

WHY TO BELIEVE WEAKLY IN WEAK KNOWLEDGE:
GOLDMAN ON KNOWLEDGE AS MERE TRUE BELIEF

Christoph JÄGER
Universität Innsbruck

Summary

In a series of influential papers and in his groundbreaking book *Knowledge in a Social World* Alvin Goldman argues that sometimes “know” just means “believe truly” (Goldman 1999; 2001; 2002b; Goldman & Olsson 2009). I argue that Goldman’s (and Olsson’s) case for “weak knowledge”, as well as a similar argument put forth by John Hawthorne, are unsuccessful. However, I also believe that Goldman does put his finger on an interesting and important phenomenon. He alerts us to the fact that sometimes we ascribe knowledge to people even though we are not interested in whether their credal attitude is based on adequate grounds. I argue that when in such contexts we say, or concede, that *S* knows that *p*, we speak loosely. What we mean is that *S* would give the correct answer when asked whether *p*. But this doesn’t entail that *S* knows *that her answer is right* or that *S* knows *that p*. My alternative analysis of the Hawthorne-Goldman-Olsson examples preserves the view that knowledge requires, even in the contexts in question, true (firm) belief that is based on adequate grounds.

1. *Weak knowledge and firm belief*

Every now and then in the history of epistemology some ingenious philosopher offers an argument designed to show that knowledge reduces, at least in certain contexts, to mere true belief. The first one who toyed with this idea is Plato (in his *Meno*). Among the most recent ones is Alvin Goldman. Plato eventually rejects the reduction *tout court*. Goldman, by contrast, has argued in various places that there is “one sense of ‘know’ in which it means, simply, believe truly” (Goldman and Olsson, forthcoming, 1; cf. also Goldman 2002b, 185f.; 2001, 164f.; 1999, 24f.). In what follows I defend Plato’s rejection against Goldman’s endorsement. Knowledge, I argue, does not reduce to mere true belief, at least not in

the kinds of context and for the kinds of reason that Goldman takes to be his main witnesses.¹

In his most recent presentation of his argument, Goldman (with Eric Olsson) begins with the claim that there are contexts in which knowledge contrasts with ignorance and in which, for a specified person and fact, knowledge and ignorance are exhaustive alternatives. For example,

Diane either knows p or is ignorant of it. The same point can be expressed using rough synonyms of ‘know’. Diane is either *aware of* (the fact that) p or is ignorant of it. She is either *cognizant of* p or ignorant of it. She either *possesses the information* that p or she is uninformed (ignorant) of it. (Goldman & Olsson 2009, 1)

The argument proceeds by way of a *reductio*. Suppose that knowledge were, in contexts of the kind in question, *justified* true belief, or justified true belief that in addition fulfilled an anti-Gettier condition. Then, if p were the case but S failed to know that p , this could be so because S failed to meet the justification condition or the anti-Gettier condition or both. Thus, since the supposition is that in this context failing to know that p means being ignorant of p , S could be said to be ignorant of p despite holding the true belief that p . However, Goldman and Olsson argue, such a result would be “plainly wrong” or at least “highly inaccurate, inappropriate and/or misleading” with regard to the notion of ignorance (Goldman & Olsson 2009, 3). We can summarize this argument as follows. Suppose that:

1. Knowledge is, in every context, justified true belief (JTB) plus some further condition X. (Supposition)
2. In certain contexts, knowledge contrasts with ignorance and these alternatives are exhaustive. (Premise)

1. Authors who opt for the view that “know” ought to be analyzed invariantly as “believe truly” include Isaac Levi (1980), Crispin Sartwell (1991, 1992), and, in the German speaking philosophical literature Franz von Kutschera (1982), Georg Meggle (1997), and, with qualifications, Wolfgang Lenzen (1980). For a sympathetic discussion of Sartwell see Ansgar Beckermann (2001). An approving discussion of von Kutschera can be found in (Beckermann 1997); for a critical discussion of von Kutschera see (Brendel 1999, chapter 1). The most detailed defense in recent German epistemological literature of the view that there are certain circumstances in which we use “know” in the sense of holding a true (firm) belief is presented by (Ernst 2002; see especially part 2). A discussion of these authors is beyond the scope of this paper, yet I believe that much of what is to follow is relevant to their arguments as well.

3. There are contexts in which ignorance contrasts with JTB plus X and these alternatives are exhaustive. (From 1 and 2)
4. One can fail to have JTB plus X regarding p but hold the true belief that p . (Premise)
5. Hence one can be ignorant of p despite having a true belief that p . (From 3 and 4)
6. (5) is false. (Premise)

(6) yields the *reductio*. The argument is that, since (2), (4), and (6) are true and the inferences from (1) and (2) to (3) and from (3) and (4) to (5) are valid, (1) is false.

The inferences in this argument are indeed valid, and neither do I want to question premise (4) or premise (6), the latter resting on the view that being ignorant of p entails lacking a true belief that p . However, why should we think that (2) is true? Goldman and Olsson's argument for this claim adopts an example from John Hawthorne (2002), which they (re)formulate as follows:

If I ask you how many people in the room know that Vienna is the capital of Austria, you will tally up the number of people in the room who possess the information that Vienna is the capital of Austria. Everyone in the room who possesses the information counts as knowing the fact; everybody else in the room is ignorant of it. It doesn't really matter, in this context, where someone apprised of the information got it. Even if they received the information from somebody they knew wasn't trustworthy, they would still be counted as cognizant of the fact, i.e., as knowing it rather than as being unaware of it. (Goldman & Olsson 2009, 1f.)

