

Conclusive Reasons to Believe

Tamar Lando
University of California, Berkeley*

Abstract

How should we understand reasons-based knowledge? According to the counterfactual or modal account of conclusive reasons, famously defended by Dretske, R is conclusive for P just in case [R would not be the case unless P were the case]. I argue that while knowing is plausibly related to having conclusive reasons to believe, having such reasons cannot be understood in terms of the obtaining of this counterfactual condition. I suggest a new theory of what it is to have conclusive reason to believe that P, which does a better job of capturing our intuitive judgments in a wide range of cases. The inadequacy of the counterfactual account has, as I argue, wider consequences for safety-based theories of knowledge, and modal accounts of ‘epistemic luck.’

1 Introduction

Here is one way in which knowledge differs from other epistemic attitudes. In knowing something, *one could not be wrong about it*. When a person knows—as opposed to merely believes, or justifiably believes—that P is the case, her epistemic credentials are “such as to eliminate the possibility of mistake,” as Dretske once said.¹ It seems right to conclude that if a person knows that P on the basis of reasons, her reasons must be *conclusive*. But what is it to have conclusive reason to believe that P?

There is a well-known counterfactual or modal answer to this question. Suppose that S believes P on the basis of R. Then R is conclusive for P if the following counterfactual holds:

(CR) R would not be the case unless P were the case.

*Email: tlando@berkeley.edu. Mailing Address: 314 Moses Hall #2390, University of California, Berkeley, CA 94720-2390. Phone: 310.991.1100

¹[Dretske, 1971]

If R and P satisfy (CR), then R counterfactually guarantees the truth of P. Indeed, if (CR) is true and R is the case, then it is correct to say not just that P *is* true, but that P *must* be true.² To have conclusive reasons, we might say, is just to believe on the basis of reasons that satisfy (CR). Even if this cannot be a fully general account of knowledge—even if there may be other ways of knowing that have nothing to do with reasons—we might still think that where knowledge does depend on reasons, (CR) specifies the crucial condition that must be met.

In what follows I will argue that the modal account of reasons-based knowledge (as I will call it) cannot succeed. In simple examples, I show that it is possible to believe P on the basis of reasons that satisfy (CR), and still *get lucky* in believing the truth about P. Given that one does not *know* what one gets right only as a matter of *luck*, the agent in these examples does not know that P. If this is right, then (CR) is not sufficient for reasons-based knowledge. So while knowledge may be a matter of having conclusive reasons to believe, we cannot understand what it is to have conclusive reasons in terms of the obtaining of such counterfactuals. Seeing where things go wrong, or where (CR) falls short, points the way to a more adequate set of conditions on reasons-based knowledge—or so I shall argue.

Talk of conclusive reasons has wider consequences. If the modal condition (CR) fails to rule out cases of epistemic luck, or luck in getting things right, what about other modal conditions advanced in the theory of knowledge more generally? Here I have in mind, in particular, the ‘safety’ and ‘sensitivity’ principles. As I argue, the very same examples adduced against the claim that (CR) rules out cases of epistemic luck also show that ‘safety’ and ‘sensitivity’ do not rule out luck. By focusing on conclusive reasons and (CR), we get a handle on a broader debate: What exactly is the place of purely external modal constraints in a theory of knowledge more generally, and can *any* such constraint rule out the presence of the pernicious kind of epistemic luck?

2 A Counterexample to the Modal Account

Let us begin with Dretske, who famously advanced a theory of conclusive reasons that relied heavily on the counterfactual condition (CR).

According to Dretske, a subject S has conclusive reason R for P just in case:

(A) R would not be the case unless P were the case.

(B) S believes, without doubt, reservation, or question, that P is the case and he believes this on the basis of R.

²As Dretske says, “we are entitled, not only to deny that, given R, not-P is the case, but also that, given R, not-P *might* be the case.” [Dretske, 1971]

(C) S knows that R is the case or R is some experimental state of S (about which it may not make sense to suppose that S *knows* that R is the case; at least it no longer makes much sense to ask *how* he knows).

To have conclusive reasons for believing P, Dretske claims, is to believe P on the basis of reasons that are in-point-of-fact conclusive: reasons that would not obtain unless the proposition believed were true.

Against this, I think it is easy to see that the three conditions stated are not jointly sufficient for reasons-based knowledge. Consider the following example:

Poker. Jessica is trying her hand at poker. She has just recently been taught the rules of the game and, despite her immense confidence, is not a very adept player. She takes note of the king of clubs lying open-face on the table, and when her neighbor, John, decides to raise she immediately concludes that John must have another king in his hand, completing a pair. (Jessica is not aware of—or does not consider—the other possibilities in the game). As it happens, and unbeknownst to Jessica, John is a rather strange player himself, and plays according to the following strategy (henceforth, Strategy): Raise whenever there is a king on the table and you have a king, and call in all other cases.

John's strategy makes the following claim true: John would not have raised unless he had a king. Moreover Jessica believes, without a trace of doubt, that John has a king on the basis of the fact that he just raised. Finally, Jessica knows that John raised because she is sitting at the table with him. Jessica satisfies (A) - (C) above, so according to Dretske's account of conclusive reasons, she has conclusive reason to believe that John has a king (and hence also *knows* that John has a king).

But this is clearly the wrong verdict. A player in similar circumstances to John might raise for any number of reasons. For example, the king on the table might finish his straight, he might have another pair in his hand unrelated to the king on the table, he may merely be bluffing, and so on. Of course *John* would not have raised under any of these other circumstances, but that is only because John is playing by a very peculiar strategy. The fact that he is playing by that strategy is itself something Jessica is entirely ignorant of. So it seems wrong to say that given what she knows, Jessica has conclusive reasons to believe that John has a king. And equally, it seems wrong to say she *knows* he has a king. If this is right, then the example forces us to give up the claim that (A) - (C) specify what it is to have conclusive reasons, as well as the claim that satisfying these conditions is sufficient for knowing that P.

3 A Response: Higher-Order Knowledge

Or does it?

Let's briefly consider a possible response on behalf of someone eager to defend (A) - (C). Here one might insist that Jessica really does *know* that John has a king. What explains our intuition to the contrary is just that while Jessica knows, she does not know *that she knows*. In the example, it is this second-order knowledge that is missing, and that we easily—or at least in this case—confuse with first-order knowledge (knowledge that John has a king).

