

**A HAPPY TRIUMVIRATE:  
VIRTUE EPISTEMOLOGY, CONTEXTUALISM,  
AND THE SAFETY PRINCIPLE**

**J. ADAM CARTER**

*Stirling University*

**ABSTRACT:** When thinking about why Gettier cases should not count as knowledge, two explanations are particularly salient: (i) the agent’s belief could have easily been wrong; (ii) luck contributes to the agent’s belief in a way that undermines our intuition that a knower is credit-worthy for her belief. In this paper, I will argue that a *virtue responsibilist* account of knowledge can accommodate both of these anti-luck concerns; to do so, it must adopt a safety condition to assuage the first concern, and allow that knowledge depend in a certain way on context to meet the second.

In the Section 1 of this paper, I examine what varieties of contextualism are theoretically compatible with the VE project, and argue that a constraint on VE is that it can only adopt varieties of contextualism that preserve semantic invariance with respect to knowledge claims.

Given this limitation, I consider in Section 2 whether VE supplemented with an invariantist variety of contextualism can preserve what I have described as our anti-luck concerns (i) and (ii). I argue that such a project will be a failure independent of the incorporation of a safety condition. A consideration of John Greco’s recent attempt to incorporate contextualism into his VE account will highlight just why safety is necessary. Duncan Pritchard has suggested recently that once safety is brought into a VE account (a move that must be made to preserve (i)) the intellectual virtues are no longer useful anti-luck machinery, given that safety will be doing all the work. I argue that although safety will do all the work needed to satisfy (i), intellectual virtues are needed to satisfy (ii); this is because, even if an agent’s belief is safe, it remains an open question whether the agent formed the belief in a way relevant to her being credit worthy for her safe belief.

In my final section, I defend my own *virtue responsibilist* account of knowledge that requires that an agent’s belief be safe, and at the same time, that to know, she must have formed this belief in a way that constitutes an achievement. This explanation will make use of a recent proposal by Kelly Becker, which suggests that safety can be indexed to belief forming processes as well as to beliefs. Creditworthiness will depend in part on whether an agent’s exhibition of intellectual virtue is a explanatorily salient feature of her forming a safe belief, a matter that will depend in part on features of context in particular cases of inquiry.

I conclude by defending my proposal against a recent criticism by Guy Axtell, which is based on a metaepistemological worry that a the very project of a VE account attempting to give a Gettier-free reductive definition of knowledge should be called into question.

### ***1.1 VE and Knowledge Invariance: Introduction***

As Crispin Wright has suggested, if a contemporary philosopher announces, unqualifiedly, that she is an “anti-realist” she will have done little more than to clear her throat.<sup>1</sup> For our present purposes, I want to point out that epistemic “contextualism” functions as a similar sort of chameleon in the philosophical lexicon, and that it is approximately as sense-opaque as metaphysical “anti-realist” is without further clarification. For this reason, determining whether a virtue epistemologist can be a contextualist is a useful question only if it is clear just what position is meant to be picked out by “contextualism”. In what follows, I distinguish between several ways that knowledge can depend on context, and will show that if VE is to endorse contextualism,

it must endorse a variety that allows knowledge can depend on context in a way that preserves invariantist treatment of the semantic value of knowledge claims.

### ***1.2 Attributer Contextualism***

The first-year philosophy student who affirms that she has hands is rightly troubled when it is pointed out that she doesn't know that she is not a brain-in-a-vat (BIV). For, if she were a BIV, then it would be false that she has hands. This 'troubling' feeling arises from the idea that knowing that one has hands requires that one know the falsity of propositions—such as the BIV proposition—that, if true, would be incompatible with one's knowing one has hands. This is the intuition that drives the closure-principle for knowledge—a principle that is both intuitive and serves as a motivator for external-world skepticism. The following is a general formulation of the closure principle as presented by Stewart Cohen<sup>2</sup> (1999):

#### **Closure Principle**

(C) If  $S$  knows  $P$  and  $S$  knows that  $P$  entails  $Q$ , then  $S$  knows  $Q$ .

If  $P$  represents the proposition “I have hands” and  $Q$  represents the proposition “I am a not a brain in a vat”, then, because the denial of  $Q$  is incompatible with  $P$ , (C) requires that knowing  $Q$  is entailed by knowing  $P$ ; hence, you can know you have hands only if you know you are not a BIV. It is easy to see why (C) motivates skeptical worries; it is both intuitive, and if true, places a daunting burden on those who want to say that they know a great deal of everyday propositions. Closure-based arguments, loosely speaking, lead us to think that we can't know the obvious unless we first know the impossible.

Despite the intuitive appeal of closure, hardly anyone has concluded from it that we don't know great deal of what we think we know. This is probably because denying that we can know we have hands (or know a great deal of what we take ourselves to know) is even less intuitive than maintaining—albeit somewhat recalcitrantly—that we do know everyday propositions in the face of the intuitiveness of closure. The net result has been that closure-based arguments for skepticism have ignited, rather than an acquiescence to external-world skepticism, a variety of anti-skeptical strategies that attempt to explain how we can know everyday propositions in light of the intuitiveness of closure.

Attributer contextualism is a prominent anti-skeptical position that embraces just this strategy. The attributer contextualist attempts to preserve a trichotomy of intuitions:

- (i) We know a great deal of everyday propositions
- (ii) We don't know we are not BIVs.
- (iii) (C) is correct.

How, save pulling a rabbit named Sophistry from a hat, can these three be endorsed under the same tent? Duncan Pritchard explains the contextualist's gameplan:

The [attributor] contextualist attempts to dissolve this dichotomy by arguing that [these] apparently conflicting intuitions can be reconciled if we are willing to allow that 'knowledge' is a context-sensitive term so that the assertions of ascription sentences in everyday conversational contexts can be true even though the assertion of those same sentences in sceptical contexts would express falsehoods.<sup>3</sup>

What the attributor contextualist wants to say, then, is that (i) is true in everyday contexts, (ii) is true in skeptical contexts, and (iii) is true in all contexts. That the endorsement of (i),(ii) and (iii) appears vulnerable to the charge of contradictoriness results from prior commitments to semantic *invariantism* with regard to the verb 'know.' This position, as John Hawthorne (2004b) has recently described it, is:

**Knowledge Invariantism:**

**(KI)** For any proposition  $P^*$  that contains the verb "know" (as it appears in knows-that constructions), *know* 'has a single, invariant, semantic value'; the truth conditions of  $P^*$  do not change 'from utterance token to utterance token.'<sup>4</sup>

If (KI) is true, then (i),(ii) and (iii) could not be non-contradictorily endorsed. More precisely, if closure is preserved, then (KI) would allow only (i) and (iii) *or* (ii) and (iii).