If someone "possesses the information" that p , does he/she believe that p ? That seems to be the intended reading, a reading also suggested in an earlier paper where Goldman presents the story as follows:

... we would want to include anyone in the room who believes or possesses the information that Vienna is the capital of Austria, even if he acquired the information in an unjustified fashion. For example, even if his only source for this fact was somebody he knew full well was untrustworthy (but he believed him anyway) he should still be counted as knowing that Vienna is the capital of Austria. This seems, intuitively, exactly right—at least for one sense of the term 'know'. (Goldman 2001, 165)²

2. Here the story is adapted from Hawthorne's discussion of the example at the Rutgers

Neither of these passages is explicit about whether the knowledge *ascriber* knows that the subjects received their information from somebody they knew wasn't trustworthy. Yet the intended reading seems to be that the ascriber is indeed aware of this. This idea also underlies the following version of the example in Goldman (2002b):

Suppose a teacher *S* wonders which of her students know that Vienna is the capital of Austria. She would surely count a pupil as knowing this fact if the pupil believes (and is disposed to answer) that Vienna is the capital of Austria, even if the student's belief is based on very poor evidence. The teacher would classify the pupil as one of those who 'know' without inquiring into the basis of his/her belief, and even in the face of evidence that it was a poor basis. (Goldman 2002b, 185f.)

Here Goldman explicitly maintains that the knowledge ascriber would count a student as one who knows the fact in question even when aware that the student's epistemic basis is poor. Is that claim right?

An initial worry is that questions such as "How many people in the room know that Vienna is the capital of Austria?" or "Which of the students know that ...?" are leading questions in contexts in which it is known that at least one of the candidates could come up with the correct answer. The formulations suggest that "none" is not among the expected replies and that such a response would be inappropriate. Suppose the speaker had put his query in a more neutral way, for example by asking: "Are there any people in the room (except for you and myself) who know that ...?", or "Which students, *if any*, know that ...?". In that case the respondent, being aware that each candidate knows that his/her source is untrustworthy, might well have replied: "no", or "none", respectively. Why is this?

If *S* knows that Vienna is the capital of Austria, *S* holds the true belief that Vienna is the capital of Austria. How firm will this belief be? Goldman and Olsson don't address this issue. According to a widely held view, however, knowledge involves firmly held belief, i.e., belief in the sense of very high, or even maximal, confidence (conviction, certainty). Note that it is this kind of belief that figures, for example, in Hawthorne's original construal of the story. He writes:

Even if someone was given the information by an informant that they knew full well they shouldn't trust (who happened to be telling the truth on this

Epistemology Conference in 2000. An early version of Hawthorne's example was published in Hawthorne (2000).

occasion), you will in this context count him as knowing what the capital of Austria was (*so long as he had the firm belief*). [Footnote Hawthorne:] Of course, someone who didn't in fact trust their informant and merely used the informant as a basis for guessing an answer—being altogether unsure on the inside—would not count. (Hawthorne 2002, 253f., emphasis C.J.)³

If someone knows full well that his informant is untrustworthy, is it plausible to assume that he generates a firm belief, in the sense of high, or even maximal conviction, in what the informant says? No; this will typically not be the case. At least for people who are minimally epistemically rational (in the situation at hand), the following propositions form an inconsistent triad:

- (1) Knowledge requires firm belief.
- (2) *S* is confronted with a piece of information *p* from a source that *S* knows isn't trustworthy (in questions of the kind at issue).
- (3) *S* knows that *p* (solely) on the basis of the fact described in (2).

In Goldman and Olsson's example, (2) and (3) are assumed to be true. Hence, if their argument is to work, they must either reject (1) or take on board the idea that we can appropriately ascribe knowledge to a person even if he is highly irrational when generating a firm belief in the proposition in question. Let us begin by considering the first option.

2. *Super-weak knowledge*

(1) is shorthand for the claim that knowledge *always* requires firm belief (conviction, subjective certainty). Accordingly, this proposition may be rejected because one thinks that knowledge never requires firm belief; or because one thinks that at least sometimes, in certain circumstances, it doesn't require firm belief. This latter claim would do the job for Goldman and Olsson. More specifically, they might retort that in order not to be ignorant in the contexts they envisage one need not hold a firm belief, but only some weaker kind of credal attitude. Weak knowledge, they may argue, requires only weak belief.⁴ So what, exactly, is weak belief?

3. A slightly different and, according to Hawthorne, amended version of the example appears in Hawthorne (2004). I shall discuss this later version below.

4. Two more radical options would be to maintain (i) that knowledge never requires any kind of belief, whether firm or weak, or (ii) that at least in certain circumstances knowledge requires

In earlier publications, Goldman has expressed reservations about modeling degrees of belief in terms of subjective probabilities (see Goldman 1986, 324–28). Yet he does endorse the view that for many epistemological purposes an approach along such lines is a “tolerable idealization” (1999, 88). In such a framework, firm belief that p is to be modeled as a doxastic state in which the subject assigns a subjective probability of 1 to p . “Weak belief”, on the other hand, is an umbrella term for credal attitudes corresponding to a confidence interval of subjective probability assignments of $0.5 < \text{Pr}(p) < 1$. Let us call weak knowledge that only involves weak belief, understood in this sense, *super-weak knowledge*. Super-weak knowledge is merely true weak belief. Would recourse to super-weak knowledge solve our problem?

To begin with, note that knowledge that only involves weak belief is a technical notion that departs significantly from our pretheoretical usage of the term “know”. Suppose Diane asks Alfred: “Do you know when the next train to Vienna leaves?” “Yes”, he replies, and presents the correct answer: “It will depart at 8:15 a.m.” “Are you sure?” “No”, says Alfred. “That I’m not. I guess 8:15 is right. But I received the information from someone who was strolling around in front of the train station. Admittedly, the guy seemed a bit drunk and did not appear to be very reliable.” Surely, a most natural reaction for Diane would be to reply: “OK, but why do you say then you *know* when the train leaves?”

