A Note on Gettier Cases in Epistemic Logic

(penultimate draft of paper in *Philosophical Studies*, published online on 1 August 2014, DOI 10.1007/s11098-014-0357-1)

Timothy Williamson

Abstract:

The paper explains how Gettier's conclusion can be reached on general theoretical grounds within the framework of epistemic logic, without reliance on thought experiments. It extends the argument to permissive conceptions of justification that invalidate principles of multi-premise closure and require neighbourhood semantics rather than semantics of a more standard type. The paper concludes by recommending a robust methodology that aims at convergence in results between thought experimentation and more formal methods. It also warns against conjunctive definitions as sharing several of the drawbacks of disjunctive definitions.

Keywords: Gettier, knowledge, justification, epistemic logic, neighbourhood semantics, thought experiments

1. *Introduction*

Edmund Gettier's 1963 paper is not just famous; it is famous for being famous. It is celebrated as a turning-point in epistemology, and partly for that reason has become a central test case in debates on philosophical method. More specifically, it is standardly taken as a convenient paradigm of recent analytic philosophy's reliance on thought experiments. Gettier is interpreted as simply eliciting case-specific judgments about his two imagined examples, that they involve justified true belief without knowledge. The question under debate is what authority, if any, such case-specific judgments or 'intuitions' should enjoy.[1]

In fact, Gettier's approach in the paper is more theoretical. He stipulates that he is using 'that sense of "justified" in which S's being justified in believing P is a necessary condition of S's knowing P', and then makes two general claims. The first is that justification is not factive: 'it is possible for a person to be justified in believing a proposition that is in fact false'. Gettier's second general claim is a closure principle for justification: 'for any proposition P, if S is justified in believing P, and P entails Q, and S deduces Q from P and accepts Q as a result of this deduction, then S is justified in believing Q'. His counterexamples then illustrate the fatal consequences of these theoretical points for justified true belief as a supposed necessary and sufficient condition for knowledge. Significantly, both of Gettier's theoretical points concern *justification*, rather than *knowledge*. That is hardly surprising, for 'S is justified in believing P' is a less ordinary, more theoretical-sounding locution than 'S knows P'. Indeed, one can argue that the way most contemporary epistemologists apply the term 'justified' to belief involves an artificial disciplinary convention, rooted in a confusion between justification and blamelessness.[2] However, for purposes of this paper, I will follow the currently standard use of 'justified' in epistemology. In particular, I will adhere to both of Gettier's theoretical principles about justification: non-factivity and closure.

Gettier provides no theoretical backing for his claims that the subjects in his examples do not know. He simply treats them as obvious. Those are the claims that have become at least mildly controversial, and they do exemplify the case-specific methodology. I regard Gettier's denials of knowledge as not just true but obvious, in a way that is neither ethnicity-specific nor gender-specific. Nevertheless, it is desirable to have some independent confirmation of them, preferably of a more theoretical kind. The point is not just to reassure ourselves that they are indeed true. By deriving them from more general considerations, we stand to gain theoretical insight into the nature of knowledge and (in the relevant sense) justification. Furthermore, we can thereby hope to rebut a milder but insidious form of scepticism, which concedes the denials of knowledge in Gettier cases but queries their significance. The fear is that they reveal mere *quirks* of knowledge, because oddities of humans' species-specific epistemological sensibility make us latch onto a gerrymandered cognitive relation that lacks deeper theoretical significance. If Gettier cases can be predicted on general theoretical grounds, that dismissive interpretation of them is unwarranted.

In previous work, I have argued that Gettier cases can indeed be predicted within the general framework of *epistemic logic*.[3] Section 2 briefly rehearses the argument. The main novelty of the present paper is in section 3, which generalizes the argument to a wider range of settings and assumptions, thereby showing its conclusions to be robust.

To many epistemologists, the intellectual style of epistemic logic is unfamiliar and unsettling. In particular, they are disturbed by its reliance on idealizations such as logical omniscience, on which one automatically knows whatever follows from what one knows. Their outraged instinct is to cite counterexamples. That is the analogue of the outraged folk physicist giving examples of planets that are not point masses. The use of idealized models is ubiquitous in the natural and social sciences. Practitioners are normally aware of the idealizations and know how to handle them. The aim is to gain insight into a phenomenon by studying how it works under simplified, rigorously described conditions that enable us to apply mathematical or quasi-mathematical reasoning that we cannot apply directly to the phenomenon as it occurs in the wild, with all its intractable complexity. We can then cautiously transfer our insight about the idealized model back to the phenomenon in the wild. The selection and application of such models requires experience, skill, and good judgment, which are not primarily matters of mathematical facility. So far, mainstream epistemology has made disappointingly little use of such techniques. The reader is urged to understand what follows in the same spirit as idealized model-building in the natural and social sciences.