How plausible is this line of response? In my view, not very. There are at least two ways to see this. Consider what *you* would say if you were sitting next to the table with Jessica, and you yourself did not know that John plays by Strategy. Jessica (you are to imagine) leans over your shoulder, and whispers to you that John has a king. "A king?" you say. "Why do you think that?" "Well, didn't you see him just raise?" she says, significantly. "There was a king on the table," she continues. "So he must have a pair." A skeptical look crosses your face. "Don't be silly," you say. "John could have raised for any number of reasons. He may have no good cards at all, may simply be bluffing." After some more back and forth, it becomes clear to you that Jessica has no further evidence to support her strange belief. Fast forward now to the end of the poker hand. Jessica is eager to prove her point, and she asks John to show her his cards. Obliging, John lays them out, and lo and behold, there is a king amongst them. Moreover, John now explains to the table that he has been playing by Strategy all along; that he *always* plays by Strategy. You look over at Jessica, to see if she knew about Strategy, but she seems just as perplexed as you are. Here is the question: Do you, at this point, say to yourself, "Aha, Jessica *knew* that John had a king all along"?

To react in this way would, I think, be very generous—indeed, far too generous. Jessica didn't *know* at the time, there simply was no *way* for her to know, given that she was not aware of Strategy. She got things right, sure—but only as a matter of luck. In the example, Jessica's ignorance of the game just so happens to coincide in the right way with John's idiosyncratic Poker strategy, producing (miraculously) a true belief. "A house of lunatics," you mumble to yourself, as you walk away from the game.

This leads us to a second way at the intuition here, by comparison with the familiar Gettier cases. In those cases, Gettier appealed to our intuition that knowledge is incompatible with a certain kind of *epistemic luck*—or luck in getting things right. If one gets things right about P only as a matter of luck, then one does not count as knowing that P is the case. (For a simple example, suppose that Jones and Smith applied for the same job, and that Smith has very good reason to think that Jones will get it. Moreover, Smith knows that Jones has ten coins in his pocket.

He believes, on this basis, that the man who will get the job has ten coins in his pocket. As it turns out, Smith himself will get the job, but he too has ten coins in his pocket.) Showing that the subject in these cases gets things right as a matter of luck was tantamount, for Gettier, to showing that the subject does not know whatever it is he is said to know.

But now it seems clear that in *Poker*, Jessica believes the truth about what cards John has only as a matter of luck. After all, Jessica has no idea that John plays by a particular strategy and merely assumes, when he raises, that his cards must complete a pair with the open-face king. That John plays by the strategy he does is, from this point of view, an amazing stroke of luck. Had John been a more ordinary player (and Jessica, we are to assume, has no reason to think he is *not* ordinary), Jessica would surely have gotten things wrong. If we accept the verdict in the Gettier cases, then we must conclude that when Jessica gets things right in *Poker*, she does not count as *knowing* that John has a king. Our judgment in the Gettier cases, in other words, extends smoothly to the verdict here. (We will return to the *asymmetries* between *Poker* and the Gettier cases in Section 5). If this is right, then it is not just that Jessica does not know *that she knows*: rather, she does not *know* in the first place.

4 Another Response: Appropriately Specific Reasons

In the final pages of “Conclusive Reasons,” Dretske himself discusses some problem cases for the claim that (A) - (C) are jointly sufficient for knowledge, and suggests a way of dealing with them. The cases bear some resemblance to the counterexample just offered, and it is instructive to look at what Dretske says in this context. Here is one of them:

Chemistry. S, upon inspecting an immersed chemical indicator, declares that the solution in which it is immersed is a base. He believes (correctly) that the indicator is the sort which turns from yellow to blue only when immersed in a base. The indicator is Thymol Blue and would not have turned from yellow to blue (in these conditions) unless the solution were a base. [...] Does S know that the solution is a base? Before answering this question the reader should be informed that there is another chemical indicator, Bromophenol Blue, which also turns from yellow to blue but only when immersed in an acid. S, however, is quite unaware of the existence of other such indicators. He merely assumes that a yellow indicator turning blue is a positive test for a base. [Dretske, 1971]

S’s ignorance on this point does not alter the fact that S satisfies the following

conditions:

- (A') The indicator's change in color is a conclusive reason for believing that the solution is a base.
- (B') S believes that the solution is a base and he believes this on the basis of the indicator's change in color.
- (C') S knows that the indicator changed from yellow to blue.

And yet Dretske says, "despite the satisfaction of these conditions, I find it (in some cases) most implausible to say [S] knows that the solution is a base."

What goes wrong according to Dretske is that although S's reason for believing P is a conclusive one, S is not exploiting those aspects of his reason *in virtue of which* it is conclusive. The fact that this indicator is Thymol blue (and not Bromophenol) is what makes it the case that the indicator's change in color is conclusive for P, but S's reason for believing P is simply that *this* indicator changed color, not that a *Thymol Blue* indicator changed color. In general, Dretske notes, the fact that an A is a B is, in some cases, conclusive reason for believing P only in virtue of the fact "that it is in particular *an A* which is B." In such cases, Dretske says, we should require that the agent's reason for belief be not simply that "*this* (something or other) is a B," but that the something or other that is a B is also an A. S need not, according to Dretske, know that his reason is conclusive for P, but he must be exploiting those features of his reason in virtue of which it is conclusive. The needed qualification on (A) - (C), Dretske thinks, can be put in terms of a further restriction on our reasons. Thus, "when one has conclusive reasons, then this is sufficient for knowing that P is the case when those reasons are properly *specific*, both with regard to what it is that displays the particular features on which one relies and on the particular features themselves." [Dretske , 1971]

I think it is easy to see, however, that even when an agent's reasons for belief are properly specific in the sense Dretske has in mind, conditions (A) - (C) are not sufficient for (reasons-based) knowledge. The reason is that while S may be exploiting the correct features of her reason—those features in virtue of which R is conclusive for P—she may do so while still getting lucky in fixing on just those features. In particular, if S is not justified in taking R to be a reason for P in the first place then in some cases although S exploits (what happen to be) the right features of her reason, she cannot be said to know that P is the case. Consider a modified version of Dretske's example.

Chemistry (Modified). S is a chemist, inspecting a chemical solution in which he has immersed a certain chemical indicator. The solution has been chosen

(at random, say) from three different chemical substances, C_1 , C_2 and C_3 . S believes (correctly) that the indicator is the sort which turns from yellow to blue when immersed in C_1 , but not in C_2 , and so upon seeing the indicator turn to blue, comes to believe that the solution is C_1 . (We can add here that S is aware of the existence of another indicator, that turns blue when immersed in C_2 , but not in C_1 – this is not important for present purposes). Nevertheless, S knows nothing at all about what happens when the indicator he is using is immersed in C_3 – a substance which is significantly different in chemical composition from both C_1 and C_2 . As it happens, the indicator remains yellow when immersed in C_3 , and for this reason, the fact that this indicator now turned to blue is a conclusive reason to believe that the solution is C_1 . However, when asked how he knows the solution is not C_3 , S replies, “ C_3 ? I’m not aware of any such solution.”

