plausible, is not obviously true.

The lack of obvious and compelling justification for KFK should not on its own be cause for concern to the many who, implicitly or explicitly, endorse KFK. After all, a similar difficulty arises when philosophers go looking for the credentials of principles that are thought to be epistemologically basic or self-evident. However, alarm bells ring when we consider an analogy with the principle of Knowledge Closure. Roughly, according to Knowledge Closure, competently deducing q from a known premise p yields *knowledge* of q . This principle was deemed uncontroversial until Dretske (1969, 1970, 1971), and later Nozick (1981) observed that certain relations and properties deemed necessary or at least important for knowledge (*is evidence for*, *is a conclusive reason for*, *is sensitively believed*) are not closed under competent deduction. For example, casual visual observation of a striped equine in a zoo pen marked 'Zebra' might constitute evidence or provide conclusive reason for believing that the animal is a zebra, but *not* for believing the entailed proposition that the animal is not a cleverly disguised mule. One can sensitively believe that the animal is a zebra without sensitively believing that it is not a cleverly disguised mule. These observations helped to spark the still-ongoing debate about the universal validity of Knowledge Closure and stripped this principle of its

² E.g., Ramachandran (2015).

indisputability. If conditions crucial or necessary to knowledge are not closed under competent deduction, it is legitimate to wonder why knowledge should be.³

By analogy, there are epistemic conditions widely thought to be necessary to knowledge such that, where q is inferred from p , their being met by a conclusion q is no guarantee that they are also met by the premise p . One such condition is truth. Where q is inferred from p , it is not a requirement on the truth of q that p is true. After all, a falsehood can entail a truth. The false proposition that Socrates ran the 100 m sprint under 10 seconds entails the true proposition that someone has done so.

Safety and sensitivity are similar to truth in this respect: as we'll discuss, the fact that the conclusion of an inference meets these conditions does not mean that the premise from which the conclusion is drawn also meets these conditions. Just as with Knowledge Closure, an attitude of uncritical endorsement is inappropriate for KFK; closer scrutiny is needed.⁴ This is especially so since, as we'll soon see, KFK has been the target of several kinds of direct attacks in recent years.⁵

An examination of the nature and an assessment of the relative merits of these attacks will be our focus in the following two chapters. Before delving into discussion, however, it will be necessary to introduce some helpful terminology.

1.2 Knowledge Counter-Closure

Let's start by giving KFK a precise formulation, bearing in mind that we are restricting our focus initially to single-premise inferences. There are

³ These considerations, if true, might however not suffice to establish that knowledge is not closed under competent deduction. See Warfield (2004) for an argument to this effect and Yan (2013) for criticism of this argument. Nonetheless, they are clearly enough to *cast doubt* on Knowledge Closure in a way that demands further investigation.

⁴ In this connection, Fitelson (2016) notes that while it is natural to expect that competent inference preserve *good* epistemic qualities, such as truth, justification and knowledge, the converse expectation that it preserve *bad* epistemic qualities, such as falsehood, lack of justification and lack of knowledge, does not enjoy similar *prima facie* plausibility.

⁵ While the recent debate on KFK was sparked by the publication of Warfield's cases of inferential knowledge from falsehood (Warfield 2005), some cases with a similar structure appeared in the literature at an earlier stage, but did not have significant impact. For example, Risto Hilpinen proposed a case of alleged inferential knowledge from falsehood, and observed 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 what is *not* known (or even true)' (1988: 164). However, Hilpinen simply mentioned this anti-KFK thought and did not provide further discussion. Peter Klein (2008: 37) and Branden Fitelson (2010a) have observed that the earliest known example of alleged inferential knowledge from falsehood is in Saunders and Champawat (1964), though, again, the full significance of that example with respect to KFK was not discussed in that article.

