

A Predominately Externalist Definition of Knowledge

Abstract: Most philosophers believe that a necessary and sufficient conditions analysis of knowledge cannot be fruitful. Against this worldview, I propose a four-condition predominately externalist (PE) definition of knowledge. In addition to the traditional 'justified true belief' conditions, a 'relevancy condition' and a 'no-defeaters condition' are added. We will examine eight case studies and explain why S knows or doesn't know in various situations, including the parked car case, barn façade case, and two lottery cases. We'll respond to the skeptical argument and the 'safety' condition. A key feature of the PE theory is that it doesn't guarantee that knowledge exists. The skeptical hypothesis is accepted, and epistemic closure is rejected, so that the PE definition doesn't assure that S knows *any* p. The definition just states the necessary and sufficient conditions for the *possibility* of knowledge. Against the skeptic's *impossibility* conclusion, with the satisfaction of the PE conditions, knowledge is clearly *possible* and obtainable.

Introduction

In this essay, I introduce a four-condition definition of 'knowledge' to succeed the traditional definition of 'knowledge' as 'justified true belief.' The methodology for explaining the significance of the predominately externalist (PE) definition of knowledge is to analyze a number of examples, in cases where knowledge is obtained, and other cases where it doesn't obtain. I assume that we can understand hypothetical cases where knowledge occurs, and that we seek the theoretical principles (and material conditions) upon which these intuitions are based. In speaking of 'intuitions' I assume that we have tentative beliefs about our normal use of epistemic concepts and the material conditions

under which they apply. I contend that knowledge can be analyzed as a 'natural kind' entity (i.e., having unity, discreteness, and essentiality) with illuminating results. The conditions for *how* and *when* 'S knows p' can be described in a fruitful way, without begging any questions against the skeptic.

What is the Function of a 'Definition' of Knowledge?

A key feature of this definition is that it seeks to show how knowledge is *possible*. The PE theory doesn't try to construct a theory of knowledge that is immune to skepticism or Gettier situations. It is agreed with Linda Zagzebski (1994) that no necessary and sufficient conditions approach will generate conditions where it is *guaranteed* that S possesses knowledge. As an alternative, this definition *describes* the material conditions for *when* it is contingently possible that 'S knows p' and explain *why*. It *explains* the Gettier problem and other epistemic problems, but *isn't* adding conditions (e.g., no-defeaters condition) with the intent to evade (or resist) Gettier situations. The definition *describes* a Gettier case as a situation where two conditions (3 and 4b) fail to be satisfied. The acceptance of the truth of the 'skeptical hypothesis' in the standard skeptical argument (below) implies that there may be *no* instances of human knowledge. As Zagzebski states, Gettier situations are inescapable and will never go away. But if knowledge exists, then these are its four conditions that answer the question, 'how is knowledge possible?' while simultaneously admitting skeptical hypotheses are possible.

The Predominantly Externalist Definition of Knowledge

I hypothesize the following Predominately Externalist (PE) definition of knowledge where S designates a subject (i.e., a person) and p designates a proposition:

'S knows **p**' if and only if:

(1) **p** is true.

(2) **S** believes **p**.

(3) **S** believes **p** upon a set of implicitly or explicitly held reasons that are substantially relevant (i.e., truth-connecting) for why **p** should be believed. In other words, the evidence and the belief forming processes (e.g., perceptual, inferential) that generate **p** are naturally or objectively why **p** should be believed.

(4) There exists no unresolved nor unconsidered undermining evidence, that would effectively lead **S** to doubt or disbelieve **p**, violating condition 2:

(4a) In situations of critical doubt, **S** must have 'adequate evidence' (i.e., strong reasons **e1**, **e2**, **e3**, etc.) for believing **p**, and **S** must be able to resolve (i.e., rule-out, discard) any actual or logical possibilities that would undermine (or defeat) the evidence possessed for **p**, and

(4b) There cannot exist undermining evidence **q** (no matter whether **S** is aware of it or not) that would significantly weaken (i.e., undermine, undercut) **S**'s belief that **p**. If there does exist evidence **q** that suggests **not-p**, and if **S** was to be aware of this evidence, then **S** must have (or acquire) evidence to dismiss (or resolve) counter-evidence **q**.

I name this a 'predominately externalist definition of knowledge' because condition 3 is an externalist condition about a person's reasons (i.e., belief-forming processes) for believing **p**. Condition 4a is an internalist condition, and 4b a mixed externalist-

internalist condition. With this definition, the externalist conditions *dominate* (or take precedence, predominate) over the internalist conditions in determining whether **S** knows.

The four primary conditions of the PE definition (that exclude sub-conditions 4a and 4b) suitably explain non-human animal knowledge. Animal knowledge is a function of how a belief arises and the surrounding material conditions. The PE definition contends that animals can know proposition **p**: if **p** is true, **p** is believed, **p** is believed upon truth-connecting reasons for why **p** should be believed, and if there are no undermining factors that would weaken belief **p**. For human knowledge, persons should additionally have reasons (or justification) for a belief if reasons are demanded by a sincere objector. Conditions 4a and 4b acknowledge the desire to assess the reasons for a belief and how a belief is known or is undermined.