Under these circumstances, it would be implausible to say that S *knows* that the solution is C_1 . Indeed, I do not think it is more plausible to say that S knows that the solution is C_1 in the modified example, than to say that S knows that the solution is a base in Dretske’s original example. There is, after all, a prominent alternative to what S believes that S cannot exclude even though S’s reason is in fact incompatible with that alternative. Seen in these terms, the original scenario and the embellished one are counterexamples to the sufficiency of (A) - (C) for the very same reason: there is in each case a possibility that the agent must rule out, but is not in a position to rule out given what she knows. But if this is right, then the insufficiency of (A) - (C) cannot be fixed simply by requiring that an agent’s reasons be properly specific with respect to what it is that displays the features on which he relies. That is because S’s reasons are *already* as specific as one would like (S’s reason is that an indicator of *this particular* type turned blue, and not simply that an indicator of any type turned blue) and yet S fails to know that the solution is C_1 .

One could respond that in this case there is no reason to think that C_3 is a prominent alternative, or at any rate an alternative that must be ruled out in order for S to count as knowing that the solution is C_1 . After all, no agent is able to rule out *all* (logically possible) alternatives to what she believes, and so given that we sometimes know things, there must be some alternatives that we can safely ignore. What makes the C_3 alternative ‘relevant’ in the sense that chemist must be able to rule it out in order to count as knowing? Here, however, it is all too easy to supplement the case with the right background circumstances in order to generate the verdict we need about relevance. If you are the sort of person who thinks that whether an alternative is relevant depends on objective factors (such as whether there was any ‘objective chance’ that the solution chosen should have been C_3),

then imagine that the solution was chosen at random from among C_1 , C_2 and C_3 , with each outcome just as likely as any other. If you are someone who thinks that what is relevant depends instead on what alternatives are contextually prominent, then imagine that the question of whether S knows that the solution is C_1 comes up in a context in which two scientists are attempting to discover, and are genuinely uncertain about, which of the three solutions, C_1 , C_2 and C_3 , this is. And so on. The crucial point is that altering these circumstances in the ways suggested does not affect the fact that the indicator's change in color is a conclusive reason for the solution's being C_1 . Change the circumstances in any of these ways—the fact that the indicator turned blue *still* satisfies the relevant counterfactual condition.³

Dretske suggests a further way in which reasons can be made appropriately specific. Consider another of his examples:

Behavior. Suppose K is behaving in such a way that it is true to say that he would not be having in that way unless he was nervous. Suppose S purports to know that K is nervous and when asked how he knows this, replies by saying, 'From the way he is behaving.' Once again, our three conditions are satisfied, or can easily be assumed to be satisfied. Yet, if we suppose that the distinctive thing about K's behavior is that he is doing B_1 while performing B_2 , then if S is relying on B_1 (or B_2) alone, we should not say that he knows that K is nervous... The fact is that the crucial aspects (those aspects which make K's behavior conclusive) are more specific than those on which S is relying in purporting to know. [Dretske, 1971]

Here Dretske's solution is to require that an agent's reasons be properly specific not just to *what it is* that displays the particular features, but to those features themselves (*i.e.*, S's reason must be specific to those features of K's behavior that make it a conclusive reason to believe that K is nervous). Thus S's reason must be that K is performing B_1 while performing B_2 , not just that, *e.g.*, K is behaving in the way he is (like *that*). But again, it's hard to see how this solves the problem. Suppose for example, that S relies on both B_1 and B_2 in believing that K is nervous. S knows that when K is happy, K performs B_1 without B_2 , or B_2 without

³One need not be a 'relevant alternatives' theorist to see the force of the fact that S is unable to rule out the C_3 alternative. Whatever account of knowledge one adopts, it seems right to say that knowing that P requires being able to rule out at least *some* alternatives to P. But once we concede this, it is just a matter of making plausible that this particular alternative—namely that the solution is C_3 —is one that S must be in a position to rule out if she is to know that the solution is C_1 . This, I submit, is something we can do just by fixing the background circumstances in the case. Moreover, we can do this without affecting the judgment that S's reason satisfies the relevant counterfactual condition on conclusive reasons. (Other accounts which appeal explicitly to the idea of ruling out alternatives include, *e.g.*, contrastivism and certain contextualist relevant alternatives accounts.)

B_1 , but never the two together. Nevertheless, S does not know that when K is sad (but not nervous) K also does not perform B_1 and B_2 together. Asked what K does when K is sad, S says, ‘I really couldn’t tell you – he’s never been that way around me.’ (Add to the picture, if you like, that K is very *often* sad, that the question of whether K is happy or sad has just come up in conversation, and the participants to this conversation are wondering whether S knows which of these is the case, and so on.) Here it seems to me that as before, although S is relying on the right aspects of K’s behavior—on those aspects that *make it* conclusive—S cannot be said to know that K is nervous. We cannot fix the problem simply by requiring that S’s reasons be *specific* to the set of features which make S’s reason conclusive. S’s reasons are already as specific as one would like (S believes on the basis of the fact that K performed B_1 and B_2 *together*), and yet S gets things right at least partly as a matter of *luck*.

What exactly do examples like *Poker*, *Chemistry*, and *Behavior* show?

It would be a mistake, I think, to conclude that S’s reason in these examples is not conclusive after all. It would be wrong to suppose, in *Chemistry* for example, that the indicator’s change in color is not a conclusive reason to believe that the solution is a base. (Or that in *Behavior*, the fact that K performs B_1 and B_2 is not a conclusive reason to believe that K is nervous.) After all, someone else *could* come to know what S fails to know on the basis of the very same reason. This person would presumably have to know that it is only Thymol Blue (and not Bromophenol Blue) that remains yellow when immersed in an acid. But if she *did* know this, then there is nothing barring her from coming to know that the solution is a base on the basis of the fact that the indicator turned to blue. In doing so she will have come to believe what she does on the basis of a reason that is in-point-of-fact conclusive. She will have come to believe P on the basis of a reason, R, that guarantees the truth of P.

The examples also do not show that there is something problematic in the very notion of one thing being conclusive for another. We know what it is for R to be conclusive for P. As we’ve said all along, it is just for R and P to satisfy the counterfactual condition (CR): R would not be the case unless P were the case. Insofar as that counterfactual is well-defined, so too is the notion of one proposition, fact, or state of affairs being conclusive for another. (These examples pose no *special* problem for conclusive reasons that is not independently a problem for these counterfactuals.) But if the examples don’t show either that (CR) does not specify what it is to be a conclusive reason, or that (CR) does not pick out a well-defined relation on propositions at all, then what *do* they show?

Here I think the answer is quite straightforward. What the examples show is that *it is possible to believe P on the basis of a conclusive reason and yet fail to*

know that P is the case. If R is conclusive for P, then R guarantees the *truth* of P (where R is the case, P must also be the case). But R cannot also guarantee that anyone who believes P on the basis of R is thereby in a position to know that P. The examples above are cases in point: cases where the right objective relationship obtains between R and P, and still in coming to believe P on the basis of R, S gets *lucky* in believing the truth about P.

