

## Being Knowledge (or: Monism of Knowledge)

Michael Della Rocca

(Note to readers: In the interest of space and time, I have omitted the last section of this paper (as well as short parts on some specific theories of knowledge I criticize). The last section (of which only the title – “Kicking away the Ladder” – remains here) is a part in which things get even more extreme, but there will, I believe, be plenty to discuss already, and of course I’d be happy to talk about the last section too should the occasion arise. Looking forward to seeing you! – Michael)

### I. Introduction

A widespread explanatory demand seems to inspire many of the pilgrims who make their way through contemporary epistemology. Expressed in one way, the demand is the question: what is it in virtue of which an individual knows that a particular state of affairs obtains or that a particular proposition is true? Or: what is it in virtue of which a given state that a knower is in amounts to a state of knowledge? Or, more grandly, what is being of knowledge? Or, most simply: what is knowledge? A fundamental task – perhaps *the* fundamental task – of epistemology is to illuminate what it is to be a state of knowledge, what makes something a state of knowledge, and what it is in virtue of which that state differs from any state of non-knowledge.

Epistemologists of very different stripes seem to embrace something like this explanatory demand. Thus, Jennifer Nagel opens her excellent introductory work on epistemology with the question, “What is knowledge?”, and she later restates the guiding question as, “What is it for you to know something, rather than merely believe it?”<sup>1</sup>

Barry Stroud, who is always deeply attuned to crucial methodological issues, sees his task and the task of epistemologists in general as one of “understanding human knowledge”. He elaborates this explanatory aim by saying that “What has come to be called ‘epistemology’ is the attempt to explain how we know the things we know”.<sup>2</sup> (p. ix). Elsewhere he continues with the same theme: in epistemology, “we want...some kind of *explanation* of our knowledge – some account of how it is possible”.<sup>3</sup> Finally, consider this general question that Stroud raises: “What exactly is knowledge, and how do human beings know the sorts of things they have to know in order to live the kind of lives they lead?”<sup>4</sup>

Similarly, Robert Nozick asks, “How is knowledge possible?” and announces that “we seek a hypothesis to explain how, even given the skeptic’s possibilities, knowledge is possible”.<sup>5</sup>

Or consider Jason Stanley’s unambiguous proclamation:

---

<sup>1</sup> Nagel, *Knowledge: A Very Short Introduction*, pp. 1, 4.

<sup>2</sup> Stroud, “Introduction”, p. ix.

<sup>3</sup> “Scepticism and the Possibility of Knowledge” p. 4, italics in the original.

<sup>4</sup> “Understanding Human Knowledge in General, p. 123.

<sup>5</sup> Nozick, *Philosophical Explanations*, pp. 167, 169.

A central part of epistemology, as traditionally conceived, consists of the study of the factors in virtue of which someone's true belief is an instance of knowledge.<sup>6</sup>

In the Preface, he says that his book's thesis concerns "what *makes* someone's true belief a case of knowledge" (p. v, my emphasis). Stanley issues a clear call to epistemologists to meet the explanatory demand.

Even Timothy Williamson, whose revolutionary views stand opposed in many respects to those of the theorists I have just quoted and who regards the notion of knowledge as primitive and unanalyzable, seeks to explain knowledge. He says that he aims to offer "a modest positive account of the concept" of knowledge (p. 33) and that a "reflective understanding of knowledge is possible" (p. 33), though not an analysis.<sup>7</sup> Williamson's account is meant to illuminate knowledge, to help us see what it is and, in this way, he also feels the pull of the explanatory demand. Rationalist that I am, "explanation" is not too strong a term for me, but perhaps it is too strong a term for Williamson or for you. No matter. I'm not going to fight about the word. It's enough to say this: Williamson certainly seeks to give an account of – to illuminate – knowledge and so, from this point of view we can speak of Williamson, as well as of the others that I have quoted, as motivated by what we might call an illuminative demand with regard to knowledge, the demand that knowledge should somehow be illuminated.

Countless other epistemologists give expression to the explanatory or illuminative demand. Each of them realizes that answering this demand enables an epistemologist to grasp the difference between knowledge and non-knowledge. Or at least that's the goal. In this paper, I examine a number of leading contemporary theories of knowledge and show how they all – while embracing the explanatory demand – fail to meet it in subtle and not so subtle ways. I will, in fact, eventually show that *any* theory of knowledge which makes a certain seemingly harmless assumption – as all or nearly all theories of knowledge do – inevitably fails to meet the explanatory demand, fails, in a way contrary to its stated aim, to illuminate knowledge. The assumption in question is that a state of knowledge is differentiated, that it somehow stands in relations. Whether a view that gives up this apparently innocuous assumption and thus adopts a non-relational view of knowledge is able to meet the explanatory demand – well, we'll just have to see....

From the conclusion that all or most standard theories of knowledge inevitably fail to meet the explanatory demand, we will be able to arrive at something like a monism of knowledge, i.e. something like the view that there are no states of knowledge differentiated from other things, rather – instead of differentiated knowledge – there is simply knowledge. To put it roughly, all is knowledge. Or perhaps even better: knowledge is. And I will further argue that, in this light, we must reconceive the entire enterprise of epistemology and, perhaps, of philosophical inquiry more generally.

To begin this examination of the failure of current theories, I want to divide such theories into two broad categories: what I will call "building block" theories and "knowledge-first" theories. As should become clear, this distinction cuts across other,

---

<sup>6</sup> Stanley, *Knowledge and Practical Interests*, p. 1.

<sup>7</sup> In a different context, Lewis draws a similar contrast between an account and an analysis in "New Work for a Theory of Universals", p. 352.

perhaps more familiar distinctions among epistemological theories, such as the internalism/externalism distinction, and the foundationalism/coherentism distinction. I do not claim that my division between building block and knowledge-first theories is exhaustive, but the distinction does cover at least a large swath of contemporary epistemology. Indeed, the positive, monistic view of knowledge that I will ultimately gesture toward is not a version of either a building block view or a knowledge-first theory, but here I'm getting ahead of myself.

Building block views seek to shed light on knowledge – to explain what it is to be knowledge – in terms of such conditions as truth, belief, and justification, and perhaps some further condition. These conditions are the building blocks of knowledge, as it were. This is the approach of the long-suffering (and long suffered!) project of analyzing the concept of knowledge that has dominated so much of analytic epistemology. But it's not the case that building block views are committed to there being an analysis of knowledge. For example, Stanley clearly holds a building block view and is explicitly not committed to knowledge being analyzable (p. 88). More generally, one can even reject the analytic/synthetic distinction and with it the notion of conceptual analysis, yet still – in building block fashion – believe that knowledge is dependent on features such as truth, belief, justification (and more) that do not presuppose the notion of knowledge. All that one needs for a building block account is the dependence of knowledge on other features. One does not need an analysis of the concept of knowledge.

The – or an – alternative to a building block view is a knowledge-first view. Such a view does not seek to illuminate knowledge by appeal to features such as belief, truth, and justification (and more) that are regarded as somehow prior to knowledge. On a knowledge-first view, knowledge is primitive or basic and not dependent on those other features. This approach seeks instead to account for other epistemic notions in terms of knowledge. Thus the nature of belief, justification, and assertion, for example, can all be accounted for in terms of knowledge (instead of vice versa). In taking this tack, the aim is to reveal important structural connections (e.g. in the case of Williamson's version of a knowledge-first view: externalism, anti-luminosity, the denial of the KK principle, and the thesis that one's evidence is one's knowledge) that serve to illuminate knowledge without analyzing it. Through charting these connections involving knowledge, we gain the illumination of knowledge that Williamson seeks. This is Williamson's way of seeking to meet the explanatory demand – not through an analysis of knowledge, but rather through a “reflective understanding” of knowledge.

Williamson's view need not be the only version of a knowledge-first strategy. For example, we can say that knowledge comes prior to belief but, unlike Williamson, have an internalist conception of belief and also insist that knowledge requires knowledge that one knows. Whether such an alternative version of a knowledge-first view can be developed is an interesting question. Williamson's version, however, is remarkably well worked out, and my point here is that it provides a useful contrast to building block views even though in the end, as I will argue, the contrast is not as significant as Williamson takes it to be.