Condition 3 and the Requirement of Relevant Reasons

Condition 3 requires that **S**'s reasons for believing **p** must be truth-connecting, non-coincidental, pertinent, or applicable (i.e., relevant) for why **p** should be believed true. A 'truth-connecting' reason (or premise) is a proposition that constitutes *evidence* from an external perspective, for *why* **p** should be believed. In this sense, the evidence 'relevant' for *why* any particular **p** should be believed true, is objective and independent of persons. Condition 3 makes no normative claim about how much relevant evidence is 'adequate' to believe **p**, and the PE definition doesn't imply that **S** needs to be conscious of (or be able to recount) the reasons for believing **p**, except in situations of critical doubt or explanation. It isn't implied that *every* proposition that is part of **S**'s evidence needs to

be true. **S** can sometimes possess some false beliefs as part of a set of (mostly) truth-connecting evidential premises. In sum, it is intuitive that for knowledge, it is necessary that **S** should believe **p** upon a set of tacitly or explicitly held reasons that are *relevant* (or *truth-connecting*) for why **p** should be believed.

A Conceptual Analysis of 'Relevance'

There are at least two senses of 'relevant' found in ordinary language. Condition 3 incorporates just one sense. This sense is illustrated by an event. The event is a house fire. Suppose that a house in a residential neighborhood is substantially damaged by a fire. The cause of the fire is initially unknown. Local fire investigators are called in to determine 'why did the fire start?' Pre-theoretically, we believe that there are *relevant reasons* for why a house catches on fire. We don't believe that a house suddenly combusts without a cause. When seeking the cause(s) of why a fire started, fire investigators want to discover whether an arsonist was involved, or if there was careless smoking, or whether there was an electrical problem, or a lit candle accidentally fell, or lightning struck and so on. In order to obtain knowledge of the cause, the investigators seek objectively *relevant evidence* and *relevant reasons* for why the house caught fire and wish to discard the consideration of any extraneous (and *irrelevant*) states-of-affairs not associated with the cause of the fire.

Let us suppose that fire investigators offer this final report that explains why the fire started: 1) an upended candle was found in the bedroom of the home, where most of the fire damage occurred, 2) the pattern of how the fire spread, and resulting damage suggests that the fire started in the bedroom near the fallen candle, 3) a resident of the

home admits having left a candle burning while leaving the home several hours earlier, 4) there is no other evidence (indicating arson, electrical problems, etc.) suggesting that the candle was not the cause of the fire, so therefore 5) **p**: the fire was caused by an unattended lit candle that fell and ignited nearby combustible materials.

Whether the premises stated by the investigators are '*relevant*' for why the fire started *is not something we decide*. Instead, if these are the relevant reasons for why the house caught fire, it is because the reasons (i.e., evidential premises) are *truth-connecting, objective, and independent* of us. The objectivity and truth-connectedness of relevant evidence is similarly presumed by physical scientists (e.g., medical researchers, chemists, physicists, etc.) seeking the causal factors for physical processes.

There is another ordinary sense of 'relevant' that is *not* being used in condition 3. It is sometimes said that what counts as 'relevant evidence' is any evidence that is 'somehow related' or 'might have some significance or probability' for why **p** is true or false. For instance, the fire investigators at the start of their inquiry will be interested in any *relevant evidence* that *may* have a bearing on why a fire started. Pragmatic interests guide investigators to what items are considered 'relevant.' In this second sense, *any* and *all* evidence (e.g., the electrical system, smoking materials, etc.) that *might* have some significance (or be truth-conducive) for why the fire started and are considered as items 'relevant' to determining the cause of the fire. With this sense of 'relevant,' *all* of the possible factors that *might* have some bearing on the cause of a house fire are considered.

We have now distinguished two senses of what 'relevant' evidence is, and make more precise (i.e., explicate) these senses, with the postulation of two definitions:

Narrow sense: The 'relevant' evidence for why **p** is true, is an explanation that is objective (and truth-connecting) for why **p** is true. The evidence 'relevant' for why **p** is true, is not something we decide.

Wide sense: The 'relevant' evidence for why **p** is true, is any evidence that might conceivably be related to the truth or falsity of **p**. The evidence that is 'relevant' for why **p** is true, is subject to pragmatics and probabilities.

This *wide sense* of 'relevant' is not the sense being used here.¹

Condition 4a: The Sufficiently Strong Evidence Requirement

Condition 4a is an internalist condition that acknowledges a normative and pragmatic contextual component to personal justification. This condition implies that it is *persons* (and not some independent objective criteria) that determines how much evidence is 'sufficient' for **S** to believe **p**. David Annis (1978) and Adam Leite (2004, 2005) argue in favor of this position, where the standards for 'strong evidence' will vary by context (e.g., upon how important it is for **S** to know **p**, as opposed to merely believing **p**). Many philosophers now agree that whether **S** knows, or is justified in believing **p**, depends in part on how much is at stake (i.e., risk) in a practical situation.