2. *Gettier cases in epistemic logic*

For the time being, a model is a *frame* <W, R, S>, where W is a set and R and S are binary relations over W, that is, sets of ordered pairs of members of W.

Informally, we interpret W as a set of mutually exclusive, jointly exhaustive worlds or states, which may but need not be metaphysically possible ('informally', because that interpretation does not constrain the formal definitions and arguments). Informally, we interpret subsets of W as coarse-grained propositions, for which mutual entailment amounts to identity. If $p \subseteq W$ then think of *p* as true at every world in *p* and false at every other world. Thus the intersection of two propositions is their conjunction and their union is their disjunction; the complement of a proposition in W is its negation. Furthermore, $p \subseteq q$ if and only if *p* entails *q*, in the sense that *q* is true at every world at which *p* is true. Similarly, $(p \cap q) \subseteq r$ if and only if *p* and *q* together entail *r*, in the sense that *r* is true at every world at which both *p* and *q* are true.

Still informally, we interpret the relations R and S as encoding the epistemology of the frame, given a cognizing subject ('the agent') and a time (the present tense). R is a relation of *epistemic accessibility*: a world *w* has R to a world *x* if and only if for all the agent knows in *w*, the agent is in *x*, that is, whatever the agent knows in *w* is true in *x*. Since knowledge is factive, everything the agent knows in *w* is true in *w*, so R should be *reflexive* (every world has R to itself); we shall only consider models meeting that condition. We define a function *K* from propositions to propositions by the following equation for all propositions *p*:

$$Kp = \{w \in W:\ \forall x \in W\ (wRx \Rightarrow x \in p)\}$$

In other words, *Kp* is true at a world if and only if *p* is true at every world epistemically accessible from that one. Informally, *Kp* is interpreted as the proposition that the agent knows *p*. Thus the agent knows something if and only if it is true at every uneliminated possibility, where the eliminated possibilities are the epistemically inaccessible worlds. These remarks are not intended as an *analysis* of knowledge in terms of epistemic accessibility, since epistemic accessibility was itself explained in terms of knowledge. Rather, they are just a recipe for informally decoding how knowledge behaves according to the frame from the relation R. For any world *w* in any frame there is a *strongest* thing the agent knows at *w*, the proposition R(*w*), defined as the set of all worlds epistemically accessible from *w*, $\{x \in W: wRx\}$. For, by the definition of *K*, for any proposition *p* and world *w*, *Kp* is true at *w* if and only if R(*w*) entails *p*: $w \in Kp$ if and only if $R(w) \subseteq p$. That convenient equivalence will be used without comment in some of the arguments to come. One obvious corollary is an unrestricted multi-premise closure principle for knowledge, for if some premises $p_1, \ldots, p_n$ entail a conclusion *q*, and the agent knows each of $p_1, \ldots, p_n$ at a world *w*, then R(*w*) entails each of $p_1, \ldots, p_n$, so R(*w*) entails *q*, so the agent knows *q* at *w*. More formally:

MPCK If $(p_1 \cap \ldots \cap p_n) \subseteq q$ then $(Kp_1 \cap \ldots \cap Kp_n) \subseteq Kq$

There is no restriction to cases in which the agent goes through a process of deducing the conclusion from the premises, competently or otherwise, nor is any allowance made for an agent's failure to recognize that the same proposition is recurring under different guises. This is the main idealization about knowledge in the models.