## II. Building Block Views

Building block views tend to take it for granted that, for a given state to be one of knowledge, that state must be a belief that is true. The question for building block theorists has generally concerned what conditions, if any, must be added to these two in order for the state to amount to knowledge.<sup>8</sup> The challenge here is to spell out the way the belief must be related to the truth in order for the belief to count as knowledge.

For some time now, the relation in question has been thought to be a justificatory relation, the relation whereby the belief is connected to the truth in such a way that the belief is justified. And so we have the justified true belief account of knowledge. But then – with one of the most consequential and cited papers in recent philosophy – along came G..., along came Ge..., along came Get... – I can't bring myself to say the name. Suffice it to say that devastating counterexamples to the justified true belief account were introduced, and philosophers went back to the ranch to try to assemble the building blocks of knowledge once again, to try to articulate the way in which a true belief must be related to the truth in order to count as knowledge.

Let me consider a potential response to these counterexamples that shall remain nameless. In these counterexamples, the belief is justified but somehow the justification is not of the right kind for the belief to count as knowledge. This suggests the following potential response: knowledge is not justified true belief, but rather is true belief that is justified *in the right way*. This may not be a response that anyone has actually given to the problem of explaining what knowledge is.<sup>9</sup> And these inadequacies should be obvious – or, if not, they will soon be. But the inadequacies in this account will be instructive. Without specifying what the right way is, we can say only that this way is right because it is the kind of justification needed for the subject to have knowledge. Thus the account would be:

S knows just in case S's true belief is justified in the way  
that is needed for that belief to amount to knowledge.

This would be a circular account of knowledge. Knowledge would be explained in terms of knowledge, and that can hardly be a satisfactory answer to the explanatory demand.<sup>10</sup>

Such a circular account is evidence of a kind of emptiness or unilluminatingness that is also found, e.g. in the iconic statement – problematic in so many ways – that is sometimes attributed to John Wayne and often repeated as a kind of parody: “A man's gotta do what a man's gotta do.”<sup>11</sup> This statement is most likely not intended, of course, as an account or explanation of any kind. In saying what he does, Wayne (or the Wayne character or the parody of the Wayne character) may be attempting to demonstrate what a man's gotta do or to embody in his laconic directness what a man's gotta do. (I'll return to this point at the very end of the paper.) But, if intended as an account of what a man's gotta do, this statement would be empty, uninformative, and unilluminating. Similarly, on

---

<sup>8</sup> See again the quote from Stanley, p. 1.

<sup>9</sup> Though, as we'll see, Goldman's initial response comes pretty darn close to this one.

<sup>10</sup> Many building block theorists acknowledge that these theories must avoid such circles. Thus, e.g., Sosa agrees that a satisfactory theory of knowledge must specify conditions in which no epistemic status plays any role at all (Sosa, *Reflective Knowledge*, pp. 175-76).

<sup>11</sup> Actually, Wayne may not have said precisely this sentence in any movie, though he said things that were close. Oddly enough, one of the first uses of the phrase in this precise wording comes from the character George Jetson in the animated television show, *The Jetsons*.

the simple account of knowledge just offered, knowledge has gotta be what knowledge has gotta be: that's what saying that knowledge must be justified in the way required for knowledge amounts to. Such a circular, uninformative account is what I call a John Wayne moment. And it would be extremely disappointing, I would think, for a philosopher to propose such an "account" when the stated, shared goal of epistemology is to provide illumination as to the nature of knowledge. Yet this is precisely what happens time and time again: as we'll see, all or most prominent accounts of knowledge amount to John Wayne moments in philosophy.<sup>12</sup>

Other building block views are more sophisticated than the skeletal view I have used as an illustration. But most or even all other such views fail in the way that my sample view does: they invoke the notion of knowledge itself in an account of knowledge (or they get into a regress), and they wind up experiencing John Wayne moments. My argument for the general conclusion will, in the first instance, be inductive: by looking at a variety of examples of prominent building block views and showing how they founder upon precisely this rock of circularity, I will offer an inductive basis for thinking that building block views of knowledge in general succumb to this problem. I will follow this up with a more or less a priori argument for precisely the same conclusion, an argument that, moreover, identifies the underlying source of this problem of circularity (or regresses). You might say that given the latter, a priori argument, the former, inductive argument is not needed. But I see the arguments as mutually supporting. Each of the arguments is controversial, and so each can use the external support that a separate argument for the same conclusion provides.

Most building block views, when they get into circles of the kind that I'm interested in do so in subtle ways.<sup>13</sup> But, in some, the circularity is right there on the surface, brazenly hanging out. This is the case with one of the earliest responses to the original counterexamples that come from Edmund L. I've-forgotten-his-last-name. This is the response found in Alvin Goldman's causal theory of knowing which is remarkably similar to the toy account I just presented. Goldman says that S knows that *p* just in case "the fact that *p* is causally connected in an 'appropriate' way with S's believing *p*" ("A Causal Theory of Knowing", p. 369). Goldman characterizes the appropriate way in terms of perception and memory both of which are, of course, ways of knowing and thus in as much need of illumination as knowledge itself is. Goldman has, at best, merely pushed the problem back one step and not resolved the question of what it is to be knowledge. It seems then that Goldman (at this stage) has no non-circular way of characterizing knowledge, and so we have a John Wayne moment in the early literature on the analysis of knowledge.<sup>14</sup>

Many other building block views take a different approach and rely crucially on the notion of relevant alternatives. As I will now argue, these views also get into circularity problems. To see the difficulties that the notion of relevant alternatives poses,

---

<sup>12</sup> Amazingly, certain philosophers appeal to such a "right way" condition in the causal theory of action: an action is a bodily movement caused by a belief and desire in the right way (Davidson, "Freedom to Act"). Attempts to spell out the right way in philosophy of action have not been successful (despite the work of Smith and Peacocke and many others), and similar attempts are no more successful in the case of epistemology as we'll now see.

<sup>13</sup> Weatherson wouldn't get into a circle. I'm not addressing people who bite the bullet and reject counterexamples. Weatherson, "What good are counterexamples".

<sup>14</sup> For this kind of criticism of the causal theory, see Nagel, p. 52. (Brewer?)

it will be helpful to split up building block accounts of knowledge into contextualist and invariantist accounts. On contextualist views, whether it is true to say that S knows a given proposition  $p$  depends on the standards for knowledge that are in force in the context of attribution, the context of the person (who may be S herself) who is evaluating whether S knows  $p$ . It is a key feature of contextualist views of knowledge that the standards for knowledge – i.e. what is required for S to know – may shift from one context of evaluation to another. Often the semantical mechanism of such shifts is regarded as a function of the allegedly indexical nature of the verb “knows”.

By contrast, an invariantist epistemologist regards the standards for knowledge as remaining constant across different contexts of evaluation. Whether it is true to say that S knows depends not on the context of attribution, but rather on the situation of S herself.

Building block views can be either contextualist or invariantist. An invariantist building block theorist would hold that whether S knows  $p$  depends on facts such as whether  $p$  is true, S believes  $p$ , and how the belief is related to the truth that  $p$ . In a given situation, whether S knows does not depend on the context in which S is being assessed for knowledge. According to a contextualist building block view, whether it’s true to say that S knows depends on the true belief, the connection between it and truth, and *also* on the context of evaluation. Different kinds of connection to the truth may be required in different contexts of evaluation in order for S to count as knowing  $p$ .

Let’s turn to the problems that the notion of relevant alternatives creates for building block views. I will begin with invariantist building block views.

Some early versions of a relevant alternatives view were invariantist.<sup>15</sup> On such accounts, S’s knowledge that  $p$  requires not just the true belief that  $p$  but also that S be in a position to eliminate alternative situations in which  $p$  is false.<sup>16</sup> Specifically, S is required to eliminate not all alternatives to  $p$ , but only those that are relevant in some way. Thus consider an example of a kind famously offered by Fred Dretske. Marion visits the zoo and sees what he thinks to be a zebra before him. Does Marion know that there is a zebra there? On a relevant alternatives view, in order to know in this situation, Marion must be in a position to eliminate certain alternatives to  $p$ , to the claim that a zebra is present, e.g. the alternative that what’s before him is not a zebra but a giraffe (or an elephant, gazelle, etc.). However, some alternatives to being a zebra are typically not – so the account goes – relevant, e.g. the alternative that the animal before Marion is a mule cleverly disguised to look like a zebra, or the possibility that there is no zebra at all before Marion and his experiences as of a zebra are, in fact, deceptively generated in Marion who is merely a brain in a vat.