1.2 Knowledge Counter-Closure

5

two reasons for this restriction: firstly, single-premise inferences provide a simpler model to work with (we will discuss multi-premise inferences in Section 4.6); secondly, this restriction will help us to engage with the relevant literature, which has similarly focused on the case of single-premise inferences. Those sympathetic to KFK should rest assured that no disservice to KFK is being done by this restriction, since this view seems *especially* plausible when the subject proceeds only from one premise. So we are not making KFK a straw man target; rather, we are focusing on its most compelling form.

The principle that encapsulates KFK for single-premise inferences is this:

Knowledge Counter-Closure (KCC): Necessarily, if (i) S believes q solely on the basis of competent inference from p, and (ii) S knows q, then S knows p.

A few details about KCC are worth noting. Firstly, clause (i) is meant to exclude both cases where S's belief that q is overdetermined and cases where one or more bases independent of the inference combine with the inferential basis to yield belief that q. For example, Maja might believe that Sarah was in town this week both on the basis of testimony from Maja's friend, who bumped into Sarah, and on the basis of inference from her independently held perceptual belief that Sarah was in town on Monday. Whether each of the two basing beliefs *independently* suffices to lead Maja to her belief that Sarah was in town this week, or whether instead they do so only jointly is not important. Both kinds of case fall outwith KCC's jurisdiction. KCC deals only with 'pure' cases where the only epistemic means exploited by the subject to reach their belief that q is the inference from p. One further important clarification: the clause 'solely on the basis of competent inference from p' in KCC's antecedent is intended strictly enough to rule out cases where one begins the inference without knowing that p but mid-inferentially acquires further evidence that makes p known. In cases of this kind, the subject may well end up with knowledge of q, but these should not be taken to be probative counterexamples to the view KCC expresses. KCC ranges only over those cases where the premise p is the sole foundation of the subject's belief that q.⁶

Secondly, KCC's name is not meant to suggest any opposition or tension with the principle of Knowledge Closure, with which it

⁶ A different way of dealing with this worry is to insist that the subject fails to know p throughout the inference. For simplicity, I will stick to the formulation provided and the strict interpretation of clause (i). See Hawthorne (2004: 33) for the parallel issue of mid-inferential knowledge-loss, which affects the formulation of Knowledge Closure.

undoubtedly shares some affinities. Rather, the ‘counter’ in its name is intended to suggest movement in the opposite direction. I am here relying on this parallel between Knowledge Closure and KCC: assuming clause (i) is fulfilled, Knowledge Closure guarantees that knowledge of the premise is accompanied by knowledge of the conclusion; assuming the same clause is satisfied, KCC guarantees that knowledge of the conclusion is accompanied by knowledge of the premise. Metaphorically, Knowledge Closure’s guarantee of knowledge acts *forward* through the inference: if you start with knowledge of the premise, you must wind up with knowledge of the conclusion. KCC’s guarantee of knowledge acts *backwards* through the inference: if you have wound up with knowledge of the conclusion, then you must have started with knowledge of the premise.⁷

The third point worth noting is that formulating KCC faces some of the same difficulties theorists are confronted with when attempting to formulate Knowledge Closure precisely.⁸ I have attempted to have KCC express KFK adequately while balancing the two desiderata of keeping KCC streamlined yet immune to obvious counterexample.

Fourthly, this principle provides some useful vocabulary. Philosophers often speak of knowledge and other epistemic conditions being closed (or not) under competent deduction to signify that they obey some closure-like principle: the fact that the premise meets the relevant condition guarantees that the conclusion does so, too. Similarly, we can speak of knowledge and other epistemic conditions being ‘counter-closed’ (or not) under competent deduction to signify that they obey some counter-closure-like principle: the fact that the conclusion meets the relevant condition guarantees that the premise does so, too. Our discussion will benefit from this shorthand.

Fifthly, the principle enjoys *prima facie* plausibility. It is buttressed by the same considerations that lend KFK its intuitive appeal.