Goldman thinks that his notion of weak knowledge captures one ordinary way of using the term “know”. Weak knowledge, he argues, corre-

neither firm nor weak belief. For example, David Lewis once remarked (1996, 556) that he would “even allow knowledge without belief, as in the case of the timid student who knows the answer but has no confidence that he has it right, and so does not believe what he knows.” *Prima facie*, this may suggest an interpretation along the lines of (ii). However, note that Lewis portrays his timid student as not being *confident* that he has it right. Hence what the student lacks is firm belief. Lewis thus seems to be using “believe” here in the sense of firm belief (conviction, being certain), and in that case his timid student example, even if it were convincing, would not in fact yield an argument for the view that knowledge, at least in certain contexts, doesn’t require any kind of belief. Second, I shall argue (in section 5) that if we say that a subject “knows the answer” to a question, even though we are aware that she doesn’t believe that her answer is true, we merely mean that she is in a position to produce words that can be used to express the “right” proposition. But this does not entail that the subject knows *that the answer is right* and hence not that she has the relevant propositional knowledge. On this topic, cf. also Colin Radford (1966), who presents examples which he thinks show that “neither being sure that P nor having the right to be sure that P , can be necessary conditions of knowing that P ” (4). Another epistemologist who opts for the view that knowledge doesn’t require firm belief is Keith Lehrer. The required kind of credal attitude Lehrer calls “acceptance” (cf., for example, Lehrer 1990, 10f.).

sponds to a strict use of “know” that “conforms to some standard, ordinary sense of the term in colloquial English (as judged by what epistemologists who attend to ordinary usage have identified as such)” (Goldman 2002b, 182). Elsewhere he writes:

Is there any ordinary sense of ‘know’ that corresponds to true belief, or have I invented it? I believe there is an ordinary sense. (Goldman 1999, 24)

However, examples as the one above illustrate that we would be disinclined to concede that *S* knows that *p*, in any ordinary sense of the term, if we assume that *S* only holds a weak (true) belief that *p*. Super-weak knowledge does not seem to meet Goldman’s ordinary usage constraint.

However, note that Goldman also says that, should it turn out that his ordinary usage view is untenable, he’d be happy to treat weak knowledge as a term of art (1999, 24). Yet a second point, which undermines this option as well, is this. In our story about Diane and Alfred the protagonists suspect, but don’t *know*, that their source is untrustworthy. In Goldman’s examples the subjects’ situation is epistemically worse (or clearer, if you wish). In these examples the candidates do know that their source is untrustworthy. But then it is hard to see why they would form *any belief at all* that the capital of Austria is Vienna (and not, for example, Graz or Innsbruck). Why would any minimally rational subject under such circumstances even assign a probability greater than 0.5 to the information in question? (We are still assuming, with Goldman, that the subject has no independent evidence for the truth of the proposition.) If someone whom I know suffers from severe schizophrenia tells me that the Martians have landed, this would not motivate me to assign a probability greater than 0.5 to this proposition. (At least so I hope.) The problem with the Hawthorne-Goldman-Olsson example thus is that a subject would not normally come to “possess the information that *p*” even in the sense of generating a *weak* true belief that *p* when *p* is stated, or in some other form presented to the subject, by someone he or she knows to be an untrustworthy informant. At least for minimally epistemically rational subjects, the following propositions form a second inconsistent triad:

- (1*) Knowledge requires belief.
- (2) *S* is confronted with a piece of information *p* from a source that *S* knows isn’t trustworthy (in questions of the kind at issue).
- (3) *S* knows that *p* (solely) on the basis of the fact described in (2).

So far I have ignored the following complication. Following Goldman's and Olsson's setup of the story, (2) says unqualifiedly that *S* knows that their informant is untrustworthy. But in what sense of "know"? In some reliabilist sense? Or in some (other) sense of justified true belief + X? Or should we perhaps read the statement in one of the weak senses of "know" just discussed (weak knowledge involving firm belief, or super-weak knowledge)? Suppose *S* enjoys merely super-weak knowledge that her informant is untrustworthy. For example, imagine that instead of a drunken loiterer Alfred asks an eight year old child, little John, about the train schedule. Alfred believes that John is fairly intelligent and that he has put considerable effort in memorizing the schedule. Still, Alfred cannot rid himself of the feeling that he shouldn't believe little John. Alfred is by no means certain that John is wrong. In fact he even believes weakly that John is right. Yet he isn't sure enough. Alfred, let us suppose, is right: John is in fact wrong about the schedule. In this story, Alfred has super-weak knowledge that John is not trustworthy (he weakly and truly believes that John is not trustworthy). Might not this be a situation in which Alfred can still (justifiably) form a *weak* belief that what John tells him is true?

No. Even if we only weakly believe that some (potential) informant is untrustworthy, when that informant is our only source of information we should not, and normally would not, assign a probability greater than 0.5 to the information (or potential information⁵) retrieved from that source. For if a given source is untrustworthy the conditional probability that the (potential) information presented by that source is false is greater than the conditional probability that this (potential) information is true.

I am not denying that there are circumstances in which we would, and should, believe a statement made by an informant of whom we have reason to believe he is untrustworthy. We would believe such an informant if we had either independent positive evidence for the proposition in question or overriding reasons for believing that our initial mistrust was unwarranted. However, in both cases the subject holds a true belief + X, where X is some fairly complex epistemic property. Situations of this kind involve an assessment of the epistemic force both of the (potentially) undermining defeater for the belief in question ("informant A, who claims that *p*, is untrustworthy") and of meta-defeaters (for example: "other, apparently reliable, informants also say that *p*", "on the present occasion A—despite

5. The term "information" is often used in a sense that rules out "false information" as a *contradictio in adjecto*. We may, however, use "information" and "informant" in a more liberal sense that doesn't have "veritistic" implications.