Crucial to attributor contextualism, then, is the rejection of (KI). "Susie knows she has hands" then would not have a truth value *simpliciter*. The truth value depends in part on the context of utterance. In everyday contexts, Susie could know she has hands. Because knowing this entails that she knows she is not a BIV, she also knows she is not a BIV, and hence, closure is preserved. Importantly, her knowledge that she has hands and is not a BIV is relative to everyday contexts, in which the standard required for knowing is easily met. In other contexts, such as ones in which sceptical scenarios are being seriously entertained, Susie would not know she is not a BIV, and so, given closure, she would not know she has hands. Closure, then, is said to be preserved in the sceptical context as well, where standards for knowing are much higher. The position described as attributor contextualism amounts to the conjunction of two claims: the denial of KI, and additionally, a claim about what features of context are relevant to determining truth values of knowledge claims. A more precise (albeit general) definition of attributor contextualism is:

**Attributor Contextualism (AC)**

- (i) KI is false.

- (ii) The standards relevant to assessing the truth values of knowledge claims are those of the attributor.

Attributor contextualists such as Cohen supplement (ii) by outlining some of the features that contribute to determining what the attributor's standards are. Importantly, these standards must not be simply 'left up to' the attributor's discretion. Here's Cohen:

How precisely do the standards for these predicates get determined in a particular context of ascription? This is a very difficult question to answer. But we can say this much. The standards are determined by some complicated function of speaker intentions, listener expectations, presuppositions of the conversation, salience relations, etc.<sup>5</sup>

Although other defenders of AC such as Keith DeRose and David Lewis have proposed different supplementary accounts of (ii), it isn't crucial that we investigate them here. This is because AC's initial claim—the denial of KI—is what puts AC at odds with the epistemic projects that I will characterize as VE. VE and AC cannot coexist in any sort of theoretical harmony. Before I defend this claim, I want to point out what I am *not* saying: I am not saying that VE and contextualism are incompatible; we are not in any position to draw such a conclusion, for we have yet to consider the VE-amenability of other forms of contextualism that do *not*, as AC does, endorse a denial of KI. I shall now consider just what it is about VE that prevents it from denying KI, and *a fortiori*, from having any hopes of compatibility with AC.

### ***1.3 Virtue Epistemology and Knowledge Invariantism***

Virtue Epistemology, in very general terms, is the project of coming to grips with central epistemic concepts such as knowledge and justification by focusing on properties of epistemic agents rather than properties of beliefs. This general definition is a vague one, but that is unavoidable. This is because there is a “great divide” among theories purporting to be VE, and the differences between theories on either side of the divide is significant enough that what I have described as the general definition of VE is about as rich as a common denominator gets.

On one side of the divide are *virtue reliabilists*; on the other side are *virtue responsibilists*. A crucial point of difference between these camps lies in their respective conceptions of intellectual virtue. Virtue reliabilists understand intellectual virtues to be faculties that need not be such that the possessor is responsible for having them. What makes a faculty an intellectual virtue, for the virtue reliabilist, is that it, as Ernest Sosa puts it, leads to a preponderance of true beliefs over false beliefs. On this view, faculties such as good eyesight and good memory qualify as intellectual virtues. Crucial to an

agent's having knowledge, on the account, is that the agent's true belief is produced by intellectual virtues of this sort. The most prominent defenders of virtue reliabilism are Ernest Sosa and John Greco, each of whom has proposed elaborate versions of the general virtue reliabilist position outlined here. Virtue Responsibilists, on the other hand, reject the virtue reliabilist conception of intellectual virtue. On the responsibilist account, intellectual virtues are understood as character traits, or as the Greeks understood them, *hexis*. Of course, genetically endowed faculties such as good eyesight will not make the list of intellectual virtues of any virtue responsibilist keeping score. Intellectual virtues are character traits which promote some epistemically valuable end, and they must be traits for which the agent is, in some relevant sense, responsible for having developed; archetypal examples include intellectual insight, intellectual sobriety, intellectual courage, openmindedness, etc.

Another point of difference between the positions is that while the virtue reliabilist takes a non-normative concept—truth—as the end relative to which faculties qualify as intellectual virtues by virtue of their promoting this end (and avoiding its converse, error), virtue responsibilists, by and large, take intellectual virtues to be components of some sort of a normative *telos*—which is intellectual flourishing. Guy Axtell, Heather Battaly, James Montmarquet are some of the primary defenders of the responsibilist view. It is Linda Zagzebski (1996), though, who has offered the most comprehensive defense of virtue responsibilism. Her responsibilist view is sometimes described as Neo-Aristotelian because of the emphasis she places on epistemic achievement as being understood within a wider framework of human flourishing.

Despite the differences that bifurcate VE, there is a unifying feature of these accounts that weds them all to KI. I will illustrate why this is so by giving an example of the following form: I will show how denying KI allows for a consequence that all forms of VE that I'm aware of must reject.

I'll begin the example by presenting a case in which both the virtue responsibilist and reliabilist would count as a clear case of knowing.<sup>6</sup> Of course, this will be a case in which an agent is motivated to achieve knowledge, and reaches it because of exhibiting *both* intellectual virtue (in Zagzebski's sense of intellectually virtuous character trait) and intellectual virtue in Greco's sense (a stable feature of a person's cognitive character).<sup>7</sup>

Suppose that U.N. bomb inspector Hans Blix is surveying the fields of Iraq in an attempt to find a W.M.D. After months of coming up empty-handed, Blix spies a suspicious looking black box half-buried in the sand. "Perhaps this is a WMD," Blix

thinks. He calls in his crew of 50 scientists to run some tests. The team, after long hours of careful analysis and sampling, detects no traces of enriched uranium emanating from the box. Blix, finally believing it is safe to open the box, finds inside a child's time capsule containing some pictures, candy and a note to "Mom and Dad." For precaution, Blix personally double checks each item in the box for enriched uranium, so that any doubt could be put to rest. No traces were found.

At the press conference the next day, Blix tells the media that the findings from yesterday were nothing but a child, Johnny's, time capsule. A reporter from Fox News Channel, however, presses Blix on the question. "Are you sure that that wasn't just a disguised WMD?" Blix looks calmly at the reporter and says: "I know that that child's box is not a WMD." And in fact, it was not.