The problem I want to focus on arises when we consider just what it is for an alternative to be relevant on an invariantist building block view. Relevant alternatives views are motivated to invoke this notion in part because they seek a way of handling

---

<sup>15</sup> Dretske and Nozick would be examples of invariantist building block folks.

<sup>16</sup> Or alternatives to knowing that  $p$ . I won’t continue to make the qualification about alternatives to knowing that  $p$ , a qualification important in some contexts, though not here. Also: the language of elimination or ruling out may problematically presuppose knowledge in a way that would render the account circular. It might be thought that an alternative is eliminated only because one *knows* it not to obtain. Lewis characterizes “elimination” in a way that may not presuppose knowledge, see Lewis, “Elusive Knowledge”, p. 424. Alternatively, one can avoid any problematic talk of elimination entirely by appealing simply to whether in nearby possible situations, S has false beliefs as to whether or not a given proposition is true. This is DeRose’s preferred, “double safety” approach (new book, chapter 7, section 5).

skepticism. A building block theory would collapse into skepticism if *all* alternatives – even the painted mule alternative or the brain in a vat alternative – were deemed relevant and needed to be ruled out in order for S to know *p*.<sup>17</sup> The account of knowledge offered by an invariantist relevant alternatives theorist is this:

S knows that *p* just in case S truly believes *p* and S is able to eliminate all relevant alternatives to *p*.

For invariantist relevant alternative theorists, safeguarding the knowledge of *p* in this way may come at the cost of denying the deductive closure of knowledge, and, in particular, at the cost of denying that one knows that the skeptical possibilities don't obtain. That may seem implausible, but Nozick and Dretske go to great lengths (unsuccessfully, in my view and in the view of many others) to defend the denial of closure. When we get to contextualism, we will see versions of the relevant alternatives view that affirm closure. In any event, the difficulties associated with the denial of closure aren't my main objection to invariantist relevant alternative views. I am focusing on a deeper worry.

And this is the worry expressed in the question: how does the relevant alternatives theorist specify which alternatives are relevant? Until that is done, the account of knowledge is not complete, the explanatory demand is not met. Sometimes, relevant alternatives theories appeal to modality and theories of counterfactuals. The giraffe counterpossibility is relevant because if *p* were false, something like the giraffe possibility (or an elephant possibility, etc.) would be the case. But the painted mule possibility or the brain in a vat possibility is not what would be the case if *p* were false. It's not that the painted mule or brain in a vat situations are impossible. They are, let's stipulate, possible. Rather, the point is that such possibilities are not what would obtain were *p* to be false. Such possibilities are – to use a metaphor from theories of possible worlds – too remote from the actual world to be relevant, whereas the giraffe possibility is not too remote. The world where a giraffe is in front of Marion is a nearby possible world, but the painted mule or brain in a vat situations obtain only in remote worlds.

Fair enough, but what makes it the case that the brain in a vat or painted mule possibilities are too remote, and the giraffe possibility is not too remote? The answer seems to be that the brain in a vat or painted mule possibilities are too remote because, if we had to rule them out in order to know, then we would not know *p*, and skepticism would be true; whereas the giraffe possibility is not too remote because requiring that it be ruled out wouldn't lead to skepticism. Thus the relevance of an alternative or the nearness or remoteness of a possibility is to be characterized in terms of what is required for avoiding skepticism. In other words, a relevant alternative is an alternative the ruling out of which is relevant for securing *knowledge* of *p*. And an irrelevant alternative is an alternative that does not need to be ruled out in order to know *p*. The relevance of an alternative is characterized in terms of knowledge.

This is explicit in many statements of the relevant alternatives view. Thus Dretske says,

---

<sup>17</sup> Dretske makes this point in "The Pragmatic Dimension of Knowledge" p. 372, where he says that for the skeptic the set of relevant alternatives is identical to the set of alternatives, or, in his terminology, the relevancy set = the contrasting set.

its being an eagle, a Mallard, or a Loon *are* members of the relevancy set since if the bird watcher could not eliminate these possibilities (sufficient unto knowing that it was not an eagle, a Mallard, or a Loon) on the basis of the bird's appearance and behavior, then he would not know that it was a Gadwall ("Pragmatic Dimension", p. 371).

Here the relevance of an alternative is characterized in terms of knowledge ("sufficient unto knowing"). Notice also that the alternative is relevant *since* (as Dretske says) if one did not eliminate it one would not *know*. An alternative is thus relevant *because* of the connection between knowing *p* and ruling out that alternative. Similarly, Dretske indicates that an alternative is irrelevant because one need not rule it out in order to know:

*for purposes of assessing someone's knowledge* that this is a table, certain alternative possibilities are simply not relevant. ("Pragmatic Dimension," p. 368, my italics).

Similarly, DeRose characterizes the notion of relevance of alternatives in terms of knowledge:

this raising of epistemic standards consists in expanding the range of relevant alternatives to what one believes, that is, the range of alternatives that one must be in a position to eliminate in order to count as knowing. ("Solving the Skeptical Problem", p. 8n11)

Notice that the notion of relevance is glossed here in terms of knowing.<sup>18</sup>

Goldman's reliabilism – which he sees as a successor to his earlier causal theory – invokes relevant alternatives in much the same spirit as Dretske's theory does. Although Goldman, unlike Dretske, is neutral between contextualism and invariantism (?), he too cashes out his notion of relevance in knowledge-presupposing ways. For Goldman, the relevance of an alternative is a function of whether the alternative is unusual enough to count as relevant *for a non-skeptic* ("Discrimination and Perceptual Knowledge", p. 787, my italics). And there are shades in Goldman's reliabilism of his earlier problematic appeal to *appropriate* causal relations when he specifies that S's propensity to form a belief must have an appropriate genesis, where the appropriateness of a genesis turns on whether it is "enough to support knowledge" (p. 789). The apple of reliabilism doesn't fall far from the tree of the causal theory.<sup>19</sup>

---

<sup>18</sup> DeRose does not endorse this account of relevance here because it is expressed in terms of the notion eliminating possibilities, a notion of which, as I indicated, DeRose is rightly suspicious. However, for DeRose, whether one knows does still turn on what one believes as to whether or not a given proposition is true in nearby worlds, and, as we will see, for DeRose, which worlds count as nearby is a function, in part, of what it takes for a subject to know.

<sup>19</sup> There may be another way that reliabilism faces a problem stemming from the notion of relevance. This is the generality problem. How are processes individuated? There are lots of different processes by which a belief may be arrived at, only some of which are reliable. If we say that all processes are relevant, we

One might think that the appeal to relevance can be saved from circularity by characterizing relevance not in terms of knowledge (as Dretske, DeRose, and others do), but merely in terms of the similarity of possible worlds to the actual world, or in terms of the nearness of possible worlds to the actual world. On this way of talking, skeptical possibilities are relatively unlike and distant from the actual world, and non-skeptical alternatives are relatively similar to or close to the actual world. And it is because of these differences in terms of similarity and closeness – and not because of considerations turning on knowledge itself – that the distinction between relevant and irrelevant alternatives is introduced. However, it's not at all clear that these claims about similarity and closeness are true. After all, the skeptical painted mule and brain in a vat possibilities are more like actuality than is the non-skeptical giraffe possibility in at least one respect: in the painted mule possibility and the brain in a vat possibility as well as in the actual world, Marion believes that there is a zebra before him. But that is not the case in the giraffe possibility. Why isn't this respect of similarity between the actual situation and the skeptical possibilities significant enough so as to make it the case that the skeptical scenarios are more similar to actuality than the giraffe scenario? I don't see how one can answer this question without appealing to what is required for Marion to know that there is a zebra before him.<sup>20</sup>

In general, we can see that many invariantist views that rely on a distinction between relevant and irrelevant alternatives invoke the notion of knowledge itself in their accounts of knowing. They say, in effect, that:

S knows *p* in part in virtue of the fact that S is able to rule out those alternatives that S needs to rule out in order to know *p*.

These accounts that invoke relevant alternatives, just like Goldman's original causal theory and just like my toy example that appeals to being justified in the right way, all say, in effect, that S knows because S does what S needs to do in order to know. Such an account clearly does not meet the explanatory demand. It amounts to what I have called a John Wayne moment in philosophy.

