

# How Believing Can Fail to Be Knowing

Murali RAMACHANDRAN

BIBLID [0495-4548 (2006) 21: 56; pp. 185-194]

ABSTRACT: This paper defends a simple, externalist account of knowledge, incorporating familiar conditions mentioned in the literature, and responds to Timothy Williamson's charge that any such analysis is futile because knowledge is semantically un-analyzable. The response, in short, is that even though such an account may not offer a *reductive* analysis of knowledge—by way of more basic, non-circular concepts—it still has an *explanatory* advantage over Williamson's own position: it explains how belief can fail to be knowledge.

Key words: knowledge, externalism, closure principle, Williamson.

## 1. A simple analysis

Here is a comparatively simple analysis of knowledge that I wish to promote and develop in this paper:

A subject,  $S$ , *knows* that  $p$  if, and only if, the following conditions obtain:

- 1) It is true that  $p$ . [*Truth*]
- 2)  $S$  believes that  $p$ . [*Belief*]
- 3)  $S$ 's rationality is not impugned by her meeting the conditions of this analysis, or by the way she meets them. [*Rationality*]
- 4)  $S$ 's belief in  $p$  does not rest significantly on any false beliefs or presumptions. [*Fact-Grounded Reasons*]
- 5) A significant body of reasons  $S$  has for believing  $p$  she would not have if  $p$  were not true. [*Fact-Dependent Reasons*]

These conditions are hardly original: conditions (1)-(3) comprise a version of the traditional *justified-true-belief* theory of knowledge; (4) echoes Clark's (1963) *well-groundedness* and Harman's (1968) *no-false-lemma* conditions; and (5) is simply Dretske's (1971) requirement that  $S$  has, in his terminology, *conclusive reasons* for believing  $p$ .<sup>1</sup> However, as we shall see, combining the conditions as above yields a surprisingly effective account.

In this section I will explicate the conditions by showing how they yield a positive verdict in a simple case. §§2 and 3 explain why conditions (4) and (5) are needed, what precisely they achieve; §5 addresses the charge, recently made by Williamson (2000) for example, that any such analysis is futile if not impossible!

---

<sup>1</sup> R is a *conclusive reason* for  $S$  believing  $p$  if  $S$  would not have R unless  $p$  were true (Dretske, 1971, p. 41).



Let us begin by seeing how the clauses work in a straightforward case of perception. *S* sees a cat in front of her and so believes truly (*p*) that there is a cat in front of her (conditions (1) and (2) are met). Her belief in *p* does not rest significantly on any false beliefs or presumptions. She believes she is seeing a cat, and she is; she presumes there is nothing unusual in the circumstances, and there isn't, etc.. Of course, she may have false beliefs about cats —e.g. that her neighbour's cat is male— or false presumptions about the circumstances —e.g. that it is colder than yesterday— but her belief in *p* does not rest on these. So, condition (4) is met.

Now, what if *p* were false, if there *weren't* a cat in front of her: would she have the same reasons for believing *p*? We are not interested here in *any old* possible situation where *p* is false. It is logically possible, for instance, that there is no cat in front of *S*, but that some devious scientists have set up a hologram of a cat in front of her; in this possible world, arguably, *S* might have precisely the same reasons for believing *p* as she has in the actual world; but, such a world is not sufficiently similar to the actual world to be relevant to our question. We are interested, rather, in not-*p*-worlds that differ as little possible in other respects from the actual world, worlds that differ *just enough* so that *p* is not true —the *nearest* not-*p*-worlds as we shall call them.<sup>2</sup> The nearest not-*p*-worlds will simply be worlds in which there isn't a cat in front of *S*, *S*'s eyesight is functioning normally, she isn't being deceived by holograms, etc.. In these worlds, she would not have a visual experience as of seeing a cat, and, consequently, would not have the reasons for believing *p* that she actually has. Condition (5) is met.

Finally, we may suppose, since we are assuming that there is nothing out of the ordinary in the situation, that the satisfaction of conditions (1), (2), (4) and (5) in this case do not impugn *S*'s rationality; in other words, that condition (3) is met. So, by the present account, *S* comes out knowing that there is a cat in front of her.

The key question is, is each condition *necessary* for knowledge? This is what the remaining sections address —by way of considering familiar test cases in the literature.

