GETTIER AND JUSTIFICATION RELIABILISM

1. Introduction

In this paper I will be arguing against the form of justification externalism (henceforth J-externalism) known as justification reliabilism (henceforth J-reliabilism). For organisational purposes and manageability, I will be distinguishing between local and global J-reliabilism and responding to each separately. Against local J-reliabilism, I will be employing a perceptual Gettier case to show that local reliability is insufficient for knowledge. Against global J-reliabilism, I construct a case in which a subject putatively has knowledge although only locally reliable in order to show that global reliability is not necessary for knowledge. The main contention of my paper is that local J-reliabilism is not sufficient for knowledge, while global J-reliabilism is not necessary for knowledge. In addition to my primary goal, I engage in a critical examination of Nozick-style reliabilism in order to cast doubt on the standard anti-luck diagnosis of Gettier counterexamples. I argue that since a certain kind of luck is involved in some putative cases of knowledge, a more sophisticated anti-luck type analysis may be called for.

2. Gettier and the Tripartite Account

If there were anything like a classical or traditional definition of knowledge, it would be something along the following lines:

JTB For any subject *S*, *S* knows that *p iff*[1]:

(i) *p* is true

(ii) *S* believes that *p*

(iii) *S*'s belief that *p* is justified

[1] The '*iff*' is to be understood as indicating necessary coextenstion, or equivalence.

JTB takes (i), (ii), and (iii) to be individually necessary and jointly sufficient for *S*'s knowing that *p*. That knowledge is composed of these three features was an assumption taken for granted by philosophers dating at least as far back as Plato.[2] However, this fact alone should not prompt us to dismiss JTB since it also appears to receive the endorsement of common sense. First, it seems counterintuitive to talk about knowing that *p* if *p* is false. Second, one cannot know that *p* if *p* is not among one's beliefs. Third, knowing that *p* requires more than simply believing that *p*. It also requires one to have some *reason* for, *evidence* in favour of, or *justification* for believing that *p*. (I will treat all these notions as synonymous.) Significantly, the third requirement of JTB has traditionally been taken to express a J-internalist intuition—namely, that one's reasons should be constituted by factors that are subjectively available to one. Each of these theses appears to be in harmony with our everyday usage and understanding of the word 'know'. Moreover, one would be hard pressed to find some additional requirement necessary for our quotidian understanding of knowledge.

Unfortunately, Gettier-type counterexamples have raised serious questions about the adequacy of JTB.[3] For example, consider the following case of veridical hallucination. Suppose *S* has strong perceptual evidence for, and comes to believe, the empirical proposition:

(a) there is a vase on the table.

S's evidence for (a) is that she is having a visual experience of there being a vase on the table. But unknown to *S*, her visual experience of there being a vase on the table is actually a drug-induced hallucination. Now, it so happens that there is in fact a vase on the table, though the real vase is being obscured from *S*'s visual field by a stack of books. The following all seem to be true of the aforementioned case:

---

[2] See: Plato [1961]: 97e-98a.
[3] See Edmund Gettier [1963].

(i) (a) is true

(ii) *S* believes that (a)

(iii) *S*'s belief that (a) is justified

Thus, all three conditions of JTB have been met. However, intuitively we would be reluctant to say that the subject in the above example knows that (a). (I will refer to the Gettier case of veridical hallucination just described with the abbreviation *GTTR*$_{VH}$.)

3. Gettier and Justification Reliabilism

J-externalists blame the failure of JTB on the J-internalist notion of justification defined purely in terms of subjectively available reasons for one's beliefs. According to this line of thought, subjectively available reasons are denizens of the 'inner' world (or the mind) and are therefore independent of, and insulated from, the 'external' world. The upshot of this picture, as Williamson [2000] notes, is that the mind and the external world are characterised as two independent variables [p. 5]. Belief is simply a function of the 'inner world' (or mind) while the truth of a given empirical proposition remains a function of the 'external world'. However, J-internalism fails to provide an account of how there could be a non-accidental connection between the two. On the J-internalist picture, any similarity between the content of the inner world and that of the external world—between belief and truth—seems to amount to happenstance or luck, or so the argument goes.