So things do not look good for invariantist building block theories of knowledge in general. Do contextualist building block accounts meet with more success? Most such views – like many of their invariantist counterparts – rely on some form of the distinction between relevant and irrelevant alternatives. S knows just in case one rules out the contextually relevant alternatives. Such a view can in a way be friendly to skepticism because it allows that in some contexts skeptical possibilities are relevant and that,

---

will land in skepticism. Only the skeptic seems to avoid an unprincipled line between relevant and not relevant processes. The skeptic may seem in this respect to be more principled. However, Comesaña makes the good point that the generality problem is a problem not just for reliabilism, but also for all or most building block views. Comesaña is insouciant about the problem because for him it is widespread. But the pervasiveness of a problem is not a reason for complacency. Bishop seems to have a better attitude: he seems to grasp the challenge to all theories of justification. [fill this out]

<sup>20</sup> More needs to be said here, of course, about just why in this situation an account of similarity that does not appeal to what is required for Marion to know is not available, but it is worth noting that Lewis, the great deployer of the notion of similarity among possible worlds, in effect reaches a similar conclusion concerning similarity in his contextualist account, as we will see shortly.

therefore, it is true to say that we don't have knowledge. But in other, more ordinary contexts, only some alternatives are relevant and other alternatives may, to use David Lewis' terminology, be "properly ignored", and we may in these contexts truly say that we have knowledge.

The problem for the contextualist – at least the problem that I want to focus on – is that the distinction between relevant and irrelevant alternatives is characterized in a way that presupposes knowledge. Thus, as we saw already, DeRose characterizes relevance for standard relevant alternatives theorists in terms of knowledge ("Solving the Skeptical Puzzle", p. 8n11). Later in the same wonderful paper, DeRose specifies what it takes to "count as knowing" (p. 36) according to his own contextualist account,

Context...determines how strong an epistemic position one must be in to count as knowing. Picture this requirement as a contextually determined sphere of possible worlds, centered on the actual world, within which a subject's belief as to whether P is true must match the fact of the matter in order for the subject to count as knowing. (p. 36)

Similarly, in an influential early contextualist paper, Stewart Cohen says:

We need to say more about the criteria of relevance. Whether S knows *q* will depend on whether any alternatives to *q* are relevant – whether the conditions are such that S's epistemic position with respect to any alternatives precludes knowledge of *q*. Thus the criteria of relevance should reflect our intuitions concerning under what conditions S does know *q*. (Cohen, "How to be a Fallibilist", pp. 101-02)

In these passages, DeRose and Cohen explicitly characterize relevance in terms of knowledge.

Similarly, in a seminal contextualist paper, Gail Stine suggests that a certain alternative is irrelevant because if it were relevant we would be prevented from giving a certain answer to the skeptic:

To allow it [a certain alternative] as relevant seems to me to preclude the kind of answer to the skeptic which I sketched in the opening paragraph of this section. (Stine, "Skepticism, Relevant Alternatives, and Deductive Closure", p. 253)

Here, the relevance or irrelevance of an alternative seems to be a function of what we know or don't know.

Finally, Lewis in a classic paper that DeRose labels "a prominent contextualist manifesto" (DeRose, *The Case for Contextualism*, p. 29), explicitly says that his "Rule of Resemblance" by which he determines which alternatives are relevant is *ad hoc* (or, more

precisely, has an *ad hoc* exception), and he does not know how to fix this situation. This is because, as he recognizes, there is no way to characterize the resemblance that is relevant to knowledge other than in terms of knowledge itself. For Lewis, what it is to know *p* is to eliminate all possibilities in which *p* is false that must be eliminated in order to secure knowledge of *p*. This is the account that, Lewis says, must be given in order to avoid “capitulation to skepticism” (Lewis, “Elusive Knowledge”, p. 430), and it “makes good sense in view of the function of attributions of knowledge (p. 430). Lewis laments the *ad hoc* nature of his account: “What would be better, though, would be to find a way to reformulate the Rule so as to get the needed exception without the *ad hocery*. I do not know how to do this” (p. 430). But apparently he does not lament the *ad hocery* enough to reject this account and try some other approach.<sup>21</sup>

A contextualist’s appeal to nearness of possible worlds to the actual world and to the similarity of possible worlds to the actual world is similarly *ad hoc*. A defender of contextualism might say that the skeptical scenarios describe possible worlds that are less similar to the actual world than are non-skeptical alternatives, and thus because of this similarity the former alternatives are irrelevant and the latter relevant. But, as I said before, it’s not at all clear that these claims about similarity and closeness are true, and I suspect that one cannot justify such claims independently of appealing to what is required for knowledge of *p*. Indeed, this point is precisely what Lewis’ lament about his *ad hoc* exception amounts to.

Jonathan Schaffer offers a distinctive form of contextualism called “contrastivism” which, despite its virtues, also offers no non-circular way to characterize the alternatives that are relevant in a given situation. For Schaffer, as for standard contextualists, the truth of simple knowledge ascriptions of the form “S knows that *p*” depends on context. But for contrastivism, the mechanism of the context dependence is different from the mechanism at work in standard epistemic contextualist theories. For the latter theories, the truth of knowledge ascriptions is, as I mentioned, often regarded as a function of the indexical nature of terms such as “knows” which are indexed to specific epistemic standards by contextual features. By contrast (!), for a contrastivist the context-dependence of “S knows that *p*” is a function of an implicit contrast condition which is supplied by context. Context thus determines that the content of the attribution is really something like “S knows that *p* rather than *q*”, e.g. “Marion knows that there is a zebra before him rather than a mule”. So the mechanism of the context-dependence of the truth of knowledge attributions is different from the mechanism in standard contextualism, but there is still context-dependence and so contrastivism counts as a form of contextualism.

Given this context-dependence, the truth of “S knows that *p*” depends on S being able to rule out certain possibilities that contrast with *p*, but does not require that one be able to rule out other such possibilities. What makes a possibility one that S needs to rule out in order to know *p*? That is, what makes a contrasting possibility one that S needs to rule out? Schaffer claims that, unlike Lewis’s account, his own way of handling the difference between relevant and irrelevant alternatives is not *ad hoc*. Perhaps, then, Schaffer may avoid the dreaded circularity other contextualist views face. But this is not the case: for Schaffer specifies which alternatives need to be addressed by appealing to

---

<sup>21</sup> For some more detailed criticism of the *ad hoc*-ness of Lewis’ account, see Vogel, “The New Relevant Alternatives Theory”.

“the role of knowledge in inquiry” (p. 267). Here Schaffer says, in effect, that to know  $p$ , one need not rule out the painted mule possibility because needing to do so would conflict with the way we employ knowledge. For Schaffer, it is “unnatural and unaccommodating” to invoke certain possibilities in certain contexts (p. 263) because doing so would undermine our knowledge and be incompatible with the role of knowledge in inquiry. Here again, we have the same kind of circularity that besets other contextualist views.

All these contextualist views rely upon a notion of relevance of alternatives that – as these views sometimes acknowledge – cannot be characterized other than in terms of knowledge. This result shows that these accounts of knowledge amount to something like this:

S knows that  $p$  only if S rules out the alternatives to  $p$  that need to be ruled out in order for S to know that  $p$ .

In other words, S knows only if S does what she needs to do in order to know. That is, S knows because S knows. And all I can say in response is: Well, sure, but such an account is hardly illuminating. We have again the failure to meet the explanatory demand. We have again a John Wayne moment in philosophy.

[Discussion of subject-sensitive invariantism and MacFarlane’s view.]

### III. A Knowledge-First Account

This critical survey of at least some building block views amounts, I believe, to a strong inductive case against building block views as failing to meet the explanatory demand. I haven’t (yet) shown that all building block views *must* fail to meet the explanatory demand because of such problems with circularity (or regresses), but one can suspect on inductive grounds that there is an underlying problem that makes such failures inevitable for building block views.<sup>22</sup> I will later give a more or less a priori argument for just this conclusion, but at this point, given the poor track record of building block views, we may, in an effort to meet the explanatory demand, be led to take seriously alternatives to a building block approach. Perhaps the most prominent such alternative is Williamson’s knowledge-first approach.

The knowledge-first approach describes a possible family of views all of which place knowledge first in a sense that I will specify. Williamson’s own view is, in principle, only an example of such a view. Nonetheless, I will, for the most part focus on Williamson’s version of a knowledge-first view.