The account of justified belief in the frame is similar. For present purposes, it is convenient to treat justified belief as a single phenomenon, rather than separating it out into justification and belief components, though nothing said here will preclude the possibility of so analysing it. S is a relation of *doxastic accessibility*: a world *w* has S to a world *x* if and only if whatever the agent believes with justification in *w* is true in *x*. Since justified belief is not factive, S will not in general be reflexive. However, S should be *serial*, in the sense that every world has S to at least one world, for a world that has S to no worlds is a world in which the agent vacuously has a justified belief in every proposition whatsoever, and so has mutually inconsistent justified beliefs, an eventuality we may exclude, at least at this level of idealization. We define a function *J* from propositions to propositions by the following equation for all propositions *p*:

$$Jp = \{w \in W: \forall x \in W\ (wSx \Rightarrow x \in p)\}$$

In other words, *Jp* is true at a world if and only if *p* is true at every world doxastically accessible from that one. Informally, *Jp* is interpreted as the proposition that the agent has a justified belief in *p*. This is not intended as an *analysis* of justified belief in terms of doxastic accessibility, since doxastic accessibility was itself explained in terms of justified belief. Rather, it is just a recipe for informally decoding how justified belief behaves according to the frame from the relation S. For any world *w* in any frame there is a *strongest* thing the agent has a justified belief in at *w*, the proposition S(*w*), defined as the set of all worlds doxastically accessible from *w*, $\{x \in W: wSx\}$. For, by the definition of *J*, for any proposition *p* and world *w*, *Jp* is true at *w* if and only if S(*w*) entails *p*: $w \in Jp$ if and only if $S(w) \subseteq p$. That convenient equivalence will be used without comment in some of the arguments to come. One obvious corollary is an unrestricted multi-premise closure principle for justified belief, for if some premises $p_1, \ldots, p_n$ entail a conclusion *q*, and the agent has a justified belief in each of $p_1, \ldots, p_n$ at a world *w*, then S(*w*) entails each of $p_1, \ldots, p_n$, so S(*w*) entails *q*, so the agent has a justified belief in *q* at *w*. More formally:

MPCJ If $(p_1 \cap \ldots \cap p_n) \subseteq q$ then $(Jp_1 \cap \ldots \cap Jp_n) \subseteq Jq$

As before, there is no restriction to cases in which the agent goes through a process of deducing the conclusion from the premises, competently or otherwise, nor is any allowance made for an agent's failure to recognize that the same proposition is recurring under different guises. This is the main idealization about justified belief in the models. Although these idealizations about knowledge and justified belief may be drastic, that can work to our advantage. For if Gettier cases occur even under these drastic idealizations, that is all the more reason to regard them as a robust phenomenon.

The justified true belief account of knowledge can be expressed as a simple equation, for all propositions p, with one conjunct on the right-hand side for truth and one for justified belief:

JTB $\quad Kp = p \cap Jp$

A trivial corollary of JTB is that knowledge entails justification: $Kp \subseteq Jp$, just as Gettier stipulated in explaining how he was using the word 'justified'. This is equivalent to the claim that at every world w, the strongest thing the agent has a justified belief in entails the strongest thing the agent knows: $S(w) \subseteq R(w)$.[4] Very roughly: since knowledge is at least as demanding as justification, there are at least as many worlds in which something must hold for it to be known in w as there are worlds in which it must hold for it to be believed with justification in w. In what follows, we may assume that, without exception, $Kp \subseteq Jp$ and $S(w) \subseteq R(w)$, for defenders of JTB must assume those relations anyway, since they follow from JTB.

We can now taxonomize worlds in such frames to describe exactly the circumstances in which Gettier cases arise, that is, in which there are counterexamples to JTB. By hypothesis, $Kp \subseteq Jp$ and $Kp \subseteq p$ (the latter because R is reflexive), so $Kp \subseteq p \cap Jp$. On these assumptions, any counterexample to JTB must be a counterexample to the converse inclusion, a case of justified true belief without knowledge, just as a Gettier case is supposed to be. Consider a given world *w* in a frame of the kind described. We divide the cases thus:

Case (i): $w \notin S(w)$. This means that the strongest proposition in which the agent has a justified belief at *w* is false at *w*, so in this case the agent has justified false beliefs. The agent also has justified true beliefs, for instance in the trivial proposition W true at all worlds. Less trivially, the agent has a justified true belief in the proposition $S(w) \cup \{w\}$. Indeed, $S(w) \cup \{w\}$ is the *strongest* proposition in which the agent has a justified true belief at *w*. For if the agent has a justified true belief in the proposition *p* at *w*, then $S(w) \subseteq p$ (because S(*w*) is the strongest proposition in which the agent has a justified belief at *w*) and $w \in p$ (because *p* is true at *w*), so $S(w) \cup \{w\}$ entails *p*. Conversely, if $S(w) \cup \{w\}$ entails *p* then the agent has a justified belief in *p* (because S(*w*) entails *p*) and *p* is true at *w*. We subdivide case (i) thus:

Case (ia): $R(w) = S(w) \cup \{w\}$. Since $S(w) \cup \{w\}$ is the strongest proposition in which the agent has a justified true belief at *w*, in this subcase what the agent knows is exactly what the agent has a justified true belief in, so *w* is not a counterexample to JTB for any proposition *p*. There are no Gettier cases at *w*. However, this looks like a rather special circumstance, a world in which the agent's justified belief and the agent's knowledge differ in content only by a single world.