5 Excursion: Sensitivity and Safety

If we are right to think that the problem in the examples has something to do with *luck*, then the failure of (CR) in cases like *Poker*, *Chemistry* and *Behavior* raises a more general question. To what extent can purely external modal constraints of *any kind* rule out the pernicious kind of epistemic luck—or the kind of luck that interferes with knowledge?⁴

Here it is natural to think of two other modal constraints on belief that have each played a prominent role in post-Gettier theories of knowledge: the ‘sensitivity’ and ‘safety’ conditions. Like (CR), these conditions impose important modal constraints on any agent who purports to know that P is the case. Moreover, they are widely taken to eliminate the kind of luck that interferes with knowledge. As Pritchard [2007] says, “If one meets the safety principle, then it is not a matter of epistemic luck that one’s belief is true.” Of course, ‘safety’ and ‘sensitivity’ are not specifically tailored to cases of reasons-based knowledge, and so, at least on the face of things, do not purport to tell us anything about what it is to have conclusive reason to believe that P.⁵ Nevertheless, these conditions do share in the broadly modal character of (CR). We do well to ask, then, in light of the examples considered above, whether Pritchard and others are right—whether, that is, ‘safety’ and ‘sensitivity’ really do eliminate the kind of luck that interferes with knowledge,

⁴It is important to be clear at the outset what we mean by luck in this context. When we say that an agent believes the truth about P as a matter of luck, we do not mean that P is brought about as a matter of luck. Nor do we mean that it was only a matter of luck that the agent was in a position to discover, or receive evidence for, P. Rather, we mean that it was only a matter of luck that the agent was *right* about P being the case. It is in this sense that luck, or ‘epistemic luck,’ interferes with knowledge. A person who gets things right about P only as a matter of luck does not count as knowing that P is the case. See [Unger, 1968].

⁵Even this disanalogy between ‘safety’ and ‘sensitivity’ on the one hand and (CR) on the other, is somewhat superficial. If we adopt methods-based versions of these principles, as we do below, then one possible ‘method’ for believing P is, presumably, to believe P on the basis of R. (One’s method, in such cases, can be stated as, ‘If R, believe that P!’) In cases where the agent’s method is to believe P on the basis of some reason or other, safety and sensitivity do seem to yield modal conditions on the agent’s reason. Indeed, in the case of sensitivity, the modal condition is (roughly): In the nearby possible worlds where not-P is true, S would not come to believe that P on the basis of R.

or the kind of luck that appears in the Gettier cases. In this section, we take a brief break from reasons-based knowledge *per se* and turn to a closer consideration of these questions.

Recall that S's belief that P is 'sensitive' just in case if P were not the case, S would not believe that P. This condition first appeared in Nozick's tracking theory of knowledge. According to that theory, if S knows that P, then S's belief that P is 'sensitive.' Well-known problems attach to this first, simple formulation of sensitivity, and in response Nozick later adopted a more sophisticated version of the principle that appeals to 'methods' of belief formation. Taking methods into account, the principle of sensitivity states that if S knows that P, then if P were not the case and S were to form a belief whether P *via the method M*, S would not believe that P. Being sensitive to the truth of P does not imply that one would *never* falsely believe that P. Rather, it implies only that in the nearby possible worlds in which not-P is true, one does not believe that P (via M).

Certainly, as Yamada himself notes, sensitivity handles easily many classic Gettier cases in which the agent gets lucky in getting things right about P. To take a familiar example, suppose that Smith and Jones have applied for the same job and Smith has every reason to believe that Jones is the man who will get it. (Smith has been assured of this, say, by the company president.) Smith also knows that Jones has ten coins in his pocket. He comes to believe on this basis that the man who will get the job has ten coins in his pocket. As it turns out, Smith himself will get the job, but unbeknownst to him, he too has ten coins in his pocket. [Gettier, 1963] We judge—correctly, I think—that Smith fails to know that the man who will get the job has ten coins in his pocket, because his belief is true only as a matter of luck. Nozick's sensitivity requirement handles this in the following way. The nearby possible worlds where it is false that the person who will get the job has ten coins in his pocket are one's in which *Smith* does not have ten coins in his pocket. But in those worlds, Smith continues to believe that the man who will get the job has ten coins in his pocket (he believes this, after all, on the basis of the fact that *Jones* has ten coins in his pocket). So Smith's belief is not sensitive and hence, on Nozick's account, does not count as knowledge. Although the 'sensitivity' requirement has fallen out of favor among epistemologists in recent years,⁶ it is still widely thought to eliminate the pernicious kind of epistemic luck, or the kind of luck that interferes with knowledge. Thus Yamada (2010) says, "If one is sensitive to the truth, it is surely no *accident* if one's beliefs turn out to be true."

A more current approach to knowledge popularized by Williamson and Sosa relies instead on the principle of 'safety,' according to which (roughly) if one knows that P one could not easily have been wrong about P. More formally, S's belief that

⁶See, e.g., [Yamada, 2011] and [Williamson, 2002].

P is safe just in case in the nearby possible worlds in which S believes that P, P is true. As with sensitivity, a more sophisticated version of safety can be stated by appealing to the agent's method of belief formation. Thus, taking methods into account, the principle of safety states that if S knows that P on the basis of the method M, then in the nearby possible worlds in which S believes that P via M, P is true. Note that safe belief does not entail sensitive belief (nor vice versa). If P could not easily have been false, then one could not easily have falsely believed that P, and so one's belief that P is safe. But that does not entail that in those distant possible worlds where P *is* false, one does not believe that P. As with sensitivity, it is commonly held that safety rules out cases of epistemic luck.⁷ As Andrew Latus says, "[T]he adoption of such a condition [*i.e.*, safety] will obviously deal with the problem of veritic luck since it defines such luck out of existence." [Latus , 2000]

How plausible are such claims? Does either 'sensitivity' or 'safety' manage to rule out epistemic luck of the kind we find in the examples considered above? For simplicity, we will focus here on *Poker*, where I think the intuition is clearest that the agent gets *lucky* in believing the truth about P. Recall that in that example, Jessica believes that John has a king on the basis of the fact that he raised and there was a king on the table. Moreover, we stipulated that John not only plays by Strategy in the current game, but always plays by Strategy. If that is the case, then the following counterfactual is warranted: If John had not had a king, he would not have raised. But if John hadn't raised, then Jessica would not have come to believe that John had a king—at least not on the basis of her current method. So in the nearby possible worlds where 'John has a king' is false, Jessica does not believe that John has a king. Jessica's belief, in other words, is *sensitive*. Is it also safe? Here we could argue as follows. Jessica's 'method' for believing whether John has a king is: When John raises and there is a king open-face on the table, believe that John has a king! In the nearby possible worlds where Jessica comes to believe that John has a king via this method, it is because John raised. But we must assume that in such worlds, John plays by Strategy (John *always* plays by Strategy—he does so in *all* nearby possible worlds). And so, in those worlds, John does in fact have a king. If this is right, then Jessica's belief is both sensitive and safe. Nevertheless, in believing what she does, Jessica gets things right only as a matter of luck and so fails to know what she is said to know. The example shows that neither 'safety' nor 'sensitivity' rules out the pernicious kind of epistemic luck after all.