The key to escaping the unfavourable implications of GTTR$_{VH}$, according to J-externalists, is to define knowledge in such a way that a subject's perceptual beliefs are connected, not to internally fixed reasons but to the external world itself. The most widely discussed (if not accepted) version of J-externalism is J-reliabilism, which asserts that for a belief that *p* to be justified it must reliably indicate the

fact that *p*, or be reliably connected to the fact that *p*, or be produced by a reliable method or process. For simplicity, I will be limiting myself to the process version of J-reliabilism:

> *Process J-Reliabilism:*
>
> For any subject *S*, *S*'s belief that *p* is justified *iff* it was formed via a reliable process (i.e., a process that tends to produce true beliefs).[4]

There are two dimensions to the notion of reliability. On the one hand, we may say that some process {M} is reliable so long as it tends to produce true beliefs with respect to some particular proposition, *p*. This would be a description of what McGinn and Goldman refer to as *local* reliability. Local reliability contrasts with what has come to be called *global* reliability. On this view, a process, {M}, is globally reliable when it tends to produce true beliefs for a range of *p*-like propositions.[5] Thus, global reliability requires that {M} be reliable with respect to not only the particular proposition *p* but also with respect to a larger range of propositions. As Goldman explains it, the notion of global reliability is a 'statistical or dispositional property of a belief-forming process or type' [p. 49]. A type of belief-forming process is globally reliable *iff* it tends to produce a high ratio of true beliefs for a range of uses of that process. For the time being I will only be concerned with local reliability. (I will return to the question of global reliability in §2.4.) We may reword the above definition of process J-reliabilism in order to capture the local nature of the reliability in question:

---

[4] See Goldman [1979] for the *locus classicus* of this view. J-reliabilism represents an externalist version of the JTB account of knowledge, and should be distinguished from TB+G accounts. According to TB+G accounts, knowledge amounts to true belief plus some codicillary anti-Gettier condition. Such accounts omit all talk of evidence or justification and instead conceives of knowledge as true belief that satisfies a causal or reliability condition. On this view reliability is what you add to true belief *instead* of evidence or justification, in order to arrive at knowledge. Examples of this version of reliabilism include Goldman's [1967] 'causal condition' approach and Robert Nozick's [1981] truth tracking account. TB+G accounts are versions of K-externalism, since they are concerned with knowledge rather than justification. J-reliabilism, by contrast, views reliability not as replacing the role of justification in an account of knowledge but rather as that which *makes* a belief justified.

[5] See Goldman [1986] p. 45.

Local Process J-Reliabilism:

For any agent $\mathcal{S}$, $\mathcal{S}$'s belief that p is justified *iff* it was formed via a process that is locally reliable. (i.e., a process that tends to produce true beliefs with respect to a particular proposition, p)

Let us assume, for the sake of argument, that knowledge is true belief acquired by a process, {M}, which is locally reliable. It remains an open question as to exactly what {M} has to be reliable for. Here are a few possibilities[6]:

(R1) For producing a true belief with regards to a particular proposition p in precisely the situation in question.

(R2) For producing a true belief with regards to a particular proposition p in situations just like the one in question.

(R3) For producing a true belief with regards to a particular proposition p in all situations worth considering.

(R4) For producing a true belief with regards to a particular proposition p in all possible situations.

Notice that (R1)-(R4) all relate to a particular proposition p, and are therefore all versions of local reliability.[7] However, each subsequent statement increases the range of possible situations in which a

[6] See Craig [1990], p. 54 for a similar list.

[7] This should not be confused with the increasing strength of the reliability requirement embodied in the shift from local reliability (i.e., regarding a particular proposition p) to global reliability (i.e., regarding all relevantly p-like propositions). A

given process, {M}, must work in order for the subject to count as knowing the particular proposition *p*. As definitions of what it means for a process to be reliable, (R1) and (R2) strike me as trivial. Both, more or less, amount to the claim that {M} is reliable when {M} is reliable (though whether we regard (R2) as trivial depends on how narrowly we understand the 'just like' in 'situations *just like* the one in question').[8] The most widely discussed version of reliability is (R3). The locution, 'situations worth considering' that features in (R3) raises important questions that bear on the debate regarding relevant alternatives and how we should go about settling which situations are close enough to the one in question. For example, are we to think of relevance as designating a similarity class (i.e., situations that resemble each other in certain fundamental respects) or as an interest claim (i.e., situations that appeal to us because of the context of discussion or particular goal we have in mind).[9] These are questions that must be settled if we are to arrive at a workable understanding of local J-reliabilism. However, they do no bear directly on our present discussion and so we may table them for now.