Case (ib): $R(w) \neq S(w) \cup \{w\}$. But $S(w) \subseteq R(w)$ by hypothesis, and $w \in R(w)$ because R is reflexive, so $S(w) \cup \{w\} \subseteq R(w)$. Hence it is not the case that $R(w) \subseteq S(w) \cup \{w\}$. But that means that the agent does not know $S(w) \cup \{w\}$ at w. Since the agent has a justified true belief in $S(w) \cup \{w\}$ at w, $S(w) \cup \{w\}$ constitutes a Gettier case at w. It has the same overall structure as Gettier's original cases, inheriting its justification from the false disjunct $S(w)$ and its truth from the unjustified disjunct $\{w\}$. Since it involves a justified true belief derived from a justified false belief, we may call it an *impurely veridical* Gettier case. Of course, 'derived' here just marks the logical relation of entailment between $S(w)$ and $S(w) \cup \{w\}$; the coarse-grained nature of the frame does not permit us to ask whether the agent has gone through a process of deducing $S(w) \cup \{w\}$ from $S(w)$, but that simplification is only to be expected.

Case (ii): $w \in S(w)$. This means that the strongest proposition in which the agent has a justified belief at w is true at w, so at w all the agent's justified beliefs are true. We subdivide case (ii) thus:

Case (iia): $R(w) = S(w)$. In this subcase, like (ia), what the agent knows is exactly what the agent has a justified true belief in, so w is not a counterexample to JTB for any proposition p. There are no Gettier cases at w. However, this too looks like a rather special circumstance, a world in which the agent's justified belief and the agent's knowledge exactly coincide in content.

Case (iib): $R(w) \neq S(w)$. But $S(w) \subseteq R(w)$ by hypothesis, so it is not the case that $R(w) \subseteq S(w)$. But that means that the agent does not know $S(w)$ at w. Since the agent has a justified belief in $S(w)$ at w, $S(w)$ constitutes a Gettier case at w. Its overall structure is more like that of a fake barn case, since it does not involve justified false beliefs.[5] We may therefore call it a *purely veridical* Gettier case.

On purely general grounds, we have thus constructed a framework in which the two main types of Gettier case found in the literature naturally arise, and in which the worlds without Gettier cases look like very special cases. However, we have to move carefully, since friends of JTB might argue that although the Gettier subcases (ib) and (iib) are formally possible, they are not genuine epistemological possibilities given the intended readings of *K* and *J* in terms of knowledge and justified belief respectively, so that the Gettier-free subcases (ia) and (iia) are not genuinely special. To counter that suggestion, we need to show that (ib) and (iib) are indeed genuine epistemological possibilities. That might seem to invite a return to the method of thought experimentation after all, thereby abandoning the attempt to provide independent corroboration of its results. But that is not so. Far more general considerations suffice.

When epistemologists try to explain what they understand by justified belief without reference to knowledge — as would be required for JTB to constitute a non-circular analysis of knowledge — they typically gloss it along the following lines. Justification may remain constant while knowledge varies, depending on factors to which the agent may have no access, such as the reliability of their perceptual faculties and the conduciveness of environmental conditions. According to an extreme version of this idea, you have exactly the same justified beliefs as a brain in a vat to whom everything appears as it does to you. Even if one rejects that version, one may allow much less extreme versions in which the agent is merely the victim of a practical joke in the bad case. Many of these less extreme versions are consistent with many externalist accounts of justification, on which it may depend on general background conditions inaccessible to the agent.