his usual untrustworthiness—appears to be trustworthy”, etc.). In such circumstances the belief condition for weak knowledge could be fulfilled. Yet, the subject’s credal attitude would not simply constitute the complement of ignorance, in Goldman’s sense of having a mere true belief. Instead, *S*’s (true) belief would be justified (and in some fairly complex way). The Goldman-Olsson-Hawthorne argument for weak knowledge thus runs into a dilemma. At least minimally rational epistemic subjects would, under the circumstances sketched in the examples, refrain from forming even a weak belief that the (potential) information they obtain is true. Since the authors do not reject the view that knowledge requires belief, the stories they offer in support of weak knowledge are therefore not coherent. On the other hand, if these stories were spelled out in such a way that it *could* coherently be maintained that the subjects generate at least a weak belief, then this would have to be on account of complex epistemic reasons they have, i.e., of reasons which override their reasons for believing that what their informant says is probably false. In the first case a minimally epistemically rational subject would refrain from forming any belief at all; in the second case the belief he does form would be justified. Either way, the subject fails to acquire even super-weak knowledge.

3. *The rationality constraint*

I have qualified my claim that the subjects in the Hawthorne-Goldman-Olsson type of example would not even form a weak belief, by adding: “at least if they are minimally epistemically rational” (in the situation at hand). What if we drop this qualification? In fact this seems what Goldman and Olsson, as well as Hawthorne, implicitly opt for in claiming that the pupils do form the belief that Vienna is the capital of Austria even though they are well aware that their source is untrustworthy. The question is whether we—if “we” refers to an average competent speaker of English—would indeed ascribe knowledge in such cases. Consider the following example, which is closely analogous to the Hawthorne-Goldman-Olsson case. Tom wants to know what the capital of Zimbabwe is. He encounters a machine that is loaded with thirty index cards displaying the names of the thirty largest cities in Zimbabwe, including the capital. When Tom types in the question: “What is the capital of Zimbabwe?” and pushes a button, the machine spits out one card at random. Tom knows that the machine works in this way, and that it contains exactly one card

with the name of the capital. He knows therefore that the information he will receive is far more likely to be incorrect than correct, he knows that the chances of receiving the correct response is 1/30. Tom thus knows that he shouldn't trust that the machine will provide him with a correct response to the question what the capital of Zimbabwe is. He pushes the button, picks the card and—solely on that basis—forms the belief that the city named on the card is the capital of Zimbabwe. As it happens, the name is correct (“Harare”). Would we say that Tom's true belief is an instance of knowledge? Clearly not. The situation, I maintain, is analogous in all relevant respects to the one where someone trusts a human informant he believes to be untrustworthy.⁶

My original argument was that if (1*) is true, then (3) entails that *S* believes that *p*, but that this is ruled out by (2) if *S* is at least minimally rational. We have now considered the option of dropping the rationality constraint. In that case (2) doesn't rule out that the belief condition for knowledge as stated in (1*) is fulfilled. Yet, if that condition is fulfilled because *S* holds an epistemically *irrational* belief, then we would not—contrary to (3)—ascribe knowledge either. Cut the pie any way you like, knowledge can't properly be ascribed.

It may be worth adding that, were we to ask Tom himself, *he* would normally deny that he knows. Similarly, the students asked about the capital of Austria would normally deny that they know. (I tested the latter kind of case, but with “Harare”, in my epistemology class.) None of the students would normally say that he or she *knows* what the capital is if he or she is aware that the source is unreliable and thus delivered, in all probability, the wrong result. It's like a lottery case. Although you don't know that you have lost, you would not normally consider yourself to know that you have won either.

4. *Weak knowledge and belief suspension about source reliability*

Before considering an alternative explanation of the examples, let us look at two more moves that may be suggested on behalf of weak knowledge. First, why not drop the even-if clause in the Hawthorne-Goldman-Olsson

6. Note that even on purely reliabilist grounds Tom's belief should not be classified as knowledge. For the process or method he employs—trusting a source that (i) he believes to be unreliable and that (ii) actually does produce far more false than true results (in all past, present, and future occasions of use, as we may assume)—is also *de facto* unreliable.

argument? This clause says that we would ascribe knowledge even if (we knew that) the subjects knew that their source is untrustworthy. Suppose it is conceded that my argument, as so far developed, is on target regarding cases in which the subjects know that their source is untrustworthy. “All right, then”, it may be responded, “so let us delete the even-if clause, and the Hawthorne-Goldman-Olsson argument stands!” In *Knowledge and Lotteries*, Hawthorne presents what he declares to be an amended version of his original example. This later version doesn’t tell us what epistemic attitude the subjects have towards their informant. Hawthorne has us imagine the following case:

I give six children six books and ask them each to pick one of the books at random. All but one contains misinformation about the capital of Austria. I ask the children to look up what the capital of Austria is and commit the answer to memory. One child learns ‘Belgrade’, another ‘Lisbon’, another ‘Vienna’, and so on. I ask an onlooker who has witnessed the whole sequence of events (or someone to whom the sequence of events is described) ‘Which one of the schoolchildren knows what the capital of Austria is?’ or ‘How many of the children know what the capital of Austria is?’ It is my experience that those presented with this kind of case will answer, not by saying ‘None of them’, but by selecting the child whose book read ‘Vienna’—even though that child was only given the correct answer by luck. (Hawthorne 2004, 68f.)

First, note that in this example the question, again, is not which one of the children—*if any*—knows, or how many of them—*if any*—know, that the capital of Austria is Vienna. It is instead which one knows, or how many know, what the capital of Austria is. So this example suffers from the same problem that has already been discussed in section 1. “Which one ... knows what the capital ... is?” as well as “How many ... know what the capital ... is?” are leading questions in a context in which it is shared knowledge that at least one of the children could present the correct answer. The formulations suggest that the response “none!” is not expected and would very likely be conversationally inappropriate. I will come back to this point in the next section.

