Sven Bernecker (Irvine)

On the Metaphysics of Knowledge

There are a number of platitudes concerning the relation between knowledge and justification, on the one hand, and truth, on the other. One such platitude is that knowledge is *factive*, that is, '*S* knows that *p*' entails '*p* is true.' Another platitude is that knowledge is incompatible with accidentally true belief in general and with *veritic epistemic luck* in particular. If it is a matter of luck that *S*'s belief is true, then *S* doesn't know that *p*. This is the point of Gettier cases. Regarding the connection between justification and truth it is generally agreed that it is conceptually possible for one to be justified (rational) in believing false propositions. This idea is known as *fallibilism*. *S* fallibly knows that *p* if and only if (i) *S* knows that *p* on the basis of some reasons *R* and yet (ii) *S*'s belief that *p* on the basis of *R* could have been either false or accidentally true. Despite widespread acceptance of fallibilism, it is widely agreed that justification is *conducive to truth*, that is, beliefs are more likely to be true if they are justified than if they are not justified.

This paper argues for an overlooked dimension in the metaphysical microstructure of knowledge. The connection between knowledge and truth is even deeper than generally acknowledged. Knowledge, I argue, supervenes not only on a specific (namely modal) relation between the proposition *p*'s truth and an agent's belief that *p*, but also on specific relations between the proposition's truthmaker and the belief's justification-maker. *S* knows that *p* only if the states of affairs referred to by *S*'s reasons for believing that *p* are identical with, causally related to, or grounded in the states of affairs that make *p* true.[1]

I Two Kinds of Gettier Cases

Gettier examples show that the traditional analysis of knowledge as justified true belief is insufficient because one can have a justified true belief that *p*, but lack knowledge that *p*. Here is an example offered by Keith Lehrer[2] that is in the spirit of Edmund Gettier's original example[3] (and that has the advantage of not resting

1 This paper draws on work published in Bernecker, Sven (2011): "Keeping Track of the Gettier Problem". In: *Pacific Philosophical Quarterly* 92, 127–152.

2 See Lehrer, Keith (1965): "Knowledge, Truth and Evidence". In: *Analysis* 25, 168–175.

3 See Gettier, Edmund (1963): "Is Justified True Belief Knowledge?". In: *Analysis* 23, 121–123.

on a confusion of the referential and attributive sense of the definite description 'the man who will get the job'):

> *Havit's Ford*. Two agents, Mr. Smith and Mr. Nogot, work in the same office. Nogot has given Smith evidence that justifies Smith in believing that Nogot owns a Ford. Imagine that Smith has seen Nogot driving a Ford, Smith has been told by persons who have in the past been reliable that Nogot owns a Ford, and so on. From this evidence Smith then infers the proposition *Someone in the office owns a Ford*. The belief that someone in the office owns a Ford is true. But, unsuspected by Smith, Nogot has lied about owing a Ford and Smith's belief is only true because another person in the office, Mr. Havit, owns a Ford. Does Smith know that someone in the office owns a Ford?

The belief that someone in the office owns a Ford is true. Smith is justified in believing that someone in the office owns a Ford, at least in senses of 'justification' that emphasize the internal or subjective: no more can reasonably be expected of Smith with respect to finding out whether Nogot owns a Ford. Moreover, Smith uses an appropriate method (deduction) for deriving the target proposition. Notwithstanding that fact, that Smith possesses a justified true belief, we would not want to say that Smith *knows* that someone in the office owns a Ford. The reason we would not want to attribute knowledge to Smith is that it just so happened that someone in the office owns a Ford, but not the person Smith thinks owns a Ford. It is a matter of sheer luck that Smith arrives at a true rather than a false belief.[4]

Besides Gettier cases like *Havit's Ford*, there are *unpossessed-defeater cases*. A paradigm unpossessed-defeater case is Alvin Goldman's fake barn example:[5]

> *Fake Barns.* Henry is driving in a part of the country where, unbeknownst to him, the inhabitants have erected a large number of fake barns, i.e., papier-mâché facades looking like barns from the highway, yet lacking back walls or interiors. From the highway, these fake barns are indistinguishable from real ones. Looking at what is in fact a real barn, Henry forms the belief that that is a barn. Does Henry know that that is a barn?

4 So-called Gettier cases have been known long before Edmund Gettier published his article 'Is Justified True Belief Knowledge?' in 1963. In Indo-Tibetan epistemology, Gettier cases have been known as early as the 8th century (see Stoltz, Jonathan (2007): "Gettier and Factivity in Indo-Tibetan Epistemology". In: *Philosophical Quarterly* 57, 394–415). In Western epistemology, Gettier cases can already be found in the works of the 15th century logician Peter of Mantua (see Martens, David B. (2011): "A Late Medieval Dispute about the Conditions for Knowledge". In: *Philosophical Papers* 40, 421–438).

5 See Goldman, Alvin I. (1976): "Discrimination and Perceptual Knowledge". In: *Journal of Philosophy* 73, 771–791.

Although Henry has a justified true belief, it seems to many that he does not know that it is a barn because he 'lucks out' in pointing to a real barn. Since he randomly chooses a barn to look at, he could have very easily picked another one that was a fake and end up with a false belief. So what is said to rob Henry of knowledge is the nearness of the potentially false belief. To know that *p* there may not be a close possible world in which one would have the same belief on the very same basis, but where the belief is false.