In the spirit just explained, let us postulate two worlds, a good case w and a corresponding bad case x, between which justification is constant, so S(w) = S(x), while the agent knows much less in x than in y, so R(x) contains many worlds (uneliminated epistemic possibilities) not in R(w). For example, for each shade of colour, it is consistent with what the brain in a vat knows that it is a brain in a vat of that shade. More specifically, we assume that R(x) contains at least *two* worlds not in R(w). Now suppose that the bad case x does not involve a Gettier case. Hence, by the taxonomy above, x falls under either subcase (ia) or (iia), so $R(x) \subseteq S(x) \cup \{x\} = S(w) \cup \{x\}$. Moreover, by standing hypothesis, $S(w) \subseteq R(w)$. Together, those two inclusions imply that R(x) contains at most one world, x, not in R(w). But that contradicts the assumption that R(x) contains at least two worlds not in R(w). Therefore x does contain a Gettier case after all. In short, JTB allows only a small difference between what the agent has justified belief in and what the agent knows; but if what the agent has justified belief in is constant between a good case and a bad case, while what the agent knows varies drastically between them, then the difference between what the agent has justified belief in and what the agent knows is not always small. Whether the Gettier case is purely or impurely veridical depends on whether x belongs to S(x) (subcase (iib)) or not (subcase (ib)). Typically, if x is only a very mildly bad case, the Gettier case will be purely veridical; if x is much worse, the Gettier case will be impurely veridical. What matters is that we have predicted on very general structural grounds that a Gettier case will occur.

The argument does not depend on the method of thought experiments. Brains in vats were mentioned only to illustrate the account of justification; much less exotic examples of blameless epistemic misfortune would do just as well. The phenomenon of blameless error of the relevant kind is far too familiar to need verification by thought experiments. Moreover, both internalists and externalists accept it. In any case, friends of the JTB analysis of knowledge obviously accept the non-factiveness of knowledge (Gettier's first theoretical claim), for if justification were factive the truth conjunct of the analysis would be redundant. They also accept the factiveness of knowledge, which follows from JTB. Nor does the support for the (idealized) closure principles for justification and knowledge MPCJ and MPCK come from thought experiments; rather, it comes from a more theoretical positive assessment of deduction as a cognitive process. Furthermore, whereas the judgment that the agent does not know is the focus of the controversy over Gettier's thought experiments, in the argument just presented the reader was at no point asked to judge that the agent in a hypothetical (or real case) does not know a given truth. Of course, any necessarily true claims in the argument, about knowledge, justification, or anything else, will *ipso facto* hold in the possible scenario of any thought experiment, but that does not mean that the thought experiment plays any role in the argument.

The argument is quite robust to perturbations of the original assumptions. Even if S(*w*) and S(*x*) are not identical, but differ only over a narrow range of cases, the argument still goes through given that R(*w*) and R(*x*) range over a significantly wider range of cases. In brief, JTB forces knowledge and justification always to stay close together while the underlying account of justification forces them sometimes to be far apart. Thus counterexamples to JTB are bound to occur. However, the argument can be made even more robust, as the next section will show.

3. Permissive conceptions of justification

The multi-premise closure principle for justification, MPCJ, fits a conception of (epistemically) justified belief as in some sense (epistemically) *obligatory* belief. In particular, it is plausible that such obligations agglomerate. If I ought to stand and I ought to salute, then I ought to do both. If one ought to believe *p* and one ought to believe *q*, then one ought to believe both, and so (arguably) ought to believe their conjunction $p \cap q$. Of course, issues arise about clutter avoidance and computational tractability, but those are just the sorts of consideration that basic epistemic logic idealizes away, and in any case they do not seem to have much to do with Gettier cases. However, on an alternative view, justified belief is *permissible* belief rather than *obligatory* belief. Unlike obligations, permissions do not plausibly agglomerate. If I am permitted to take the ice cream and permitted to take the cheese, it does not follow that I am permitted to take both. A treatment of the lottery paradox has been derived from such a permissive account of justified belief: for each ticket one is permitted to believe that it will lose, but one is not permitted to believe that they will all lose.[6] On a more extreme version of the view, one is permitted to believe that there is a God, and one is permitted to believe that there is no God, but one is not permitted to believe a contradiction. A natural probabilistic implementation of the permissive conception is that one is permitted to believe a proposition if and only if its probability on one's evidence reaches some fixed threshold strictly between 0 and 1, say 99% or even 50%. For any such threshold, there will be a pair of propositions each of which reaches the threshold while their conjunction does not. One of the two contemporary philosophers whom Gettier cites for the JTB analysis, A. J. Ayer, provides just such a permissive account, with 'S has the right to be sure that P' in place of 'S is justified in believing that P' and 'S is sure that P' in place of 'S believes that P' (Ayer 1956, pp. 33-4).