If we can grant that Blix was motivated to know whether the box was a WMD or not, and that he reached the truth of the matter because of exhibiting intellectual virtue in both the reliabilist's and the responsibility's senses, then the virtue epistemologist (of either breed) is going to rule that Blix knows it is not a WMD.

Let us suppose that as the words "I know that the child's box is not a WMD" leave Blix's mouth, two members of the media, George Edward and Keith, form beliefs of their own which are as follows:

George Edward: "*Blix knows that the child's box is not a WMD.*"

Keith: "*Blix does not know that the child's box is not a WMD.*"

Because the attributor contextualist denies KI, it is possible that what both say of Blix is true, simultaneously. The attributor contextualist will point out that Keith has spent his day reading Descartes and watching the *Matrix*, and has thus been immersed in the serious entertainment of sceptical scenarios. (I shall add that Keith has also been reading a science fiction book about time-capsules cleverly disguised as WMDs.) This, the AC theorist will argue, leaves Keith's standards much higher than George Edward's, and hence, relative to George Edward's standards, Blix knows, but relative to Keith's standards, Blix doesn't know. The AC theorist can get away with such a move because the view is not committed to "know" as having an invariant truth value.

Because the VE theorist (both camps) rules that Blix knows, the VE theorist will have to say that George Edward got it right and Keith got it wrong. That follows from the fact that VE rules that Blix knows. Attributor contextualism, then, by denying knowledge invariantism, renders itself an unavailable semantic view for the VE theorist to endorse.

It is true that VE provides, in various forms, accounts of the nature of knowledge such that, when some agent S knows a proposition P, the context in which S comes to believe P is relevant to whether S knows P. This sort of context-dependence certainly isn't unique to VE. What is important to note is that the AC theorist literally treats the verb "know" like we might be inclined to treat the verb 'tall'. As John Hawthorne points out, it would be quite puzzling to think of a sentence "Michael Jordan is tall" to be true or false *simpliciter*. Whether it is true or false will depend on features such as (for example) whether he is on a Subway in China or whether he is competing in the world's tallest man context. 'Flat' functions like tall in this respect; as Hawthorne says: "Assuming that what is flat for an airstrip is different from what is flat for a hockey field, [a pilot and a hockey player] may both utter 'That is flat' and yet one speaks truly while the other speaks falsely."<sup>8</sup> Because VE attempts to give conditions under which an agent, in fact, has knowledge, the VE theorist simply does not think of "know" the same way (i.e. like 'flat' or 'tall') as the attributor contextualist does and, hence, is prevented from employing the attributor contextualist manoeuvres that rely on the rejection of KI. Now that we've seen that VE, given its commitments, can't coherently deny KI, we shall turn now to a consideration of what contextualist options are still on the table for VE.

#### **1.4 VE and Contextualism in Harmony with Knowledge Invariantism**

Ill equipped to explain away closure-based arguments for scepticism by manoeuvres that require the denial of KI, the VE theorist had better find some other way. At least, the VE theorist must have a way to preserve the intuitive closure principle whilst also explaining how it is that we know much of what we take ourselves to know. This requirement is, I take it, one that any adequate theory of the nature of knowledge not succumbing to radical scepticism must meet.

I think that VE has a way to meet this requirement; also, I will propose, VE is better suited to do so by incorporating features of contextualism than it would by going it alone.

At this point, I want to distinguish between two ways that knowledge can be contextual, neither of which involves endorsing a rejection of KI. This distinction is drawn by Sosa (1988):

Knowledge is said to be twice contextual...it is said to be contextual in respect of how sensitive one must be to [sic] evidence if it is to be effective in determining whether one knows that p. It is said to be contextual in what, for short, we shall call *the first way*. It is said to be contextual in a second respect, in respect of which "alternatives" to the proposition that p count as *relevant*. Knowledge is thus said to be contextual in *the second way*.<sup>9</sup>

Sosa's distinction brings us to what I think is an important obfuscation about contextualism. What is generally meant by "knowledge is contextual" is that knowledge in some way *depends* on features of context. For the AC theorist, whether sentences with "knows-that" constructions are true or false in the mouth of a particular utterer *depends* on contextual standards relative to the utterer. A consideration of Sosa's point shows that this is not the *only* way that knowledge can "depend" on context; whether one *knows* a particular proposition can depend on the degree of sensitivity to evidence required to know in a certain context; this is Sosa's first way. In the second way, whether one *knows* a particular proposition can depend (at least in part) on which propositions incompatible with the agent's belief need to be ruled out, which itself can be determined by the context in which the belief is formed. Neither of the two ways Sosa mentions in which knowledge is contextual implies the context-dependence relation espoused by those who reject KI.

To illustrate this, just consider that the extent to which Blix should be sensitive to his evidence, and which alternatives he should rule out in order to know the box is not a WMD would be different if, for example, his informants had disclosed to him that time-capsule boxes were made of some new-age material which can contain uranium undetectable without further, more sophisticated tests. The VE theorist can grant that such dependence on context is needed to know, and will maintain that either Blix did or did not know (but not both) in the context in which he investigated.

Alas, VE can't use contextualism as a crutch to preserve what I described as the trichotomy of intuitions in the way that the AC theorist does. However, denying an invariantist truth value for "know" is not the only way out of the sceptical dilemma engendered by closure. A consideration of how VE can utilize contextualism in the ways Sosa mentions will illuminate a different contextually-driven response to the skeptic—one which harnesses not the skeptic-battling semantic equipment of AC, but rather, skeptic-battling anti-luck equipment without which the VE theorist is (probably) doomed.

### ***2.1 Virtue Epistemology, Safety and Luck***

One of the most promising features of the virtue-based project in epistemology is that it provides resources, where other views do not, for preserving intuitions we have about knowledge and its incompatibility with luck. We must be careful here, though, not to take this intuition too far; as Jon Kvanvig rightly points out, knowledge is not achieved solely through works, but also partly through grace.<sup>10</sup> This is a somewhat

profound way of presenting an obvious feature of knowing: namely, there must be some cooperation of the world involved, and not all facets of this cooperation are within the control of the knower.

Duncan Pritchard (2003) has recently distinguished between the special sort of luck that is simply ineliminable—which he calls *reflective luck*—from the sort that, upon learning that a belief is true because of it, we are prepared to deny knowledge. The latter sort, which he calls *veritic luck*, is the sort of luck that leads us to rule Gettier cases as ones in which the agent in question does not know.

