Philosophy and Phenomenological Research
Vol. LVII, No. 4, December 1997

Easy Possibilities

R. M. SAINSBURY
King's College London

1. Introduction

If you know, you couldn't easily have been wrong. Call this the Reliability Conditional. Its modality relates to the knowing, rather than to what is known. A true proposition which might easily have been false can be known, and a true proposition which could not easily have been false can be believed without being known. If a machine functioned reliably (on a given occasion), it couldn't easily have gone wrong (on that occasion). Reliability and easy possibility are closely related, so an account of knowledge in which reliability plays an important part is an account that connects closely with the notion of what could easily happen.

This note explores whether the Reliability Conditional is consistent with claims about knowledge that Williamson (1994a) has made in elaborating his epistemic theory of vagueness. Having argued that all intelligible expressions draw sharp boundaries, he suggests that we can explain ignorance of borderline cases for vague predicates in terms of "small differences of meaning". The essence of the explanation transposes to concepts. Suppose some borderline object is red and I believe this. The envisaged explanation of why this belief does not amount to knowledge is that I could easily have been wrong because I could easily have had a slightly different concept of *red*, one which ruled the object not red. I could easily have had a different concept because my practice (which is at least in part fixed by that of my conceptual community) could easily have been slightly different.

On the other hand, Williamson wants to hold that I do have knowledge of the meaning of my vague terms. Suppose that this knowledge can be properly expressed in this kind of way: "red" means *red*. Suppose also that I could easily have had a concept different from *red*, say *red**. Doesn't this mean that I could easily have been wrong in believing that "red" means *red*, for I could easily have been employing the maverick concept *red**, in which case I would have falsely believed that "red" means *red**?

An affirmative answer appears tempting; and, if correct, it would testify to an incoherence in Williamson's position. However, I shall argue that this appearance is delusory: the notion of what is "easily possible", as applied to

knowledge through the Reliability Conditional, turns out to comport well with various demands made by Williamson's form of epistemic theory, in particular with the combination of knowledge claims described in the previous two paragraphs.

2. Tracking

The Reliability Conditional is that if you know, you could not easily have been wrong. No equivalence is claimed: the absence of easy possibility is only claimed to be a necessary condition for knowledge, and I bracket all questions about sufficient conditions.

Is "you could not easily have been wrong" equivalent to "you tracked the truth"? Nozick characterizes tracking in terms of counterfactuals: you track the truth of p iff both (a) and (b) hold:

(a) if it were not the case that p, then you would not believe that p.

(b) if it were the case that p, then you would believe that p.[1]

That you could not easily have been wrong does not entail that you track the truth. You believe that you are not a brain in a vat. You could not easily have been wrong: to have been wrong, you would have had to be a brain in a vat, and it seems that this is not easily possible. So the Reliability Conditional seems not to preclude your current belief that you are not a brain in a vat from counting as knowledge. But, at least according to Nozick, this belief does not track the truth, for it does not satisfy the condition that if it were not the case you would not believe it. On the contrary, if it were not the case that you are not a brain in a vat, that is, if you were a brain in a vat, you would still continue to believe that you are not (setting aside doubts posed by semantic externalism).

Suppose I know that p and believe that I do not mistakenly believe that p. If I were to mistakenly believe that p, I would not believe that I mistakenly believed that p, so tracking would preclude my knowing that I do not mistakenly believe that p, even when it allows that I know that p. On the other hand, under the envisaged conditions, I could not easily be wrong in believing that I do not mistakenly believe that p; so the Reliability Conditional does not preclude my knowing this.[2]

If you might have wrong, it is because it might have been that your belief is false. In this context. "your belief" does not designate rigidly. One possibility is that your belief retains its content but is false. Perhaps (a) addresses this case. Another possibility is that you believe something else, and that

---

1 Compare Nozick (1981) p.172, 176. This is not his finished account of knowledge, which essentially involves reference to methods of belief formation.