The second point is that in the present form of the story the children’s attitudes towards their informant is significantly underdescribed. In this version, the assumption seems to be that they do not mistrust the person who distributes the books. But this still requires a case distinction. If they don’t believe that their source is untrustworthy they either (i) believe that it is trustworthy or (ii) suspend belief as to whether it is trustworthy

or untrustworthy. The way Hawthorne presents his example suggests that—contrary to what was envisaged in Goldman and Olsson’s version of the school example—it is (i) that he has in mind. (“One child *learns* ‘Belgrade’, another ‘Lisbon’, another ‘Vienna’ ...”) However, if that is the idea, the story clearly fails to yield a good case for weak knowledge! For if *S* acquires the true belief that *p* on the basis of her belief that their source is *trustworthy*, *S* doesn’t acquire a “mere true belief”, but a true belief that is (however weakly) justified. The belief is at least “subjectively justified”, as we may say, for example in the sense of its being internally rational for the subject to form that belief. On certain deontological accounts of justification one could also declare the subject to be epistemically blameless when she forms the belief on the basis of thinking that her informant is trustworthy. Such notions are of course internalist notions of justification. Yet, the kind of positive epistemic status under consideration is not confined to internalism. The process of belief formation may plausibly be described as a process or method of the type “trusting a teacher (an informant) who usually provides her pupils (hearers) with correct information”. At least in that case the children also enjoy justification in some standard process reliabilist (and externalist) sense. Hence, if it is assumed that *S* believes that the source of their information is *trustworthy*, then the knowledge one might be inclined to ascribe to *S* would—contrary to what is required for weak knowledge—not constitute the complement of ignorance.

What we are left with, therefore, are cases in which *S* is told that *p* (or is provided in some other way with the information that *p*), but neither believes nor disbelieves that the source is trustworthy. Here the answer is analogous to the one I have given above. If *S* suspends belief as to whether his informant is trustworthy, then if we knew this and were asked whether *S* *knows* the information he has obtained, we would, other things being equal, assume that *S*’s epistemic reaction displays some basic level of epistemic rationality. Accordingly, we would assume that *S* suspends belief as to whether what he or she was told is true, and hence not normally ascribe knowledge.⁷ So much for the suggestion of dropping the even-if-clause in the Hawthorne–Goldman–Olsson argument.

7. It may be worth emphasizing that this conclusion has no implications for non-epistemic forms of rationality. The fact that from an epistemic point of view the proper attitude for *S* is belief suspension regarding *p* does not of course entail that it may not be rational for *S* in some non-epistemic sense to act *as if* he/she believed that *p*, or to act *on the assumption* that *p*. You are lost in the mountains. A fellow mountaineer who appears to be familiar with the territory

There is yet another move that may be suggested on behalf weak knowledge. So far we have, with Goldman and friends, considered only cases in which at least the knowledge *ascriber* knows that the source on which *S* bases his/her answer is untrustworthy. What if we drop *that* constraint? Suppose John, who knows nothing about the teaching situation in a certain class, enters the classroom and witnesses the teacher asking: “What is the capital of Austria?” Only Lisa replies “Vienna”. When John is asked which of the children, if any, knows that Vienna is the capital of Austria, might he not appropriately reply “Lisa”? If so, it may be suggested, it doesn’t alter the situation if we add that unbeknownst to John the teacher has distributed various books with only one (received by Lisa) containing the correct information, and that the children are well aware that their sources are untrustworthy.⁸

The response to this is that a knowledge ascription would *not* in fact be appropriate in this case. John’s answer would be false and, I maintain, he would accordingly, in normal circumstances, retract his claim that Lisa knows when he is informed that it was by sheer luck that she got hold of the right information. The reason is that, under normal circumstances, John would assume some basic epistemic rationality on Lisa’s part, which precludes her from holding either the firm or the weak belief that Vienna is the capital of Austria if she knows that her source is untrustworthy. And if she generates that belief nonetheless, her “epistemic behavior” would display a high degree of epistemic irrationality, which would again preclude her from being correctly classified as a knower.

From what has been said so far I conclude that, as it stands, the Hawthorne-Goldman-Olsson argument for weak knowledge fails. However, I don’t wish to deny that their examples highlight an interesting and important phenomenon that calls for explanation. I argued that, if the questions in the stories were framed in a more neutral way (“Which of the students, *if any*, knows that ...?”), it is doubtful whether a typical addressee would indeed simply pick the student who utters the right words. However, I don’t wish to dispute that if the question is phrased in one of the ways these authors envisage, there may be circumstances in which the addressee

tells you that the only chance to reach the valley before nightfall is to take the trail to your left. A signpost which appears to be well maintained by the local mountaineer’s club directs you to the trail on your right. Suppose your evidence for the truth of either suggestion is evenly balanced. Even so, you had better not contemplate until sunset about which trail you should take.

8. For this point I am indebted to an anonymous referee of the *Grazer Philosophische Studien*.

is indeed inclined to “count those pupils” who reply “Vienna”. The question remains: Why would this be so?

Unfortunately, it is always much easier to criticize an analysis of a phenomenon than to come up with a plausible alternative. I don't have the space here to present a full account of what I think might constitute such an alternative. Yet in the final section of this paper I shall at least outline an explanation of the examples that does not invoke weak knowledge.