## 2. Fact-grounded reasons

Conditions (1)-(3) are generally taken to be necessary. So, we may begin by considering why these alone are not jointly sufficient for knowledge. Gettier (1963) offers counterexamples to their sufficiency; the examples have *S* justifiably believing a true proposition *p*, but on the basis of a false belief. Here is a version from Lehrer (1965):

NOGOT (*why (1)-(3) are not jointly sufficient for knowledge*)

A colleague in *S*'s office, Mr. Nogot, has given evidence which leads *S* to justifiably believe (*r*) that Mr. Nogot owns a Ford, from which *S* correctly deduces (*p*) that someone in this office owns a Ford. But, unbeknownst to *S*, Nogot has been shamming and *p* is only true because another person in the office, Ms. Havit, owns a Ford.

---

<sup>2</sup> So, *pace* Lewis (1971) and Nozick (1984), counterfactual conditionals of the form "If *A* were the case, then *B* would be the case" are interpreted as affirming that the nearest *A*-worlds are *B*-worlds.

$S$  clearly does not know that  $p$ . Yet,  $p$  is true and  $S$ 's belief in  $p$  is justified —or, what is relevant for present purposes,  $S$ 's rationality is not impugned by her believing  $p$  for the reasons she does. So, (1)-(3) are not jointly sufficient for knowledge.

An obvious moral, noted by Clark (1963) and Harman (1968), is to discount cases where the belief in  $p$  is based on false beliefs; this is what condition (4), the *fact-groundedness* condition as I shall call it, does.  $S$  does not come out knowing that someone in the office owns a Ford, because her belief that this is so is based on her false belief that Nogot owns a Ford.

By using the phrase “rest significantly on” rather than just “rest on” in the formulation of (4), I want to allow for the possibility of knowledge in cases where some false beliefs play a negligible part in the formation of  $S$ 's belief in  $p$  (see also Harman 1968). Suppose  $S$  has seen (and recognized) Jones in London on many occasions, but has mistaken someone else for Jones in London last Saturday. Condition (4) does not rule out  $S$ 's knowing ( $p$ ) that *Jones has been in London*. For,  $S$ 's belief that  $p$  does not rest *significantly* on her false belief that Jones was in London last Saturday.

The appeal to *presumptions* in the condition is intended to handle the following sort of example from Goldman (1976):

BARN (*why the demand for no-false-presumptions*)

$S$  is driving in a region of the countryside where, unbeknownst to her, there are many papier-maché facsimiles of barns (easily mistaken for genuine barns from the roadside —indeed, we may suppose that  $S$  has mistaken many of them for barns in the course of her journey). As it happens,  $S$  stops in front of a genuine barn and believes truly ( $p$ ) that *there is a barn in front of her*.

Intuitively,  $S$  does not know ( $p$ ) that *there is a barn in front of her*. Yet,  $p$  is true, and, arguably,  $S$ 's belief in  $p$  does not rest on any false beliefs she has. One might think that her belief or reasoning rests on her believing (falsely) ( $q$ ) that *there are no fake barns in the neighbourhood*; but,  $S$  has not even entertained the possibility of there being fake barns; so, the argument runs, it is incorrect to say she has beliefs about them: she neither believes nor disbelieves  $q$ . However, even if one accepts this, it seems correct to maintain that  $S$ 's belief rests on false *presumptions*, e.g. the presumption that there are no barn look-alikes around. Thus, (4) secures the desired result that  $S$  doesn't know that  $p$ .

Here is another example, adapted from Harman (1968), where false presumptions rather than false beliefs are involved:

DICTATOR (*why the demand for no-false-presumptions*)

The dictator of a country dies and his death is reported in a newspaper.  $S$  reads the report and comes to believe ( $p$ ) that *the dictator is dead*. However, all the leading newspapers, including this one, are soon pressurized by the military to dismiss the story as malicious and unfounded rumour.  $S$  does not read the denials and so continues to believe (truly) that  $p$ . Her compatriots, how-

ever, are persuaded by the later denials and believe otherwise —as *S* would have done had she read the denials.

There is some reluctance to say that *S* knows (*p*) that the dictator is dead, but, as in the BARN example, it might be argued that her true belief does not rest on any false beliefs she has. However, it may well rest on a false *presumption*, such as the presumption that the newspaper would not lie about such an important matter.<sup>3</sup> Condition (4) would in that case discount this as a case of knowledge. (Obviously, if *S*'s belief does *not* rest on this or any other false presumption, then she may still come out knowing *p*.)<sup>4</sup>

Some analyses rely on counterfactual conditions —conditions like our (5), for example— in order to handle cases like BARN and DICTATOR.<sup>5</sup> If the foregoing illustrations of condition (4) are sound, however, these problem-cases can be resolved without appeal to counterfactuals. So, the question arises, what *does* condition (5) achieve?