2 Cf. Hughes (1996) p. 319 and Sosa (1996) p. 276.

other content is false. The relation between this case and tracking condition (b) is not entirely clear. The simplest thought is provided by familiar semantics for counterfactuals, which makes the truth of both antecedent and consequent sufficient for truth. Under this ruling, (a) and (b) can hold even when you could easily have been wrong. Suppose you come to believe a mathematical truth in some random way (for example, you simply guess at the sum of two large numbers, and by good fortune you are right). It is clear that you could easily have been wrong, that is, you could easily have guessed otherwise. Yet it appears that (a) is true in some trivial way, and (b) holds under the familiar semantics.

If one abandons the familiar semantics for counterfactuals, one may still detect some discrepancy between tracking and not being easily able to be wrong. Presumably (b) fails if, were it the case that *p*, you would simply fail to believe it, which is not a case of being wrong, but a case of being ignorant. So if we can make sense of your truly believing that *p*, yet it not being the case that were *p* true you would believe it, we can make sense of this in a situation in which the latter holds because of your failing to believe it rather than your believing anything false; arguably, this is consistent with your actual belief being one concerning which you could not easily have been wrong. For example: I believe that Mary is married because I notice her wedding ring. She is married, but she hardly ever wears the ring, and indeed didn't intend to wear it on this occasion (she left it on by accident). I have no curiosity about her marital status, so if I had not noticed the ring, I would have formed no belief about whether or not she was married. Arguably, this falsifies (b),[3] so that I don't track the truth. Yet, arguably, I could not easily have been wrong (given the uniformly prevailing convention in my culture that only married persons wear such a ring). The Reliability Conditional does better than tracking: I do know that I am not a brain in a vat and that Mary is married and that I am not mistakenly believing that *p*; and I do not know the guessed mathematical truth.

The problem raised by the case of Mary is widespread and general.[4] It can be just an accident, a piece of good or bad fortune, that I come to know that *p*. I happen to know that Mary is here, because I happened to turn my head (despite being deeply engrossed) just as she came in; so I saw her, but I might easily not have done. Much of our knowledge is in this way fortuitous. Further, I can know that *p* even if it is just an accident that *p* is true, and much of our knowledge is of accidental truths. Both kinds of happen-

---

3 It certainly does so if a necessary condition for the truth of a counterfactual is: "there be no sufficiently close possible world in which the counterfactual has a true antecedent and a false consequent" (Plantinga 1997, p. 144). Given Mary's disinclination to wear her ring, a world in which she is married but I don't believe it (because my only clue is absent) is close.

4 Thanks to Ernest Sosa for helping me get clearer about this.

stance are consistent with knowledge; the tracking condition, but not the Reliability Conditional, is apt to rule them out. So Nozick was clearly right to think that one could not give an account of knowledge just in terms of tracking. Introducing the further notion of *methods* has to be done with care. The method by which I come to know that Mary is present is perceptual and it leads to knowledge. The method by which someone whose environment contains many fake barns comes to believe truly that a (real) barn is present is also perceptual, yet it does not lead to knowledge. We must say they use different methods, one reliable the other not. Perhaps the Mary-spotter uses the method: infer from Mary-appearances to the presence of Mary. This is a reliable method in the Mary-believer's circumstances. Perhaps the barn-spotter uses the method: infer from barn-appearances to the presence of a barn. This is an unreliable method in the barn-believer's circumstances.

The Reliability Conditional takes the cases in its stride, precluding knowledge of the presence of the barn (because I could easily have been wrong), not precluding knowledge of the presence of Mary (because I could not easily have been wrong). Explaining why it takes them in its stride may well involve reference to methods. But methods might be better thought of as an explanation of the truth of the Reliability Conditional than as a corrective to tracking.