1.3 A Schema for KCC Failures

We are now in a position to sketch a generic profile of potential KCC failures. These will involve a subject *S* who competently infers a conclusion *q* from a premise *p*, where:

- (I) *S*’s epistemic standing with respect to *p* is insufficient for knowledge, since some particular constraint on knowledge – call

⁷ This parallel between Knowledge Closure and KCC glosses over the fact that the former is restricted to deductive inference, whereas I have formulated KCC as ranging over both ampliative and non-ampliative inference. This difference is not important in this context.

⁸ See David and Warfield (2008) for probing discussion.

1.3 A Schema for KCC Failures

7

this X – is not satisfied. Besides failure to meet X, however, S's epistemic standing with respect to p is excellent – so much so that failure to meet X is the only obstacle standing between S and knowledge of p: all other things being equal, if X were met then S would know that p.

- (II) S believes that q on no basis other than competent deductive inference from p. Thus, other epistemic routes to belief in q are either unavailable or, if available, are not exploited by S.
- (III) S's epistemic standing with respect to q is extremely good: not only does it share all the good epistemic qualities of S's standing with respect to p, but it also satisfies X, i.e., the crucial constraint that S's standing with respect to p fails to meet.

Cases that meet the profile described in (I)–(III) are arguably cases of inferential knowledge from an unknown premise, or *knowledge from non-knowledge*. As we will see, there are various candidate conditions that can play the role of X, and which thereby account for S's failure to know the premise of its inference. Because the emerging dialectical situation will differ depending on this choice, several candidate conditions for X demand different treatment.

We will start in Chapter 2 with the kind of case that has received the greatest attention in the literature: inferential knowledge from falsehood. Discussion of various cases of alleged *knowledge from unknown truth* is the focus of Chapter 3. The viability of the cases discussed in these two chapters suggests exploring a view that denies KFK. In Chapter 4 I outline what I take to be a plausible and coherent version of such a view, making clear along the way that its endorsement is not as radical as might seem initially. My overall argumentative strategy is to explain that cases of knowledge from non-knowledge differ significantly from epistemically suspect cases of 'easy knowledge' and 'transmission failure', to show that plausible replacements to KFK are in the offing, and to highlight that some further unpalatable theoretical consequences that appear to result from denying KFK are, on reflection, in fact avoided.

CHAPTER 2

*Inferential Knowledge from Falsehood***2.1 Introduction**

In this chapter we will discuss examples of alleged inferential knowledge from falsehood in the recent literature. If the examples are *bona fide*, then we have a first kind of counterexample to KCC and the default view it expresses. In particular, these cases aim to show that the following principle, entailed by KCC on the plausible assumption that knowledge is factive, is false:

No-False-Premise: Necessarily, if (i) S believes q solely on the basis of competent inference from p, and (ii) S knows q, then p is *true*.

We'll also examine in some detail the arguments put forward in defence of the *No-False-Premise* principle, and therefore of KCC, by those who think that the proposed cases fall short of their target and that inferential knowledge from falsehood cannot exist. I will focus in particular on Martin Montminy's (2015) response.

The discussion in this chapter constitutes a first step in setting up a comparison among various types of challenge to KCC, which will be one of the foci of Chapter 3.

2.2 Warfield on Inferential Knowledge from Falsehood

The recent literature on inferential knowledge from falsehood (hereafter KFF) was sparked by Ted Warfield's description of several cases of putative knowledge obtained by single-premise inference from a false premise (Warfield 2005).¹ The cases he presents are structurally similar. Here are two of them:

¹ Peter Klein (2008) also proposed and discussed similar cases. I will present Klein's discussion in Section 2.7.

2.2 Warfield on Inferential Knowledge from Falsehood

9

Fancy Watch: I have a 7 p.m. meeting and extreme confidence in the 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: 'It is exactly 2:58 p.m.; therefore, I am not late for my 7 p.m. meeting' ... I know my conclusion, but as it happens it is exactly 2:56 p.m., not 2:58 p.m. (Warfield 2005: 408)

Border: CNN breaks in with a live report. The headline is 'The President is speaking now to supporters in Utah'. I reason: 'The President is in Utah; therefore he is not attending today's NATO talks in Brussels'. I know my conclusion but my premise is false: the President is in Nevada – he is speaking at a 'border rally' at the border of those two states and the speaking platform on which he is standing is in Nevada. The crowd listening to the speech is in Utah. (ibid.)