One might wonder at this point whether in specifying Jessica's method in this way, we accurately represent her epistemic state in *Poker*. Would it not be more accurate to describe her method as: When *anyone* at the table raises and there is a king open-face on the table, believe that that player has a king! But if we do this

⁷See, *e.g.*, [Latus , 2000], [Pritchard , 2007], and [Yamada , 2011].

then Jessica's method is bound to produce false beliefs—if not about John, then about the other players in the game.

There are two ways to respond. First, even if we allow that this restatement of Jessica's method is more accurate, it does not impact the judgment that her belief that 'John has a king' is safe. Indeed, in the nearby possible worlds where she comes to believe that proposition, it is because *John* raised, and John, as we know, only raises when there is a king in his hand. So Jessica's belief about John is safe, and yet she is lucky to get things right. But second, by supplementing the example with the right background circumstances, we can, I think, generate the verdict that our first formulation of Jessica's method is the right one. Suppose, for example, that Jessica believes that John is the only player at the table who *really* knows how to play Poker—the others are just too stupid to get it, or too inconsistent, and therefore it is not worth hypothesizing anything about what cards they have from their behavior in the game. Here it seems right to say that Jessica's method really is specific to John. 'Believe *of John* that if he raises and there is a king on the table, he has a king!' Using this method, Jessica is still lucky in getting things right about John's cards. She is lucky, as we might say, in adopting the method she does in the first place, given that she does not know that John plays by Strategy. (Another way to put this: She is lucky in adopting what happens to be a truth-conducive method).

6 A Diagnosis

It seems that whether we focus on modal conditions that pertain only to reasons, or ones that pertain directly to belief, the negative moral is roughly the same. There is on the one hand the objective relationship that obtains between our reason, R (or method, in the case of knowledge generally) and the proposition P, and there is on the *other* hand, our own epistemic position *vis à vis* the proposition P. The point of the counterexamples is that these two can come apart. We can believe P on the basis of reasons that happen to be conclusive—reasons that bear the right objective relationship to P—and yet not have what it takes to know that P is the case. If this is right, then to know that P is not simply a matter of what we might call the objective strength of our reasons. To see things this way is to neglect our relationship to them *as reasons*—to ignore, as it were, the grounds on which we take them to be reasons in the first place.

To put the point somewhat more precisely: There is, on one side of things, the possibilities that are left open by our reason, or by our 'method' of belief formation—those possibilities that 'might' be the case, given R.⁸ But there is also the way that we as agents are able to reason with respect to those possibilities. The

⁸Equivalently, those possibilities that are 'counterfactually compatible' with R.

examples show that the two are not one and the same. Suppose that S believes P on the basis of R, and that R is counterfactually incompatible with the alternative, Q, to P. Then, given R, it is correct to say not just that Q is not the case, but that Q *could* not be the case. Nevertheless, it does not follow that in believing P on the basis of R, S is in a position to rule out the alternative Q. Whether or not S is able to do so will depend, in general, on what *other* things S knows, on S's broader epistemic state. We see this directly in the examples. In *Chemistry (Modified)*, the indicator's change in color is counterfactually incompatible with the possibility that the solution is C_3 . But this does not mean that in believing that the solution is C_1 on the basis of the indicator's change in color, S is able to rule out the C_3 possibility. As we saw, she is not. Asked how she knows that the solution is not C_3 , S says, "I don't know of any such solution." Seen in this light, what counterfactuals like (CR) or modal conditions like 'sensitivity' and 'safety' guarantee is only one half of the equation: They guarantee that certain possibilities are incompatible with one's reason or method. But they cannot also guarantee that any agent who uses that reason or method is able to rule out those very same possibilities. If knowledge requires ruling out certain alternatives—if, in other words, it places conditions on how the agent reasons with respect to these possibilities, and not just on the truth of various modal claims—then the purely external constraints we examined above will not do the job.

All of this, of course, leaves us in the position of having to say something more definite about what exactly goes wrong with a person who believes P on the basis of an in-point-of-fact conclusive reason and yet fails to know that P is the case. I suggested above that what is missing in the examples we considered is something to do with the agent's relationship to her reason *as a reason*. It is time now to stick my neck out a little further. If, against Dretske, we simply help ourselves to the notion of justification in this context then I think it is correct to characterize the examples as ones in which the agent is not *justified* in believing P, but believes P on the basis of an in-point-of-fact conclusive reason. Furthermore, the reason that the agent is not justified in believing P is that she is not justified in taking R to be a reason for P in the first place. Consider again the scenario described in *Poker*. Here Jessica believes that John has a king on the basis of the fact that he raised when there was a king open-face on the table. In order to be justified in believing that John has a king—indeed, in order to be justified in taking John's raising to be a reason to believe he has a king—Jessica would have to know, or justifiably believe, that John plays by Strategy. This fact would have to play some role in her belief calculus. Given that it does not, her conviction that John has a king is bound to strike us as ungrounded—the outcome of mere omission, ignorance, or a failure to conceive of the real possibilities in the game. This, even though her reason is in-point-of-fact conclusive (John would not have raised unless he had a king). The

same is true, if to a lesser degree, in the case of *Chemistry (Modified)* and *Behavior*. Given that, *e.g.*, S does not know what would happen to the indicator if immersed in C_3 , and given that this possibility is a ‘live’ one, S is not fully justified in taking the indicator’s change in color as a reason for believing the solution is C_1 .⁹

The point, I think, is more forcefully brought out when we compare examples like *Poker* to typical Gettier cases. In both cases, the agent gets things right only as a matter of luck, and in so doing, fails to know what she is said to know. But in Gettier cases, the agent *is* justified in believing what she does, and gets lucky in that what she believes turns out to be true for other reasons. In *Poker*, on the other hand, the agent is *not* fully justified, and gets lucky in that the proposition she believes turns out to be true *for the very reason she believes it*. The examples show that this second scenario is not as paradoxical as it sounds. One can get lucky not just in believing the right thing, but in believing it for the right reasons.

In discussing Peter Unger’s account of knowledge, Dretske says that “absolute justification”—or the kind of justification required for knowledge—requires being able to rule out the ‘relevant’ alternatives to the proposition believed.