3. KK

Arguably, there are cases in which it could easily have been that it could easily have been that the machine went wrong, even though it could not easily have been that the machine went wrong. This would suggest that the operator "it could not easily have been that it is not the case that" does not iterate; so the modality is not one apt to support the KK principle (if you know, you know that you know). This machine has just functioned reliably: it worked correctly and couldn't easily have gone wrong. But when the engineer was making it, he had a choice between installing one of two kinds of switch, and he had no good reason to choose one kind rather than another. As good fortune had it, he chose the one that performed well under the machine's operating conditions. But he could easily have chosen the other one (his choice was arbitrary—perhaps he flipped a coin), and the other switch would have made the machine unreliable. As things are, it is not easily possible for the machine to have gone wrong; uncontroversially (if the story is accepted) it could easily have been that it was easily possible for it to go wrong. Arguably, the story also sustains the counterexample to iteration: on that occasion it was not easily possible for the machine to have gone wrong, but

on that occasion it was easily possible for it to have been easily possible for the machine to have gone wrong.[5]

Reliability has a feature structurally analogous to a "margin for error" principle. A reliable machine will still produce the appropriate output even if the operating conditions had been slightly different (in likely ways). Suppose x and x' are two similar inputs, with x closer than x' to what is a paradigmatically "normal" operating condition. Suppose that the machine responds reliably to x. Then it produces the correct or designed output and could not easily have gone wrong. So it would have produced the correct output in response to x'. But does it follow that it would *reliably* have produced the correct output in response to x'? We had best not say that this follows (though of course it may also be true in many cases), for then the machine's reliability would extend endlessly outwards from the paradigm cases of normal operation in a sorites-like manner: we would be forced to say that it is reliable under all conditions, since any can be reached by a series of small steps from the paradigm ones in which it is, by hypothesis, reliable. When a machine operates well at the outer limits of its range, it may manifest its reliability, that is, its reliable operation in more central cases, without the near-the-limit operations being actual instances of reliable operation.

One moral is that reliability is a relativized concept: good operation is not required under every possible circumstance, but only under some contextually specified range of circumstances. Indeed, it is hard to make sense of the idea of a machine which can operate reliably under *all possible* circumstances, since these include ones with, for example, very different laws of nature, and ones with all sorts of freak demonic interference.

4. Hidden snags

Suppose a mechanism, M_1, is reliable within the values 4–7 of some parameter. Idealizing, this may mean that it delivers the right output when the values are 3–8; but is not reliable for the values 3 and 8 themselves, since it produces a wrong output for values 2 and below and for values 9 and above. Suppose now that M_1 is coupled to M_2, which reliably detects the values of the parameter in question, and allows M_1 to operate only if they fall within the range for which M_1 produces the right output (i.e. 3–8). The coupled mechanism, $M_1 + M_2$, is reliable even for values 3 and 8, for if a value marginally different from these (in the unfavourable direction) had been realized, there would have been, not an incorrect output, but no output (or, if you prefer, an output correctly saying "M_1 cannot reliably handle this situation"). Coupling one mechanism with another which is sensitive to the first's limitations can extend its reliability.

---

[5] As Christopher Hughes pointed out, the interaction between tense and modality leads to many complications. The alleged counterexample cannot be regarded as definitive.

The reliability of human cognitive mechanisms can likewise be extended by an appreciation of their limitations. We are often in states of the following kind: if a way of arriving at a belief had been apt to produce error in the actual circumstances of a specific use, we would have known this (at least tacitly: the belief-forming mechanism in question would be inhibited). Features of our situation (following Sosa, we may call them "hidden snags") which make us not in this state may defeat knowledge, by shrinking (sometimes devastatingly) the range of our reliability. This explains why many people feel that the barn-spotter does not know that there is a barn before him. The local fake barns are hidden snags, depriving a normally reliable way of telling whether there's a barn before one (cursory visual inspection) of its reliability (while leaving in place the reliability of other more close-up and searching ways of spotting barns).

Those who say that hidden snags need to be actual and not merely possible if they are to defeat knowledge are certainly right within a reliabilist perspective. Although the perspective suggests that, typically, knowledge requires absence of hidden snags, it does not suggest that it requires knowledge of the absence of hidden snags.

5. Determinism

"Could easily" does not need to invoke indeterministic notions. A determinist can properly affirm: "You should not have driven so fast—you could easily have crashed". The claim could be defended in terms of small differences in relevant parameters, for example: "You would have crashed if the road had been a little more slippery". There is no need to suppose that the slipperiness of the road is other than completely determined by previous circumstances, or that there is a causal possibility that the actual slipperiness could (with no changes elsewhere) have produced the crash. The point, rather, seems to be that if the road had been a little more slippery, you wouldn't have known. You were rash because you engaged in a course of action which, for all you could know in advance, would have led to disaster.