How can we generalize the structural argument against JTB to cover permissive conceptions of justification? We can no longer appeal to the definition of the *J* function in terms of the doxastic accessibility relation S, since it automatically validates MPCJ. In this new setting, there may be no such thing as the strongest proposition in which the agent has a justified belief at a given world. However, we can still model justification as a function *J* from propositions to propositions, and propositions as sets of worlds, in the setting of so-called *neighbourhood semantics*.[7] Moreover, a weaker principle of *single*-premise closure for justification remains plausible on the permissive conception:[8]

SPCJ If $p \subseteq q$ then $Jp \subseteq Jq$

For instance, if *p* reaches the probabilistic threshold, and *p* entails *q*, then *q* reaches the probabilistic threshold. More generally, if one is permitted to believe a proposition, one is plausibly permitted to believe anything it entails, and the permissive conception is consistent with an idealization to an agent who does indeed believe those entailed propositions. Thus we assume SPCJ in what follows.

Having abandoned the definition of *J* in terms of doxastic accessibility for present purposes, we may also be wary of the structurally parallel definition of *K* in terms of epistemic accessibility, on grounds of fairness to defenders of JBT, since the close relation it requires between justification and knowledge may suggest that the two definitions should stand or fall together. We shall therefore not appeal to that definition of *K*, although we still model knowledge as a function *K* from propositions to propositions. We can no longer appeal to MPCK, the principle of multi-premise closure for knowledge, since it goes with the definition of *K* in terms of epistemic accessibility. In this new setting, there may be no such thing as the strongest truth the agent knows at a given world. We can still maintain a principle of single-premise closure for knowledge, in parallel with justification:

SPCK If $p \subseteq q$ then $Kp \subseteq Kq$

Indeed, SPCK follows from SPCJ given JTB (just as MPCK follows from MPCJ given JTB), so friends of JTB should accept SPCK if they accept SPCJ. We still have the principle that knowledge entails justification, for the same reason as before.

We can now reconstruct the argument for Gettier cases just on the basis of these weaker assumptions. As before, we suppose a good case *w* and a less good case *x*. Just as before, we assume that justified belief is constant between *w* and *x*, but we now have to express the assumption thus: for every proposition *p*, $w \in Jp$ if and only if $x \in Jp$. We cash out the assumption that the agent knows much less at *x* than at *w* thus: for some proposition *p*, the agent knows *p* at *w* and does not know anything nearly as strong as *p* at *x*. Here a proposition *q* counts as 'nearly as strong as *p*' if and only if there are only a few worlds where *p* is false and *q* is true: *q* excludes all but a few of the worlds that *p* excludes. (These are stipulative definitions of 'knows much less' and 'nearly as strong', not assumptions.) For example, in the good case one may know that one's car has not been stolen, while in the bad case for many natural numbers *n* it is consistent with everything one knows that one's car was stolen exactly *n* minutes ago, so one knows nothing nearly as strong as that one's car has not been stolen. Thus the assumption that one knows much less much in the bad case than in the good case is not very demanding.

Suppose that the agent knows *p* at *w* and does not know anything nearly as strong as *p* at *x*. Since knowledge entails justification, the agent has a justified belief in *p* at *w*. Since justified belief is constant between *w* and *x*, the agent also has a justified belief in *p* at *x*. Of course, *p* may be false at *x*, even though it was true at *w* (because it was known at *w*). However, *p* entails $p \cup \{x\}$, so by SPCJ the agent also has a justified belief in $p \cup \{x\}$ at *x*, and $p \cup \{x\}$ *is* true at *x*. But $p \cup \{x\}$ is nearly as strong as *p*, since there is at most one world, *x*, where *p* is false and $p \cup \{x\}$ is true.[9] By hypothesis, therefore, the agent does not know $p \cup \{x\}$ at *x*. Thus we have our Gettier case at *x*: justified true belief without knowledge of $p \cup \{x\}$. As before, the argument does not depend on the method of thought experiments. It too is robust to minor perturbations of its premises.

For permissive conceptions of justification we can still distinguish between impurely veridical and purely veridical Gettier cases, although in slightly different terms from before, since even if the agent has some justified false beliefs at *x*, they need not include one that entails *p*. Instead, we distinguish the two sorts of case relative to the proposition *p* as well as the world *x*, where (as above) the agent does not know $p \cup \{x\}$ at *x*:

Case (i). The agent has a justified false belief at *x* in some proposition *r* that entails *p* ($x \notin r$; $x \in Jr$; $r \subseteq p$). Hence *r* entails $p \cup \{x\}$. Then the agent's justified true belief without knowledge at *x* in $p \cup \{x\}$ is an impurely veridical Gettier case.