Two primary concerns that lead us to flinch at the thought of awarding knowledge in Gettier cases are:

- (i) the agent could easily have been wrong;
- (ii) Gettier cases are at tension with the thought that knowledge is an achievement on behalf of the knowing agent—an achievement which requires that knowledge is creditable to the agent—credit which does not appear awardable in Gettier cases.

Defining knowledge in a way that eliminates the first concern is a project that has as its champion the *safety principle*; defining knowledge in a way that eliminates the second concern has been a central task of the virtue epistemologist. I will, for the moment, set the safety principle aside to get clear just how it is the VE is a useful anti-luck theory for its ability to deal with the second concern.

Both camps of VE employ the idea that knowledge requires true belief reached through intellectual virtue. The idea relevant to combating our second luck-based concern is that if a belief is reached through the agent’s intellectual virtue, then it is not by accident that the agent believes what she does. Requiring that knowledge be reached through virtue is supposed to eliminate the sort of “accidentality” by which Gettier-style beliefs are true and, at the same time, explain how knowledge is something for which the knower is credit-worthy.

In a recent paper, Pritchard<sup>11</sup> casts doubt on the prospects of VE being successful in this task. I will examine these concerns shortly, but I want to first mention where I think his worries are right. He thinks that VE is simply not up to the task of giving a Gettier-free definition of the nature of knowledge without bringing in the *safety principle*—a principle he thinks is needed for VE to not count as knowledge cases in which an agent who forms a true belief through intellectual virtue *could have easily been wrong*. Put another way, VE focuses on alleviating what I described as the second luck-based concern but, to its detriment, cannot alleviate the first concern. The end result, Pritchard thinks, is that once the safety principle is brought in to supplement the first

concern, the machinery of “intellectual virtue” seems to be tagging along for no good reason, as safety seems to do all the work.

Although I don’t agree that endorsing the safety principle “swamps” the theoretical usefulness of intellectual virtues in a VE safety-endorsing account, I think Pritchard is right to think that VE accounts are doomed to rule as knowledge beliefs that could have easily been wrong; this fatal flaw surfaces in the form of Gettier susceptibility for these accounts. I have argued elsewhere<sup>12</sup> that Zagzebski’s account isn’t up to the task of avoiding Gettier problems, and more recently, have come to the conclusion that no VE account, without the help of extra anti-luck resources, is even *capable* of being Gettier immune.

My thinking now is that, bluntly speaking, there is no good way for VE accounts to handle fake barn cases. Such cases can always be constructed so that the agent *would not have formed some true belief* had it not been for an exhibition of intellectual virtue and, at the same time, we need merely “barn façade” the situation so that it is clear that the intellectually virtuously formed true belief could have easily been wrong. The VE theorist can explain until blue in the face about how the fake barn cases are not ones in which the true belief was formed *because* of the intellectual virtue; however, in cases in which the formation of the intellectually virtuous belief depends counterfactually on the agent’s displaying intellectual virtue, it is simply false to say that the belief did not arise because the agent displayed the virtue. If x would not have happened if y hadn’t happened, then it would be overtly dubious to try to deny that “x because of y” or, furthermore, to deny that y is an explanatorily salient feature of x’s coming about<sup>13</sup>. VE theorists who deny that intellectually virtuous true beliefs in fake barn scenarios are true because of intellectually virtuous believing can only do so by committing the mistake I’ve just mentioned.

I will pause here to summarize: VE appears to be a useful anti-luck strategy because it requires a condition—intellectually virtuous believing—which seems to by definition rule out the kind of accidentality of beliefs for which we don’t want to give the possessor credit when true. This strategy, though, is only speciously promising; the implementing of the virtue conditions still leaves open the possibility that beliefs reached because of intellectual virtue could have easily been wrong. VE, then, cannot achieve the anti-luck task to which it is motivated without adding some feature to the theory that eliminates the possibility that the virtuously formed true belief could have easily been wrong.

Recall that we still have two types of VE-compatible contextualism left on the table. VE appears desperate, and so we shall turn now to the question of whether the remaining forms of contextualism can patch up what has been exposed as VE's Achilles heel. Additionally, to attend to Pritchard's concern, we must see if the heel can be patched without patching so much as to make useless the runner.

### ***2.2 First try: Greco's Agent-Reliabilist contextualism***

In his forthcoming draft "What's Wrong With Contextualism", John Greco attempts to make exactly the sort of move our discussion to this point has suggested is needed for VE, and for exactly the reasons mentioned. Greco's agent reliabilist brand of virtue reliabilism fits the bill of a VE account that attempts to eliminate non-credit-worthy belief as knowing, but does so in a way that nonetheless invites fake-barn-style Gettier cases.

Against this background, Greco has recently called to the aid of contextualism to save the day. The sort of contextualism he incorporates into his view is in line with the sorts that Sosa mentions, which preserve KI while allowing for knowledge to depend importantly on context in other respects. On the view he proposes, knowledge inherits the "context-sensitivity of causal explanations."<sup>14</sup> Here is an example of how this is supposed to work:

Suppose Gary, with excellent eyesight, sees a sheep-looking figure in a field. He forms the belief that there is a sheep in the field, a belief which is formed by the faculty of his keen and reliable eyesight. Unfortunately for Gary, what he spotted was a dog that looks very much like a sheep. However, there is a sheep hiding in far-away bush in the field, and so Gary's belief is true.

Greco claims that while Gary's faculty of good eyesight contributed to his forming the true belief that there is a sheep in the field, his success was not related in the right way to his ability. On Greco's view, cases of knowledge require that intellectual abilities be a *salient* factor in the explanation of why the belief is true.<sup>15</sup> In the case of Gary, his abilities are not explanatorily salient in his forming a true belief, and so he does not know. What counts as salient in a causal explanation, Greco thinks, depends on context. He writes:

Putting all this together, we get the result that knowledge attributions are sensitive to context because they involve causal explanations, and causal explanations are sensitive to context. Knowledge attributions inherit the context-sensitivity of causal explanations.<sup>16</sup>

To get a better idea how causal explanations are supposed to depend on context on his view, consider a case of a car wreck at an oddly designed intersection. Suppose a

driver, confused by the intersection, runs through a light and hits another car. The police rule that the driver caused the wreck. Suppose further, that the city planners, discussing a large proportion of wrecks at that intersection, rule that the poorly designed intersection is the cause of wrecks such as the one mentioned. Greco thinks that they can both be correct and that what is explanatorily salient as the cause of a wreck depends on the context in which the wreck is discussed, and hence, we should think of causal explanations as context dependent. To reiterate, because knowledge turns on causal explanations, knowledge thus inherits context sensitivity.