¹ Many philosophers are accustomed to the wide sense of 'relevant.' Gail Stine (1976) says that an alternative is 'relevant' only if there is some reason to think that it is true. Stewart Cohen (1988) says that "factors pertaining exclusively to **S**'s own evidence affect whether alternatives become relevant.... an alternative **h** is relevant, if **S** lacks sufficient evidence (reason) to deny **h**..." (p. 103). John Hawthorne (2004) says that "What counts as relevant depends on the interests and intentions of the user" (p. 55).

In Keith DeRose's (1992) bank case and Stewart Cohen's (1999) airport case, whether **S** *knows* 'the bank is open' or 'the time of plane arrival' will depend upon **S**'s personal circumstance and pragmatic normative standards for possessing 'sufficient' evidence.

Condition 4a requires **S** to have at least some minimal reasons for a belief when questioned sincerely by someone seeking the basis of a belief. If there are actual undermining factors present, and if **S** was made aware of them, then **S** should be able to resolve any possibilities that imply $\sim p$ with a high degree of probabilistic and psychological certainty. In practice, **S** only needs to respond to counter-possibilities in the context of a properly motivated challenge. A challenge must be earned by finding specific reasons for questioning either the truth of the belief or **S**'s reasons for holding it.

Condition 4a endorses fallibilism. Fallibilism suggests that '**S** can know **p** upon strong reasons, but **S**'s strong reasons for believing **p** do not guarantee the truth of **p**.' It implies that **S** must *fallibly rule-out and resolve* actual and logical possibilities that could undermine (or defeat) the premises for believing **p**. With condition 4a, **S** must tacitly or explicitly *discard* improbable undermining possibilities, and *assume* them false, without *always* knowing them false. Skeptical possibilities are subjectively ruled out (without denying them as 'possibilities') in a pragmatic and normative fashion. Two statements of 'fallibilism' in terms of 'justification' and 'evidence' are as follows:

- (a) For some **p**, it is possible for **S** to know that **p** even if **S** could have exactly the same personal justification for believing **p** when **p** is false.
- (b) For some **p**, it is possible for **S** to know that **p** even if one's evidence for **p** does not make certain the truth of **p**.

We can further assume that an evidential proposition **e** can be (1) true and relevant for why **p** should be believed, or (2) true and coincidental, and irrelevant for why **p** should be believed or (3) false and irrelevant for why **p** should be believed. This third possibility allows that **S**'s body of *evidence* (or reasons) for believing **p** might contain some *false* propositions. Again, persons sometimes possess minor or insignificant false beliefs as 'evidence' (or reasons) for believing **p**.

Condition 4b: The No-Defeaters Requirement

The first sentence of condition 4b maintains that in order for **S** to have knowledge, there cannot exist additional evidence not possessed by **S**, such that if this evidence were possessed by **S**, this evidence would undermine **S**'s personal justification, causing **S** to seriously doubt **p**. Whether **S** knows depends on whether there actually exists true propositions that when considered as evidence would undermine **S**'s belief in that context. The second sentence states that *if* there does *exist* undermining evidence **e2** (no matter whether **S** actually considers it), *then* **S** needs to be able to resolve the counterevidence with evidence **S** already has, or seek additional evidence (if the proposition is of importance) in order to maintain knowledge of **p**. This is a subjunctive conditional about *what S's belief would be* if contrary-to-belief evidence **e2** were to be introduced. The truth value of a subjunctive conditional is determined by the content of its components and is somewhat speculative. The truth of such a conditional is judged by imagining circumstances that are close to the ways things actually stand, and how things would go in alternative circumstances. But as Max Baker-Hytch and Matthew Benton (2015) argue, there is no uniform way of precisely stating what undermining evidence is.

Condition 4b implies that if there exists undermining evidence, even if it is *misleading* and *inaccessible* to **S**, and if **S** is (or would be) *unable* to discard this undermining evidence so as to retain a firm belief that **p**, then 4b is unsatisfied and **S** does *not* know **p**. For many philosophers it is unacceptable that spurious undermining evidence could block what are normally instances of knowledge, and the no-defeaters condition is rejected. For example, in the Tom Grabit case, **S** might be denied perceptual knowledge that ‘Tom stole a book’ because of a judicial oath by a mother about a (false) twin brother. On the contrary intuition, I argue that inaccessible misleading evidence might (sometimes) *prevent* knowledge. The problem of misleading evidence isn’t a counterexample to the PE definition, but this can’t be argued here.

The Two Senses of 'Justification'

An important feature of the PE definition of knowledge is that the term 'justification' is eliminated. *We can describe what 'knowledge' is, without using the word 'justification' at all.* But why eliminate the word 'justification' from the definition of knowledge? The crucial reason for eliminating the word 'justification' from the definition of knowledge is that *the term is ambiguous.* The term conflates two concepts. As suggested by Mylan Engel, Jr. (1992) and Michael Williams (2001), the word 'justified' is capable of being understood in more than one way. When speaking of epistemic justification, we can speak either about a *person S* being justified in holding a belief, or about a *belief p* being justified. The first sense is about what it takes for a person **S** to be justified in believing a proposition, and the second sense is about whether a belief **p** has the property of being justified from an external perspective.