The application of local J-reliabilism to the Gettier case of veridical hallucination described earlier is quite straightforward. Hallucinating a vase is not a reliable method for arriving at true beliefs about vases. In the case of the subject who comes to believe (a), *there is a vase on the table*, based on a hallucination of a vase, the belief that (a), though true, is not reliably formed. Thus, according to local J-reliabilism, the GTTR$_{VH}$ subject does not have a justified belief. However, below I will attempt to construct a Gettier case against local J-reliabilism. Since (R3) is the most widely accepted version of local J-reliabilism, my proposed Gettier case will take it as its target. I will then consider a strengthened version of local J-reliabilism, one that appeals to Nozick-style counterfactual conditionals.

---

similar list can be constructed on behalf of global reliability by adjusting the notion of relevant likeness to include a progressively wider range of relevant *p*-like propositions.

[8] This is not to suggest that (R1) and (R2) are trivial vis-à-vis a definition of justification. In fact, as a definition of justification (R1) is no less demanding than (R4) since it implies one is only justified if one's belief that *p*, in a particular instance, is factive.

[9] Which situations count as 'similar' also raises issues relating to what features of a given situation should count as salient vis-à-vis a subject's justification.

Consider the following perceptual Gettier case. Bob enters a classroom and sees what he takes to be a vase in the middle of the floor. Based on this perceptual experience, Bob forms the following belief:

(a*) There is a vase in the room

Bob believes that *there is a vase in the room* because he thinks that he sees that *there is a vase in the room*. However, what Bob takes to be a vase is actually a computer generated hologram. Now, there is in fact a real vase in the room, but the actual vase is hidden behind an opaque black curtain at the back of the room. Furthermore, the computer that generates the hologram is equipped with special sensors that detect the presence of the vase behind the curtain. This then prompts the computer to project the hologram of the vase in the centre of the room. In short, the hologram in the centre of the room is causally connected to the vase hidden behind the curtain at the back of the classroom. Moreover, the computer's programming is perfectly reliable, so that once there is a vase in the room, it will project the hologram, and it only projects the hologram when there is a vase in the room. In sum, as far as the computer program is concerned, the presence of the actual vase in the back of the room is a necessary and sufficient condition for the hologram of the vase in the centre of the room. Unfortunately (or perhaps fortunately), Bob lacks any of this background information, and forms his belief that there is a vase in the room base solely on the hologram in the centre of the room. All of the following seem true in the above case:

(i*) (a*) is true

(ii*) Bob believes (a*) is true

(iii*) Bob's belief that (a*) is formed via a reliable process

Ex hypothesi, (iii*) is true since the computer is programmed to only project the hologram of a vase when there is an actual vase in the room. Thus, Bob's belief that there is a vase in the room is locally reliable since the process by which the belief was formed would, given the computer's programming, tend to produce true beliefs with respect to the particular proposition (a*). However, I believe the above scenario represents a bona fide Gettier case since, though Bob's belief is justified (i.e., reliably formed) and true, we would not want to say that he has knowledge.[10] I will refer to the Gettier case against local reliability just described with the abbreviation *GTTR*$_{LR}$.

It seems clear that local J-reliabilism, when understood in terms of (R1) is insufficient for resisting GTTR$_{LR}$. According to (R1), a process, {M}, is reliable so long as it is reliable with regards to the situation in question. The computer program described in GTTR$_{LR}$ clearly meets this stipulation. (R2), by contrast, is concerned not only with the situation in question, but also with situations 'just like it'. Thus, (R2) seems to introduce the idea of some sort of unspecified counterfactual condition for a process being reliable. This notion takes on a much more robust form in (R3) where a range of possible situations seems to be implicated.

4. Gettier and Nozick-Style Justification

We may further unpack the idea of counterfactuals, hinted at in (R2) and suggested by (R3), in terms of Nozick's [1981: 179] subjunctive conditionals:

(N1) If *p* were not true, then *S* would not believe that *p*

(N2) If *p* were true, then *S* would believe that *p*[11]

[10] I believe Bob does not have perceptual knowledge that there is a vase in the classroom because he does not see the vase. On this score, I am in agreement with Paul Snowdon's observation that seeing that *p* puts one in a position to have demonstrative thoughts about *p*. If Bob attempted to have demonstrative thoughts about the vase he would fail to do so since he would point to the hologram rather than the real vase. As such, Bob does not see the vase.