I have two bones to pick with this suggestion. First off, I think we should reject his idea that the semantics of causal explanation requires a contextual treatment. Such a view appears to me to be metaphysically indefensible. For example, take a paragon case of causation: two snooker balls collide, the first knocks into the second, causing it to move. We can give several causal explanations for the second ball moving. On the one hand, an event-theorist about causal relata, who maintains that causation is a relation between events, will point out that the event of the first ball coming into contact with the second ball explains the event of the ball moving. Event theorists, such as David Lewis (1973), will make recourse to counterfactual analyses of causation to give precise conditions under which one event is a cause of another. For example, on a simple counterfactual analysis, event *a* is a cause of event *b* if and only if, had *b* not occurred, *a* would not have occurred.<sup>17</sup>

An important idea behind this kind of project is that causation is a phenomena that admits of a reductive account, one which supposes an invariantist treatment of causation. Many objections have been presented to simple counterfactual analyses, and Lewis actually rejects his early simple counterfactual analysis for a more elaborate counterfactual account in which a cause must be a “difference maker” in its effect. Counterfactual analyses of causation which take events as casual relata certainly come in quite a few varieties in an attempt to preserve our intuitions about particular cases<sup>18</sup>; however, it is a supposition among such views that causation is to be understood within the context of invariantist semantics. Others who part ways with Lewis on the score of what counts as causal relata—for example, by claiming as Hugh Mellor and David Armstrong do—that facts are causal relata—yet nonetheless preserve invariantist semantics with regard to causal claims.

For Greco to deny invariantist semantics about causation, then, puts him directly at odds with precedent. *Prima facie*, this appears to be a strike against his view, and at

minimum, leaves him in need of an explanation at the level of the metaphysics of causation that can accommodate what he wants to say.

However, I think Greco has a response available that will get around this problem, though. Rather than to treat causal claims themselves as contextual, he could maintain invariantist semantics for causal claims while allowing that the *extent* to which an agent's virtuous believing must be a cause of the belief in order to get credit will depend on features of context. On such a view, it wouldn't be a matter of context *whether* virtuous believing is a cause of a belief, but what would depend on context is whether the virtuous believing is a cause of a true belief in a way that is relevant to award the agent credit. Making this move would alleviate my earlier worry, but would require a careful articulation of just how we are to understand credit earning while maintaining that the truth values of causal explanations do not vary from utterance token to utterance token.

But even this project, I think, will run into a more serious concern, which is that of Pritchard's. Consider that even if Greco incorporates invariantist contextualism in his view in a way that explains how credit is earned, and in a way that doesn't treat causal explanations as themselves contextual, his view is still susceptible to the criticism that his view can't adequately explain knowledge in a way that eliminates our first worry about luck—namely that: if you know, then you couldn't easily have been wrong. Greco, then, in supplementing his account of credit, will have to make some reference to how an agent's holding a *safe* belief is relevant assessing to assessing credit. Pritchard claims that once the safety principle is added to a VE account, however, the role of virtues will then become extraneous for the reason that safety is doing all the anti-luck work that virtues were intended to do in the first place. To see why he thinks this, consider Pritchard's formulation of the safety principle:

**Safety Principle:**

For all agents,  $\phi$ , if an agent knows a contingent proposition  $\phi$ , then, in most nearby possible worlds, that agent only believes that  $\phi$  when  $\phi$  is true.<sup>19</sup>

Adding a safety requirement to Greco's view would then alleviate anti-luck concerns to which his view is susceptible, but in a way that Pritchard thinks will leave his virtue-based account undermotivated. Here's Pritchard:

But if safety is to be part of the view anyway, then why is there any need to appeal to the specific virtue-theoretic aspects of the position in order to deal with other cases of epistemic luck that are neutralized by this anti-luck condition? That is, why not simply opt for a safety-based theory of knowledge and make no essential reference, in one's definition of knowledge at any rate to the epistemic virtues and cognitive faculties?<sup>20</sup>

It is true that if Greco requires that beliefs counting as knowledge meet the safety condition, then there is no reason to appeal to virtues in any attempt to explain how agents who have knowledge couldn't easily have been wrong.

I want to draw attention at this point to the distinction I made in the previous section regarding two separate anti-luck concerns we have. I describe these as “anti-luck” concerns because they are what I take to be the most salient concerns we have when thinking about why Gettier cases should not count as knowledge. First, we point out that the agent could have easily been wrong. Secondly, we point out that knowledge is a sort of achievement on the part of the knower, and that Gettier cases count as knowledge beliefs that the agent didn't form in the right sort of way so as to be credited for this achievement.

If Greco requires that the safety condition be met, he would address the first concern in a manner which would trump any usefulness virtues would have at meeting this task. However, requiring that beliefs be safe would not exhaust the usefulness of virtues in an account wishing to preserve our second anti-luck concern.

It isn't clear that we can conclude from the fact that an agent's belief is safe that the agent formed the belief in a manner that would render her creditworthy for this achievement. Put another way, an agent could hold a belief such that it couldn't easily have been false, and yet, it would remain an open question whether the agent deserves credit for holding such a belief. And so, to the extent that we take it as intuitive that knowledge is some sort of achievement, we should reject the idea that safe-belief alone is capable of assuaging both our first *and* second anti-luck concerns.

Yet a further problem surfaces: even if requiring that a belief be safe would not trump the usefulness of virtues on Greco's account toward the task of preserving *both* anti-luck conditions, Greco will nonetheless be faced with a further problem that I think is intractable for his view, and additionally, intractable for any virtue-reliabilist account wishing to address the anti-luck concerns I've mentioned.

The problem is this: Greco's reason for appealing to contextualism in the first place was motivated by his goal of providing an account of how virtues are explanatorily salient in causal explanations of knowledge claims. However, once Greco brings in the safety principle—which he must to alleviate the first anti-luck concern—his view must then explain how an agent's believing virtuously is an explanatorily salient feature of explanations for why an agent's true beliefs are safe. As I see it, no virtue reliabilist account, Greco's included, is capable of such an explanation.