Let us define these distinctions:

Personal justification (sense PJ): 'S is justified in believing **p**' if **p** is acquired as an immediate non-inferential belief, or **S** has reasonably (non-recklessly) acquired strong evidence and used good inferential reasoning for believing **p**.

Belief justification (sense BJ): '**p** is a justified belief' if **p** is believed from inductive evidence (or deductive, abductive reasoning) that is relevant (i.e., truth-connecting, pertinent, non-defective) for why **p** should be believed true.

Let us continue to elaborate.

Personal justification (PJ) is where *persons are judged*, given their internal set of evidential premises. When evaluating whether a person is justified in believing **p**, we are asking whether they are warranted in believing **p**, given the evidence they possess. In cases of perception, 'non-inferential' beliefs are normally taken to be spontaneously justified (e.g., seeing a chair in a classroom). Otherwise, **S** is usually judged to be justified, or warranted in believing **p** if **S** is diligent in assessing available evidence, and if good inductive, abductive, or deductive reasoning are used.

Belief justification (BJ) is about a proposition **p**, and whether the *belief* is justified. Being knowledge seekers, we wish to have beliefs based upon relevant, truth-connecting evidence for why **p** is true. If **S** believes a proposition **p** upon a set of premises (or reasons) that are not relevant for why **p** should be believed, then **S** doesn't possess a justified belief. Whether a *belief is justified*, is *external* to **S**'s subjective conception of her epistemic situation. If **S** possesses a justified belief, this implies that there is no defeating (i.e., rebutting) facts that render **p** false.

While the PE definition excludes the term 'justification,' condition 3 implies the 'BJ sense' and condition 4a implies the 'PJ sense' of justification. Let's examine a case study that involves Mary's belief **p**: 'Some people in my workplace own a Ford truck.'

(#1) Mary believes that 'Someone in my office owns a Ford truck.'

Suppose that Mary works at a small office. Suppose that Mary has strong evidence for **p**: 'Some people in my workplace own a Ford truck.' Mary believes that five people in her workplace own Ford trucks. She enumerates her evidence: Tom owns a truck, Jennifer owns a truck, Fred owns a truck, and so on. These five fellow employees have been long-time Ford truck owners, and Mary knows them well. In fact, Mary saw two of these co-workers park Ford trucks in the parking lot this morning. We will examine three different material situations.

Situation #1: Unknown to Mary, one of the five workers sold his truck yesterday, and presently owns no vehicle. Does Mary still know that *some* people in her workplace own a Ford truck? Does one false belief (i.e., premise) make it false that Mary knows **p**? From intuition and the rules of logic, Mary still knows **p**. She knows that someone in my workplace owns a Ford truck, even if one of her evidential premises is false. Mary's strong evidence can consist of four true (and relevant) premises for why **p** should be believed without every single proposition being true or relevant for why **p** should be believed. In order for condition 3 to be satisfied, Mary's reasoning need not be flawless, but it must be substantially truth-connecting for why **p** should be believed.

Situation #2: Suppose unknown to Mary, four of her co-workers suddenly sold their Ford trucks, so that only one remaining co-worker owned a Ford truck. Even if four

out of five evidential premises are false, *condition 3 would still be satisfied* because *Mary would still possess objectively (weak) relevant evidence* for the claim that 'some people at my workplace, own a Ford truck.'

But does Mary *know p* despite the fact that most of her reasons for believing *p* are false, and irrelevant for why *p* is true? For most of us, if Mary believes this true *p*, with just one of her five premises being true, we are inclined to say that Mary no longer knows *p*, based upon a violation of condition 4b. If Mary was made aware that four co-workers recently sold their trucks, and if Mary hadn't confirmed the status of the fifth person's ownership recently, she would likely have immediate doubt whether anyone in her workplace owned a Ford truck. Mary would likely deny knowing *p* if she became aware that four out of her five reasons for believing *p* were false. In this case, even though *p* is true, *S* believes *p*, *S* has strong evidence to believe *p*, and *S* has a relevant reason for believing *p*, *S* does not know *p* because of undermining counterevidence not considered. Because Mary is unaware of undermining (but not defeating) counterevidence, knowledge conditions 1, 2, 3, and 4a are all satisfied, but condition 4b is not.

Situation #3: Now consider an even more drastic possibility. Suppose that unknown to Mary, all five of her co-workers suddenly sold their Ford trucks, so that *none* of her known remaining co-workers owned a Ford truck. Suppose too, that earlier in the day, unknown to Mary, her boss hired a new employee. This new co-worker not yet acquainted with Mary, in fact owns a Ford truck. Even though Mary possesses a personally justified (based on her evidence) true belief that *p*: 'Some people in my workplace own a Ford truck,' Mary doesn't know *p*. In this case, even though *p* is true,

and Mary believes **p**, has strong evidence to believe **p**; she no longer has reasons that are relevant (i.e., substantially related) for believing **p**. Mary does not know **p** because knowledge conditions 3 and 4b are both violated, which is a standard Gettier situation.