Gettier cases and unpossessed-defeater cases have in common that they describe situations where knowledge is destroyed by epistemic luck. What distinguishes these kinds of cases is the specific role played by the lucky occurrence. In Gettier cases, the lucky occurrence functions *beneficially* in the sense that if the lucky occurrence were absent, then (all else being equal) the subject would not have a justified true belief. The subject would lack the truth or the belief or the justification. *Havit's Ford* is like this. If Havit did not happen to own a Ford, then (all else being equal) Smith's belief that someone in the office owns a Ford would be false. In unpossessed-defeater cases, however, the lucky occurrence functions as an *unseen threat* to the subject's having a justified true belief. If the lucky occurrence were absent, then (all else being equal) the subject would not be in any real danger of *not* having a justified true belief; instead, we would have a normal case of knowledge. *Fake Barns* is like this. If there did not happen to be fake barns in the vicinity, then (all else being equal) Henry would truly and justifiably believe that what he is looking at is a barn. So with the absence of the strange occurrence (viz., the fake barns) and with all things being equal, Henry would have a belief that is true and justified in the normal way. This is why Stephen Hetherington calls unpossessed-defeater cases 'dangerous' Gettier cases and ordinary Gettier cases 'helpful' ones.[6]

The distinction between Gettier cases and unpossessed-defeater cases is not sharp. Bertrand Russell's famous stopped clock case,[7] for instance, belongs to both categories:

> *Stopped Clock*. Suppose Bert looks at what he takes to be a reliable clock, sees that it reads eight o'clock, and so on that basis believes that it is eight in the morning. It is true, let's suppose, that it is eight a.m. Suppose further that the clock is actually broken, but that it stopped the night before at exactly eight p.m. Does Bert know that it is eight a.m.?

6 Hertherington, Stephen (2001): *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*. Oxford, 72–75.

7 Russell, Bertrand (1948): *Human Knowledge: Its Scope and Limits*. London, 170.

Even though Bert has a justified, true belief about what time it is, we judge that the subject lacks knowledge. Why? Because the belief, while justified, is true by virtue of luck alone. Had Bert looked at the stopped clock a few minutes before or after eight a.m., which could easily have happened, he would have acquired a false belief.

The *Stopped Clock* example is both a Gettier case and an unpossessed-defeater case. It is a Gettier case because the lucky occurrence (the fact that the clock stopped twelve hours ago) functions beneficially in the sense that if the lucky occurrence were absent, Bert would not have a justified true belief. But the lucky occurrence (the fact that the clock stopped) also functions as an unseen threat to Bert's having a justified true belief.

There is near universal agreement that subjects in Gettier cases do not know. In the case of unpossessed-defeater cases, however, the situation is less clear. There is some controversy as to whether subjects in unpossessed-defeater cases know.[8] A number of epistemologists judge that Henry (in *Fake Barns*) *knows* that what he is looking at is a barn. Whether or not it is reasonable to attribute knowledge to Henry depends crucially on the specification of his belief-forming process. If the belief-forming processed employed by Henry is described as, say, 'seeing a barn in fake-barn country,' then his belief is only accidentally true. But if the belief-forming process is described as, say, 'seeing a barn in an area within fake-barn country where there are no fake barns,' then Henry's belief is not only true by virtue of luck alone and it would be reasonable to grant him knowledge. Since it seems to be entirely up to us how we describe the belief-forming process employed by Henry, it also seems to be up to us whether or not we want to grant Henry knowledge. This is known as the 'generality problem.' In light of the generality problem, it is not surprising that a number of epistemologists attribute knowledge in unpossessed-defeater cases. In what follows, I will concern myself only with genuine (or helpful) Gettier cases.

8 Among those who think that Gettier cases 'and' unpossessed-defeater cases are compatible with knowledge are Hetherington (see Hetherington 2001) and Sartwell (see Sartwell, Crispin (1991): "Knowledge is Merely True Belief". In: *American Philosophical Quarterly* 28, 157–165). Heathcote (see Heathcote, Adrian (2006): "Truthmaking and the Gettier Problem". In: Stephen Hetherington (ed.): *Aspects of Knowing*. Amsterdam, 151–168, especially 166.) and Lycan (see Lycan, William G. (2006): "On the Gettier Problem Problem". In: Stephen Hetherington (ed.): *Epistemology Futures*. Oxford, 148–168, especially 161–163) share the Gettier intuition, but claim that unpossessed-defeater cases are compatible with knowledge.

II Two Approaches to the Gettier Problem

According to the orthodox *covariationist reading*, what prevents Smith (in *Havit's Ford*) from knowing someone in the office owns a Ford is the fact that, given Smith's evidential situation, it is just a matter of luck that this proposition is true. There is a number of possible scenarios which are epistemically indistinguishable vis-à-vis Smith's evidential situation, but in which the proposition is false because, say, Havit has recently traded his Ford for a Toyota. So on the covariationist reading, Smith does not know that p because, given his reasons for believing p, Bp fails to covary with p through a sphere of possibilities. Gettier cases are instances of what is called 'veritic luck.' A belief is veritically lucky if it is true in the actual world, but in some close possible worlds, in which the subject forms the same belief on the basis of the same evidence or via the same method of belief formation, the belief is false. The covariationist reading is widely accepted.

On the *identificationist reading* of the Gettier problem, what prevents Smith from knowing p is the fact that his reasons for holding p true have nothing to do with what makes p true. Smith clearly has reasons for believing that someone in the office owns a Ford, namely that Nogot has claimed to own a Ford. Yet, if we were to explain why this belief is true – what makes it true –, we would refer not to Nogot, but rather to Havit. The truth-maker for Smith's belief is disjoint from the state of affairs in which the justification is grounded. The justification does not direct us to what accounts for the truth of the belief. Smith's reasons for believing that it is true that someone in the office owns a Ford misidentify the actual truthmaker of the proposition. He does not know because he takes the target proposition to have a truth-maker other than it has.

According to the covariationist interpretation, Gettier cases result from a failure of the belief in p, the truth of p, and the evidence E for believing p to covary in close possible worlds. On the identificationist interpretation, however, the crux with Gettier cases is not covariation-failure in close possible worlds, but identification-failure in the actual world: the subject's reasons for holding the belief true misidentify the belief's truth-maker. Whether an epistemic situation is a Gettier case on the identificationist reading cannot simply be read off from the truth-values of Bp, E, and p in a range of possibilities, but demands comparing the states of affairs that make p true with S's reasons for thinking that p is true.