As I noted at the outset, Williamson, like most or all other epistemologists, seeks to meet the explanatory demand. He seeks to give some kind of account of knowledge which at least illuminates the notion of knowledge and enables us to distinguish knowledge from ignorance and from non-knowledge generally. But the way in which

---

<sup>22</sup> Others note – for reasons different from mine – that there is an inductive case against building block views. See Williamson pp. 30, 63, 91. And see Cassam who quotes Williamson while making a similar point, “Given that attempts to analyse the concept of knowledge have succumbed to a ‘pattern of counterexamples epicycles’ (Williamson, p. 31), it is not unreasonable to conclude that all such attempts suffer from a common underlying defect” (Cassam, “Can the Concept of Knowledge be Analysed?”, p. 21). See also Craig, *Knowledge and the State of Nature*. (?)

Williamson endeavors to meet the explanatory demand exhibits considerable – dare I say, true – grit: it undertakes a reversal of the procedure found in building block approaches. Instead of understanding knowledge in terms of other epistemic notions such as belief and justification, we are, for Williamson, to understand other epistemic notions in terms of knowledge. The idea is that, in so doing, we illuminate knowledge itself even while treating knowledge as a primitive. In Williamson’s version, this illumination comes in the form of a characterization of knowledge as the most general factive or truth-entailing mental state. This account is explicitly not an analysis (p. 36). For Williamson, the notion of knowledge is prior to the notion of a factive mental state or at least the notion of a factive mental state is best understood in terms of knowledge (see p. 39 where the notion of a factive mental state operator (FMSO) is understood in terms of knowledge.) But we can see how knowledge is connected to other mental states, and thus we shed light on knowledge. Not only is knowledge characterized as the most general factive mental state, but there is also the illuminating outline of structural connections or the lack of structural connections with other notions. Among such connections or non-connections are the denial of the KK principle, the denial of luminosity, and the claim that one’s evidence is identical to one’s knowledge. What is important, for Williamson, is that these connections do somehow follow from the nature of knowledge. Thus, for Williamson, much light can be shed on knowledge once we give up the analytic project and begin our account with knowledge treated as basic.

How successful is Williamson at meeting the explanatory demand? Let’s look more closely at what he takes to be the factors that, apparently, illuminate knowledge. The crucial claim for Williamson is, as we’ve seen, that knowledge is the most general factive mental state. On the basis of this characterization, Williamson draws further conclusions, all of which are intended to shed light on knowledge.

Let’s focus first on the connection between knowledge and being the most general factive mental state. Williamson recognizes that, in the end, this way of meeting the explanatory demand is circular. For him, being a factive mental state is best understood in a way that is not independent of knowledge itself. Being a factive mental state cannot be understood independently of knowledge – if it could, then Williamson’s characterization would not only be illuminating but would threaten to be an analysis of knowledge after all, an outcome Williamson would not welcome. So it must be the case that for Williamson the notion of a factive mental state is best understood through the notion of knowledge itself. Williamson indicates as much when he summarizes what it is for an operator to be a factive mental state operator (FMSO). He offers “three principles”:

- (18) If Phi is an FMSO, from ‘S Phi’s that A’ one may infer ‘A’.
- (19) ‘Know’ is an FMSO.
- (20) If Phi is an FMSO, from ‘S phi’s that A’ one may infer ‘S knows that A’. (p. 39; Williamson’s numbering)

Here Williamson indicates that being a factive mental state is to be understood in part in terms of the notion of knowledge. He makes a similar claim when he says,

If we could isolate a core of states which constituted ‘pure mind’ by being mental in some more thoroughgoing way than knowing is, then the term ‘mental’ might be extended to knowing as a mere courtesy title. On the conception defended here, there is no such core of mental states exclusive of knowing. (p. 6)

For Williamson knowledge is one of the core states in terms of which we can best understand the notion of the mental.

In this light, we can see that Williamson’s characterization is ultimately circular: the notion of knowledge is illuminated through – explained in terms of – its connection to the notion of being a factive mental state, but this notion in turn is illuminated through the notion of knowledge. How is this roundabout performance illuminating? In the end, this seems to be just another John Wayne moment in epistemology. Knowledge has gotta be what knowledge has gotta be. Knowledge is understood in terms of knowledge.

Similarly, the other purportedly illuminating connections Williamson draws between knowledge and other structural epistemic features, such as the denial of the KK principle and the identification of evidence and knowledge, also presuppose or are understood through the notion of knowledge itself. So once again we are faced with a circle of dubious value as far as illumination is concerned.

Here’s another way to make what is essentially the same point. Although for Williamson knowledge is a primitive notion and has no analysis, he nonetheless claims to be able to draw connections between the nature of knowledge and features such as being a general factive mental state, etc. Since knowledge is thus primitive, how can we find the hooks in it, as it were, which enable us to make these connections between the nature of knowledge and these other features? The question we must ask – and that Williamson does not and cannot answer – is this: *how* is it that these purportedly illuminating connections hold? Without specifying this “how”, we do not have an informative account of knowledge. All we know – or apparently know – is that the nature of knowledge *somehow* dictates that knowledge is the most general factive mental state and *somehow* dictates that knowledge has the other structural connections Williamson outlines. But without insight into these “how’s”, no genuine advance in understanding knowledge has been achieved.

One might think that this kind of charge of emptiness is unfair to Williamson. After all, Williamson is merely attempting to shed light on knowledge through the explanatory role it plays or can play. This approach, one might say, is a specific version of a perfectly legitimate attempt to clarify a phenomenon by outlining its theoretical role, something we do whenever we Ramsify in our theory-building. Thus in a functionalist philosophy of mind, for example, we might elucidate belief as that mental state which, together with certain desires and other beliefs, leads to certain actions. And we might understand desires as those mental states that, together with beliefs and other desires, lead to action, etc. Here belief is understood in terms of its connection to desires and actions both of which are understood in terms of belief. Yes, one might argue, this account of belief may be circular at some level, but it’s still illuminating as long as the circle is large enough. Similarly, one might say, knowledge, on Williamson’s view, is understood in terms of its connection to (in particular, its priority over) belief and justification (etc.)

which are understood in terms of knowledge. In this case too, there is circularity, but the circularity is OK and there is illumination as long as the circle is large enough.

However, I think that in each of these cases the circularity is problematic and, indeed, undermining of any genuine illumination.<sup>23</sup> In the belief case, if we are in the dark as to what belief is and if we seek illumination as to what belief is, then being told that belief is understood in terms of desire which is understood in terms of belief is not going to help us with our illumination deficit. How can we make use of these connections for the purpose of illuminating belief if these connections can, in the end, be understood only in terms of the very notion – viz. belief – that we are trying to illuminate? Similarly, in the case of Williamson’s appeal to the theoretical role that knowledge plays in an economy of other notions: if the connections to other phenomena – such as to justification, luminosity, etc. – can, in the end, be understood only in terms of knowledge itself, then how can these alleged connections relieve us of our illumination deficit when it comes to knowledge? We gain illumination here only if we already understand what knowledge is, but, of course, we do not already understand what knowledge is – otherwise why would we be embarked on this inquiry at all?

What Williamson offers here is what might be called a dormitive virtue explanation of knowledge, a dormitive virtue response to the explanatory or illuminative demand. In poking fun at Aristotelian explanation in terms of natures or powers, Molière (actually, one of his characters) quips that opium has the power to put one to sleep because it has a dormitive virtue. This explanation is unilluminating because it purports to shed light on the power of a thing simply by appealing to the nature of that thing. What’s missing is a way of indicating *how* the opium has this power. In this light, the insight of Molière and of much of the early modern critique of Aristotelian natural philosophy is that mere appeals to natures are not illuminating. What is required is a way – beyond a simple invocation of the nature of opium – of saying *how* it is that the opium performs this wonderful task. That is, we require an *independent* way of saying how opium does this job. Williamson’s account is, obviously, much more sophisticated, but at bottom the purportedly illuminating step is one in which he seeks to illuminate knowledge through knowledge itself. Without spelling out, in independent terms, *how* it is that a given state gets to be a factive mental state (let alone the most general such state), we don’t really have an understanding of what it is to be knowledge. More generally, although there may be the connections Williamson alludes to, if all we know or think we know is *that* there are these connections without seeing in independent terms *how* they follow from the nature of knowledge, then we are none the wiser when it comes to understanding knowledge. Thus we have no more insight into knowledge than Molière’s characters have into the workings of opium.<sup>24</sup>

So not only is Williamson’s account – because of its circularity – yet another unilluminating John Wayne moment, it is also – because it rests on a brute appeal to the nature of knowledge as acting in a certain way we know not how – an unilluminating Molière moment in philosophy.<sup>25</sup>

---

<sup>23</sup> Yes, I’m doubling down.