At first blush, it seems that the subjects in these cases know their conclusion even though this is based solely on a false premise (respectively, *it is exactly 2:58 p.m.* and *the president is in Utah*). If so, then inferential knowledge from falsehood is possible, and the *No-False-Premise* principle is undermined.

The main idea behind all the cases proposed by Warfield is that the subject's premise fails to be true, but not by much, and in any case not by an extent capable of compromising the subject's believing the conclusion truly on its basis. To use Warfield's metaphor, some falsehoods lead to truths reliably, and other falsehoods lead to truths unreliably. For example, if in *Fancy Watch* the subject had performed the inference on the basis of the false premise *it is exactly 2:58 p.m.* to the conclusion *it is not 2:57 p.m.* then the latter belief, while true, would have been too luckily true to count as knowledge. In cases of KFF, by contrast, the stability of the 'path to the truth' of the conclusion is unaffected by the falsehood of the premise because the small margin by which the subject's belief is inaccurate is insignificant to the truth of the conclusion. The truth of the conclusion does not seem hostage to accidentality or luck in the way that the belief *it is not 2:57 p.m.* would be if inferred from the same false premise.

According to Warfield, this is a mark of distinction between cases of KFF and cases where, indisputably, a false premise leads to a conclusion that is true but *not* known. Recall that the first response to Edmund Gettier's (1963) examples of justified true belief that intuitively fail to constitute knowledge suggested that this knowledge failure was explained by the presence of a false lemma in the subject's inference to the justified true belief at issue (Clark 1963). For example, suppose Smith has strong but misleading evidence that Jones will get the job as well as strong

evidence that Jones has ten coins in his pocket. Smith infers the justified true belief *the person who will get the job has ten coins in his pocket*. In fact, unbeknownst to him, it is Smith who will get the job and he also happens to have ten coins in his pocket. Smith's inferred conclusion is not a case of knowledge from falsehood. A key difference between this case and the cases of KFF proposed by Warfield is that the false premise *Jones will get the job* makes the path to the truth of the conclusion too unstable. A plausible way of expressing this instability is by saying that, while true, Smith's belief *the person who gets the job has ten coins in his pocket* could very easily have been mistaken. By contrast, it seems that in *Fancy Watch* the subject's belief *I am not late for my 7 p.m. meeting* could not easily have been mistaken, given the substantial time remaining until the meeting and the very slight inaccuracy of the watch (whose inaccuracy would in any case lead the subject to be early rather than late). If taken at face value, then, Warfield's cases show not just that KFF is possible, but that the 'no-false-lemmas' constraint on knowledge espoused by some theorists as an explanation of knowledge failure in Gettier's original cases does not generally apply to all cases where a subject infers a justified true belief from a falsehood.

2.3 Warfield against the 'Proxy Premise' Strategy

Warfield's main effort in his discussion is to show that retaining the *No-False-Premise* principle in the light of his proposed cases of KFF is implausible. To resist these alleged counterexamples, defenders of *No-False-Premise* must argue either that S does not know her conclusion, or that her belief in the conclusion is not based solely on inference from the false premise. Warfield observes that 'clear and widely shared intuitions about the cases' (2005: 408) prevent us from arguing that in these cases the subject lacks knowledge; clause (ii) of the antecedent of *No-False-Premise* seems secure. Those wishing to defend *No-False-Premise* from these cases are better off arguing that clause (i) is unfulfilled, by claiming that the known conclusion is not believed solely on the basis of inference from the false premise, and that some other, true premise is somehow doing important work.

The most promising resistance strategy, which is in fact the one pursued by all of Warfield's critics, involves commitment to the view that in the relevant inferences, some true proposition 'in the neighborhood' of the false premise plays the relevant epistemic role in its stead. Let's call the true proposition that in each case could be argued to be acting in lieu of the