Granted that the identificationist reading of Gettier cases is correct, the question arises as to what is the right sort of relation between the satisfaction of the justification condition and the satisfaction of the truth condition. Later, I will

offer a response. Right now, however, my point is that, regardless of how the relation sketched by the identificationist reading of the Gettier problem is spelled out, it is (or can be) distinct from the relation sketched by the covariationist reading. The latter relation is one of truth-values across a range of possible worlds; the former is one of identification in the actual world.

I focus on Gettier cases as opposed to unpossessed-defeater cases because the identificationist reading does not apply to the latter. Henry who drives through fake-barn country and happens to look at a real barn does *not* misidentify the state of affairs that renders his belief true. Henry's reason for believing that something is a barn (viz., that it looks like a barn) *is* suitably connected to its truth-maker (viz., that it is a barn). So the identificationist reading of the Gettier problem does not apply to unpossessed-defeater cases. Unpossessed-defeater cases are examples of covariationist Gettierization without identificationist Gettierization. What I intend to show in this paper, among other things, is that the converse is possible as well. There are cases of identificationist Gettierization without covariationist Gettierization.

The epistemic defect referred to by the identificationist reading is usually the reason for the epistemic defect referred to by the covariationist reading. When the belief in *p* fails to counterfactually covary with the truth of *p*, this is usually because the subject misidentifies *p*'s truth-maker. Yet, it is possible that the defect referred to by the identificationist reading is present while the defect referred to by the covariationist reading is absent. Since the epistemic defects identified by the two readings of the Gettier problem can come apart, a theory of knowledge may work for Gettier cases due to covariation-failure but not for Gettier cases due to identification-failure. Tracking accounts of knowledge are a case in point.

III Truth-Tracking

Given the orthodox covariationist reading of the Gettier problem, the obvious way of blocking the Gettierization process is to rule out possible situations in which the agent has the same belief as in the actual situation but in which the belief either goes wrong or gets defeated. This is precisely the strategy adopted by truth-tracking accounts of knowledge. Truth-tracking comes in different flavors: the *sensitivity/variation condition*, the *adherence condition*, and the *safety condition*. While a number of objections have been raised for each of these conditions, the general idea of truth-tracking is widely accepted. Let's start with sensitivity.

The best-known sensitivity-based account of knowledge is due to Robert Nozick.[9] Nozick suggests that *S* knows that *p* only if (i) *p* is true, (ii) *S* believes that *p*, (iii) if *p* were not true, *S* would not believe that *p* ($\sim p \mathrel{\Box\!\!\rightarrow} \sim Bp$), and (iv) if *p* were true, *S* would believe that *p* ($p \mathrel{\Box\!\!\rightarrow} Bp$). Condition (iii) is the sensitivity/variation condition and condition (iv) is the adherence condition. A belief in *p* is sensitive if and only if *S* would not believe *p* if *p* were false. A belief in *p* is adherent if and only if *S* would believe *p*, if *p* were true. A belief that fulfills conditions (i)–(iv) is one that, in Nozick's expression, 'tracks the facts' that make it true.

Sensitivity-based accounts of knowledge differ concerning the range of error-possibilities on which the truth of *p* must covary with *S*'s belief in *p*. According to Nozick's original account, only the nearest ~*p*-worlds are of relevance for determining whether a belief is sensitive and adherent. On Keith DeRose's[10] and Mark Heller's[11] contextualist versions of the sensitivity account, the set of ~*p*-worlds differs from context to context. In some contexts a belief must track the truth only in the nearest ~*p*-world, in other contexts the same belief must covary with the truth through a wide range of nearby ~*p*-worlds.

Sensitivity theorists are convinced that their account is immune to Gettierization. To see why, consider again the case of Havit's Ford: Smith's belief that someone in the office owns a Ford fails to be sensitive to the truth in the sense that, in the closest possible world in which the proposition is false, he will continue to form the same belief in the same way as he formed his belief in the actual world. The sensitivity condition is not met, for it is not the case that if the proposition had been false, Smith would have responded differently to his environment and, in particular, would not have believed the proposition.

Despite its intuitive plausibility, there are a number of objections to the sensitivity condition. The principal objection is that sensitivity leads to the breakdown of the closure principle.[12] One can know everyday propositions (such as that one

9 See Nozick, Robert (1981): *Philosophical Explanations*. Cambridge, 172–196.

10 DeRose, Keith (1995): "Solving the Skeptical Problem". In: *Philosophical Review* 104, 1–52.

11 Heller, Mark (1999): "The Proper Role for Contextualism in an Anti-Luck Epistemology". In: *Philosophical Perspectives* 13, 115–129.

12 This objection to the sensitivity condition has been leveled by Kripke (see Kripke, Saul A. (2011): "Nozick on Knowledge". In: Saul A. Kripke: *Philosophical Troubles: Collected Papers* Vol. 1. New York, 162–224). Sosa (see Sosa , Ernest (1999): "How to Defeat Opposition to Moore". In: *Philosophical Perspectives* 13, 141–154, especially 141–142 and 149.) and Williamson (see Williamson, Timothy (2000): *Knowledge and its Limits*. Oxford, 116–117). In Bernecker (see Bernecker, Sven (2012): "Sensitivity, Safety, and Closure". In: *Acta Analytica* 27, 367–381) I argue that the argument, to the effect that sensitivity is not closed, is invalid. For sensitivity to violate closure, it must be possible to sensitively believe *p* and to sensitively believe that *p* entails *q* but not meet the sensitivity condition with respect to *q*. It is not hard to come up with cases where

has hands) in virtue of possessing a sensitive belief in these propositions, know that they entail the denials of skeptical hypotheses (like the brain-in-a-vat hypothesis), and yet fail to know the denials of skeptical hypotheses, in virtue of lacking a sensitive belief in these propositions. For in the nearest possible worlds, in which a skeptical hypothesis is true, one continues to believe that one is not the victim of this hypothesis. Those who want to hold on to the closure principle, while accepting the basic insight of truth tracking, have replaced the sensitivity condition with its contrapositive – the *safety condition*. The safety condition reads: *S* would believe that *p* only if *p* were true ($Bp \Box\!\!\rightarrow p$).[13]