<sup>24</sup> In a similar way, Cassam criticizes Williamson’s minimalism as leaving knowledge ultimately unexplained (p. 30). Stroud?

<sup>25</sup> In the same way, I would respond to Samuel Elgin’s important suggestion – not Williamson’s – that knowledge, though it lacks a full analysis, has a partial analysis (say, in terms of justification, truth, and

[Williamson, Reflective Equilibrium, and “Good” Circles.  
IV. A More or Less A Priori Argument

The knowledge-first and building block approaches are – for all their differences – in the same boat: each of them (or the examples of these approaches that we have examined), though in different ways, fails to meet the explanatory demand that drives contemporary epistemology. And because these views insist on a demand which they themselves do not meet, these views also suffer from a kind of internal incoherence. For these theories, there is no illuminating account of knowledge, and, ultimately or even *ab initio*, what knowledge has gotta be just is what knowledge has gotta be. The Duke would be proud.

There is thus an inductive case now not only against building block views but also against knowledge-first views, and so we see the failure of most ordinary theories of knowledge. But, as I’ve mentioned, there could be, for all I’ve said so far, building block or knowledge-first views other than the ones I’ve discussed that do meet the explanatory demand. So it’s not clear that this kind of problem is inevitable. Further, and more importantly, it’s not clear what the mechanism of the problem that building block and knowledge-first views face is, what the source of this problem is. I would like to give now a more or less a priori argument that simultaneously addresses both of these issues. First, the argument shows that building block views and knowledge-first views in general and indeed any view which, like these, makes a certain seemingly innocent assumption inevitably fail to meet the explanatory demand. Given that ordinary views insist on the explanatory demand, this failure – this inevitable failure – is damning indeed. Second, this more or less a priori argument has the virtue of isolating the common culprit – the varmint as John Wayne might say – behind the failures of building block and knowledge-first views. Appropriately enough – since all the views which I am about to undermine trade on the explanatory demand – my argument will turn on the notion of explanation.

In making the explanatory demand or in insisting that knowledge be illuminated, epistemologists tend to search for what might be called the knowledge-making relation. They ask, in effect, what is the relation between the state that is in fact a state of knowledge – call the state “K” – and some item or items that make K a state of knowledge. Different theories of knowledge have different proposals as to what the knowledge-making relation is. On many accounts, the relevant relation is a relation between K and other items such as other states that make it the case that the knower’s beliefs are justified in the right way or, more generally, connected to the truth in the right way. (This would be the approach of building block theories of knowledge.) Or, the relation may be a relation between K and a certain nature – K’s own nature – which is that in virtue of which K is a state of knowledge. In this case, the relation would be a kind of internal relation. (This would be the approach of knowledge-first theories.) More

---

belief). (See Elgin, “Knowledge is a Justified True Belief”.) Issues of unacceptable circularity arise here if the notion of justification can be understood only in terms of knowledge. Further, given that knowledge has no full analysis, it’s not clear that we can be confident that the nature of knowledge as a whole dictates that knowledge requires justification. As before, we need to see not only *that* this connection holds but also *how* it holds. But, given the lack of a full analysis, this connection must remain opaque. Once again, the account fails to provide the insight that we seek and fails to meet the explanatory demand.

simply and more generally, the relevant relation may be a relation of distinction between K and states of non-knowledge, and the epistemological work in this vein seeks to shed light on what differentiates K from states of non-knowledge in general.

In this light, we can see that most accounts of knowledge proceed upon the seemingly harmless presupposition that what it is for something to be a state of knowledge requires a distinction – a relation – between the state of knowledge, K, and some items. The epistemologist thus tends to presuppose that we are trying to understand what it is to be a *differentiated* state of knowledge, a state of knowledge that somehow stands in relations. This focus on differentiated or relational knowledge will be significant in what follows.

In identifying the knowledge-making relation, R, there are three exhaustive and mutually exclusive options:

- (1) R is primitive. That is, the relation, R, between K and items is not grounded in or explained by anything. There is nothing in virtue of which K stands in R to the items.
- (2) R is not primitive, but is, rather, ultimately grounded in or explained by or ultimately holds in part in virtue of a relation to something other than K.
- (3) R is not primitive, but is, rather, ultimately fully grounded in or explained by or ultimately holds in virtue of the nature of K alone. That is, it follows from the nature of K alone that K is a state of knowledge.

I say “ultimately grounded” (in (2) and (3)) in order to contrast such grounding with what might be called intermediate grounding. It may be that something *x* is explained by some item *y* and *y* is fully grounded in *z*. In this case, *x* would be grounded in *y* but not ultimately grounded in *y*. Instead, *x* would be intermediately grounded in *y* and ultimately grounded in *z*. This distinction between ultimate and intermediate grounding will soon be important.

R – the knowledge-making relation – must fall into one of the three categories just outlined. Which option, if any, will be the lucky winner?

The first option can be eliminated simply because relations are, by their nature, grounded in something. They are, as is widely agreed, grounded in, at least, one or more of their relata. A relation exists at least in part *because* of one or more of its relata. Or – to put things in a not quite equivalent way<sup>26</sup> – whether a relation obtains depends on the existence of its relata. Relations are thus grounded and are not primitive. Thus option (1) is not really an option.

Unlike option (1), option (2) is widely endorsed among epistemologists and, indeed, is definitive of what I have called building block theories. According to this option, the crucial relation is grounded in part in certain items other than K, items to which K is related in such a way as to render K a state of knowledge. The other items in question may be K’s causes (specified in a certain way) or, additionally or alternatively, the other item may be mental states which the knower is in and which are related to K in such a way that K amounts to knowledge. In general, those engaged in the project of analyzing knowledge take up option (2): what it is to know is to be broken down, analyzed, into individually necessary and jointly sufficient conditions – conditions that

---

<sup>26</sup> Ref. to Vertanen.

detail the kind of relation that K must bear to other items in order for K to amount to knowledge.

However, while the lack of an explanation of R offered by a proponent of option (1) is *immediately* unsatisfactory in this context, the explanation that a proponent of option (2) offers is *ultimately* unsatisfactory. To see why, I will now offer a Bradleyan argument against the reality of relations in general, but an argument keyed to the specific case of knowledge and the knowledge-making relation. I will return to the general anti-relational argument later in the paper.

If the knowledge-making relation, R, is, in part, grounded in something other than K and, in part, grounded in K, then the fact that K amounts to a state of knowledge holds, in part, because K stands in a relation of partial grounding to R. The other relatum or relata of the knowledge-making relation also stand in a relation of partial grounding to R. Thus, not only is R – the knowledge-making relation – grounded in part in K and in part in the other relatum or relata, it is also grounded in part in a further relation between, e.g., K and R itself. This is a relation of partial grounding between K and R.

To see why the demand for this further relation emerges, note that R – the knowledge-making relation – and indeed relations in general as such and by their nature, seem to be dependent entities. Relations are not, as I indicated in eliminating option (1), groundless. Rather they are, by their nature, dependent on their relata. Thus, the relation by its nature demands not only that its relata be in place, but – because a relation is by its nature dependent on its relata – the relation also demands that there be a relation of partial grounding between the relation and each relatum individually. Because, in general, a thing depends on any item built into its nature, and because, as I’ve just claimed, a relation of partial grounding between R and K is built into R’s nature, it follows that R depends on the relation of partial grounding. Call this further relation R’.

We can fancifully illustrate this point as follows. Let’s say that I am relation R and I don’t yet exist, but I’m a pushy not-yet-existent relation and so I demand of God that God create me. (I am, as Leibniz would say, one of the striving possibles.) In demanding that I come into existence, I demand not only that there first be K (and perhaps another relatum of the relation R), but I also demand the path – the relation R’ – between K and me, i.e. between K and R. (I may also demand that there be a path between the other relatum and me.) So the relation, R, depends not only on K, but also on R’, the relation of partial grounding between K and R.