To illustrate this, I will return a final time to our story of Blix and his knowledge that the black box is not a WMD. An account that rules Blix's belief as knowledge will have to preserve both that he couldn't have easily been wrong, and that his knowing this is in some sense a credit-worthy achievement. Suppose Greco tries to explain this. Endorsing the safety principle will allow him to explain why Blix couldn't have easily been wrong—namely, in most nearby possible worlds in which he forms the belief in the same way, his belief will be true. However, once Greco tries to explain why Blix's knowledge should accrue him credit, he must explain how Blix's intellectual virtues are explanatorily salient in his coming to form a safe belief. However, on Greco's account of intellectual virtue, in which stable, reliable faculties of one's cognitive character are what he takes to be intellectual virtues, these sorts of faculties simply would *not* be explanatorily salient in any explanation we give for why Blix's true belief is safe. To see why this is so, consider that if we ask the question "Why did Blix form a safe belief in this case?" what is explanatorily salient is that Blix exhibited dispositions that led him to be sensitive to what evidence is relevant to knowing whether the box was a WMD, and additionally, that he evaluated this evidence in an intellectually virtuous way. In other words, Blix was sensitive to features of the context of his inquiry which, in turn, led him to form a safe belief in a way to which we can ascribe him credit.

Faculties that Greco rules as virtues, such as good memory, good eyesight, etc., would, though instrumental to Blix's forming a safe belief, not be explanatorily salient in an explanation of why he did so. What is explanatorily salient is features of Blix's intellectual character that he is responsible for having developed and employed in this case. This fact points to the following conclusion: A VE account that endorses a safety condition can preserve both of our anti-luck concerns only if the account can explain why the agent's exhibition of intellectual virtue is explanatorily salient—in a way sufficient to explain how the agent is credit-worthy—for forming a safe true belief. An VE account can only do *this* if it endorses a view of intellectual virtue that allows for virtues to actually *be* explanatorily salient in such an explanation. Virtue reliabilist accounts think of virtues in a way that would preclude them from being explanatorily salient in an explanation for why an agent's true belief is safe. Intellectual virtues can be explanatorily salient in an explanation of why an agent's true belief is safe only within a picture of intellectual virtue that thinks of intellectual virtues as the sorts of dispositions that really are explanatorily salient in explanations of why an agent forms a true safe belief, and these dispositions are ones which lead an agent to be sensitive to what

evidence is relevant in a particular context, and to evaluate this evidence appropriately. The virtue *responsibilist* account, and not a virtue reliabilist account, understands intellectual virtues in a way that is capable of such an explanation.

### ***3.1 Virtue responsibilism, context and safety in harmony***

The discussion to this point has led to the conclusion that quite a bit of promise lies in the direction of incorporating safety into a virtue responsibilist account. As we've learned from considering Greco's view, an incorporation of safety into a VE account capable of preserving both anti-luck intuitions must incorporate safety in such a way that the agent forms a safe belief *because* of having formed the belief through intellectual virtue. Additionally, a promising way of doing this is to, as Greco proposed, require that the agent's intellectual virtue be an explanatorily salient feature of the explanation for why the agent forms a safe belief. We don't want to, as Greco did, endorse an invariantist treatment of causal explanations in our effort; however, as I mentioned in the last section, such an endorsement isn't necessary to the task.

So what is our picture to look like? I'll attempt to answer this by adding our needed components one at a time. We begin with a very general virtue responsibilist framework:

**VR1:** S knows some contingent proposition  $p$  iff  $p$  is true and S believes  $P$  because of having exhibited intellectual virtue.

I am going to set aside at this point whether this account requires a 'motivation' component for the reason that I want to keep this initial proposal very general. VR1, as stated, cannot preserve our first and second anti-luck concerns, of course, and so the next step will be to add the safety condition so as to eliminate the possibility that S knows  $p$  and could have easily been wrong. Hence:

**VR2:** S knows some contingent proposition  $p$  iff (i)  $p$  is true, (ii) in most nearby possible worlds, S believes  $p$  only when  $p$  is true, and (iii) S believes  $p$  only when  $p$  is true in most nearby possible worlds *because* S exhibits intellectual virtue in forming the belief  $p$ .

VR2 now has no problem meeting our first anti-luck concern. However, to understand how it will meet the second—namely, that when S knows, it is an achievement creditable to S (an achievement that isn't entailed simply by S holding a safe belief)—we must explain more carefully how condition (iii) in the definition is satisfied. Satisfying (iii) will require that we must be able to draw conclusions from how S forms a belief in a particular case to how S would form beliefs in most nearby worlds. I think we can do this if we think of safety the way that Kelly Becker (2006) has recently suggested, in which safety is indexed to belief formation processes—a move he makes in defending his

process reliabilist account against what he takes to be an extension of Pritchard's criticism of VE to his own account. If Becker is right, and we can talk about the safety of belief formations processes as well as the safety of beliefs, then we are well on our way to showing how (iii) in VR2 is to be satisfied. We shall say then that S believes p only when p is true in most nearby possible worlds because S exhibits a safe belief forming process, the exhibition of which depends on S's exhibiting intellectual virtue in forming the belief. Requiring that S exhibits a safe belief forming process depends on S's exhibition of intellectual virtue is a key move here because it requires that the safe belief forming process be one for which the agent is responsible for having developed. We can rightly give credit to an agent who forms a safe belief in this way, where as we could not if the requirement were only that the safe belief be formed by a safe belief forming process—one which would leave it an open question as to whether the agent deserves credit for such a process. For example, "Swampman" might form safe beliefs from safe-belief forming processes, and yet our second anti-luck concern would not be addressed because it wouldn't be clear that Swampman's resultant beliefs are achievements deserving of credit.

The account I am proposing is also one in which knowledge depends on context. Contextual dependence is required by the third condition. I have said thus far that the agent's safe belief forming process must arise from the agent's having exhibited intellectual virtue. This, I've said, is important if the agent is to be creditable for forming safe beliefs. We shall bring on board here Greco's notion of explanatory salience to establish how intellectual virtue and resultant safe belief forming processes are to hook up in the right way. My suggestion here is that the exhibition of intellectual virtue must be explanatorily salient in a causal explanation of why it is that the agent's belief forming process is safe. Of course, we do not need to do this by denying a semantic invariantist treatment of causal explanations, as Greco did. We can maintain that, while the causal relationship between the intellectual virtue and the safe belief forming process is not such that claims about this relation have truth values that vary from utterance token to utterance token, nonetheless, the extent to which the intellectual virtue must be instrumental in engendering the safe belief forming process *in order for the agent to get credit* is a relation that will depend on features of context. Note that this is a distinct claim from the claim that whether or not the intellectual virtue and the safe process stand in some particular causal relation is such that explanations can have variant truth values. Rather, I suggest that we maintain that the causal relation between the intellectual virtue and the

safe belief forming process “is what it is”, and whether “what it is” is relevantly explanatorily salient in our assessment of credit to the agent will depend on features of context that Sosa mentions, such as how sensitive one must be to available evidence.