The safety condition handles Gettier cases in essentially the same way as the sensitivity condition. Each condition excludes a range of close possible worlds, in which the agent forms his belief in the same way as he does in the actual world, but forms a false belief instead of a true one. Thus, given the standard covariationist diagnosis of Gettier cases, the tracking account eliminates *all* Gettier cases.

IV Knowing Necessary Truths

Given the identificationist reading of the Gettier problem, Gettier cases do not depend on the belief in question being false in some nearby possible world. One's belief that *p* could conform to the truth in *all* close possible worlds and still one could be a victim of Gettierization in the sense that one is radically mistaken regarding *p*'s truth-maker. The covariation of *Bp* with *p* does not ensure that the subject's reasons for holding *p* true identify what accounts for the truth of *p*. The mistake of tracking accounts of knowledge is to suppose that the adequacy of epistemic reasons can be specified purely in terms of the covariation between the *truth* of *p* and the subject's reasons for holding *p* true – that one need not take into consideration the *content* of *p* vis-à-vis the subject's reasons. Since tracking accounts of knowledge allow for the

---

one's belief in *p* is sensitive but one's belief in *q* is not. The problem, however, is applying the sensitivity condition to the belief in the entailment. The counterfactual conditional 'if it were not the case that *p* entails *q*, one would not believe that *p* entails *q*' is vacuously true. To avoid having to attribute vacuous knowledge whenever someone believes in a necessary truth, the sensitivity condition is applied only to contingent truths. But if beliefs in necessary truths do not meet the sensitivity condition, the argument for the non-closure of sensitivity crumbles.

13 Among the proponents of a safety-based account of knowledge are Pritchard (see Pritchard, Duncan (2005): *Epistemic Luck*. Oxford, 161–73), Sosa (see Sosa 1999 and Sosa, Ernest (2000): "Skepticism and Contextualism". In: *Philosophical Issues* 10, 1–18) and Williamson (see Williamson 2000, 123–128).

subject's reasons for believing something to misidentify the circumstances underlying the truth of the belief, there are 'intractable' Gettier cases, i. e., Gettier cases that tracking accounts cannot handle.

The shortcoming of safety-based and sensitivity-based accounts of knowledge can be illustrated by a slight variation of the story of Havit's Ford:

> *Prime Lie:* Nogot has given Peter evidence that justifies Peter in believing that Nogot owns a Ford. Imagine that Peter has seen Nogot driving a Ford, Peter has been told by persons who have in the past been reliable that Nogot owns a Ford, and so on. Using disjunction-introduction Peter forms the belief that Nogot owns a Ford or the 100th prime number is 541. Peter has plenty of evidence for the first disjunct, but only the second disjunct is true and Peter has no evidence that it is true. He has simply guessed that the 100th prime number is 541. Does Peter know 'Nogot owns a Ford or the 100th prime number is 541?'

Peter's belief that Nogot owns a Ford or the 100th prime number is 541 is true in *all* possible worlds. But just because Peter's belief *cannot be false* does not mean that it is automatically knowledge. Since one can believe necessary truth on the basis of silly reasons, counterfactual dependence is not the appropriate way to handle Gettier cases for necessary truths. What then 'is' the reason for Peter not knowing that Nogot owns a Ford or the 100th prime number is 541? In my view, Peter does not know it, because the fact that makes the disjunctive propositions true – a necessary fact as it happens – is not properly linked to his reasons for holding it true. Peter's reason for holding the disjunction true is not suitably related to (and in this case is totally independent from) the features that render it justified. The belief is supported by the wrong kind of reasons.

Although the story of *Prime Lie* shows that safety-based and sensitivity-based accounts of knowledge have problems excluding certain kinds of Gettier cases, it is not prudent to rest the case for identificationism solely on this counterexample. First, safety-based and sensitivity-based accounts of knowledge are not designed to handle necessarily true and necessarily false propositions. It is common to restrict the scope of safety and sensitivity to contingent propositions.[14] Second, tracking theorists can sidestep counterexamples such as *Prime Lie* by asserting that safety or sensitivity, respectively, are only necessary for knowledge, but not sufficient. The reason a belief in a necessary truth may not qualify as knowledge is that it fails to meet some other necessary knowledge

14 Necessary truths are not the only truths to cause problems for truth-tracking accounts of knowledge. Contingent truths whose negations are nomologically impossible are just as problematic. Truth-tracking needs to be restricted to *fully contingent* propositions, that is, propositions that are neither logically, nor nomologically, nor metaphysically necessary.

condition, such as a virtue-theoretic condition[15] or the adherence condition.[16] The adherence condition excludes Gettierized beliefs in necessary truths. When someone believes a necessary truth for a silly reason there are close possible worlds in which *p* is true, but he does not believe it. I will come back to the adherence condition in section 5. Third, identificationism is committed to there being a state of affairs for every knowable proposition and to that state of affairs acting as the proposition's truth-maker. The problem with necessary truths is that it is not clear that they are made true by particular states of affairs, for they are true regardless of what facts obtain. But if necessary truths do not have truth-makers, identificationism does not get a foothold, for there is nothing for the reason in support of a belief in a necessary truth to identify.[17]

There are cases of identification-failure without covariation-failure that do not rely on necessary truths. To drive home this point, consider a variation of Gettier's second example:[18]