In contrast to a 'Gettier case' in situation #3, the circumstances of situation #2 is better termed as a 'Harman case.' Gilbert Harman (1973, pp. 142-150) presents a series of examples where **S** possesses belief **p** and has strong evidence for believing **p**, and **p** is true, and despite the fact that **S**'s evidence is relevant for why **p** should be believed, there exists evidence (sometimes misleading) that if accessible to **S**, might lead **S** to doubt (and not know) that **p**. Most philosophers conflate Harman cases (4b violated) as Gettier cases (3 & 4b violated) oftentimes under the concept of 'luck,' Gettier cases involve 'intervening luck' and Harman cases involve 'environmental luck':

(1) In Gettier cases, **S** is personally justified (with strong evidence) in believing **p**, but **S**'s reasons for believing **p** are *irrelevant* to the truth of **p**. There exist other reasons (or intervening circumstances) that explain *why* **p** is true. The truth of **p** is just coincidentally (i.e., accidentally) true, given **S**'s *evidence*.

(2) In Harman cases, **S** is personally justified in believing **p**, and **S**'s *reasons* for believing **p** are relevant to the truth of **p**, but there exists evidence (not accessed by **S**) that would weaken (i.e., undermine, undercut) **S**'s belief that **p**. The truth of **p** is just coincidentally (i.e., accidentally) true, given **S**'s *environment*.

In both kinds of cases, **S** lacks knowledge, and is just 'lucky' to have a personally justified true belief that **p**. But rather than consolidate these two kinds of cases under the concept of 'luck,' as Duncan Prichard (2009) does, we will keep them as distinct.

(#2) Do You Know Where Your Car is Parked?

Suppose that you have legally parked your car on the side of a street thirty minutes ago, in a large city in the United States. You are asked, 'Do you know where your car is?' You respond **p**: 'Yes, I know that my car is parked at the corner of Maple and Nelson, about six blocks away.' We are inclined to say that you know where your car is, if your car is stationary where you parked it, and if your memory is good, and if you have some premises (or implicit memories) for recalling where your car is now parked.

However, it is true that thousands of cars are stolen every day in the major cities of the United States. Do you know that your car has not been stolen? Many people would respond that they would *not know* whether their car had been stolen. Is it consistent to both *know* where your car is, and admit that you *don't know* whether it has been stolen or not? The PE definition shows how this is consistent.

Material Case #1: S knows where his car is. This is an example of a normal situation.

- (a) **S**'s car is exactly where **S** believes it is parked: The car is parked at the corner of Maple and Nelson Streets.
- (b) **S** believes that the car is parked at the corner of Maple and Nelson Streets.
- (c) From memory of where he parked the car, **S** has premises describing (or memory for finding) where his car is. **S** acknowledges that there is a small statistical chance that his car has been stolen, or in some other way moved, but **S** has no significant doubt that the car is at the corner of Maple and Nelson.
- (d) There exist no undermining facts (e.g., the existence of a nearby car thief, or another car crashing into it and moving it, etc.) that would weaken **S**'s belief **p**.

In such a case, we affirm that the conditions for S's knowing where his car is parked are obtained. Although S doesn't possess infallible evidence (premises) that eliminate the logical possibility of the car being moved, the material conditions for S's knowing **p** can be described, as we are assuming that **p** (i.e., the car is parked at Maple and Nelson) is true. In this example, S is not required to *know* that there are no undermining facts (e.g., a car thief has spotted the car) nor defeating facts (e.g., a car thief is now moving the car), there just exist no facts that undermine S's premises for **p**.

Material Case #2: An example of additional evidence that may violate condition 4b is the fact that an auto thief has actually entered into S's car and intends to take it. If S were to be made aware of this fact about the car and its environment, this *might* constitute unresolved undermining evidence for S that brings doubt about where his car is.

However, if S was confronted with this evidence **q**, and S believes that the car with its sophisticated anti-theft devices, are sufficient to stop any would-be thief, this fact when brought to S's attention, might not constitute undermining evidence for S (since anti-theft devices are counterevidence). So, condition 4 and the other three conditions would all be in place, and S continues to know where his car is parked.

Material Case #3: An example of a fact that would eliminate S's knowledge of where his car was parked, would be if a joyrider stole S's car and drove it around the block and quickly parked it in the same spot at Maple and Nelson, and then fled. If it is true that the car is parked where S believes it is parked, but S is unaware of the joyrider; then S would not know where his car was since conditions 3 and 4b are violated. S's premises for

having a true belief where his car is parked would be coincidental to where his car is. (This is a Gettier case).

Material Case #4: Examples of defeating facts that would eliminate S's knowledge of where his car is parked, is if a thief moved the car to another location, or if S's car was hit by another vehicle and pushed it away from its original location. Not only would conditions 3 and 4b be violated; condition 1 would be violated, since **p** would be false.