The social or pragmatic dimension to knowledge, if it exists at all, has to do with what counts as a relevant alternative, a possibility that must be evidentially excluded, in order to have knowledge. It does not change the fact that to know one must be in a position to exclude all such possibilities. It does not alter the fact that one must have, in this sense, an optimal justification - one that eliminates every (relevant) possibility of being mistaken. [Dretske , 1981]

But what is it to rule out alternatives to what one believes? Dretske is, as far as I can tell, never fully explicit about this, and I think we can now see that his language in the passage is indicative of a real ambiguity. Is it that relevant alternatives must be somehow “evidentially excluded,”—note the passive tense—or that the agent himself must be in a position to “eliminate” all such possibilities? We have chosen

⁹In fact, there are two ways to think about cases like *Chemistry (Modified)*. According to the first, S is not justified *at all* in taking the indicator’s change in color to indicate that the solution is C_1 , and is only justified in taking it to indicate that the solution is *not* C_2 , or is *either* C_1 or C_3 . According to the second, S is *partly* justified in taking the indicator’s change in color to indicate that the solution is C_1 , but is not *fully* justified in this, given that the solution could have been C_3 . (Those who are attracted to the second reading would perhaps motivate it by saying that the indicator’s change in color makes the C_1 alternative more likely than it was before the test.) I do not think we need to choose between these two readings. It is enough to note that either way, there is *something* to do with S’s justification that goes sour here—something in the way of her justification that makes it not adequate to a knowledge attribution. I will express this by saying that the agent is not *fully* justified in taking R to be a reason for P.

above to think about ruling out alternatives as a *competence* of some kind on the part of the agent—as having to do with the way in which an agent reasons about those possibilities—and I think this not only gets the requirement on knowledge right, but captures our intuitive judgments about when an agent is, and is not, able to rule out certain possibilities. (On this understanding of ‘ruling out,’ Jessica is *not* able to rule out the possibility that John has a straight, and S is *not* able to rule out the possibility that the solution is C_3). If this is right, then we have here a deeper way of understanding what goes wrong in the examples presented above. Justification requires ruling out relevant alternatives, where this is most plausibly understood as a competence on the part of the agent. But, as the examples show, one can rely on in-point-of-fact conclusive reasons without having any particular competence—without, that is, being in a position to reason with respect to these possibilities in the way that knowledge demands. And so, one can rely on what turn out to be conclusive reasons for P without either being justified in believing P, or knowing that P is the case.

7 *Having Conclusive Reasons*

“Does this mean that *having conclusive reasons* is not sufficient for reasons-based knowledge?” you ask. This sounds implausible. After all, in saying that someone has conclusive reason to believe P, we seem to assert that her grounds are *as good as they come*—that her grounds leave no room for reasonable doubt. But then how could it be true that in having conclusive reasons for P, a person may yet fall short of knowing that P is the case? What more could we reasonably demand of any putative knower?

If *having* conclusive reasons to believe P is just believing P on the basis of a reason that is in-point-of-fact conclusive, then in view of the examples above, we are forced to say that having conclusive reasons for P is *not* sufficient for knowing that P is the case. One can rely on in-point-of-fact conclusive reasons, and yet fall short of knowledge, as we’ve already seen. But to reverse the point, if *having* conclusive reasons *is* sufficient for knowledge, then what it is to have conclusive reasons cannot be (simply) to believe P on the basis of reasons that are in-point-of-fact conclusive.

Consider the following purportedly valid argument. (I think that something very much like this is at the heart of Dretske’s analysis of knowledge in “Conclusive Reasons,” although Dretske himself does not make this explicit.)

(P1) S knows that P if S has conclusive reason to believe that P.

(P2) A reason R is conclusive for P if and only if R would not be the case unless

P were the case.

(C) S knows that P if S believes P on the basis of R and R would not be the case unless P were the case.

The examples in Sections 2 and 4 show that the conclusion of the argument is incorrect. S can believe P on the basis of a reason that satisfies the counterfactual condition in (C), and yet fail to know that P is the case. But then we must either reject one of the two premises, or reject validity.¹⁰ I've argued already that I see no reason to reject (P2). This specifies correctly the relationship that has to hold between R and P if R is to be conclusive for P under the given circumstances. (P1), on the other hand, states that having conclusive reasons is sufficient for knowledge, and this too seems—at least on one way of hearing this phrase—intuitively plausible. When we assert that a person has conclusive reason for believing P, we claim that her epistemic justification with respect to that proposition leaves nothing to be desired. Given the things she knows, there is no room for doubt that P is the case. It would be strange if for all that, this person could fail to know that P. But if (P1) and (P2) are true, and (C) is not, this leaves us with only one option: rejecting validity. The inference to (C) is not truth-preserving.

If this is right then, looking at the two premises, the explanation must lie in the fact that to have conclusive reasons for believing P is not the same as believing P on the basis of reasons that are in-point-of-fact conclusive. In other words, (P1) and (P2) are concerned with different things. (P1) is concerned with what it is for a given agent to *have* conclusive reasons, as this relates to whether or not that agent *knows* that P is the case. It is a claim about the strength of the agent's epistemic position vis-à-vis the proposition P. (P2) is concerned, on the other hand, with what it is for a given reason to be in-point-of-fact conclusive for P, where that is a relationship that obtains independently of what any given agent knows or believes. Simply put, the argument is invalid because these two notions are different. To *have* conclusive reasons for P, is not simply to believe P on the basis of reasons that are in-point-of-fact conclusive.

We began with the claim that knowing that P on the basis of reasons is a matter of having conclusive reasons to believe that P. If the argument of this section convinces, then we can still accept this claim. Where knowledge is based on reasons, those reasons must be conclusive; and likewise, where a person has conclusive reasons to believe that P, her belief counts as knowledge. What we should reject is the claim that what it is to have conclusive reasons to believe that P is to rely on

¹⁰'Validity' as I use it here is just necessary truth-preservation. An argument is valid if it is impossible for the premises to be true and the conclusion false. I do not mean to suggest that the argument is valid in virtue of its *logical form*—or that anyone takes this to be the case.

reasons for P that are in-point-of-fact conclusive. In particular, we should reject any account according to which ‘having conclusive reasons for P’ is reduced to the obtaining of a certain counterfactual relationship between the proposition P and a person’s reason, R.

8 Justification in Taking R to be a Reason

But what is it, then, to *have* conclusive reasons to believe that P?

Presumably, an accurate account of this more proprietary sense of conclusive reasons will steer clear of the counterexamples above. It will not predict that S knows that the chemical solution is C_1 but is not able to rule out the possibility that it is C_2 , or that S knows that K is happy, but is not able to rule out the possibility that K is sad. One easy way to ensure this would be to simply specify that having conclusive reasons for P is (or requires) being able to rule out some set of alternatives to P. This, however, would be quite unilluminating. It would leave the connection between the more proprietary sense of *having* conclusive reasons, and reasons themselves, entirely unexplained. Indeed, on such a view, it would look as though having conclusive reasons has nothing to do with *reasons*.