> *Leaving Barcelona:* Nogot has given Paul evidence that justifies Paul in believing that Nogot owns a Ford. Imagine that Paul has seen Nogot driving a Ford, Paul has been told by persons who have in the past been reliable that Nogot owns a Ford, and so on. Using disjunction introduction Paul forms the belief that Nogot owns a Ford or Brown is in Barcelona. Paul has plenty of evidence for the first disjunct, but only the second disjunct is true and Paul has no evidence that it is true. He has simply guessed the whereabouts of Brown. But let's further assume that it is no mere accident that Brown is in Barcelona. Brown is so constituted (psychologically, financially, and otherwise) that it is extremely unlikely that he would ever leave Barcelona. Thus, Paul's belief to the effect that either Nogot owns a Ford or Brown is in Barcelona is true in nearly all (if not all) close possible worlds. Does Paul know *Nogot owns a Ford or Brown is in Barcelona*?

Intuitively Paul fails to have knowledge, despite having a contingently true belief that meets the safety version of the tracking condition: he would believe that Nogot owns a Ford or Brown is in Barcelona only if it were true, for *Brown is in Barcelona* is true in each and every of the *close* possible worlds. Identificationism, on the other hand, offers a compelling explanation of why Paul does not know. The reason Paul fails to know that Nogot owns a Ford or Brown is in Bar-

15 Cf. Sosa.

16 Cf. Nozick 1981.

17 This objection crucially depends on particular states of affairs not being able to act as truth-makers for necessary truths. But why should we not say that, for instance, the necessary truth $p \vee \sim p$ is made true by either some particular state of affairs, making p true, or by some particular state of affairs making $\sim p$ true?

18 This example is adapted from Hiller, Avram/Neta, Ram (2007): "Safety and Epistemic Luck". In: *Synthese* 158, 303–313, here: 307–8.

celona is that his reflectively accessible grounds for believing this proposition misidentify its truth-maker.

Proponents of the safety condition and the sensitivity condition hold that for a safe or sensitive belief to qualify as knowledge, it must stem from a reliable method of belief-formation. There are at least two ways of conceiving of reliable methods:

> (a) Reliable method for believing that *p:* if the method yields the belief that *p* in close possible worlds, *p* is true.
> (b) Reliable method for believing that *p:* if the method is applied under relevantly similar conditions in close possible worlds, it yields only true beliefs.

Definition (b) is clearly more demanding than definition (a). Definition (a) requires only that the belief that *p* may not be false if it is formed on the basis of the same process or method in a close possible world. Definition (b) requires that a reliable method yields no false beliefs in close possible worlds. The liberal definition (a) is endorsed by Nozick and Pritchard; the stringent definition (b) is endorsed by Goldman, Sainsbury and Sosa.

The advantage of definition (b) is obvious: it allows proponents of the safety condition to account for our intuition that the protagonists in *Prime Lie* and *Leaving Barcelona* do not know; they do not know because they fail to satisfy the safety condition. The safety condition is not satisfied because the belief-formation method employed in both cases – guessing – could have easily generated false beliefs. In *Prime Lie*, Peter simply guesses that 541 is the 100th prime number. And in *Leaving Barcelona*, Paul guesses that Brown is in Barcelona.

Given that (b) allows the proponent of truth-tracking theories to eliminate so-called 'intractable' Gettier cases, what – if anything – prevents him from adopting this account of reliable belief-forming methods? The problem with (b) is that it is too stringent. If knowledge requires that one employs belief-forming methods that yield no false beliefs in close possible worlds, then hardly any of the belief-forming methods used in everyday life are reliable and hence knowledge becomes a very rare commodity indeed.

V Knowledge and Preemption

Not all Gettier cases due to identification-failure involve logical or contingent necessities. Here is a case of identificationist Gettierization without covariation-failure that manages without the use of necessary truths:

> *Double Trouble:* Two independently operating snipers aim with the same kind of gun at one and the same spy. The bullet from sniper *A* arrives first and kills the spy by piercing his heart. The bullet from sniper *B* arrives a split-second thereafter, and would have been sufficient for killing the spy in the same manner, only the spy was already dead. Since the bullet of sniper *B* hits the spy's body with the same speed, at the same angle, and in the same place as the bullet of sniper *A*, the bullet of sniper *B* does not cause any additional damage to the spy's body. Both bullets travel through the spy's body (creating an exit wound), fall through the slits of a manhole cover, and disappear. The investigating sheriff is aware of sniper *B*, but ignorant of the existence of sniper *A*. The evidence the sheriff gathers justifies him in believing that the spy died due to a bullet from sniper *B*. On the basis of this evidence, the sheriff infers the true proposition *The spy died due to a sniper's bullet*. Does the sheriff know what he justifiably and truly believes?

This is a case of preemptive causal overdetermination with the sheriff being aware of the preempted sufficient condition for the spy's death (the bullet from sniper *B*), but not the causally effective one (the bullet from sniper *A*). According to identificationism, the sheriff does not know that the spy died due to a sniper's bullet since his reasons for believing this proposition point to its potential rather than its actual truth-maker.

Whether proponents of the truth-tracking account of knowledge are committed to attributing knowledge to the sheriff depends in part on the underlying notion of a reliable belief-forming process. Granted the demanding notion of a reliable belief-forming method, the sheriff was not using a reliable method of belief-formation when he inferred from 'the spy was shot' to 'the spy was shot by sniper *B*.' The belief-forming method used by the sheriff is unreliable because it easily generates false beliefs. In fact it generates a false belief in the actual situation. So a tracking theorist can maintain that the sheriff does *not* know that the spy died due to a sniper's bullet, for the sheriff's belief that the spy died due to a bullet from sniper *B* is *unreliably* formed and hence *unjustified*. Yet, as I already mentioned at the end of the previous section, the demanding notion of a reliable belief-forming process leads straight into skepticism. Given this notion of a reliable belief-forming method, we possess very few, if any, justified beliefs.