This example may suggest that "could easily" is consistent with determinism only when the possibility itself issues from ignorance. The impression might be reinforced by the following contrast. The bridge will break under a load of 80 tons but not under a load of 79 tons. I drove my 79-ton truck across and got safely to the other side. Safely? Case 1: my weigh bridge registered 79 tons but is accurate only plus or minus 5%. Then it was rash of me to cross because my truck could easily have weighed enough to break the bridge, so the bridge could easily have broken. Case 2: my weigh bridge was dead accurate, and I really knew all the facts. Then my crossing was not rash: the bridge could not easily have broken. The second case involves the same small differences as the first, but in this case they are discriminated by the subject. This may suggest that easy possibilities are to be explained in terms

of small differences which are significant for the outcome but which could not easily be known.

This impression is just an artifact of the examples so far. If we turn to mere mechanisms, we may reasonably expect a more general feature from which the appeal to knowledge emerges in special (e.g. action-involving) cases. Moreover, the relevant feature is consistent with determinism. A machine is reliable if it is so structured as to produce a good frequency of successful outcomes under a reasonably wide range of conditions (a range determined contextually). In particular, a small variation in input conditions should not normally be enough to tip it into malfunction. The action-involving cases in which ignorance is crucial are just special cases of this. You were lucky to get round the corner without a crash: if you go on behaving like that you'll kill yourself. We're cashiering you for crossing the bridge with that load: a rash fellow like you will undoubtedly cause a disaster sooner or later. The system consisting of you and your vehicle is not reliable, because there are closely similar circumstances under which the desired outcome (getting round the corner without crashing, getting across the bridge) will not be held causally in place by the system's structure. A determinist need find nothing amiss, since we do not have to envisage alleged possibilities in which the inputs are just the same and the outputs different.

6. Worlds

It is easily possible for me to be wrong in believing that p (even if it is true that p) iff at some world "close" to the actual world the actual episode of forming the belief that p (or the counterpart of this episode) is one in which a false belief is formed. In some cases, this is because the same belief is formed, but is false at the close alternative world. In other cases, it is because a different, and false, belief is formed at the close world.

Our notion of a "close" world is not an input to our understanding of knowledge, reliability or easy possibility. Rather, we must fashion the notion of closeness to do proper justice to these notions. For example, no intuitions about closeness should be allowed to rule either that the demon's world is close or that it is not: argument is needed, and the concept of closeness must be so developed as to reflect their outcome.

The initial idea is that worlds are close when there is only a small difference in the value of some parameter. It is because 79 is close to 80 that it was dangerous to cross the bridge with a 79-ton truck. But not just any parameter is relevant. No doubt two worlds that are now completely different may have been close at the time of their Big Bang, for a small difference in parameter values then can lead to endlessly large differences later. When assessing the reliability of machinery, we have already used the notion of intended operating conditions and intended function. A normal computer is not unreliable because it will not operate under 50 meters of water, but a

Rolex which failed to operate under these conditions would have to be returned as defective. We need some notion like that of intended function to specify what counts as correct operation. A screwdriver is not an unreliable screwdriver just because it breaks when used as a lever.

A reliabilist conception of knowledge requires it to be reliably delivered by a mechanism whose job it is to produce true belief.[6] It will accordingly be enough for a difference to count as small that it would not produce any significant difference in the mechanism's operation. Some differences in an aspect of the agent's situation which he would not have noticed count as not having affected his belief-forming mechanism, and so count as small, even if relevant to the truth or falsehood of the belief he actually formed. To use one of Williamson's examples: I believe truly that there are just so many people in a crowded stadium, but I could easily have been wrong because there could easily have been a person more or less: that is, a person more or less would not have affected the mechanisms which issued in that belief, and the same belief would have ensued. So a world in which there is a person more or less and I form the same belief is close. We must be careful to avoid rash generalization, or we will find ourselves counting the difference between our world and a demon world as small, on the grounds that the subject would not have noticed the difference. The talk of "noticing" is misleading, as it suggests a comparison by the subject, whereas what should properly be at issue is whether there would be a significant difference in the subject's overall behaviour. In difficult cases, it will be impossible to assess this other than from the perspective of a theory. (Would my "behaviour" have been "significantly different" in a demon world?) A separate issue is that we may apply different standards in different contexts: a reliable Skoda may break down more often than a reliable Honda, and the demands we place on whether a witness "really knows" may be related to the severity of the crime.