A more promising strategy would be to allow that believing that P on the basis of in-point-of-fact conclusive reasons is at least a necessary condition on having conclusive reasons for P. We could then go on to ask what further conditions need to be added in order to yield an adequate account of what it is to have conclusive reasons in the proprietary sense. This at least has the advantage of allowing for a tight connection between having conclusive reasons to believe, and reasons themselves. But what extra conditions are needed? In the counterexamples we considered above, it seems appropriate to say that while the agent’s reason is in-point-of-fact conclusive, it is not a conclusive reason *for her*. What is left out here isn’t anything about the reason itself (after all, it is in-point-of-fact conclusive), or anything about the agent’s state of belief (the agent believes, without a trace of doubt, that P is the case). Rather, it is something to do with the agent’s relationship to the reason *as a reason*. In view of the discussion of justification in Section 6, it seems natural to say that an agent has conclusive reasons to believe that P only if she is *justified* in taking her reason *to be a reason* for P in the first place.

Consider, then, the following account:

(HCR) S has conclusive reasons for believing that P if:

- (1) S believes P on the basis of R;
- (2) R would not be the case unless P were the case;

- (3) S knows that R is the case;
- (4) S is (fully) justified in taking R to be a reason for P.

(HCR) avoids the counterexamples we worried about above. In *Poker*, Jessica is not justified in taking the fact that John raised to be a reason for believing he has a king. Her belief, as we said, is grounded in her own ignorance of the game; someone who was more competent at poker could easily point out why, in general, someone's raising is *not* a reason to infer that that person has a pair. Because Jessica does not satisfy condition (4), she does not have, according to (HCR), conclusive reason to believe that John has a king. This is plainly the right verdict. Moreover, if we take seriously the notion of justification, then I think similar arguments apply in both *Chemistry (Modified)* and *Behavior*. In *Chemistry (Modified)*, for example, the fact that the indicator turns to blue when immersed in C_1 but not C_2 allows S to rule out the possibility that the solution is C_2 . But supposing, as we did, that C_3 is a prominent alternative and that S knows nothing about what happens to the indicator when immersed in C_3 , S is not fully justified in taking the indicator's change in color to be a reason for believing the solution is C_1 . Again, there are two ways to look at the situation (see Note 9). On the first, S is not at all justified in believing that the solution is C_1 , given the fact that she cannot rule out the C_3 possibility. On the second, she is partly justified in believing that the solution is C_1 , but fully justified only in believing that it is *either* C_1 or C_3 . Whichever way we look at things, S fails to meet condition (4) in (HCR), and so, on that account, does not have conclusive reason to believe that the solution is C_1 .

Although this is not the place to defend a theory of justification, there is one important negative point worth making. We cannot here understand what it is for S to be justified in taking R to be a reason for P in terms of a demand that S have a *conclusive reason* for taking R to be a reason for P. Quite apart from any worry this would give rise to about infinite regress in (HCR), a quick look at our intuitions in the Gettier cases shows that this would be overly demanding. Take, for example, the case of Smith and Jones who apply for the same job (see Section 6). We judge, in that example, that Smith is justified in believing that Jones will get the job that both men applied for. Moreover, we are told that Smith believes this on the basis of the fact that he has been told as much by the company president. But supposing, as we do in the example, that Jones will not in fact get the job, Smith's reason is not in-point-of-fact conclusive. Indeed, Smith has this reason despite the fact that Jones will *not* get the job (and so, presumably, the relevant counterfactual, "Smith would not have heard from the company president that Jones would get the job unless Jones were to get the job," is false.) If we want to hold on to our intuition that Smith is justified in believing that Jones will get the job (and so, too, that he is justified in believing that the man who will get the job has ten coins in his pocket),

then we cannot understand justification as requiring having a conclusive reason to believe. This is the case not just for this particular Gettier example, but for the Gettier examples quite generally.

As I've indicated, these considerations seem to point in favor of adopting (HCR) as a necessary condition on having conclusive reasons in the proprietary sense that we are after. But is (HCR) also sufficient? Unfortunately, this is not the case; the sufficiency claim falls prey to examples similar in many respects to the Gettier cases. Consider the following scenario:

KFC. Augie drives to his local KFC, intending to pick up some chicken for dinner, but as he goes past, he sees that the electric 'CLOSED' sign is illuminated. He concludes that KFC is closed, and with grave disappointment moves on to the Jack in the Box down the street. Does Augie have conclusive reason to believe that KFC is closed? Before deciding, the reader should be informed that for the past three days, the electric 'CLOSED' sign has been turned off at this particular branch. The reason for this is that the manager of the store took a bribe from the head of a crime syndicate allowing the crime boss to break in at night and steal the contents of the register without setting off the building alarm. The secret arrangement between the crime boss and KFC manager was that the manager would give the go-ahead signal by illuminating the 'CLOSED' sign on the night of the break-in. Of course, the manager would never give the go-ahead signal at a time when the restaurant was still open.

The fact that the 'CLOSED' sign is illuminated is an in-point-of-fact conclusive reason to believe the restaurant is closed. After all, the manager would not have given the go-ahead signal at a time when the restaurant was still open for business. Moreover, Augie believes that the restaurant is closed on the basis of the fact that the sign is illuminated, and he is, presumably, justified in taking this to be a reason in the first place. (Here we rely on the reader's intuitive judgments about what is and is not justified, and not on any theoretical commitments regarding the nature of justification.) So Augie satisfies (HCR). Nevertheless, it seems wrong to say that Augie has conclusive reason to believe that KFC is closed. Why? Although he relies on what is an in-point-of-fact conclusive reason, and is justified in taking this reason to be a reason, it's conclusiveness is located in facts that are, as it happens, entirely unknown to Augie. Augie takes the illuminated sign to be a reason for believing the restaurant is closed only because he thinks the sign is performing its usual function. This example is what we might call an 'amplified' Gettier case—a case in which, although Augie is justified in believing that KFC is closed, and moreover believes it on the basis of a reason that is in-point-of-fact conclusive, he still does not have conclusive reason to believe that KFC is closed (and hence, does

not *know* that KFC is closed).

The example shows that while incorporating condition (4) in (HCR) moves us significantly closer to an adequate account of what it is to have conclusive reasons to believe, it does not quite close the distance.

9 Conclusive Reasons to Believe

Let us summarize, briefly, where we have gotten to up to this point. We saw that purely external modal constraints like (CR) cannot guarantee anything about the way in which the agent reasons with respect to possibilities that she must rule out in order to count as knowing that P. In the examples we considered, although the agent relied on a reason R that satisfied (CR), she was not (fully) justified in taking R to be a reason for P. In the previous section, we attempted to fix the problem by requiring that the agent not only believe P on the basis of an in-point-of-fact conclusive reason, but that she be justified in taking her reason to be a reason in the first place. But we saw that this was not enough: In *KFC*, Augie is justified in taking his reason to be a reason, but nevertheless fails to know.