Another way for proponents of the tracking-account of knowledge to respond to the case of *Double Trouble* is to claim that the sheriff *knows* that the spy died due to a sniper's bullet. He knows this since the basis for his belief is a sufficient condition (a lethal bullet). Whether or not the sufficient condition is preempted by another sufficient condition does not undermine the belief's positive epistemic status. The sole function of epistemic reasons is to ensure that it is not a lucky coincidence that the subject holds a true belief; it is not re-

quired that the epistemic reasons identify the circumstances that underlie the belief's truth.

I disagree. If truth-conduciveness were the sole function of epistemic reasons, as suggested by the objection at hand, the original Gettier case would turn out to be nothing but a pseudo-problem. There would be no sense in which Smith (in *Havit's Ford*) is in the least bit justified in believing that someone in the office owns a Ford. The reason is that Smith's evidence for holding the belief true is not at all connected to its truth. But if Smith is not justified in believing what he believes, then the example fails to show that justified true belief is insufficient for knowledge. Unless there is more to justification than truth-conduciveness, *Havit's Ford* does not show what it is supposed to show: that the justified-true-belief analysis of knowledge is insufficient. According to identificationism, knowledge not only requires the belief that p, the truth of p, and the evidence E for believing p to covary in close possible worlds; the evidence E must also identify the features accounting for the truth of p. To qualify as knowledge, a belief must track the facts *for the right reasons*.

Yet, another way for proponents of the tracking-account of knowledge to respond to the case of *Double Trouble* is to resort to Nozick's adherence condition ($p \Box\!\!\rightarrow Bp$). In fact, adherence can handle *all* of the so-called 'intractable' Gettier cases. Consider *Prime Lie*. Peter's belief that Nogot owns a Ford or 541 is the 100th prime number is not adherent, for in one of the worlds closest to actuality in which the disjunction is true, Peter has no misleading evidence that Nogot owns a Ford, so does not believe that Nogot owns a Ford, and hence does not infer the disjunction. In *Leaving Barcelona*, Paul's belief that Nogot owns a Ford or Brown is in Barcelona also fails to meet the adherence condition. In a close possible world, in which Paul does not get the misleading evidence whereupon Nogot owns a Ford, he does not endorse the disjunctive proposition that Nogot owns a Ford or Brown is in Barcelona, even though the proposition is still true. And in *Double Trouble* the adherence condition is also violated. There is a close possible world in which the spy dies due to a sniper's bullet (namely sniper *A*'s bullet) without the sheriff believing it (because there is no evidence of sniper *B* having fired a shot).

Even though adherence excludes 'intractable' Gettier cases, proponents of the truth-tracking account of knowledge would be ill-advised to make use of adherence. Adherence may be sufficient for knowledge, but it is not necessary. To see this, consider a doorbell with a short circuit. Whenever the doorbell rings inside the house, someone is outside pressing the button. But sometimes pressing the button does not result in the bell ringing. Given this scenario, whenever the bell rings, you know that someone is at the door. But when the bell does not ring, you cannot be sure whether someone is at the door. Now the question is wheth-

er, by hearing the bell ring, you can come to know that there is someone at the door. Since the adherence condition is not fulfilled, Nozick seems to be committed to answer in the negative. Intuitively, however, you know that someone is at the door when the bell rings. Thus, knowledge is not adherent. So adding adherence to our knowledge conditions means excluding 'intractable' Gettier cases, but it also means excluding genuine cases of knowledge.

VI False Evidence

Some of the earliest attempts to explain the difference between knowledge and justified true belief emphasize the role that falsehoods play in yielding cases of justified true belief that are not knowledge. It was suggested that Gettier cases can be ruled out by stipulating that one belief can justify another only if it is true. Thus, in *Havit's Ford*, Smith is justified in believing that Nogot owns a Ford. However, Smith is not justified in believing that someone in the office owns a Ford. This is because the epistemic justification for the latter belief is grounded on the former belief which is false.[19]

Even though the original no-false-premise approach was soon shown to be both too strong and too weak, a number of epistemologists continue to think that knowledge must not rest on any essential false assumptions/presuppositions/implicit beliefs. On this view, knowledge can be based on evidence and background assumptions, some of which are false, so long as no essential element of the reasoning is false. Any false evidence and assumption on which the justification rests must be dispensable. An assumption E is essential for S to be justified in believing p on the basis of another belief q if and only if: S is justified in believing p on the basis of q only if S believes E.[20]

19 Opinions differ on whether the no-essential-false-assumption approach is successful in ruling out not only Gettier cases but also unpossessed-defeater cases such as *Fake Barns*. Levin (see Levin, Michael (2006): "Gettier Cases Without False Lemma?". In: *Erkenntnis* 64, 381–392, especially 390) and Lycan (see Lycan 2006, 157–158) maintain that in dangerous Gettier cases there are no identifiable false tacit assumptions, wherefore the no-essential-false-assumption approach does not get a foothold. I disagree. Henry falsely assumes that if something looks like a barn from the highway, then it is a barn. Moreover, he falsely assumes that there is nothing unusual about the part of the county he is driving through.

20 Advocates of the no-essential-false-assumption approach are Feldman (see Feldman, Richard (2003): *Epistemology*. Upper Saddle River, 37), Harman (see Harman, Gilbert (1973): *Thought*. Princeton, 46–50 and 120–124), Lehrer (see Lehrer, Keith (1974): *Knowledge*. Oxford, 219–220), Levin (see Levin 2006), and Lycan (see Lycan 2006, 156–157).