[11] Counterexamples such as the Grandma case have prompted Nozick [1981, p. 179] to revise (N1) and (N2), limiting them to the same method (i.e., vision):

Together, (N1) and (N2) allow a subject's belief to track the truth of a given proposition, which has earned Nozick's version of reliabilism the title: the truth-tracking account. Significantly, Nozick's truth-tracking account is actually a type of TB+G reliabilism and is therefore not a genuine version of J-reliabilism (See footnote 4 above). However, we can, for the time being, ignore this fact by thinking of (N1) and (N2) as the necessary and sufficient conditions for justifiably believing that *p*.[12] We may say that (N1) and (N2) are true, as counterfactual statements, *iff* in possible worlds near to the actual world, if *p* is false, *S* does not believe that *p*, and if *p* is true *S* believes that *p*.[13] Nozick recommends that we assess (N1) and (N2) by reference to what is the case in all nearby possible worlds and not just the actual world. Roughly, a world may be described as 'nearby' if it is only slightly different from the actual world and 'distant' if it is radically different. Given the above account of nearby worlds, the following reply to $GTTR_{LR}$ seems available to Nozick. Let us suppose that the set of nearby possible worlds include ones in which the computer crashes or malfunctions. In such nearby worlds, although there is a vase in the room, there is no hologram of a vase. Since in such a world the subject would not believe that (a*) although (a*) were true, (N1) has not been satisfied.

There are a number of reasons why the above reply to $GTTR_{LR}$ is unsatisfactory. First, there does not seem to be any principled reason for holding that the set of nearby worlds include ones in which the computer crashes or malfunctions. For example, suppose the computer in question is of a particularly

---

(N1*) If *p* were true, *S* (using M) would believe that *p*

(N2*) If *p* were not true, then *S* (using M) would not believe that *p*

(N1*) and (N2*) restricts us, not only to a particular proposition *p*, but also to a particular method M for arriving at *p*. Since I am combining Nozick's conditionals with process reliabilism, the restriction to a single method is already built into the present analysis.

[12] Some may find Nozick's conditionals implausible as a criterion for justification since, *inta alia*, it makes justification factive. However, factive justification figures in certain prominent contemporary philosophical system, such as that of John McDowell, who writes: 'seeing that *p* constitutes false-hood excluding justification for believing that *p*' (McDowell [2002], p. 98). Thus, the mere fact that a system implicates factive justification is not enough to dismiss it from consideration. (Significantly, McDowell's account of justification is different from that discussed here since it relies on a safety principle rather than sensitivity.) In any case, my purpose for construing Nozick's conditionals as a criterion for justification is aimed at exposing a certain flaw I see in the anti-luck diagnosis of Gettier. (See below for details.) I am not here suggesting that it is a position actually represented in the literature.

[13] I appeal to the argot of possible worlds only as a linguistic convenience and nothing more. For further clarification of the semantics of subjunctive conditionals see Stalnaker [1968]; Lewis [1973] and Nozick [1981].

high quality (we may suppose that it is a Mac rather than a PC). Would a world where the computer crashes or malfunctions still count as nearby? This question points towards a certain capriciousness in the determination of nearby worlds. Consider the following often cited example:

(A) If Nixon had pressed the red button there would have been a nuclear war.

(B) If Nixon had pressed the red button the connecting wire would have broken.

On at least one interpretation, the world in which (B) obtains is more similar to the actual world in that it, like the actual world, has not undergone a nuclear war. However, the possibility of the wire breaking does not seem salient given the context suggested by the content of the two sentences. Thus, intuitively, (A) seems like the acceptable counterfactual, while (B) seems both arbitrary and unwarrantedly optimistic. The question that confronts the nearby worlds type reply to $GTTR_{LR}$ is why should we consider the computer malfunctioning salient vis-à-vis the determination of nearby worlds, especially given that we have already stipulated that the computer in question is not one prone to malfunctions. Of course it would best serve Nozick's purposes to do so, but such a motivation would seem, at best, ad hoc.

But even more crucially, Nozick seems to have good reason for limiting the set of nearby worlds to ones in which the computer is working properly. For example, while not common, people do occasionally go blind due to damage or disease. Now, if the set of nearby worlds is taken to include ones in which a perfectly healthy subject suddenly goes blind, then vision would no longer count as a reliable method. But this is simply implausible. Now conceivably, there could be a computer which is no less reliable in its operation, than a human sense organ. Thus, if worlds in which a healthy subject goes blind are excluded from the set of nearby worlds, then the same should hold for worlds in which a perfectly functioning computer malfunctions.