We can put the difficulty we are up against roughly as follows. So long as the conclusiveness of S's reason R is located in facts that S herself does not take into account—so long as the very circumstances which make R conclusive for P are not the ones that in some sense *explain* S's taking R to be a reason for P in the first place—S does not count as knowing that P is the case. S's grounds for P strike us as inadequate for knowledge. S may be entirely *reasonable* in believing what she does in the way she does. After all, who could object to Augie's taking the illuminated 'CLOSED' sign to signal that the restaurant is in fact closed? And yet it seems right to say that Augie in this case gets lucky. The restaurant *was* closed, his reason *was* conclusive, but only in virtue of facts that Augie himself was not aware of: facts which played no role in Augie's coming to believing what he did.¹¹

Put in slightly more abstract terms, the problem is that there can easily be a gap between the facts in virtue of which the agent is *justified* in taking his reason to be a reason, and the facts in virtue of which his reason is in-point-of-fact conclusive.

¹¹It is interesting to note here that both 'sensitivity' and 'safety' predict that Augie *does* know that the restaurant is closed. Had the restaurant been open, the 'CLOSED' sign would not have been illuminated, so Augie would not have believed that the restaurant was closed—at least not on the basis of his current method. Thus Augie's belief is sensitive. On the other hand, in the nearby possible worlds where Augie believes the restaurant is closed on the basis of the illuminated 'CLOSED' sign, it really is closed. So Augie's belief is safe. There are, of course, many nearby possible worlds in which Augie does not believe the restaurant is closed (because the 'CLOSED' sign is not illuminated), but in which the restaurant really is closed. Neither 'sensitivity' nor 'safety' is sensitive to this fact.

In *KFC*, the fact that makes it the case that Augie is justified is just that ‘CLOSED’ signs in general function in the way that they do. But the facts that make the illuminated sign an in-point-of-fact conclusive reason for believing that the restaurant is closed are entirely different. They have to do with the secret arrangement between restaurant manager and crime boss, and their agreed-upon go-ahead signal. The scenario sketched in *KFC* is not a one-off example. Putting things in these terms shows that there is a formula for constructing all kinds of scenarios that fit the same pattern. Simply take an agent who believes that P on the basis of a reason R, which generally *is* conclusive for P—and which any reasonable believer would take to be a reason to believe that P. Then rig the background circumstances so that in this particular case, R does not bear the usual relationship it does to P. Now alter the circumstances again, so that for entirely different reasons, R is nevertheless in-point-of-fact conclusive for P. *Voilà*: You have a case structurally similar to *KFC*, in which the agent fails to know although he is justified in taking his reason to be a reason, and although his reason is in-point-of-fact conclusive.

If the diagnosis of the failure of (HCR) is correct, then what is needed is presumably some condition that ensures a link between the facts in virtue of which the agent is justified in taking her reason to be a reason, and the facts that make it the case that her reason is in-point-of-fact conclusive. Ideally, these facts should coincide. In other words, the very facts that the agent takes into account in taking her reason to be a reason should be the ones in virtue of which her reason is in-point-of-fact conclusive. This suggests, I think, the following account of what it is to have conclusive reasons for P:

(HCR*) S has conclusive reasons for believing that P if:

- (1) S believes P on the basis of R;
- (2) R would not be the case unless P were the case;
- (3) S knows that R is the case;
- (4) S is (fully) justified in taking R to be a reason for P.
- (5) The facts in virtue of which R is conclusive for P are ones that S takes into account in taking R to be a reason for P (they *explain*, in this sense, S’s taking R to be a reason for P).

What should S’s epistemic position be with respect to those facts and circumstances which she ‘takes into account’ in relying on her reason, R? In my view, the requirement here should be quite stringent. If S is to come to know that P is the case in this way, we must require not just that S believes, or justifiably believes, that these circumstances obtain. Rather (and here I am afraid I will lose some readers) we must require that S *knows* that these circumstances obtain. Although this

requirement is strong, it need not lead to a vicious regress. So long as there are other ways of knowing not based on reasons (in particular, so long as one can have perceptual knowledge of certain facts or states of affair), then S need not satisfy conditions (1) - (5) with respect to those facts or circumstances which she 'takes into account' in believing what she does. In going this route, we put ourselves in roughly the same position as Dretske is in with respect to conditions (A) - (C) in his account of reasons-based knowledge (see Section 2). As Dretske says, the reference to knowledge itself in condition (C) can be eliminated by recursive application of the definition of 'know' until one arrives at a reason R "about which it does not make sense to say that an agent knows that R." (Here we should replace "knows that R" with "knows that R on the basis of reasons.")

Certainly, much work remains to be done on just what it is to 'take into account' a given circumstance in coming to believe what one does. This, I take it, is the real challenge for any more fully developed theory along the lines of (HCR*). In the meantime, the conditions as stated point us in what I think is a fruitful direction. Where propositional knowledge is based on reasons, those reasons must be conclusive. Rightly understood, having conclusive reasons is not simply a matter of the objective relations that hold between the agent's reason and the proposition she believes on that basis. What matters is our relationship to our reason *as a reason*—the facts and circumstances that we take into account in taking our reason to be a reason in the first place.

References

- DRETSKE, F. (1971). *Conclusive Reasons*. *Australasian Journal of Philosophy*, vol.49, no.1, 1–22.
- DRETSKE, F. (1972). *Contrastive Statements*. *Philosophical Review*, vol.81, no.4, 411–437.
- DRETSKE, F. (1981). *The Pragmatic Dimension of Knowledge*. *Philosophical Studies*, vol. 40, no.3, 363–378.
- GETTIER, E.(1963). *Is Justified True Belief Knowledge?* *Analysis*, vol. 23, 121–123.
- GOLDMAN, A.(1996). *Discrimination and Perceptual Knowledge* *Australasian Journal of Philosophy*, vol.131, no.1-3, 287–301.
- LACKEY, J.(2008). *What Luck is Not*. *Australasian Journal of Philosophy*, vol. 86, no. 2, 255–267.

- LATUS, A. (2000). *Moral and Epistemic Luck*. ***Journal of Philosophical Research***, vol. 25, 149–172.
- LEWIS, D.(1996). *Elusive Knowledge* ***Australasian Journal of Philosophy***, vol.131, no.1-3, 287–301.
- NOZICK, R.(1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- PRITCHARD, D. (2004). *Epistemic Luck*. ***Journal of Philosophical Research***, vol. 29, .
- PRITCHARD, D. (2007). *Epistemic Luck*. Oxford: Oxford University Press.
- SOSA, E. (1986). *On Knowledge and Context*. ***The Journal of Philosophy***, vol.83, no. 10, 584–585.
- SOSA, E. (2004). *Relevant Alternatives, Contextualism Included* ***Philosophical Studies***, vol.119, 35–65.
- UNGER, P. (1968). *An Analysis of Factual Knowledge* ***Journal of Philosophy***, vol.65, 157–170.
- WILLIAMSON, T.(2002). *Knowledge and its Limits*. Oxford: Oxford University Press.
- YAMADA, M. (2011). *Getting it Right by Accident* ***Philosophy and Phenomenological Research***, vol.83, No. 1, 72–105.