With these four case situations, S knows **p** if all four PE conditions are satisfied, as in case #1. Whether S knows where his car is parked is contingent upon S's belief-forming mechanisms (e.g., perception, memory) and the material conditions that surround S. It seems intuitive that *when S knows p*, S cannot always possess relevant reasons to know the falsity of every proposition that is inconsistent or undermining to belief **p**.

These material cases illustrate our intuitions about *when it is true* that 'S knows where his car is parked.' They jointly illustrate an intuitive principle:

The Axiom of Knowledge Contingency: *When S knows p*, S is not required to *know* that there are no unconsidered facts (or ruled-out possibilities), that defeat **p**, or unconsidered facts that would undermine S's belief, leading to doubt that **p**, there *just are* (i.e., exist) *no unconsidered facts* (i.e., true propositions) that defeat (or substantially undermine) S's premises for believing **p**.

This intuition is consistent with the following epistemic principles found in the literature:

- (1) The *acceptance* of *fallibilism* where S can know **p** upon strong evidence, based upon assumptions against counter-possibilities that S takes for granted, without knowing these counter-possibilities to be false.

(2) The *denial* of the *KK principle* that says, 'For any **p**, if one knows that **p**, then one knows that one knows it.' It is intuitive that in order for **S** to know **p**, **S** does *not* need to *know* that all defeating counter-possibilities are false; they just are false. To know **p**, you are not required to 'know that you know **p**.' Knowing that one knows would require **S** to be omniscient of all pertinent facts related to a proposition, including an externalistic perspective of one's own material situation.

(3) The *denial* of the *epistemic closure principle* which states 'If **S** knows **p**, and **S** knows **p** entails **q**, then **S** is in position to know **q**.' The standard principle of propositional closure (*modus ponens* across true propositions) is truth-preserving, but the principle of epistemic closure (*modus ponens* across knowledge states) is not truth-preserving. This will be explained below.

(#3) Henry and the Barn

An example popularized by Alvin Goldman (1976): While driving his car in an apparently normal rural setting, Henry migrates into an area with 99 barn facades where there is just a single genuine barn. Henry happens to see the genuine barn and says, 'there is a barn.' Most philosophers agree that although Henry possesses a personally justified true belief, he doesn't have knowledge. Why doesn't Henry know that he sees a barn?

The PE definition suggests that the fact that Henry has a true belief based upon truth-connecting and pertinent evidence is not sufficient for Henry to know that 'There is a barn.' The reason why Henry doesn't know is that there is undermining evidence that Henry is unaware of in this situation. If Henry was to become aware that the other 99 nearby objects were barn facades, then Henry would acknowledge that he couldn't

distinguish a real barn from a barn facade. Henry would concede that his belief that he was presently seeing a barn had weakened considerably. He would admit that he does *not* know whether he is seeing a barn or not. Condition 4b is violated, a Harman case.

Let us contrast this explanation with Zagzebski (1994). She states that Gettier problems such as ‘Brown in Barcelona’ and ‘Henry and the Barn’ arise when it is only by *chance* that a justified true belief is true. The concepts of ‘luck,’ ‘accident,’ ‘easy,’ and ‘chance’ are incorporated into her epistemology. Zagzebski believes that since **p** might *easily* have been false in these cases; these true beliefs are not instances of knowledge. Her explanation of why Henry lacks knowledge, is that it is an ‘accident’ that Henry’s visual facilities *normally reliable* in this sort of situation, are *not reliable* in this particular situation and it is a ‘second accident’ that Henry is looking at a barn. She says that an accident of ‘bad luck’ is cancelled out by ‘good luck.’

Zagzebski characterizes ‘process reliabilism’ as an externalist theory that **S** is justified in believing **p** when **p** is formed in a reliable truth-conducive manner. But she says that there is *no guarantee* that such justified beliefs are true, and that a breakdown in the connection between a reliable belief-forming process and the truth is possible. When this happens, even if **p** is true, **S** doesn’t have knowledge. This intuition is correct.

But in criticism of Zagzebski, is it *true* that Henry’s visual facilities, normally reliable in this sort of situation, are *not reliable* in this situation? On the contrary, Henry’s visual facilities *are reliable* (i.e., well-functioning and relevant) to this situation. It is *not* an ‘accident’ that Henry is *seeing* a barn. A better explanation is that the circumstances of a hostile environment (99 visually indistinguishable barn facades) is

what provides the coincidental (i.e., statistically improbable) 'good luck' of seeing a real barn. Henry's lack of knowledge is explained as a 4b violation although PE 3 is satisfied.

Zagzebski's argument against 'internalist' and 'externalist' theories of knowledge has been influential as part of the widespread dismissal of a necessary and sufficient conditions approach to knowledge. In contrast, from the perspective of the PE theory, the historical attempt to specify the conditions (including specification of 'positive,' 'undermining,' 'misleading,' and 'defeating' evidence) required to *guarantee* 'S knows p' was a mistake. A conditional approach need not attempt to stipulate Gettier-proof or Harman-proof conditions for attaining knowledge. It is admitted that these possible situations will never go away. What is needed is an *explanation* of these cases.