### 3. Fact-dependent reasons and the closure principle

There are two main achievements. Firstly, (5) secures the correct verdict in the following sort of example, where *S* believes *p* purely on statistical grounds (see Dretske 1971, p. 44 ff.):

LOTTERY (*why (5) is required*)

*S* has been given a national lottery ticket for next Saturday's draw. Today, *S* believes (*p*) that she will not win the lottery next Saturday; and she is right: she does not win the lottery on Saturday.<sup>6</sup>

What is relevant to us is that *S*'s belief is true, perfectly rational, and does not rest on any false beliefs or presumptions. If conditions (1)-(4) are taken to be jointly sufficient for knowledge, *S* comes out *knowing* that she will not win the lottery. This seems wrong. Condition (5) reverses the verdict, because *S* would have the same reasons for believing *p* even if *p* were false.

---

<sup>3</sup> We cannot take the presumption to be that the newspaper is not lying *in this particular case*, because this would be tantamount to presuming *p* itself; *S*'s belief in *p* would not then be *justified* or *supported* by such a presumption, and could not, therefore, be said to *rest* on it.

<sup>4</sup> I mention this because in my experience intuitions vary about whether *S* knows that the dictator is dead—some people remain resolutely convinced that *S* *does* know. The differing intuitions perhaps reflect that the example is under-described. Anyhow, I do not think too much weight should be placed on what verdict an analysis of knowledge delivers in this example.

<sup>5</sup> E.g. Nozick's (1984) "tracking" account requires this: *if p were not true, then S would not believe p*; *indefeasibility* accounts (e.g. Lehrer and Paxson 1963; Sturgeon 1993) require something like this: *there is no fact q such that if S were aware of q, then she would no longer be justified in believing p*.

<sup>6</sup> This example is particularly problematic for defeasibility accounts (see Audi 1998, pp. 217-18).

The second achievement is that in appropriate cases (5) defies the principle that knowledge is closed under known entailment:

(CP) If  $S$  knows that  $p$ , and knows that  $p$  entails  $q$ , then  $S$  knows that  $q$ .

Admittedly, this principle is *prima facie* plausible; indeed, one may wonder how extension of our knowledge by *inference* is to be accommodated otherwise; however, there appear to be everyday counterexamples. E.g.:

LAPTOP (*counterexample to closure principle; why (5) is required*)

$S$  knows ( $p$ ) that her laptop is at home. She also knows that  $p$  entails ( $q$ ) that the laptop has not been stolen in the last half hour. However, she surely does not *know*  $q$ . Knowing such a thing would appear to require more evidence than she has.

Condition (5) delivers the desired verdicts on  $p$  and  $q$ . In the normal run of things, the nearest not- $p$ -worlds (the nearest worlds in which  $S$ 's laptop is not in her home) will be worlds in which she brought it to work, or lent it to a friend, or sent it off for repair, etc.. In such worlds,  $S$  would not have the reasons she actually has for believing  $p$ —so (5) is met. By contrast, the nearest not- $q$ -worlds (worlds in which  $S$ 's laptop was stolen in the last half hour) will be worlds in which  $S$  has the same reasons for believing  $q$  as she has in the actual world—condition (5) is therefore *not* met. Thus,  $S$  can come out knowing  $p$  without knowing  $q$ .<sup>7</sup>

It might be objected here that once  $S$  entertains the possibility of  $q$  she would no longer lay claim to, or feel confident about, knowing  $p$ . I agree, but the present analysis does not take  $S$ 's views or claims as to whether she knows to be crucial to whether she does know—(though, of course, these views and claims may well play a part in fixing whether she meets the conditions of the analysis). If the claim is simply that  $S$  stops knowing  $p$  when she entertains the possibility of  $q$ —a view endorsed by Lewis (1996) for example—I disagree. Such a stance makes knowledge far too fragile. One would have thought that knowledge should in some sense be more *robust* than mere belief; yet, on this view, knowledge is irrevocably unstable, and epistemologists come out knowing far less about their environment than unquestioning people!

*In any case*, the LAPTOP example makes no mention of  $S$  'entertaining' the possibility of  $q$ , or, for that matter, of her 'entertaining' the truth of  $p$ , that her laptop is at home. One may take the knowledge ascriptions in LAPTOP purely from a third-person point of view. Even though  $S$  is not thinking about her laptop at present, it may be entirely appropriate to maintain that she does know it is at her home (i.e.  $p$ ).