Case (ii). The agent has no justified false belief at *x* in a proposition *r* that entails *p* (in particular, therefore, since *p* entails itself and is justified at *x*, *p* is true at *x*). This is a purely veridical case Gettier case.

Whether a given Gettier case is purely or impurely veridical depends on the details of the case and on how 'justification' is understood.

In general, the idea that justification can be constant across corresponding good and bad cases provides the natural motivation for the first of Gettier's two general theoretical claims mentioned in section 1, that one can be justified in believing falsehoods. Notably, Gettier's second general theoretical claim is a principle of *single*-premise closure for justification, corresponding to SPCJ rather than MPCJ, though with the proviso required in Gettier's unidealized setting, that the agent accepts the conclusion as a resulting of deducing it from the premise. Thus the materials of the present argument are quite similar to Gettier's general theoretical claims.

The upshot of this section is that the underlying structural objection to JTB generalizes from an understanding of justifications as obligations to an understanding of them as permissions.

4. *Concluding reflections*

Natural formal models of knowledge and justified belief (in a non-factive sense) provide robust evidence against JTB, independently of thought experiments in any distinctive sense, but in a way closely related to Gettier's original arguments. Experimental philosophy will not save JTB. Nor does the idea that the word 'knowledge' picks out non-factively justified true belief because it is the most natural candidate roughly to fit our use of the word look promising on the evidence of the models.[10] For if justified belief is basic, then on JTB the strongest things we know at *w* will typically be odd disjunctions of the strongest things in which we have justified belief at *w* with the singleton of *w* itself, which would make knowledge a far from natural relation.[11]

None of this means that formal models make thought experiments redundant. Rather, the mutual confirmation of the results of the two methods should increase our confidence in each method.

We might also draw another general moral for philosophical method. Philosophers are used to the idea that disjunctive definitions tend not to pick out theoretically useful distinctions, because they do not carve at the joints. By contrast, conjunctive definitions such as JTB have stood under no such cloud. After all, when two things fall under a disjunctive definition, they may be quite dissimilar, because they fall under different disjuncts, whereas when two things fall under a conjunctive definition, they must be quite similar, because they both fall under the same conjuncts (all of them). However, we saw that in epistemic logic what fall under the JTB definition are all and only disjunctions of a disjunct believed with justification and a true disjunct, a point already hinted at in Gettier's counterexamples. The conjunctive definition has a disjunctive obverse. That is not very surprising, given the logical duality of conjunction and disjunction (interchanging ‘T’ and ‘F’ throughout the truth-table for either operator yields the truth-table for the other); the distinction between a conjunction and its negation is equivalent to the distinction between the disjunction of the negated conjuncts and its negation.[12] Faced with a conjunctive definition, our first thought should be to doubt that it has enough unity to give us a theoretically useful distinction.[13]

Notes

1 Much of the debate concerns alleged experimental findings of ethnic or gender variation in judgments about Gettier cases, following Weinberg, Nichols, and Stich 2001. For a recent defence of the method of cases see Nagel 2012, and for a recent experimental study that did not find such bias see Nagel, San Juan, and Mar 2013. For a different kind of scepticism about Gettier cases see Weatherson 2003. My discussion of Gettier cases in Williamson 2007 also concentrated on case-specific judgments.

2 See Williamson 2013a, 2013b, 201X.

3 See Williamson 2013a, 2013b, 2014.

4 Proof: Suppose $Kp \subseteq Jp$ for all propositions p; but for any world w, $w \in K\mathrm{R}(w)$, so $w \in J\mathrm{R}(w)$, so $\mathrm{S}(w) \subseteq \mathrm{R}(w)$. Conversely, suppose $\mathrm{S}(w) \subseteq \mathrm{R}(w)$ for all worlds w; but if $w \in Kp$, then $\mathrm{R}(w) \subseteq p$, so $\mathrm{S}(w) \subseteq p$, so $w \in Jp$; thus $Kp \subseteq Jp$.