(#4) Lottery Paradox

Suppose that there is a lottery where 100,000,000 tickets are issued. S purchases one ticket, and the lottery will be held later this week. S is well-aware of the statistical probability that her ticket won't win, and on this basis, dismisses the small probability of the ticket being a winner, stating p: 'My ticket is a losing ticket.' Does S know p?

The widely accepted response is that S does *not* know that her ticket is a loser. The PE definition explains why. The reason that S doesn't know p, despite the high probability that p is true, is that S's evidence for believing p isn't connected to any particular ticket (including her own). With a fair drawing S can never be in a position to possess relevant premises for believing p is true. Knowledge condition 3 is not satisfied. Even after S examines the statistical evidence, and resolves and discards the small chance that the ticket may win, and has a personally justified belief, S cannot know p.

(#5) Another Lottery Paradox: 'S knows she will never be a multimillionaire.'

Another lottery paradox has been developed by John Hawthorne (2004) and others. Its form is described by Elke Brendel and Christoph Jäger (2004, p. 148):

Let us assume **S** bought a ticket in a fair lottery and the chances of this ticket winning are very low-- 1: 10 million. If **S** is the lucky winner, she will get ten million dollars. Although there is overwhelming statistical evidence for the belief that **S**'s ticket will lose, many people share the intuition that **S** nevertheless does not know that her ticket will lose. Let us assume furthermore that, given **S**'s meager income and her lack of rich relatives, **S** *claims to know* that she will never be a multimillionaire. Now we have a problem: **S**'s knowing that she will never be a multimillionaire seems to imply her knowing that she will not win the lottery-- which contradicts the intuition that **S** fails to know that she will lose.

Thus, we have the following:

- (1) **S** knows that she will never be a millionaire.
- (2) If **S** knows she won't ever be a millionaire, she is in a position to know she won't win a major lottery prize. (Epistemic closure).
- (3) Conclusion: **S** is in position to know she won't win a major lottery prize.

One simple resolution to this paradox is to state that premise (1) is false, and **S** *doesn't* know **p** ('she will never be a multimillionaire'). Since **S** has played the lottery, and is awaiting the result, it is possible that **S** may win the lottery, and so **S** doesn't know she'll never be a millionaire. (This is plausible and is ultimately correct).

But this response might not seem entirely intuitive. It might seem plausible that **S** can know **p** (i.e., 'she will never be a multimillionaire') based upon her strong reasons: (1) I have a low-paying job with no hope or desire for advancement, (2) I have no rich relatives from which to expect inheritance, nor pending legal settlement case, nor any reason for expecting financial windfall, and (3) I know that the chances of my winning the lottery are infinitesimal. (4) Therefore, I *know* I will never be a multimillionaire.

These three reasons certainly constitute strong evidence (and personal justification) to support **S**'s belief that "she will never be a multimillionaire." But are these reasons 'relevant' for **S**'s claim to *know* that she *won't* be a multimillionaire before the drawing? *Not* with the *narrow sense* of 'relevant.' **S cannot** possess relevant evidence (to satisfy condition 3) to *know* that her ticket is a loser.

However, given **S**'s financial circumstance, **S** can have 'relevant' reasons (in the *wide sense* of relevant) to believe (but not know) that she will never be a multimillionaire. With a wide sense of 'relevance,' **S**'s immediate concern is with financial reality and probabilities that lead **S** to believe that she won't be a multimillionaire. With an infinitesimal probability of winning the drawing, **S** concedes that the potential to win the lottery is 'irrelevant.' She discards the existing possibility of winning the lottery as being 'relevant' (in the wide sense) to her financial life.

This paradox is explained by recognizing that there are *two* ordinary language senses of 'relevant.' In one sense, **S cannot** have relevant (truth-connecting) evidence to *know* she won't win the lottery and become rich, but in another sense, **S** has 'relevant' (wider, related) probabilistic evidence that there is little reason to believe that winning the

lottery would become true. **S**'s claim to *know* that she'll *never* be a multimillionaire is *false*. In sum, both premises of the paradox are false, and so is the conclusion.

(#6) The Rubbish Chute Case

In determining whether '**S** knows **p**,' the contingencies of **S**'s external environment take precedence in deciding whether **S** knows (or not). A simple illustration of this can be drawn from an example from Ernest Sosa (1999, p. 145):

On my way to the elevator, I release a trash bag down the chute from my high-rise condo. Presumably, I know my bag will soon be in the basement. But what if, having been released, it still (incredibly) was not to arrive there? That presumably would be because it had been snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance.

In response to this case study, we ask does **S** *know* **p**: 'My bag is in basement?' If so, why? If not, why not? With the PE definition of knowledge, like the parked car case, whether **S** knows **p** is in part determined by one's material situation. We are inclined to say that you know that your bag is in the basement, given that upon releasing the bag into the chute, it functions in a normal way for the bag to reach the basement. **S** typically knows **p** (my bag is in the basement) in a well-maintained high-rise.