My 'intractable' Gettier cases obviously violate the no-essential-false-assumption condition. Peter in *Prime Lie* and Paul in *Leaving Barcelona* falsely believe that Nogot owns a Ford. The falsity in their overall sets of reasons is ineliminable. Similarly, the sheriff's justification in *Double Trouble* essentially rests on the false assumption that there was no further sufficient condition for the spy's death. So a critic could argue that the reason Peter, Paul and the sheriff fail to know is because they violate the no-false-essential-false assumption condition, not because they fail to meet the identificationist condition.

Given that cases of identificationist Gettierization involve essential false assumptions, does this mean that identificationism amounts to the same thing as the no-essential-false-assumption approach? Is requiring that a reason for believing *p* identifies *p*'s truth-maker tantamount to demanding that the inference to *p* does not rest on an essential false assumption? The answer is 'no.' Even though identificationism and the no-essential-false-assumption theory rule out the same cases, the former is not only distinct from the latter, but has a clear explanatory advantage over it.

The no-essential-false-assumption approach is successful in ruling out Gettier cases (including cases of identificationist Gettierization), but it does not shed light on the nature of knowledge. Since the no-essential-false-assumption approach offers no explanation as to why knowledge is incompatible with false assumptions, it is explanatorily unsatisfactory. What is more, the no-essential-false-assumption approach puts the cart before the horse: evidence should be a guide to (non-accidental) truth rather than truth being a requirement for evidence. Identificationism, by contrast, gives us insight into the nature of justification and knowledge: knowing *p* involves having a true belief and properly basing the belief on reasons that identify *p*'s truth-maker.

VII Identificationism

Knowledge requires an adequate connection between the state of believing that *p* on the basis of reasons and the truth-maker for *p*. But what kind of connection? The mistake of the truth-tracking account of knowledge, as we saw, is to suppose that the adequacy of epistemic reasons can be specified purely in terms of truth-covariation, and that one need not take into consideration the content of *p* vis-à-vis the subject's reasons. How should truth-tracking theories be complemented so as to rule out Gettier cases due to identification-failure?

Prima facie one might be tempted to demand that the reason-providing belief stands in a *semantic* relation to the target belief. However, the problem with this proposal is that in Gettier cases such a semantic relation is indeed in place.

In *Havit's Ford*, for example, there is a semantic relationship between the reason-providing belief that Nogot owns a Ford and the target belief that someone in the office owns a Ford. This semantic relation is known as *hyponymy*.

Alternatively, one could try to supplement the truth-tracking account of knowledge by demanding that the reason-providing belief stands in an *explanatory* relation to the target belief. Given this proposal, it must be possible to explain the likely truth of the target belief on the basis of the assumed truth of the reason-providing belief. Despite its intuitive appeal, this proposal may turn out to be circular. At least sometimes, the notion of explanation already presupposes the notion of knowledge: to explain something is to lay out the conditions in virtue of which the subject knows. Given this kind of explanation, it is circular to explicate the notion of knowledge in virtue of the notion of explanation. Moreover, since explanations are context-sensitive there is the worry that even in Gettier cases there is a context in which the justification-maker and the truthmaker *are* explanatorily connected. In other words, the worry is that the explanatory approach to knowledge is not in a position to discriminate good cases of knowledge from Gettierized ones.

The knowledge-constituting connection between the state of justified believing and the truthmaker for the proposition believed is neither of a semantic nor of an explanatory kind but *metaphysical* in nature. According to identificationism, *S* knows that *p* on the basis of reasons *R* only if (i) *p* is true, (ii) *S* believes that *p*, (iii) *S* tracks the facts that make *p* true (by satisfying the sensitivity or safety condition), and (iv) *S*'s reasons *R* for believing *p* identify *p*'s truthmaker. *S*'s reasons *R* for believing *p* identify *p*'s truthmaker if and only if *R* refer to states of affairs that are identical to, causally related to, or grounded in, the states of affairs that make *p* true.

The most straightforward cases of knowledge are those in which the epistemic reasons for believing *p* represent *p*'s truthmaker itself. But there are also cases of knowledge in which the connection between the epistemic reasons and the truthmaker is less direct. For example, I am justified in believing, on the basis of testimony, that there are kangaroos in Australia, because I have positive evidence in favor of the truthworthiness of the attester. Yet, the states of affairs that speak in favor of the attester being trustworthy are disjoint from the states of affairs that make the proposition true. This suggests that the satisfaction of the justification condition need not be identical with the satisfaction of the truth condition. The connection between reasons and truthmaker may be that of causation or grounding.

Note that identificationism does not require that the truthmaker be *the cause* of the states of affairs referred to by the subject's epistemic reasons. This would make it impossible to know future events. Identificationism requires only that

there be a causal relation between the justification-maker and the truthmaker. This allows for the justification-maker to cause the truthmaker or both having a common cause.

Identity and causation at the level of states of affairs are not the only kind of relations that constitute knowledge. Besides identity and causation, grounding is another knowledge-constituting relation. 'Grounding' is the collective term for non-causal dependence relations among states of affairs. Examples of grounding are the relations that hold between neurophysiological and mental facts, between categorical and dispositional properties, between facts about parts and their arrangement and facts about wholes, as well as between non-social and social facts. There is considerable controversy over what, if anything, unites these diverse examples of non-causal dependence.[21] Yet, I think it is reasonable to operate under the assumption that there is one kind of distinctively metaphysical relation that unifies these examples. And the idea proposed here is that in the case of mathematical, logical, and inductive knowledge, the justification-maker must be grounded in the truthmaker.