It may be that R’ comes for free with K and the other relatum, i.e. it may be that the relation R’ between K and R comes automatically, as a “free lunch” to use Armstrong’s term,<sup>27</sup> as soon as K and the other relatum or relata are in place. In such a case, R may not depend *ultimately* on R’. But, even so, R depends non-ultimately on R’: R depends intermediately on the path from K to R. Or, to put the point in yet other terms: R holds at least in part because each of K and the other relatum or relata *do their part* in making it the case that R holds. And, K’s part is to partially ground R, to be related to R by R’, the relation of partial grounding. So given the commitment to the grounding of the knowledge-making relation, R depends on at least K and R’ (and also on some other relatum and, perhaps, some relation of partial grounding between that other relatum and R).

---

<sup>27</sup> David Armstrong, *A World of States of Affairs*, pp. 12-13. Bennett (“By Our Bootstraps”) adopts this strategy for avoiding a vicious regress.

Thus I disagree with Karen Bennett when she says, “Quite generally, when A grounds B, the ground of B is not A plus the grounding relation – the grounds of B are simply A” (“By Our Bootstraps,” p. 31). On the contrary, I say, if standing in a grounding relation is part of the nature of the grounded (as it is in the case of relations and as it would be for grounded items in general), then the grounded is grounded in part in the relation of grounding between the grounded and ground.

Given that R depends in part on R', our endeavor to understand the knowledge-making relation leads us to inquire: in virtue of what does the relation, R', of partial grounding between R and K hold? This question concerns a metaphysical dependence relation between R' and other items. There is also a parallel epistemological point: we cannot be said to have understood the knowledge-making relation, R, until we understand on what it is grounded, and so we cannot understand R until we understand R'. Thus we cannot understand R until we understand on what R' is grounded. So we ask: on what is R' grounded? Well, R' is a relation (of partial grounding) between K and R and, as such, it depends not only on K and on R but also on a relation of partial grounding between K and R' itself. (R' also depends on a relation of partial grounding between R' and R). Call the relation of partial grounding between R' and K, R'' – i.e. R double prime. Notice that in order to explain R and thus in order to explain R', it now turns out that we must explain R''. Our understanding of what makes K a (differentiated) state of knowledge is not complete until we have a grasp of the knowledge-making relation, R. But we cannot be said to grasp R until we grasp the relation, R', of partial grounding on which R depends. And we cannot be said to grasp R' until we grasp a further relation, R'', of partial grounding, and so on ad infinitum.

As my presentation of it already indicates, this regress is vicious. The initial relation R *depends* on a further relation, R', which in turn *depends* on a further relation, R''. Before one can explain R, one must *first* explain R'. But *before* one can explain R', one must *first* explain R'', and so on ad infinitum. One's explanation of R can thus never be completed and since, the explanatory goal – essential to the theory of knowledge – demands that we account for the knowledge-making relation, our explanatory goal cannot be reached, and, as Hume says in a related context, “such a discovery...cuts off all hope of attaining satisfaction” (*Treatise*, p. 266). Thus option (2) is, in the end, incompatible with the explanatory aims that all theorists of (differentiated) knowledge embrace.

Before moving on to option (3), let me address some potential objections to the line of argument I have just developed.

Many philosophers agree that this kind of regress is vicious.<sup>28</sup> But some disagree.<sup>29</sup> These are complicated issues worthy of further exploration. In any event, the problem I have isolated doesn't turn on whether there is a regress here and so doesn't turn on whether that regress is vicious, for there is another way to show that the grounding question we have asked leads to an intolerable problem. So forget about the regress and turn to another old favorite philosophical problem, a problem at work throughout this paper: circularity. For once we ask the grounding question, we can see that in addition to

---

<sup>28</sup> For example, Vallicella, “Relations, Monism, and the Vindication of Bradley's Regress”, Lewis, “Tensing the Copula”, and Leibniz (see the discussion in Della Rocca, “Violations of the Principle of Sufficient Reason (in Leibniz and Spinoza)”, p. 150).

<sup>29</sup> See Wollheim, p. 114, Gaskin, *The Unity of the Proposition*, Quine, “Responses” in *Theories and Things*, p. 183, Bennett, “By Our Bootstraps”.

generating a regress (vicious or not), our question also reveals a circle. Thus, as we saw, R is grounded in part at least in K and R'. But R' is grounded in (*inter alia*) its relata, and so R' is grounded in part in K and R. Putting these two grounding claims together and given the transitivity of grounding, we reach the result that R is grounded in part in R itself. This is because R is grounded in part in R' which is grounded in part in R. Thus R is grounded in part in R, and so we have a circle.<sup>30</sup>

Whatever one thinks of the regress in this case, it seems that in this context, the grounding of R – the metaphysical explanation of R – in terms of R itself is unacceptably unilluminating. Thus R – which stands for any relation here – can only be circularly grounded and cannot be made intelligible.

It is also important to note that the core of this argument doesn't turn on treating relations as objects in their own right. Bradley who purveys the kind of argument I have just made is sometimes criticized because in generating the regress, he invokes (e.g. at *Appearance and Reality*, p. 126n) relations as *things* that require other *things* to be related to them. One might argue that the above argument just misses the point that relations are precisely not things and so the regress (or circle) cannot even get going.<sup>31</sup>

However, this argument does *not* depend on treating relations as objects. For one way to see why this is so, let's focus not on objects as grounds and grounded, but on facts as grounds and grounded. So, ask not: what *object* or objects does relation R depend on; ask: on what *fact* or facts does the fact – call it F – that K is related to other items depend on? In asking this question, we do not presuppose that there are things that are relations. Just as a relation depends by its very nature on its relata, so too the relational fact, F, depends at least in part on the fact that K exists. And just as the relation between K and other items depends on the relation of partial grounding between K and R, so too the fact that K and other items are related depends in part on the fact, F', that the fact that K exists partially grounds the fact – F – that K is related to other items. Figuratively: the pushy fact that K is related to other items demands that there be a path of partial grounding from the fact that K exists to the fact that K is related to other items.

But what fact is this further fact, F', grounded in? F' is grounded in the fact that K exists and in the fact – F'' – that K exists partially grounds the fact that F' obtains, etc. This result is every bit as viciously regressive or every bit as circular as the result in the previous argument that seemed to treat relations themselves as objects.

An appeal to relational facts as opposed to relations as objects can avert another potential objection to the line of argument I am advancing. If one takes relations as objects, then a question arises: are these objects universals or particulars? It might be thought that if the relations are particulars – something like a trope, for example K's relation to another item – then there may be a problematic regress. But, it might also be thought, if we treat relations as universals, then there may be no problematic regress. For consider the relations of partial grounding that were at work in the regress R', R'', R''', etc. If relations are universals, then they can be wholly present in multiple places or times. Thus it might seem that the relations of partial grounding are really the same

---

<sup>30</sup> For a challenge to the transitivity of grounding, see Schaffer, "Grounding, Transitivity, and Contrastivity." For a good defense, see Litland.

<sup>31</sup> For discussion of this charge, see Hylton, *Russell, Idealism and the Emergence of Analytic Philosophy*, p. 48; Wollheim, *F.H. Bradley*, p. 114, Candlish, *The Russell/Bradley Dispute*, p. 168. Russell makes the charge in *Outline of Philosophy*, p. 263.

relation that is multiply instantiated, i.e.  $R' = R'' = R'''$ , etc. But if such identities hold, then it's no longer clear that there is any problematic regress.

Even if it is the case that invoking relations as universals would bypass the regress, the regress re-arises once we turn (again) to relational facts instead of relations as objects. For facts are not universals. They are not capable of multiple instantiation, i.e. they are more like tropes. Thus the fact that R is partially grounded in R' and the fact that R' is partially grounded in R'' are, though similar, nonetheless distinct facts, and so any regress here would involve a proliferation of distinct entities (the relational facts) with the earlier items in the series dependent on later ones.

The demise of option (2) signals the demise of the project of analyzing knowledge and the demise of the building block project in epistemology more generally. Thus, I agree with Williamson's conclusion that this project fails. However, Williamson's reasons for welcoming this demise are different from mine. Williamson's argument, as we've seen, is largely an account of what he sees as the theoretical advantages of adopting his knowledge-first approach which rejects such an analysis.<sup>32</sup> By contrast, my reasons for ushering the building block project out the door do not stem from any purported advantages of Williamson's knowledge-first approach. Indeed, as I've argued for the case of Williamson's approach and as I'm about to argue for the case of knowledge-first views in general, I deny that there are any such advantages at least insofar as satisfying the explanatory aim is concerned.