5. *Outline of an alternative explanation of the examples*

I have argued that a question such as “Which of the students, *if any*, knows what the capital of Austria is?” would not, in contexts of the type Goldman discusses, typically be answered by mentioning the candidate who comes up with the correct answer. I also claimed that such a question, if it lacks the “if any” qualification is a leading question when it is shared knowledge between speaker and hearer that some candidates, as we may say, “possess the correct answer”. More precisely, when both (i) know that the other knows this and (ii) assume of their interlocutor that he knows that the other knows this, a question such as the unqualified “Which of the students knows what the capital of Austria is?” suggests that the answer “none” is not expected. Instead, it invites the hearer to mention the candidate who has uttered the right word(s). This description of the case is, I think, intuitively highly plausible. For example, compare the case to a multiple choice exam that asks “Which of the following five propositions are correct?” The supposition clearly is that at least one of the propositions listed *is* correct. If you are the examiner and—as customary for example in British universities—your questions had to be checked by the Faculty's exam board before you were allowed to use them, this way of phrasing the question would certainly not pass if there were no correct answer among the options offered. So there is initial evidence for the view that the which-question in the Hawthorne-Goldman-Olsson example suggests to the hearer that, in the context in question, at least one person should be counted even if it is known that every candidate knows that their source is untrustworthy. Can this point be substantiated on a more theoretical level? More specifically, are there theories of meaning and communication within which this can be fleshed out? Suppose this were the case. Even so, it may be argued, what is wrong with leading questions? Such devices are included among our standard conversational practices. So why should not

Goldman point to such usage and take *it* to show that there are contexts in which “knowing that *p*” simply means “believing truly that *p*”?

Let us consider Hawthorne’s 2004 version of the example. Exactly one child in the room, suppose it is Lisa, has been so lucky as to receive a book that contains the correct information. Suppose again that the hearer is aware that Lisa knows her source is untrustworthy, and someone asks an onlooker, “Which of the children knows that Vienna is the capital of Austria?” To begin with, note that Hawthorne himself indicates that things are somewhat fishy here. In a footnote to the 2004-passage quoted above he concedes:

I note in passing that a few informants claimed to have slightly different intuitions as between ‘Which one of the schoolchildren knows what the capital of Austria is?’ and ‘Which one of the schoolchildren knows that Vienna is the capital of Austria?’. (Hawthorne 2004, 69)

However that may be, let us suppose that the addressee does respond “Lisa” when asked either of these questions. What he means, I suggest, is in neither case that, strictly speaking, Lisa knows *that the capital is Vienna*. Instead, what he means is a proposition that could also, and more appropriately, be expressed by saying: “Lisa knows, or possesses, the correct answer” or, still more appropriately: “Lisa is acquainted with, and disposed to utter, a word that can serve to give the correct answer”. If the respondent doesn’t use any of these sentences, this is because the context and the ways the questions are put conversationally license, in response, the simple utterance of “Lisa”. What is crucial is that Lisa’s possessing the right answer, in the mere sense of being acquainted with the right word, does not entail that she knows *that her answer is right* or that she knows *that the capital of Austria is Vienna*. Compare the case once more to Tom and the card machine. What we would say when Tom receives the “Harare” card is that, due to a lucky accident, he has got hold of the right name. This may, in certain circumstances, also be expressed by saying that he knows, or possesses, the right answer. But since it was by sheer luck that Tom came to possess the right answer, and since Tom knows this, we would *not* normally say that he knows that his answer is right and hence knows that the capital of Zimbabwe is Harare.

If a subject is able to *give the correct answer*, it may be asked, why would she not be able to infer *propositional knowledge of the correct answer* from this ability?⁹ The subject cannot infer this because she doesn’t know *that*

9. This question has been raised by an anonymous referee.

she can give the right answer. If she knows that her source is untrustworthy, she doesn't know—indeed does not even believe—that she was in fact lucky enough to get hold of the right information. In fact, if she is minimally epistemically rational, she even believes that what her book (or the teacher, the index card, or whatever) tells her is probably false. If the story goes such that, when asked, the subject is nonetheless disposed to utter certain words that can be used to give the correct answer, then this may plausibly be explained by the fact that she sees this as her only chance (however small) to hit the truth. My proposal is thus that when in the cases under consideration we concede that the subject knows we speak loosely, assuming for the conversational purposes at hand that what we say is precise enough. To corroborate this interpretation, I shall now take a closer look at the speech acts performed in the example.

Consider Paul Grice's famous Cooperative Principle (CP). This principle says, "contribute what is required by the accepted purpose or direction of the conversation" (cf. Grice 1989, 26f.). One of CP's "supermaxims" concerns what Grice calls "conversational manners" and prescribes: "Be perspicuous!" Grice invokes CP in an attempt to explain implicature, which—as analyzed by Grice—is a phenomenon that pertains to assertive utterances. (Very roughly, to implicate that p is the case is to mean or imply that p is the case by saying that something else is the case.) However, in its general form stated above CP applies to non-assertive utterances as well. Gricean implicature is only one type of *indirect speech act*, and even though Grice's notion of implicature may not be directly applicable to non-assertive utterances, questions can be used as well to perform indirect speech acts. By asking a question one can, we may say, conversationally *imply* that something is the case without *saying* that it is the case.

According to John Searle, indirect speech acts are speech acts in which "the speaker communicates to the hearer more than he actually says by way of relying on their mutually shared background information, both linguistic and nonlinguistic, together with the general powers of rationality and inference on the part of the hearer" (Searle 1975, 60f.). More specifically, the machinery needed to explain the indirect part of indirect speech acts includes "a theory of speech acts, certain general principles of cooperative conversation (some of which have been discussed by Grice [...]), and mutually shared factual background information of the speaker and the hearer, together with an ability on the part of the hearer to make

inferences" (ibid.). A standard example is the question "Can you reach the salt?", as uttered for example during dinner. In such a context the speaker normally utters this sentence not, or not merely, to ask a question, but also as a request to pass the salt. A key question regarding such indirect speech acts is, as Searle notes, "that of how it is possible for the hearer to understand the indirect speech act when the sentence he hears and understands means something else" (Searle 1975, 60).