Identificationism is committed to the idea that every knowable proposition has a truthmaker. This idea is known as the truthmaker principle: Necessarily, for all propositions *p*, if *p* is true, there exists something that makes *p* true. The truthmaker principle works well for propositions expressed by positive existentials ('*a* exists') and true contingent positive characterizations ('*a* is *F*'). Yet, the truthmaker principle seems to run into some problems when applied to propositions expressed by true contingent negative characterizations ('*a* is not *F*'), negative existentials ('there are no *F*'s'), general truths ('all *F*'s are *G*'s'), truth of mere possibility ('*a* is possibly *G*'), necessary truths ('2 + 2 = 4'), subjunctive conditionals ('if *a* were *G*, then *a* would be *H*'), and (if we assume presentism) past truths ('*a* was *F*').[22] Even though this is not the place to defend the truthmaker principle, I would like to at least sketch how the truthmaker principle can be defended by adopting *truthmaker pluralism* (the thesis that there is more than one kind of thing that can make a proposition true) and *truthmaking pluralism* (the thesis that there is more than one way for a proposition to be made true). Insofar as the truthmaker pluralist[23] takes truth to con-

21 See Correia, Fabrice/Schnieder, Benjamin (eds.) (2012): *Metaphysical Grounding: Understanding the Structure of Reality*. Cambridge.

22 These issues are discussed in Horvath, Joachim (2015): "Taking the Metaphysics of Knowledge Seriously. A Response to the Paper of Sven Bernecker. In: Markus Gabriel/Wolfram Hogrebe/Andreas Speer (eds.): *Das neue Bedürfnis nach Metaphysik – The New Desire for Metaphysics*. Berlin.

23 For instance Lynch, Michael (2009): *Truth as One and Many*. Oxford; Sher, Gila (2004): "In Search of a Substantive Theory of Truth". In: *Journal of Philosophy* 101, 5–36.

sist in different relational properties in different domains (correspondence, coherence, superwarrant, etc.), he will need an account of truthmaking tailored to each domain. For example, if truth consists in coherence for ethical truths, then the truthmaker pluralist will need a conception of truthmaking in which propositions make other propositions true. Now given pluralism about truthmaking, the truthmaker principle can be tweaked so as to handle the problem cases mentioned above. For example, we can say that *x* is a truthmaker for a negative characterization *p* if and only if (i) *p* is about a proper part of *x* and *x* is one of the things that (ii) prevents the existence of any falsemaker *y* for *p* and that (iii) necessitates that *p* is true. And we can specify the truthmaker of a negative existential as follows: For any negative existential *p*, *p* is made true by the world as a whole if and only if *p* would be false if at least one entity existed in virtue of which *p* were false.[24]

Identificationism is a version of *epistemic externalism*. The knowing subject need not have cognitive access to the fact that there is a causal, identity or grounding relation connecting the states of affairs referred to by the reason-providing belief or experience, on the one hand, and the states of affairs that make the target belief true, on the other. Adding an internalist awareness requirement to the account of identifying reasons strikes me as psychologically implausible. Just as one does not have to be aware that one's beliefs track the facts for them to have positive epistemic status, one does not have to be aware of the fact that the proposition's truthmaker and the facts referred to by one's reasons stand in an identity, causal or grounding relation. One can know *p* in virtue of believing *p* on the basis of reasons that identify *p*'s truthmaker, while being ignorant of the fact that one's reasons identify *p*'s truthmaker.

Externalist identificationism takes reflective epistemic luck cases – such as BonJour's clairvoyance example[25] or Pritchard's chicken sexer example[26] – to be fully compatible with knowledge. Another consequence of the marriage of epistemic externalism with identificationism is that the causal relations between the justification-maker and the truthmaker may be deviant. After all, there is no such thing as a deviant causal chain *per se*. A causal chain is only deviant relative to our expectations. But given epistemic externalism, the subject need not have views regarding the metaphysical microstructure of knowledge in order to know.

It should be obvious that the identificationist account of knowledge can rule out *all* Gettier cases, regardless of whether the Gettierization is due to covariation-failure or due to identification-failure. The reason the sheriff in *Double Trou-*

24 For pluralism about truthmaking see Griffith, Aaron M. (2013): *Ways of Truthmaking: A Pluralist Theory of Truthmaking*. Dissertation. Irvine.

25 BonJour, Laurence (1985): *The Structure of Empirical Knowledge*. Cambridge, 42.

26 Pritchard 2005, 43.

ble does not know that the spy died due to a sniper's bullet is that the sheriff's evidence points to sniper *B*, but it is the presence of sniper *A* that makes the belief true. Since sniper *A* and sniper *B* operate independently from one another, there is no adequate connection between the justification-maker and the truthmaker. If the case were changed so that sniper *B* fires a shot if and only if sniper *A* has fired a shot, the epistemic assessment of the situation would be different. Since now there is a causal relation between the justification-maker and the truthmaker, the sheriff gets to know that the spy died due to a sniper's bullet.

VIII Conclusion

If the only goal of an epistemological theory were to come up with a condition that rules out all kinds of Gettier cases, there would be no need to look beyond the no-essential-false-assumption approach to knowledge (provided there is a satisfactory account of what it is for a false assumption to be essential). The no-essential-false-assumption approach is sufficient to rule out both ordinary and 'intractable' Gettier cases. But given that we are in the business of developing *theories* of knowledge, we cannot content ourselves with claiming that what it means for epistemic reasons to be truth-conducive is that they are true. As was explained before, the no-essential-false-assumption approach gets things backwards: evidence is a guide to (non-accidental) truth rather than truth being a requirement for evidence. Identificationism is as good as the no-essential-false-assumption approach in ruling out Gettier cases but, in addition, it sheds light on the nature of knowledge: *S* knows that *p* on the basis of reasons *R* only if (i) *p* is true, (ii) *S* believes that *p*, (iii) the belief in *p*, the truth of *p*, and the reasons *R* for *Bp* covary in close possible worlds, and (iv) the states of affairs referred to by *S*'s reasons *R* for believing that *p* are identical with, causally related to, or grounded in, *p*'s truthmaker.[27]

27 For helpful comments on an earlier draft of this paper I am grateful to Forrest Fleming and Joachim Horvath.