Even given a suitable notion of closeness, we cannot gloss "could not easily have been otherwise" thus: for every world in which it is otherwise, there is a closer one in which it is not. Suppose machines of a certain kind are reliable only if they will work within the range 0–10 of some parameter. This machine has just functioned correctly, with the parameter value 6. However, it will go wrong at values 8 or higher. That is, it could easily go wrong (within its intended operating range) so it is not reliable. Yet for every world in which it goes wrong (with the parameter values 8 or higher) there is a more similar one (because closer to 6) at which it goes right.

---

6 This is to be distinguished from "requires it to be delivered by a reliable mechanism". A reliable mechanism may not operate reliably at the outer limit of its normal range.

7. Vagueness

That I do not know of a borderline case that it is red, even if I truly believe it, is overdetermined: I could easily have formed a false belief about a different object, for example if an indistinguishable or barely distinguishable non-red shade had occupied my attention at that moment; I could easily have come to the opposite, and so false, view about this very object; and, finally, the possibility Williamson stresses: I could easily have had a slightly different concept, one which, drawing the borderline in a slightly different place, rules the object not to fall under the concept.

This third possibility requires a little elaboration if we are to grasp its knowledge-defeating character. Suppose I stipulate that, in a certain examination, students are to count as "excellent" just on condition that they attain a mark of 72 or more. I could easily have made a different stipulation, for example, that students are to count as "excellent" just on condition that they attain a mark of 73 or more. Does it follow that, when I apply the epithet "excellent" to a student who has scored 72, I do not know that the student counts as excellent, because I might easily have used a concept which, had it played the relevant part in my belief, would have ensured that the belief was false? Obviously not. Knowledge is defeated if the shift of concept functions as a hidden snag, making a way of arriving at a belief we have actually used unreliable without our knowing this (or without inhibiting belief-formation by that method). If we had made a different stipulation governing the local use of "excellent", setting the minimum at 73, we would know this, and this knowledge would have inhibited the passage from 72 to excellent; so our actual method of forming the judgement, by noting that the student scored 72 or more, is reliable. By contrast, it is possible for the extension of *red* in the thoughts of ourselves and our community to differ from its actual extension in a way I do not detect. This possibility renders unreliable any actual method of reaching the belief that something close to the borderline is red. The first difference changes my belief-forming behaviour, whereas the second does not.

The possibility in question does not attach to every concept. Consider one which is sharp (or let us so suppose), like being a foot long. It would not be easily possible for me to use, in place of *foot*, the concept *metric foot*, without this bringing with it many related differences, despite the fact that there is only a small difference between a foot and a metric foot.[7] If I had the *metric foot* concept in place of that of *foot*, my measuring practices would be different (for example, I would happily measure with a centimetre rule, whereas I actually find the conversion from centimetres to feet too hard); I would associate the price of timber with different numbers, depending as I took the price to be per foot or per metric foot; in many circumstances, judgements I would actually formulate with the words "that is a foot long" would not be made;

7 One metric foot is 300mm; a foot is about 305mm.

and so on. So an episode of coming to believe truly that something is four feet long could not easily have been an episode of coming to believe falsely that it is four metric feet long. Knowledge using non-vague concepts is not impaired. Likewise, knowledge applying vague concepts to clear cases is not impaired. A difference in my concept *red* so great that a clear case of red no longer falls under the different concept is a difference that could not easily occur, since its occurrence would bring with it large scale changes in my behaviour.