I think that a satisfactory treatment of (iii) will require a more rigorous presentation than what I have provided here; however I hope to have provided at least a promising sketch for how such a treatment must proceed to accommodate both anti-luck concerns.

#### ***4. Strengths and Problems***

What is promising about the account proposed here is that it offers a way to accommodate not just one, but two key intuitions we have about knowledge and luck. Luck can infest a theory twice over: it can lead beliefs to be counted as knowledge which could have easily been false, and it can contribute to a beliefs being formed in a way that inclines us to deny credit to the agent for having formed the belief.

Incorporating virtue responsibilism with safety in a way that credit for knowledge will depend on context—i.e. context will determine what is relevantly salient about the relationship between an agent’s exhibiting intellectual virtue and believing  $p$  through a safe belief forming process—has the resources for doing justice to both anti-luck concerns. Additionally, because credit is a normative notion on this account, we have resources for explaining the value of knowledge that wouldn’t be available to accounts which either fail to explain why knowledge is an achievement of the agent, or attempt to foist such an explanation by claiming that credit accrues agents by virtue of the agent’s genetically endowed processes, faculties or “cognitive” character—features for of an agent for which the agent need not be responsible, and features which are not explanatorily salient in an explanation of the agent’s safe belief formation.

I turn now to what I think are some difficult objections for this account. First off, it appears incapable of explaining *a priori* knowledge. Given that for any *a priori* belief  $\phi$ ,  $\phi$  is necessarily safe. One might object that one could know  $\phi$  even if, for example, she were intellectually vicious, and wished to hold all false beliefs. Even such an agent, upon reflecting upon  $\phi$  would know  $\phi$ , and so, the account I have presented is incomplete. At least, this is how I see the objection going.

No appeal I can give to the relationship between intellectual virtue and safe belief forming process will conclude that intellectual virtue is explanatorily salient in the agent’s forming a safe belief forming process. But perhaps, there is a different sort of response available. To give this response, we need only reflect on what actually motivates our

second anti-luck concern. We do think that knowledge of contingent propositions is some sort of achievement on behalf of the agent, for which she is credit worthy. But do we think that an agent is credit worthy for her *a priori* knowledge? Perhaps not. A way to think about this is as follows: The anti-luck concerns I mentioned are ones that we have when thinking about Gettier cases—cases in which a theory rules as knowledge a belief in which veritic luck undermines either the safety of a belief, or the fact that the belief is creditable to the agent. I am not aware of any “a priori Gettier cases” and am not such a case would even be possible. This is because, when one knows an *a priori* proposition, it is never by accident that the belief formed is true.

If we in fact do not think of *a priori* knowledge as the sort that, when the believer believes it, she is credit worthy, then it would not be a strike against a view that it doesn't explain this fact. However, a consequence nonetheless is that my account must be restricted in its domain so that it is an account of *contingent* knowledge.

A second objection to the view proposed here can be found in a recent paper by Guy Axtell (2006), in which Axtell defends a virtue responsibilist account in a manner much different than the one given here. In his manuscript “Two For the Show”, Axtell takes as a starting point for discussion the a move Pritchard makes in *Epistemic Luck*, which is to reject the epistemic luck platitude (ELP)—a platitude that has as an assumption the mistaken idea that *all* forms of luck are eliminable. Axtell thinks Pritchard is right to think that mistaken conceptions of the extent to which luck should be eliminated from an account of knowledge arise from dogmatic adherence to ELP; however, Axtell thinks that there is a second platitude that epistemologists endorse and which is equally, if not more, problematic for their accounts. This platitude is the Platonic Definition Platitude (PDP) which is the platitude that knowledge admits of necessary and sufficient conditions in terms of which it is definable. Axtell's view is that, once we bring PDP on board our knowledge-defining ship, we are embarking on a journey that will inevitably end at the bottom of the sea. He compares the endorsement of PDP to making the decision to play the board game *Jumanji*—a fictional game in which the only rule is that, once you start playing, you play until you finish, even though the game itself is a tortuous one. His idea here is that, once we buy in to PDP, we will inevitably be sucked into the “Gettier Game”—a game which he thinks is unwinnable. I am sympathetic to Axtell's disenchantment of the Gettier game; in fact, Robert Shope's *The Analysis of Knowing: A Decade of Research* amounts to a recorded history of failed attempts of those who dared play the game. What Axtell proposes is that, rather than to

attempt to solve the Gettier problem, we should not play the game in the first place. More specifically, we should analyze knowledge independent of the PDP.

His suggestion for an alternative draws from recent ideas by Heather Battaly, which require that we think about intellectual virtues as *thin* rather than *thick* concepts. In doing so, he thinks, we can offer a virtue responsibilist account of knowledge that has an anti-luck (*toucheic*) condition as well as a condition that the belief be creditable to the agent in the right way. In this respect, his view is similar to the one I have proposed here. However, Axtell thinks that by defining knowledge in terms of intellectual virtues that are understood *thinly*, the result will be an account of knowledge that does not admit of necessary and sufficient conditions, as PDP-endorsing accounts do. He thinks that the virtues, by nature, are concepts more akin to “religion” than (for example) “bachelor” in the sense that there is no description of intellectual virtue that is applicable to all particular instances of intellectual virtue. And so, because knowledge is defined in part in terms of *thin* virtues, it will itself not admit of necessary and sufficient conditions. I take it that his account is one that attempts to give conditions of verisimilitude; he thinks that a stronger project would suck us back into the Gettier game.

My response to Axtell’s suggestion here is two-fold. First, if he is correct that the Gettier game is unwinnable, then I agree that it is time to start thinking outside of the box and consider ways of analyzing knowledge independent of PDP. However, making this move should be our last resort. A reductive account of knowledge would be, in principle, more illuminating and helpful to us than a mere account of verisimilitude; we should resort to the latter only if it is clear that the Gettier game is unwinnable. We ought not pay the price of amputating our legs until we are sure that what diseases them is untreatable.