---

<sup>7</sup> The closure principle is sometimes challenged by reference to radical sceptical hypotheses, such as e.g. (RASH) the hypothesis we are brains in vats on Alpha Centauri with all our sensations and experiences being fed into us by alien, research-active cognitive scientists (see e.g. Nozick 1984).  $S$  may know ( $p$ ) that her laptop is at home, and also know that this entails the denial of RASH; yet, supposedly, she does not know RASH to be false. I think less radical examples, such as LAPTOP, are more persuasive.

Likewise, even if  $S$  may never considered the possibility of her laptop being stolen, it may be entirely appropriate to also hold that she also knows that it would not be at home if it had been stolen within the last half hour. But, importantly, we can maintain these two things while also denying that she knows her laptop has not been stolen in the last half hour.

Two achievements of condition (5) have been mentioned: its verdicts on LOTTERY and LAPTOP; however, the same is achieved by a simpler condition on knowledge imposed by Nozick (1984):

(Noz) If  $p$  were not true, then  $S$  would not believe  $p$ . (I.e. the nearest not- $p$ -worlds are worlds in which  $S$  does not believe  $p$ ).

So, why not settle for (Noz)? Here's why not:

#### OUTING

$S$  has believed for sometime, purely on the basis of national statistics, ( $p$ ) *that someone in her office is gay*. Recently, her belief has been confirmed: her colleague, Mr. Make-No-Bones-About-It, has confided to  $S$  that he is gay.

Intuitively,  $S$  now knows that  $p$ . But even if  $p$  were false, that is, even if *no one* in  $S$ 's office were gay,  $S$  would still have believed  $p$ . So, (Noz) is not met. (5), however, *is* met: there is a main reason  $S$  has for believing  $p$  (namely, her colleague's confession) that she would not have if  $p$  were false. Thus, (5), but not (Noz), permits the correct verdict in the OUTING example.

One final point concerning the roles of conditions (4) and (5). It might be thought that so long as we have (5), condition (4) is unnecessary. What, after all, is wrong with the analysis that takes conditions (1), (2), (3) and (5) to be individually necessary and jointly sufficient for knowledge? The only verdict that is affected among the examples we have considered is the verdict on DICTATOR; but, as has been noted (fn. 4), intuitions diverge in this case. The following variation on the NOGOT example shows that (4) is still required:

#### DEVIIOUS NOGOT

Mr Nogot is mad. He decides that he will fool his gullible colleague  $S$  into believing ( $p$ ) *that someone in the office owns a Ford* IF AND ONLY IF someone in the office does own a Ford. (Let us suppose that Nogot has the means to find out if someone does.) Nogot finds out that Havit owns a Ford and embarks on his plan. The rest is history —  $S$  comes to believe truly ( $p$ ) *that someone in this office owns a Ford*.

Clearly, conditions (1)-(3) are met in this case. As for (5):  $S$  would *not* have the reasons she actually has for believing  $p$  if  $p$  were not true, so (5) is also met. Hence, without condition (4), we would end up with the counterintuitive verdict that  $S$  knows that  $p$ . (4) reverses this verdict because  $S$ 's belief in  $p$  does rest on a false belief, the same false

belief as in the original NOGOT example, namely, the belief *that Nogot owns a Ford*. So, condition (4) is not redundant.

That concludes the case for conditions (1)-(5). Now let us consider what the point of all this is.

#### 4. *How one can fail to know*

The issue is pertinent because there is an increasing trend to dismiss the task of providing an analysis of knowledge as wrong-headed. Williamson's (1995, 2000) has been a recent influential voice. He takes *knowing* to be a mental state—more accurately, a propositional attitude—in its own right, along with *believing*, *desiring*, *fearing*, etc.; and it is semantically un-analysable. In particular, knowing cannot be analysed as a special form of believing; i.e. knowledge is not simply belief with a special pedigree, belief that meets certain conditions. Clearly, this threatens the kind of account proposed here. A full discussion of the issue would take us too far afield; I will make a few remarks directed against Williamson.