Must an epistemic theorist insist on the third source of ignorance? I presume that the motivation is that not all true belief involving vagueness need be perceptual. Suppose that 34 grains (of a given kind) are the fewest that can make a heap. As Williamson stresses, if this is true it is necessarily true. Perhaps someone comes to believe it not on any perceptual basis, but in a way that strikes him as the same as the way in which he comes to know truths of reason. How does the epistemic theorist show that, despite truth and this apparently favourable aetiology, the belief is not knowledge? Appeal to less than perfect perceptual powers seems irrelevant. Appeal to a potential conceptual instability will cover any case, perceptual or not, of purported knowledge in which a vague concept is applied to a borderline case.[8]

Will it cover too many cases, including knowledge of meaning? Earlier I envisaged the claim that one does not know that "red" means *red* because one could easily have had a different concept, *red**, in which case one would have judged falsely that "red" means *red**. There seems also to be another instability: the word "red" could easily have meant something different, e.g. *red**; if it had, then my judgement that "red" means *red* would have been false since "red" would mean *red**. Why do these possibilities not show that I do not know the meaning of my vague expressions?

One way to answer depends on an externalist idea. I am party to a conceptual and linguistic practice; there are some possible differences in the practice which simply entail that my concept or language is different, regardless of how similar things may seem to me in the different situations. We cannot say in any detail what such differences are like (if we could, we would understand in detail how meaning supervenes on use, which we do not); but they are easily possible in that their obtaining would involve only minute behavioural shifts in the community, not ones with any significant impact on belief-forming mechanisms. The easy possibility is thus a social shift which drags my concept with it. What could not so easily happen is that my concept would get out of line with the community's. The difficulty is not causal, but issues from the externalist view: what counts as the precise exten-

[8] Williamson comments that he was concerned that vagueness should be a distinctive source of ignorance; thus distinct, for example, from inescapable ignorance due to limitations on accuracy of measurement. Cf. Williamson (1994b).

sion of a subject's vague concepts is fixed on a community wide basis, so, in close worlds, idiosyncrasies in his usage will not count as manifestation of an idiosyncratic concept. In close worlds, a subject's concepts co-incide with those of his fellows, whether theirs are the same as or different from their (and his) actual ones.

If my concept cannot easily get out of line, then I cannot easily come to believe that "red" means *red** when this is false. The only easy way to have the concept *red** it is to have it because the others do; but if that's the concept we all have, "red" does mean *red**. Running the point in the other direction: if, as could easily have been, "red" meant *red**, then my actual judging that "red" means *red* would have been a judging that "red" means *red**, and so would have been true.[9]

Externalism is not the only way to attain the structure of this answer (though it seems to be the way Williamson would prefer). Perhaps I have a mechanism which reliably aligns my concepts with the ambient ones. Its reliability ensures that neither of the two possibilities which seemed to threaten knowledge of meaning would be easy; so the threat would peter out. The truth, I presume, requires a combination of both kinds of factor. There will be some externalist, thus constitutive, determinations; but these will be possible only if some causal mechanism reliably keeps me in touch with the concepts, language and topics of discourse of my fellows, for this mechanism will be relevant to *which* other speakers and objects help constitutively to determine my concepts and meanings.

9. Objections

1) I know that I do not have arthritis in my tibia, but an objection to the proposed account is that it entails that I do not. The inflammation I am suffering is similar in kind to that associated with arthritis, but the tibia is not a joint. However, I could easily have been wrong because it could easily have been that my community had no use for the concept of arthritis, and used instead the concept *twarthritis* which applies to any arthritis-like inflammation. If this had been so, my belief-forming episode would have issued in the falsehood that I do not have twarthritis in my tibia. But I know that I do not have arthritis in my tibia, so the account of knowledge is incorrect.

Being conceptually normal, I actually share my fellows' concepts and would do so if they had been different: I am actually an *arthritis* user and potentially a *twarthritis* user. In the latter case, it cannot be assumed that the (counterpart of my actual) relevant belief-forming episode would have delivered the false belief that I do not have twarthritis in my thigh. Actually, my belief that I do not have arthritis in my thigh comes from my doctor's expla-

[9] This retracts a suspicion voiced in my review of Williamson: Sainsbury (1995) p. 595–97.

nation that arthritis does not affect joints. There can be no such explanation in the twarthritis world, since, in that world, there is no concept *arthritis* to enter into such an explanation. If my twarthritis belief comes from my doctor, it is more likely to be the true belief that I have it in my thigh than the false belief that I do not. Likewise for other aetiologies of my actual belief.