I think that before jumping the PDP ship, we must be quite confident that no reductive account of knowledge is capable of preserving the anti-luck conditions—the assuaging of which has been the target of the Gettier gamesmen. The account I provide here, however, draws on mistakes the Gettier gamesmen have made in the past. Consider that Gettier solutions were being entertained for over 20 years before the first discussions of intellectual virtue were introduced in the literature. Similarly, the benefits of endorsing contextualism into an account are, relatively speaking, in their infancy.

I think that we should attempt to utilize these anti-luck tools in every way possible before giving up on a reductive definition. Again, this is because the reductive

definition, if achievable, is more illuminating and useful than what we get if we settle for less.

#### REFERENCES

- Axtell, Guy. (2007) "Two for the Show: Anti-luck and Virtue Epistemologies in Consonance," Synthese special edition on epistemic luck (D. Pritchard, ed.), forthcoming.
- Axtell, Guy. (2007) "Virtue Theoretic Responses to Scepticism", Forthcoming in J. Greco (ed.) *Oxford Handbook on Skepticism*, OUP.
- Carter, J. Adam. (2004) "Zagzebski and the Gettier Problem" (manuscript).
- Collins, John; Paul, L.A. and Hall, Ned. (2004) *Causation and Counterfactuals*, Cambridge, Mass.: MIT Press.
- Cohen, Stewart. (1999) "Contextualism, Skepticism and the Structure of Reasons," *Nous*, Vol. 33, Supplement: Philosophical Perspectives, 13, Epistemology.
- Greco, John. (2006). "What's Wrong With Contextualism?" (manuscript)
- Hawthorne, John. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press.
- Kvanvig, Jonathan L. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.
- Lewis, David. (1973). "Causation," *The Journal of Philosophy*, Vol. 70, No. 17, Seventieth Annual Meeting of the American Philosophical Association Eastern Division
- Lewis, David. (2004) "Causation as Influence", published in Collins, J. *Causation and Counterfactuals*, Cambridge: MIT Press.
- Pritchard, Duncan. (2005). *Epistemic Luck*, Oxford: Oxford University Press.
- Pritchard, Duncan. (2007) "Virtue Epistemology and Epistemic Luck (Revisited)" *Metaphilosophy* (forthcoming)
- Jonathan Schaffer (2000) "Trumping Preemption," *Journal of Philosophy* 97, pp. 165-81.
- Sosa, Ernest. (1988). "Knowledge in Context, Skepticism in Doubt: The Virtue of Our Faculties," *Philosophical Perspectives*, Vol. 2, Epistemology.
- Wright, Crispin. (1992) *Truth and Objectivity*, Cambridge, Mass: Harvard University Press.

#### NOTES

---

<sup>1</sup> Wright, C. (1992). *Truth and Objectivity*, Cambridge, Mass.: Harvard University Press, p. 1.

<sup>2</sup> Cohen, S. (1999) 'Contextualism, Skepticism and the Structure of Reasons' *Nous*, Vol. 33, Supplement: Philosophical Perspectives, 13, Epistemology. p. 62.

<sup>3</sup> Pritchard, D. (2005). *Epistemic Luck*, Oxford: Oxford University Press, p. 55. In this passage, Pritchard is describing the attributor contextualist strategy of dissolving the tension between the infallibilist intuition that we don't know the denials of skeptical hypotheses and the anti-skeptical intuition that we know a great deal of everyday propositions. I have presented the reconciling dilemma for the contextualist as between the two intuitions Pritchard mentions plus a third, which is that closure should be preserved. Pritchard seems to take the closure-endorsing requirement as part of the dilemma as well given that he points out elsewhere that, because closure is so intuitive, it is a strike against a theory if it fails to endorse it (i.e. see, for example, his criticism of Dretske's attempt to do so) *Ibid.* Sec. 2.4.

<sup>4</sup> Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press, p. 53.

<sup>5</sup> Cohen, S. (1999) 'Contextualism, Skepticism and the Structure of Reasons' *Nous*, Vol. 33, Supplement: Philosophical Perspectives, 13, Epistemology. p. 61.

<sup>6</sup> The views I have in mind here as representative of the two camps are Greco's (agent reliabilist) and Zagzebski's (1996).

<sup>7</sup> The motivational condition I mention here is presented to satisfy Zagzebski's criteria.

<sup>8</sup> Hawthorne, J. (2004). *Knowledge and Lotteries*, Oxford: Clarendon Press, p. 55.

<sup>9</sup> Sosa, E. (1988). "Knowledge in Context, Skepticism in Doubt: The Virtue of Our Faculties," *Philosophical Perspectives*, Vol. 2, Epistemology, p. 143.

<sup>10</sup> Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*; Cambridge: Cambridge University Press.

<sup>11</sup> Pritchard, D. "Virtue Epistemology and Epistemic Luck, Revisited" ( )

<sup>12</sup> Carter, J.A. "Zagzebski and the Gettier Problem" (manuscript).

<sup>13</sup> This is the move that Greco ("What's Wrong With Contextualism-(manuscript)) makes.

<sup>14</sup> Greco, J. "What's Wrong With Contextualism" (manuscript) p. 5.

<sup>15</sup> *Ibid.* p. 4.

<sup>16</sup> *Ibid.* p. 5.

<sup>17</sup> Lewis, D. (1973). "Causation," *The Journal of Philosophy*, Vol. 70, No. 17, Seventieth Annual Meeting of the American Philosophical Association Eastern Division. (Oct. 11, 1973), pp. 556-56

<sup>18</sup> For example, simple counterfactual analyses of causation branch into nuanced varieties in attempts to get right standard problem cases, such as ones introduced by Jonathan Schaffer (2000) in "Trumping Preemption," *Journal of Philosophy* 97, pp. 165-81

<sup>19</sup> Pritchard, D. (2005) *Epistemic Luck*, Oxford: Oxford University Press, p. 71. Pritchard has since revised his account of safety in light of criticisms from Greco that a reading of "all nearby possible worlds" would be too strong and "most nearby possible worlds" will be too weak. In response, Pritchard presents an intermediate account of the safety principle in "Virtue Epistemology and Epistemic Luck, Revisited" which addresses this objection.

<sup>20</sup> Pritchard, D. (200 ) "Virtue Epistemology and Epistemic Luck, Revisited," p. 7.