Basically, Searle's (generally plausible) answer is that the hearer (*H*) infers the relevant facts about the speaker's intentions by invoking some general principles of conversation and shared factual background knowledge. For example, *H* will interpret the question "Can you reach the salt?" as a request to pass the salt by the following kind of reasoning. (What follows is a simplified and slightly modified sketch of Searle's account in (1975, 73f.) and (Searle 1979), where he explains indirect speech acts partly in terms of Gricean conversational implicature). "S has asked me whether I have the ability to pass the salt. I may assume that he is cooperating in our conversation and thus that his utterance has some aim or point (principle of conversational cooperation). The context is not such as to indicate any theoretical interest on S's part in my ability to pass the salt; for clearly, S knows that I have that ability (background knowledge). Hence his utterance is probably not meant just as a question, but has some further illocutionary point. People often use salt at dinner; there is no salt within S's reach, but I can reach the salt (background knowledge). Therefore, since there is no other plausible illocutionary point, S is probably requesting that I pass him the salt." Searle also notes a number of general facts about sentences used to perform directive indirect speech acts. The most important ones are the following (cf. 1975, 67–69). (i) Such sentences do not have an imperative force as part of their meaning. (ii) They are not ambiguous between an imperative and a nonimperative illocutionary force. (iii) Yet such sentences are standardly used to issue directives. (iv) They are idiomatic, but (v) don't constitute idioms (they don't work, for example, like: "This is where the rubber hits the road"). (vi) When such sentences are uttered as requests, they retain their literal meaning and are uttered with, and as having, their literal meaning. (vii) Even when they are uttered with the illocutionary point of a directive, the literal illocutionary act is also performed.

I have suggested that the question posed in the Hawthorne-Goldman-Olsson example ("Which of the students know ...?") is an indirect speech act. Assuming, details aside, that Searle's account is on the right track the

task is to provide an analysis, and explanation, of the Hawthorne-Goldman-Olsson case that is analogous in the relevant respects to the story about indirect speech acts just outlined. I think that such an account can indeed be given. Uttering the sentence, “Which of the students know what the capital of Austria is?” constitutes, not the complex indirect speech act of asking-a-question-plus-making-a-request, but the complex indirect speech act of asking-a-question-plus-making-an-assertion. According to this proposal, what the speaker means when he/she utters this sentence in contexts of the kind in question could also be expressed by uttering (more awkwardly): “Some of the students—even though they know they received their information from an untrustworthy source—possess the correct answer to the question what the capital of Austria is. Which ones?” In order to see that this proposal matches the fundamental tenets of Searle’s account of indirect speech acts, I shall now state some facts that correspond to the features (i)-(vii) above. I will then sketch the reasoning by which a hearer can indeed understand the speaker’s utterance *as* the indirect speech act the latter performs.

Here are some relevant facts about the sentence, “Which of the students know what the capital of Austria is?”, as uttered in a context of the kind envisaged by Goldman. (i) The assertive force of that sentence is not part of its meaning. Witness the fact that its literal utterance can coherently be supplemented with “bracketing” its assertive force, as in: “Which of the students know what the capital of Austria is? (Note that the answer may be ‘none!’).” Compare the multiple-choice test that asks: “Which—if any—of the following propositions is correct?” Here, too, what would otherwise be an assertive component of the indirect speech act performed by uttering the unqualified “Which of the following propositions is correct?”, is explicitly cancelled. Or consider the following children’s trick question: “You have a box filled with a hundred pounds of stones and another one filled with a hundred pounds of feathers. Which one is heavier?” Depending on their age and stage of education, the children who respond will either go for the box of stones or spot the catch and reply, with a smile: “neither”. However, in the latter case they will not typically accuse the questioner of having asked an incorrect question and argue that the *meaning* of the sentence “Which one is heavier?” doesn’t allow for the response “neither”. Yet, why does a corresponding utterance work, for some subjects and in some contexts, as a trick question? Because in the contexts in question it does produce the assumption that one of the boxes is indeed heavier.

(ii) The sentence is not ambiguous between the illocutionary force of a question and that of an assertion. This seems intuitively clear. Moreover it may also be pointed out, with Searle, that the onus of proof would seem to be on those who are inclined to maintain that the sentence is ambiguous. For “one does not multiply meanings without necessity” (Searle 1975, 67f.).

(iii) The sentence can standardly be used to *conversationally imply* an assertion. Several arguments for this have already been laid out in this paper. Remember, for example, the fact that sometimes it will be appropriate explicitly to neutralize the assertive force of the question by qualifying it, as in: “Which of the students, *if any*, knows ...?” In certain contexts this qualification is necessary precisely because otherwise the question *would* be taken to imply that there is at least one candidate who knows the correct answer.

(iv) The sentence is clearly idiomatic; but it is

(v) not an idiom.

(vi) When the sentence is uttered (also) to make the assertion in question (i.e., “There is at least one student who possesses the correct answer”), it still has its literal meaning.

(vii) When it is uttered with the illocutionary point of an assertion, the literal illocutionary act is also performed. After all, the assertion involved is made by way of asking a question. Moreover, the utterance may subsequently be reported by reporting the literal illocutionary act (cf. Searle 1975, 70).

Having described these features of indirect speech acts, let us finally reconstruct the reasoning through which the hearer will normally understand the speaker’s utterance as the indirect speech. The hearer’s reasoning, I suggest, will roughly proceed along the following lines:

“S has asked me which of the students know what the capital of Austria is. Both S and I know—and S knows that I know and that I know that S knows—that the candidates know that they have received their information from an untrustworthy source. Yet unbeknownst to one of the candidates (Lisa) she got hold of the correct answer. Neither the speaker nor I have reason to assume that Lisa is not minimally epistemically rational. So Lisa—despite in some moderate sense possessing the correct answer—fails to *know that* Vienna is the capital of Austria because she doesn’t even believe this. I may assume, however, that S is cooperating in the conversation and trying to be perspicuous. So, had S considered

it an option for me to respond: ‘None (of the students knows that it is Vienna)’, *S* would have indicated this—for example by adding: ‘... if any’ to his question. The only candidate who can be counted in this context is Lisa. Therefore, when *S* is asking ‘Which of the students know what the capital of Austria is?’, *S* probably means this: ‘Some of the students—even though they know they received their information from an untrustworthy source—possess the correct response to the question what the capital of Austria is. Which ones?’”