2) I know that there is a glass of water in front of me, but an objection to the envisaged strategy is that it entails that I do not. I could easily have been on twin earth, in which case it would have been twater rather than water in the glass and had I formed my belief in the way I actually did it would have been false. So although I believe truly that there is a glass of water in front of me, I could easily have been wrong, and so I do not know.

To firm up the example, let us suppose that twin earth and earth are mutually inaccessible, so that there is no question of an inhabitant of one finding himself on, or traveling to, the other. This does not settle whether I could or could not easily have been on twin earth, for that question, as the notion has been developed here, is not about ease of access but about the extent of hypothetical changes in my behaviour.

Suppose it is allowed that I could easily have been on twin earth, that is, that my belief-forming mechanism would count as operating in essentially the same way on twin-earth as on earth: in both cases, the same kinds of inputs yield the same kinds of outputs. This involves thinking of inputs as the same, and so not thinking of earth inputs as responses to water and twin earth inputs as responses to twater. Perhaps the inputs should be thought of as water-appearances, to which I would be subject on either planet. If the output is thought of as a belief whose content subsumes both water and twater, then I would not have formed a false belief on twin earth, so this is not an example of an easy possibility of error. The combination of views required for the counterexample sees the outputs as the same through both being judgements that there is a glass of water in front of me.

As far as I can see, the counterexample, once it has been allowed to develop to this point, can be blocked only by denying that the twin output would have this same content; in other words, only by some form of semantic externalism. The externalist will say that the supposition that my belief-forming episode occurred on twin earth amounts (given inaccessibility) to the supposition that I was always there, and so had the concept *twater* where I actually have the concept *water*; so that episode (or its twin counterpart) would have issued in the true belief that there is a glass of twater before me.

If this is right, then something said in §2 needs to be reconsidered. As the notion of easy possibility has been developed here, my being a brain in a vat cannot be ruled not to be easily possible on the grounds that it would require difficult technology to arrange. Ease is determined by how great the behavioural changes would be. Externalistically considered, there would indeed be great differences between the vat world and the actual world: depend-

ing on precise details of the set-up, I would likely have access to few if any of the contents I actually have, and little or none of my behaviour could be regarded as seeing or responding to barns, or whatever. But if this is the ground for saying that a vat world is not easily possible, and so is no threat to knowledge as constrained by the Reliability Conditional, then the vat world does not pose the envisaged problem for Nozickian tracking, since it verifies the counterfactual that if I were a brain in a vat I would not believe I was not one. This is because, if I were a brain in a vat, I would not, given externalism, have the conceptual resources required in order to believe anything about brains or vats.

The upshot is that if the easiness of a possibility has been rightly characterized here, then, given standard views about what is known, the Reliability Conditional leads to some form of externalism. This destination is not, of course, one that would be unwelcome to Williamson.[10]

REFERENCES

Plantinga, Alvin (1997). "Warrant and accidentally true belief." *Analysis* 57.2, 140–45.

Sainsbury, R.M. (1995). "Vagueness, ignorance and margin for error." *British Journal for Philosophy of Science* 46, 589–601.

Sosa, Ernest (1996). "Postscript to 'Proper Functioning and Virtue Epistemology'." In Kvanvig, Jonathan L. (ed.) *Warrant in Contemporary Epistemology: Essays in Honor of Plantinga's Theory of Knowledge*, 271–80, Lanham MD and London: Rowman and Littlefield.

Williamson, Timothy (1994a). *Vagueness*. London: Routledge.

Williamson, Timothy (1994b). "Definiteness and Knowability." *The Southern Journal of Philosophy* 33 (Supp.), 171–91.

[10] Thanks to Timothy Williamson for commenting on the review mentioned in footnote 9, and to Keith Hossack, Christopher Hughes and Ernest Sosa for helpful discussion of the present paper.