Williamson's view and knowledge-first views generally fall under my option (3), the only remaining option. According to this option, the knowledge-making relation, R, that K stands in is not primitive, nor does K stand in this relation because of something external to K. Instead a state of knowledge is a state of knowledge because it follows simply from its nature that it stands in the knowledge-making relation. R is thus, in one sense of the term, an internal relation, a relation that stems simply from the nature of a thing. It might be thought that the situation would look brighter for option (3) than for option (2) because, with (3), R is grounded in K *alone*. Here there is no other relatum external to K, and so, it seems, we do not have to appeal to any notion of partial grounding in the way that we had to with option (2) in a way that seemed to invite a regress of relations.

However, even with option (3), we have to appeal to a further relation in spelling out the grounds of relation R. If we don't appeal to a further relation as a ground of R, if we just say that R is grounded in K and leave things at that, then we will be saying that K stands in the knowledge-making relation R simply because that is the nature of K. But while such a pronouncement is nice, it is nothing more than another dormitive virtue kind of explanation, an empty appeal that is the mere shell of the kind of substantive explanation we were seeking. As I've already indicated, if we've learned anything from the early modern critique of Aristotelian explanations, it's that mere appeals to natures are *not* explanatorily illuminating. So if we are to preserve option (3), we need to say not just *that* R – the knowledge-making relation – is grounded in K's nature alone, but we also need to say *how* it is grounded in K's nature alone.

In terms of Williamson's view, the challenge is to spell out this relation by specifying *how* it is that a certain mental state is a factive mental state. Williamson acknowledges the need to spell out how R is grounded in K: "Of course, something needs

---

<sup>32</sup> See Ichikawa and Steup, "The Analysis of Knowledge", section 10.

to be said about the nature and significance of this matching [between mind and world].”<sup>33</sup> However, to unpack this “how”, we need to specify another relation –  $R^*$  – in virtue of which  $R$  is grounded in  $K$  alone. But since  $R^*$  is now the knowledge-making relation, we are led to ask the same Bradleyan question – “What is  $R^*$  grounded in?” – and we face the same Bradleyan challenges as before. Turning once again to Williamson’s view in particular which is a version of (3), we can see that, Williamson fails to meet the Bradleyan challenge by failing to explain how  $K$  is a general factive mental state. Williamson leaves it a mystery how  $K$  is state of knowledge.<sup>34</sup>

This argument against option (3) turns on a rejection of internal relations, a rejection of relations that stem simply from the nature of an object. Like the rejection of external relations in eliminating option (2), this rejection of internal relations is also Bradleyan in spirit. As is often overlooked, Bradley rejects internal relations as well as external relations.

The Bradleyan argument I just gave which eliminates options (1)-(3) is a more or less a priori argument for the result that our survey of leading theories of knowledge had given us inductive reason to suspect was the case. In reaching this conclusion, I am not doing anything so tame as raising the skeptical possibility that, for all we know, all or most of our beliefs may be false. Rather, I’m raising the deeper and prior question of whether we have a grasp of what it would be for us to be in a state of knowledge. That is, I am asking whether knowledge itself – the nature of knowledge – as we conceive it, or think we conceive it, is intelligible. And what I have argued is that, for all standard theories of knowledge – either building block or knowledge-first – knowledge itself is unintelligible. Knowledge, it may seem, like Parmenides’ non-being, cannot be and cannot be thought.

## VI. The Varmint and the Ascent

The problem that we are now facing has arisen only on the presupposition shared by Williamson, Goldman, DeRose, Stanley, and all or most epistemologists that states of knowledge stand in relations. In other words, the difficulty arises on the assumption that the knowledge we are trying (and failing) to understand is differentiated knowledge, i.e. a state of knowledge that is distinct from cases of non-knowledge or distinct from items such as justifications or mere beliefs or causes of states of knowledge, or, again, a state of knowledge that stands in an internal relation to its own nature. It is only on such a presupposition that we faced the exhaustive options (1)-(3) for dealing with the knowledge-making relation, options each of which is unacceptable. Indeed, I contend that it is the alleged relationality of knowledge that is the culprit – the varmint – responsible for the problems of circularity and regress to which each of the leading theories of knowledge have fallen prey.

But if, for this reason, we get rid of the problematic assumption that there are states of knowledge that stand in relations, that is, the assumption that there is relational knowledge, then we no longer face these problems. For these problems arise only on the assumption of the relationality of knowledge. The problems that relational knowledge

---

<sup>33</sup> *Knowledge and its Limits*, p. 40.

<sup>34</sup> In the same way, I would argue that Anscombe’s view of action which can be seen as, in some respects, an action-first view leaves it a mystery how an action is an action. Ref. to Levy, “Intentional Action First”.

faces do not arise for non-relational, undifferentiated knowledge. For the proponent of non-relational knowledge, although there are no states of knowledge as distinct from other states of knowledge and as distinct from things that are not knowledge, there may simply be knowledge. And if there is to be knowledge at all, it cannot occupy a place differentiated from and related to things: we cannot intelligibly place knowledge *within* reality in a place differentiated from other places. Thus, if there is to be knowledge at all, then, as we might put it, all is knowledge. The reality that epistemologists have been trying – unsuccessfully – to capture with their building block theories and knowledge-first theories which appeal to relational, differentiated knowledge is better captured by a view that rejects relations (even internal relations) between states of knowledge and things. This is a move that I have elsewhere called a Parmenidean Ascent: a rejection of a certain kind of distinction and, by transcending this distinction, an attempt to capture the reality that was only imperfectly expressed by appealing to such a distinction.

This view according to which there is simply knowledge, there is simply relationless, structureless knowledge, can be seen as a monism of knowledge. What we reach is something like a version of what Schaffer calls existence monism – the only thing is the cosmos itself – but a version of existence monism that is expressed in the idiom of knowledge. There is at most one thing, and this thing is a state of knowledge which does not contain within itself any multiplicity. Actually, I think that this view may be an even stronger form of existence monism than existence monism is ordinarily understood to be. This is because existence monism may allow for a distinction between the one fundamental concrete thing that is the cosmos itself and its non-concrete properties. The view I outline here – this monism of knowledge – would reject any such distinction between knowledge – the one thing – and its non-concrete properties.

I have argued that all is knowledge. But it might seem that that conclusion does not follow: for even if relational knowledge is problematic, all that follows is that there is no relational knowledge. That conclusion is compatible with either the claim that all is knowledge or with the claim that all is non-knowledge. Wouldn't this latter claim be an equally good way to get rid of relational knowledge? Perhaps, but my reason for preferring the claim that all is knowledge over the claim that all is non-knowledge is that non-knowledge is inherently relational. Non-knowledge is, as it were, defined in terms of what it is not, viz. knowledge. So non-knowledge – by its very nature – is relational. And thus, in a Parmenidean vein, I reach the conclusion that there cannot be, as it were, undifferentiated non-knowledge. However, there is, perhaps, no similar incoherence in undifferentiated knowledge. In particular, undifferentiated knowledge would bring with it none of the problematic varmint – relations – that have plagued us. We are looking, in other words, at a kind of pure knowledge. Undifferentiated knowledge – unlike non-knowledge – is purely positive and not at all relational and, as such, it is free from the paradoxes of relationality.

Actually I'm a bit reluctant to call this view a kind of monism at all, because it may be that to number a thing as *one* (as monism seems to do) presupposes that that thing stands in relations. It may be that there cannot be one thing without there being some kind of relation involving that thing. Thus because this thing that is knowledge is non-relational – as I have just argued – it may not be proper to call it “one” (or to refer to it as “this thing”). And so monism may not be the correct label for my view after all. (For similar reasons – having to do with relationality – I think Spinoza recognizes that God is not

one and so Spinoza may not strictly be a monist either.)<sup>35</sup> But regardless of this worry concerning one and relationality, the view I have offered is like a monism of knowledge in that it affirms that there is knowledge while denying that there is differentiated knowledge. There is thus simply pure, relationless, structureless knowledge.

## VII. Kicking Away the Ladder

---

<sup>35</sup> Della Rocca, "The Elusiveness of the One and the Many in Spinoza: Substance, Attribute, and Mode."