The point of the above example, then, is that any account which included worlds in which the computer malfunctions as nearby seems to prove too much. This is a fact that Nozick himself seems to be sensitive to, as evinced in his treatment of the Jesse James case.[14] Here is the Jesse James example, as described by Craig [1990]:

> Jesse James, the reader will recall, is riding away from the scene of the crime with his scarf tied round his face just below the eyes in the approved manner. The mask slips, and a bystander, who has studied the 'wanted' posters, recognises him. The bystander now knows, surely, that it was James who robbed the bank. But Nozick has a problem: there is a possible world, and a 'close' one, in which James' mask didn't slip, or didn't slip until he was already past the bystander; and in that world the bystander wouldn't believe that James robbed the bank, although it would still be true that he did. So Nozick's condition [N2] is not satisfied, and he is threatened with having to say that the bystander doesn't know that it was James, even though the mask did slip. So his analysis looks like ruling out something which is as good a case of knowledge as one could wish for [p. 22].

Nozick is already equipped with a reply to the 'Jesse James' counterexample, akin to that employed in the Grandma case (see footnote 11 above). In brief, he may simply argue that in worlds in which the mask did not slip, the method employed by the bystander would be different. Thus, given the version of Nozick's counterfactuals revised to include the subject's method, worlds in which the mask did not slip would not be included in the relevant nearby possible worlds. However, to the extant that this reply is effective in preserving the bystander's knowledge in the Jesse James case, it is also effective at preserving Bob's 'knowledge' that there is a vase in the room. In $GTTR_{LR}$, the computer constitutes part of the method by which Bob arrives at his belief that there is a vase in the room. Thus, by the lights of (N2*), all worlds in which the computer has crashed, malfunctioned or missing altogether, a different method is being used and such worlds would *eo ipso* fail to count as nearby.

[14] See: Nozick [1981] p. 193.

5. Gettier and the (Anti) Anti-Luck Diagnosis

The traditional diagnosis of what goes wrong in Gettier-type scenarios is in terms of luck. While I believe there is much to be said for the standard anti-luck diagnosis, I do not think it is the sole explanation of what can go wrong in Gettier-type scenarios. I believe this claim is supported by the fact that a subject whose belief was formed by a reliable process can still be Gettiered. If a subject's belief is reliably formed, then prima facie, it cannot simply be a matter of luck that it is true. However, it may be objected that there is still a species of luck involved in Bob's belief that (a) since it is only a matter of luck that the hologram he sees is reliably linked to a real vase via a computer. Thus, we may say that Bob is lucky that his belief is reliable. Moreover, one may insist that it is this higher order luck (if you wish to put it like that) that is responsible for Bob's belief being Gettiered.

While tempting, I believe the above defence of the anti-luck diagnosis is deficient. Even if we grant that it is matter of luck that Bob's belief is reliably formed, this cannot be the explanation of why his fails to constitute perceptual knowledge. This is because there are cases implicating such (higher order) luck which still seem to constitute legitimate cases of perceptual knowledge. For example, consider a subject, Alan, who is out on a leisurely drive in an unfamiliar countryside. He comes to a fork in the road, which we may describe as {ROAD1} and {ROAD2} respectively. Alan cannot decide with of the two roads to take, and so he decides to flip a coin. By luck, {ROAD1} wins the coin toss and so Alan goes down {ROAD1}. Let us suppose further that both roads terminate at a different aviary. The aviary at the end of {ROAD1} contains only Robbins and Woodpeckers. The aviary at the end of {ROAD2} contains only Robbins and Sparrows.

Now suppose that Alan is able to reliably distinguish between Robbins and Woodpeckers, so that he never confuses the two. However, Alan cannot distinguish between Robbins and Sparrows and is prone to identify both species as Robbins. We may sum up this idea by saying that Alan's ability to recognise Robbins by sight is indexed to environments containing only Robbins and Woodpeckers but