Even though he takes knowing to be a mental state, Williamson is an *epistemic externalist* insofar as he maintains that one need not be in a position to know whether one knows.<sup>8</sup> He draws an analogy between *semantic externalism* and his position. According to the semantic externalist, one cannot have the thought *that gold is precious* without appropriate acquaintance with *gold*; in a world without gold, even a physical duplicate of me could not have thoughts about *gold*. So, *gold-beliefs depend* on the existence of *gold*, but they are mental states nevertheless. Likewise, *knowing* that *p depends* on the truth of *p*, but this does not preclude knowing from being a mental state: "If the content of a mental state can depend on the external world, so can the attitude to that content" (Williamson 2000, p. 6).

But there is a crucial disanalogy. The fact that I have beliefs about gold whereas my twin (my physical duplicate in a world without gold) does not has a *causal explanation*—at any rate, it is partially explained by the different *causal interactions*. Now, consider one of Williamson's own examples:

##### *The dead President example*

Consider... the situation of a generally well-informed citizen N.N. who has not yet heard the news from the theatre where Lincoln has just been assassinated. Since Lincoln is dead, he is no longer President, so N.N. no longer knows that Lincoln is President (knowing is factive). (Williamson 2000, p. 23).

According to Williamson, when N.N. stops knowing (L) that Lincoln is President, there is a change in her overall mental state. But, clearly, this change is not the result of any of N.N.'s causal interactions with the external world; intuitively, the change in L's truth value does not impact on her consciousness at all. It therefore seems odd to

---

<sup>8</sup> A characteristic feature of externalist accounts of knowledge is that the conditions they specify for *S*'s knowing that *p*, in addition to *S*'s truly believing *p*, are conditions that *S* need not be able to tell are met by introspection; consequently, one need not be in a position to know whether one knows. So, I do not dispute the externalist credentials of Williamson's theory.

hold that the epistemic change corresponds to a change in *mental* state as Williamson has it. So, I think Williamson's comparison with semantic externalism does not really support his position.

Moreover, the epistemic change in N.N. is not accompanied with changes in other propositional attitudes: her beliefs, desires and fears, for example, are *entirely unaffected*. By contrast, one would expect a change in belief, say, to leave some mark on other propositional attitudes —one's desires, fears, expectations or even knowledge. The comparative inertness of knowing in such cases provides another reason, I venture, to question the claim that knowing is a mental state.

Williamson is on safer territory when he claims: "On quite general grounds, one would not expect the concept *knows* to have a non-trivial analysis in somehow more *basic* terms (2000, p. 31). His point here is that previous attempts to analyse other concepts suggest that we should not expect to be able to analyse knowing in terms that can be understood independently of the concept of knowing. The analysis is likely to be circular, and hence not fully reductive. (This point holds even if the conditions the account takes to be necessary and sufficient for knowledge *are* necessary and sufficient.) So in what follows I will assume the material adequacy of the conditions.)

The point is well taken —and I fully concede that the account I have proposed may well invoke concepts that are inextricable from the concept of knowing. For example, the concept of *believing* may not be fully understood but by contrast with *knowing*; and we may not have an independent grasp of when one's belief *rests significantly* on other beliefs. So, I do not pretend to have provided a non-circular, reductive analysis of knowledge. But, I contend that such an account nevertheless *has a point*.

To begin with, even if the account is ultimately circular, it unpacks important connections between knowledge and other concepts, and it reveals certain *constraints* on knowing. This surely enhances our understanding of the phenomenon of knowledge. Most importantly, though, the account specifies (albeit indirectly) the ways in which one may *fail* to know. Condition (1) is in fact redundant; it is entailed by condition (5): if it is true that S wouldn't have the reasons she actually has if *p* were not true, then, clearly, *p* must actually be true. So, assuming the material adequacy of the account, there are three, and only three, different ways in which believing can fail to be knowing —roughly: the subject's rationality is impugned, her belief rests significantly on falsehood, or the reasons behind it do not depend on the truth of the believed fact. Thus, we have explanations of why the subject S does not know the propositions in question in examples such as LOTTERY, BARN, DEVIOUS NOGOT. These explanations are more direct and less obscure than Williamson's own: for example, he resorts to an un-analysable, unorthodox, externalist notion of *evidence* in order to explain the absence of knowledge in some cases (see e.g. 2000, ch. 9). And, presumably, given his take on the un-analysability of knowledge, he cannot rule out the possibility of cases where there is *no explanation whatsoever* for one's knowing or not knowing a certain fact.



So, I contend that even if knowing is semantically irreducible as Williamson maintains, the necessary and sufficient conditions sought (or proffered) by traditional accounts may have an *essential explanatory* role.