5 The case was first published in Goldman 1976, which acknowledges Carl Ginet for the example.

6 See Kroedel 2012.

7 See Hughes and Cresswell 1998, pp. 221-3.

8 Chellas 1980, p. p. 234, calls modal logics with this rule *monotonic*.

9 Of course, if *p* is true at *x* then $p \cup \{x\} = p$.

10 See Weatherson 2003.

11 Artemov 2008 analyses Gettier's arguments and related considerations in the framework of *justification logic*, a refinement of epistemic logic in which the structure of justifications can be explicitly represented in the formal language. For present purposes, the austere framework of unrefined epistemic logic is preferable, because it assumes less and makes the comparison between knowledge and justified true belief more perspicuous. Nevertheless, justification logic is an intriguing resource for epistemologists to exploit.

12 Another problem for strictly conjunctive analyses is that they disallow *compensation* between how a putative instance scores on the various dimension relevant to the conjuncts. To put the point schematically, let being *F* depend on doing well on *n* dimensions, with compensation between dimensions. Suppose that we analyse what it is for *x* to be *F* as a conjunction of *n* conjuncts, where the *i*th conjunct is that $t_i < x_i$, where x_i is how well *x* does on the *i*th dimension, t_i is the required threshold for that dimension, and < is the relevant ordering relation. Given compensation

between dimensions, we should have cases like this: *a* is *F* and *b* is *F*, where $b_i < a_i$ but $a_j < b_j$ (*b* compensates for doing worse than *a* on dimension *i* by doing better than *a* on dimension *j*), but *c* is not *F*, where $c_i = b_i$ and $c_j = a_j$ (*c* does not compensate for doing worse than *a* on dimension *i* by doing better than *a* on dimension *j*; for simplicity, assume that on each other dimension *a*, *b*, and *c* are equal). But this cannot happen on the conjunctive model. For since *b* is *F*, it satisfies the *i*th conjunct, so $t_i < b_i = c_i$, so *c* satisfies the *i*th conjunct too; since *a* is *F*, it satisfies the *j*th conjunct, so $t_j < a_j = c_j$, so *c* satisfies the *j*th conjunct too; since *c* equals *a* and *b* on all the other dimensions, it also satisfies all the other conjuncts; thus *c* is *F* on the conjunctive analysis.

13 This article develops half of my talk at the 2013 'Gettier Problem at 50' conference in Edinburgh; Williamson 201X develops the other half. I thank Allan Hazlett and audiences there and at the Universities of Michigan, Oxford, and Virginia for helpful comments.

References

Artemov, Sergei 2008: 'The logic of justification', *The Review of Symbolic Logic*, 1, 477-513.

Ayer, A. J. 1956: *The Problem of Knowledge*. London: Macmillan.

Chellas, Brian 1980: *Modal Logic: An Introduction*. Cambridge: Cambridge University Press.

Gettier, Edmund 1963: 'Is justified true belief knowledge?', *Analysis*, 23, pp. 121-123.

Goldman, Alvin 1976: 'Discrimination and perceptual knowledge', *The Journal of Philosophy*, 73, pp. 771-791.

Hughes, George; Cresswell, Max 1996: *A New Introduction to Modal Logic*. London: Routledge.

Kroedel, Thomas 2012: 'The lottery paradox, epistemic justification, and permissibility', *Analysis*, 72, pp. 57-60.

Nagel, Jennifer 2012: 'Intuitions and experiments: a defense of the case method in epistemology', *Philosophy and Phenomenological Research*, 85, pp. 495-527.

Nagel, Jennifer; San Juan, Valerie; Mar, Raymond 2013: 'Lay denial of knowledge for justified true beliefs', *Cognition*, 129, pp. 652-661.

Weatherson, Brian 2003: 'What good are counterexamples?', *Philosophical Studies*, 115, pp. 1 -31.

Weinberg, Jonathan; Nichols, Shaun; Stich, Stephen 2001: 'Normativity and epistemic intuitions', *Philosophical Topics*, 29, pp. 429-460.

Williamson, Timothy 2007: *The Philosophy of Philosophy*. Oxford: Blackwell Wiley.

Williamson, Timothy 2013a: 'Gettier cases in epistemic logic', *Inquiry*, 56, pp. 1-14.

Williamson, Timothy 2013b: 'Response to Cohen, Comesaña, Goodman, Nagel, and

Weatherson on Gettier cases in epistemic logic', *Inquiry*, 56, pp. 77-96.

Williamson, Timothy 2014: 'Very improbable knowing', *Erkenntnis*, forthcoming.

Williamson, Timothy 201X: 'Justifications, excuses, and sceptical scenarios', in Fabian Dorsch and Julien Dutant (eds.), *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*, forthcoming.