not to environments containing Robbins and Sparrows. Since Alan decided to take {ROAD1} he ultimately visits the aviary with only Robbins and Woodpeckers. Alan sees a Robin and pointing to it says, 'there goes a Robin!' Since Alan can reliably distinguish between Robbins and Woodpeckers, intuitively we would say that Alan knows that the bird is a Robin. However, it is only a matter of luck that Alan finds himself in an environment to which his ability to recognise Robbins is indexed. After all, if the coin-flip had gone the other way, Alan would have gone down {ROAD2} and would have found himself in an environment in which he was not able to reliably identify Robbins by sight. Thus, we may say that it is only a matter of luck that Alan's belief is reliably formed. Nevertheless, the fact that Alan's belief that the bird before him is a Robin is lucky in this respect does not undermine our intuition that he has perceptual knowledge. I conclude that in the Alan case, just as in the Bob example, it is only a matter of luck that the subject's belief is reliably formed. However, while intuitively the former has perceptual knowledge, the latter does not. Thus, the presence of luck cannot account for the epistemic difference between the two cases.

6. Gettier and Global Justification Reliabilism

The stronger notion of global reliability offers us a way to respond to GTTR$_{LR}$. Recall, global reliability requires that the process by which *S* forms her belief that *p* be reliable with respect a range of *p*-like propositions. However, since the subject in the GTTR$_{LR}$ is prone to mistake holograms of vases for real vases, then the subject's reliability does not extend to other cases in which a hologram is not causally connected to an actual vase. Thus, from a global reliabilist perspective, the subject is not reliable. Thus, the stronger notion of global reliability succeeds where the weaker notion of local reliability fails. Unfortunately, global reliability's strength (pun intended) may prove to be its greatest weakness, since it is doubtful that we consider someone justified in believing that *p*, only if they are reliable regarding a range of *p*-like propositions.

More often than not, any process, {M}, that is reliable with regards to a particular proposition *p*, would also be reliable with regards to a number of *p*-like propositions as well. Even so, it is far from established that this is a conceptual requirement of what it means for some process, {M}, to be reliable. In other words, it may be that while the extension of our concept of reliability is typically global in scope, the intention is nevertheless only local in nature. But it is only the intention of the concept that concerns us. The question, then, is this: is *S*' knowing that *p* contingent on her being reliable regarding a wider class of *p*-like propositions? There appear to be a least some cases in which there could be knowledge that *p* when the relevant class of '*p*-like' truths doesn't reach beyond *p* itself, assessed at other times. This typically occurs when the subject stands in a unique relationship to the fact in question (such as perceptual cases), or may be presumed to be uniquely familiar with it. Edward Craig [1990] advances an argument to this effect:

> An infant may know what its name is without knowing what anyone else's name is, and without knowing its age, and it may know its forename without knowing its surname. Being confident that someone knows their name doesn't help us to judge that they are valuable informants on many other questions, if any. If we try hard enough we will probably be able to imagine circumstances in which we might reasonably trust a potential informant on the question whether there was a table in the next room, but not on the question whether there were chairs there, or anything else about the furniture. (In fact one doesn't even have to try very hard: suppose that he has never been in the room, but has just seen a table being carried through the door.) So if we are looking for strictly necessary conditions...we shall plump for localism and confine ourselves to the individual proposition at issue. [p. 57]

However, more needs to be said if we are to find Craig's examples persuasive. For example, a defender of global reliability may counter as follows: The above argument seems to rest on the mistaken assumption that for a process, {M}, to be reliable regarding *p*, it must have actually been reliably employed in relation to a number of *p*-like propositions. Craig argues that the fact that we would count an infant as knowing its own name although it does not know anyone else's shows that local reliability is enough. But the fact that an infant may know its name and no one else's does not mean that the process,

{M}, by which it acquired this knowledge is only locally reliable. It may very well be the case that {M} would also allow the baby to learn other names, but the infant simply hasn't had the opportunity (as of yet) to implement the method towards this end. Recall that according to Goldman, reliability is merely a dispositional feature of {M}. The question, then, is not whether {M} has been reliable with regards to a number of *p*-like propositions, but rather *would* it be. The same goes for Craig's example of the person who sees the chair (and only the chair) being taken into the room. Let us assume that the method by which Craig's subject comes to know that there is a chair in the room is something like the following:

{M*} seeing α being taken into the room (where α stands for some object or other)

Now, it so happens that the Craig subject has only been able to implement {M*} once, and so only has knowledge of the single proposition, *there is a chair in the room*. However, presumably, if the subject were to see a table or a television being taken into the room, he would also, using {M*} come to know that there was a table or television in the room. The fact that the subject has not had the opportunity to implement this method (as of yet) is beside the point.