### 5. A closing suggestion regarding closure

In §3 we rejected the unrestricted validity of the closure principle:

(CP) If  $S$  knows that  $p$ , and that  $p$  entails  $q$ , then  $S$  knows that  $q$ .

But, a great deal of our knowledge *is* acquired by inference —e.g. mathematical and scientific knowledge. So, some qualified version of (CP) must be correct. I suggest the following:

(CP\*) If  $S$  knows that  $p$  and that  $p$  entails  $q$ , then, *so long as the corresponding beliefs of  $S$ 's do not rest on the belief or presumption that  $q$ ,  $S$  knows that  $q$ .*<sup>9</sup>

In the LAPTOP example,  $S$ 's belief in  $p$  (that her laptop is at home) *rests* on her presumption ( $q$ ) that it has not been stolen in the last half hour. So, (CP\*) does not, in this instance, lead to the conclusion that  $S$  knows  $q$ . But consider a different consequence of  $p$ : ( $r$ ) *that her laptop is at home or in Paris*;  $S$ 's belief in  $p$  does not rest on her belief (or presumption) that her laptop is at home or in Paris. So, (CP\*) sanctions the conclusion that  $S$  knows  $r$ .

The question arises whether the validity of (CP\*) follows from conditions (1)-(5) of the account proposed here. For, if it does not follow, then our inferential knowledge may not come out as knowledge by that account. We need not struggle with attempting to derive (CP\*) from (1)-(5). I propose we simply modify our account as follows:

A subject,  $S$ , *knows* that  $p$  if conditions (1)-(5) obtain, *or if* she comes out knowing  $p$  by (CP\*); otherwise, she does *not* know  $p$ .

It may be a bit trickier to accomplish, but one may still extract from this revised account the precise ways in which believing can fail to be knowing.<sup>10</sup>

---

<sup>9</sup> Davies (1998) proposes an analogous qualification for the transmission of *epistemic warrant*: “Epistemic warrant cannot be transferred from  $A$  to  $B$ , even given *a priori* known entailment from  $A$  to  $B$ , if the truth of  $B$  is a precondition for our warrant for  $A$  to count as a warrant.” (1998, p. 351).

<sup>10</sup> I have presented versions of this paper at the Universities of Sussex, Manchester and York, and at Birkbeck College, London; I thank the participants for the useful discussions. Thanks also to Michael Morris and Kathleen Stock for more detailed comments. Research supported by the research Fund of the School of Humanities (Sussex), and HUM2004-02330.

## REFERENCES

- Audi, R. (1998). *Epistemology*. London: Routledge.
- Clark, M. (1963). "Knowledge and grounds: A comment on Mr. Gettier's paper", *Analysis* 24, pp. 46-48.
- Davies, M. (1998). "Externalism, Architecturalism, and Epistemic Warrant", in C.J.G. Wright, B.C. Smith and C. Macdonald (eds.), *Knowing Our Own Minds*. Oxford: Clarendon Press.
- Dretske, F. (1971). "Conclusive Reasons", *Australasian Journal of Philosophy* 49, pp. 1-22. Reprinted in G. Pappas and M. Swain (eds.), *Essays on Knowledge and Justification*. Ithaca: Cornell University Press, 1978, pp. 41-60.
- Gettier, E. (1963). "Is justified true belief knowledge?", *Analysis* 23, pp. 121-23.
- Harman, G. (1968). "Knowledge, inference and explanation", *American Philosophical Quarterly* 5, pp. 164-73.
- Lehrer, K., and D. Paxson, Jr. (1969). "Knowledge: Undefeated Justified True Belief?", *Journal of Philosophy*, 66, pp. 225-37.
- Lewis, D. (1996). "Elusive Knowledge", *Australasian Journal of Philosophy* 74, pp. 548-67.
- Nozick, R. (1984). *Philosophical Investigations*. Oxford: Clarendon Press.
- Sturgeon, S. (1993). "The Gettier problem", *Analysis* 53, pp. 156-64.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.

**Murali RAMACHANDRAN** (BA Nottingham, MA Oklahoma) was a Mathematics teacher before becoming a Philosophy lecturer. He is presently Reader in Philosophy at the University of Sussex. His main areas of research are philosophical logic, metaphysics, and, more recently, epistemology. He has published papers on the theory of descriptions, counterfactual analyses of causation, counterpart theory, and contingent identity.

**ADDRESS:** Department of Philosophy, School of Humanities, University of Sussex. Falmer, Brighton BN1 9QN, UK. E-mail: muralir@sussex.ac.uk.