H certainly does not have to go through any conscious process of inference to derive this conclusion. Instead, he may simply “hear” the speech act as involving the assertion. Admittedly, this is not an uncomplicated story; but that holds as well for Searle’s explanation of the—arguably simpler—indirect speech act “Can you reach the salt?”. (Searle’s exposition takes about two pages; cf. Searle 1975, 73–75.)

With the above conclusion, *H* is in a position to give an appropriate response. If he/she is cooperative, the response will be “Lisa”. And while this is, in the given context, a *grammatical* ellipsis for: “Lisa knows that Vienna is the capital of Austria”, what *H* means, on the basis of his/her understanding of *S*’s utterance, is that Lisa possesses, in the moderate sense of being acquainted with, the correct answer to the question of what the capital of Austria is. *H*’s response is an indirect speech act as well. More precisely, it constitutes an *elliptical indirect speech act*. In this case however the illocutionary force of the indirect component corresponds to the illocutionary force of the direct component: both speech acts are assertive.

In *Knowledge in a Social World* Goldman says that he believes “there is an ordinary sense of ‘know’ in which it means ‘truly believe’” (24). I believe—and I believe I believe truly—that the above discussion casts considerable doubt on Goldman’s argument for this view. However, as noted above, Goldman also writes that, should his ordinary-sense claim turn out to be untenable, he will be “prepared to proceed cheerfully with ‘weak knowledge’ as a term of art (or technical term)” (*ibid.*). If my arguments are on target, Goldman’s ordinary-sense claim is problematic. This does not debar him from switching to the term-of-art view. Indeed, we have seen over the last ten years how the notion of weak knowledge can facilitate pursuing novel and important epistemological projects, projects concerning which Goldman has once more presented pioneering, insightful, and inspiring work. Hence my conclusion

is not that the notion of weak knowledge should be dismissed root and branch. Yet I should like to recommend: Let us believe weakly in weak knowledge.¹⁰

LITERATURE

- Beckermann, Ansgar 1997: “Wissen und wahre Meinung”. In Wolfgang Lenzen (Hg.), *Das weite Spektrum der analytischen Philosophie—Festschrift für Franz von Kutschera*. Berlin, New York: De Gruyter, 24–43.
- 2001: “Zur Inkohärenz und Irrelevanz des Wissensbegriffs. Plädoyer für eine neue Agenda in der Erkenntnistheorie”. *Zeitschrift für philosophische Forschung* 55, 571–593.
- Brendel, Elke 1999: *Wahrheit und Wissen*. Paderborn: mentis-Verlag.
- Davis, Wayne 1998: “*Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*”. Cambridge: Cambridge University Press.
- Ernst, Gerhard 2002: *Das Problem des Wissens*. Paderborn: mentis-Verlag.
- Goldman, Alvin I. 1986: *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- 1999: *Knowledge in a Social World*. Oxford: Clarendon Press.
- 2001: “Social Routes to Belief and Knowledge”. *The Monist* 84, repr. in id., 2002a, 164–181.
- 2002a: *Pathways to Knowledge: Private and Public*. Oxford: Oxford University Press.
- 2002b: “What is Social Epistemology? A Smorgasbord of Projects”. In id., 2002a, 182–204.
- Goldman, Alvin I., and Erik J. Olsson 2009: “Reliabilism and the Value of Knowledge”. To appear in *Epistemic Value*, ed. by D. Pritchard, A. Millar, and A. Haddock, Oxford University Press, 19–41. (Page references are to the manuscript version of this paper.)
- Grice, Paul 1989: *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Hawthorne, John 2000: “Implicit Belief and A Priori Knowledge”, *Southern Journal of Philosophy* 38, Spindel Conference Suppl., 191–210.
- 2002: “Deeply Contingent A Priori Knowledge”, *Philosophy and Phenomenological Research* 65 (2), 247–269.

10. For valuable comments and discussions I am indebted to Wayne Davis, Alvin Goldman, Michael Gorman, Katherine Munn, Gerhard Schurz, and, especially, an anonymous referee of *Grazer Philosophische Studien*.

- Hawthorne, John 2004: *Knowledge and Lotteries*. Oxford: Clarendon Press.
- Kutschera, Franz von: *Grundfragen der Erkenntnistheorie*. Berlin, New York: De Gruyter.
- Lehrer, Keith 1990: *Theory of Knowledge*. London: Routledge.
- LeMorvan, Pierre 2005: "Goldman on Knowledge as True Belief". *Erkenntnis* 62, 145–155.
- Lenzen, Wolfgang 1980: *Glauben, Wissen und Wahrscheinlichkeit*, Wien, New York: Springer.
- Levi, Isaac 1980: *The Enterprise of Knowledge*. Cambridge, Mass.: MIT Press.
- Lewis, David 1996: "Elusive Knowledge". *Australasian Journal of Philosophy* (74), 549–567.
- Meggler, Georg 1997: *Grundbegriffe der Kommunikation*, 2nd ed. Berlin, New York: de Gruyter.
- Radford, Colin, "Knowledge—By Examples". *Analysis* 27 (1966), 1–11.
- Sartwell, Crispin 1991: "Knowledge is Mere True Belief". *American Philosophical Quarterly* 28, 157–165.
- 1992: "Why Knowledge is Mere True Belief". *The Journal of Philosophy* 89, 167–180.
- Searle, John 1975: "Indirect Speech Acts". in *Syntax and Semantics*, ed. Peter Cole and Jerry L. Morgan, New York, San Francisco, London, 59–82.
- 1979: "Indirect Speech Acts". in id., *Expression and Meaning: Studies in the Theory of Speech Acts*, Cambridge: Cambridge University Press.