Although Craig's objection is ultimately unsuccessful in impugning global reliabilism, I believe the problem lies in Craig's example rather than his strategy. If we could construct an example in which a subject intuitively has knowledge, despite the fact that the process by which she acquires her beliefs lacks global reliability, then perhaps we may succeed where Craig fails. I believe such an example is available. Consider a Royal Air Force officer who is using a second-rate radar to identify enemy airplanes. The radar system is designed to bleep only when there is a B-52 within a hundred yards of the military base. However, the radar cannot detect any other airplanes besides B-52s. Moreover, when it detects a B-52, it cannot tell exactly how near or far it is, the bomber's velocity or trajectory, or any other subsidiary information apart from the fact that the bomber is within a hundred yards of the base. In short, the radar

is reliable with regards to the proposition: 'there is a B-52 within a hundred yards of the military base', but not with regards any other propositions. Thus described, we may say that the radar lacks global reliability. However, despite the radar's lack of global reliability, we would still say that when the radar bleeps, the Royal Air Force officer knows that there is a B-52 within a hundred yards of the base. This suggests that a subject may arrive at knowledge even though the process by which she acquires her beliefs lacks global reliability. Global reliability therefore fails to constitute a necessary condition for knowledge.

7. Conclusion

The definitive challenge to JTB are Gettier counterexamples which seem to present cases of justified true belief that fail to constitute knowledge. According to the J-externalist, the key to resisting Gettier is to eliminate the element of luck involved in Gettier-type situations by establishing an objective connection between justification and truth. The standard means by which such a connection is achieved is by insisting that a belief is justified only if it is formed via a reliable process. On this view, the Gettier subject lacks reliability and therefore is not justified. However, local reliability is not sufficient for responding to all Gettier cases since the fact that a particular belief was formed via a reliable process may itself be a matter of luck. Moreover, I have suggested that since luck features in certain putative cases of knowledge, a more sophisticated alternative to the straightforward anti-luck approach is called for. Global reliability appears to succeed where its local counterpart fails. However, while global reliability seems sufficient for resisting Gettier, it does not appear to be necessary for knowledge. We therefore find ourselves confronted with an all too familiar epistemological dilemma, how to find a justification condition that is both sufficient for resisting Gettier, and necessary for knowledge. I have argued that J-reliabilism fails to provide a satisfactory answer to this question.

REFERENCES

Alston, W. [1989], *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.

Armstrong, D. M. [1973], *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.

BonJour, L. [1992], 'Externalism/internalism', in J. Dancy and E. Sosa (eds.), *A Companion to Epistemology*. Oxford: Blackwell.

Craig, E. [1990], *Knowledge and the State of Nature*. Oxford: Clarendon Press.

Feldman, R. [1974], 'An Alleged Defect in Gettier Counter Examples', *The Australasian Journal of Philosophy*, 52: 68-9.

Gettier, E. [1963], 'Is Justified True Belief Knowledge?', *Analysis*, 23, 6: 121-3.

Goldman, A., [1967] 'A Causal Theory of Knowing', *The Journal of Philosophy*, 64, 12: 357-372.

Goldman, A. [1979], 'What is Justified Belief?' in G. Pappas (ed.) *Justification and Knowledge*. Dordrecht: D. Reidel.

Greco, J. [2004], 'Externalism and Skepticism', in R. Shantz (ed.) *The Externalist Challenge: New Studies on Cognition and Intentionality*. New York: de Gruyter.

Harman, G., (1970], 'Knowledge, Reasons, and Causes', *The Journal of Philosophy* 67: 844-55.

Harman, G. [1973] *Thought*. Princeton, NJ: Princeton University Press.

Lewis, D. [1973], *Counterfactuals*. Oxford: Basil Blackwell.

McDowell, J. [2002], 'Knowledge and the Internal Revisited', *Philosophy and Phenomenological Research*, 64, 22-30.

Pollock, J. [1986], *Contemporary Theories of Knowledge*. Savage, MD: Rowman and Littlefield.

Stalnaker, R. [1968], 'A Theory of Conditionals', *American Philosophical Quarterly*, monograph no. 2: 98-112.

Sturgen, S. [1993], 'The Gettier Problem', *Analysis*, 53,3: 156-164.

Vogel, J. [2000], 'Reliabilism Leveled', *The Journal of Philosophy.* 97, 11: 602-623.

Williamson, T. [2000], *Knowledge and Its Limits*. Oxford: Oxford University Press.