But what about the possibility (or the actuality) of a snag caused by the initial improper installation of a portion of tube, or an obstruction from a new squirrel's nest? If conditions 1, 3, and 4b are all false (i.e., the bag is snagged in the chute), then **S** fails to know **p**. If **S** is aware of the possibility of a malfunction (in another situation), and for some reason cannot dismiss this counter-possibility **q**, then **S** doesn't know **p**. In the

ordinary case of knowledge, **S** fallibly discards undermining possibilities (e.g., a snag), and there *just are* (i.e., *exist*) no other unconsidered facts (i.e., true propositions) that defeat (or substantially undermine) **S**'s premises for believing **p**. This is analogous to knowing a sports score by reading a newspaper (e.g., there being no misprint). With Gettier and Harman cases, it has been demonstrated that 'personally justified true belief' isn't sufficient for knowledge, and that *externalist conditions* (3 and 4b) are needed.

(#7) Does S know that she has purchased five pounds of grapes?

Suppose that **S** is at a respected grocery store purchasing some grapes. The scale reads 5 pounds. Given the possibility that the scale might be malfunctioning (and the weight is only 4.5 pounds), does **S** know that she has five pounds of grapes?

Like most of us, **S** cannot easily detect the difference between 4.5 and 5 pounds of grapes, but she trusts her grocer. But if **S** assumes the scale is reliable, and believes herself to know that she has five pounds, then **S** is *alleged* to possess this entailment:

- (1) **S** knows that she has five pounds of grapes.
- (2) If **S** knows that she has five pounds of grapes, then she knows that the scale isn't malfunctioning.
- (3) Therefore, **S** knows that the scale isn't malfunctioning.

But **S** *doesn't* outwardly *know* that the scale isn't malfunctioning! On the PE definition, *if* **S** knows that she has five pounds of grapes, then **S** has dismissed the possibility of malfunction (as part of her background evidence) and if there is no material malfunction (satisfying PE 4b), then **S**'s reading a reliable scale satisfies PE 3, and **S**'s true belief ('I have five pounds of grapes') is knowledge. All conditions are satisfied. But does **S**'s

presumed knowledge that she has five pounds of grapes *entail* that she knows the scale isn't malfunctioning? Intuitively, it *doesn't* seem that **S** knows that the scale *isn't* malfunctioning; it is just *contingent* that it isn't. Of course, **S** might verify the weight of the grapes with a reliable scale at her home, and then know that the store's scale was accurate, and if this is the case, then this just confirms that **S** already knew **p**.

Can I Know that I'm not a BIV?

The PE definition has major implications for the radical skeptical argument. With the *denial* of the epistemic closure, it is acknowledged that the '*skeptical hypothesis*' in the argument is *true*. The skeptical argument makes reference to 'I' assuming a single conscious thinker, and **o** that designates an ordinary proposition (e.g., I have two hands):

(#1) I do not know that 'I am not a brain-in-a-vat.' (Skeptical hypothesis).

(#2) If I do not know that 'I am not a brain-in-a-vat,' then I do not know **o**.

(#3) Therefore I do not know **o**.

With a narrow sense of 'relevant evidence,' I argue that **S** *cannot* have relevant (i.e., truth-connecting) evidence to deny the truth of premise #1. If **S** cannot possess relevant evidence to know 'I am not a BIV' is false, then premise #1 above is entailed true.

The truth of the 'skeptical hypothesis' can be argued:

(1) In order to know the BIV possibility is false, **S** *must* possess relevant reasons for believing the BIV possibility is false.

(2) For **S** to have 'relevant' reasons for believing that the BIV possibility is false, those reasons must be truth-connecting (i.e., germane, material, applicable, pertinent) for why the BIV possibility should be believed false.

(3) As an assumed material condition of the BIV possibility, **S** cannot possess relevant (i.e., truth-connecting) reasons to believe that the BIV possibility is false, because any objective and independent evidence for testing the falsity of the BIV possibility is cognitively inaccessible. The phenomena of one's experiences are the same, no matter whether one is or isn't a BIV, so **S** can never have adequate evidence (or reasons) to deny the physical possibility of being a BIV.

(4) If **S** *cannot* have relevant reasons for believing that the BIV possibility is false, then 'I do not know that I am not a BIV' is true.

(5) Therefore, the hypothesis is true (i.e., 'I don't know that I am not a BIV').

In this deductive argument: Premise 1 assumes that PE condition 3 must be satisfied, as necessary for knowledge. Premise 2 defines what a 'relevant' reason is. Premise 3 states the nature of the BIV possibility. Premise 4 states that if **S** is unable to possess relevant reasons for believing the BIV possibility is false, then 'I do not know that I am not a BIV.' The truth of the skeptical hypothesis is entailed by the truth of the four premises. Given that the skeptic's premise #1 is proven true, it is implied that premise #2 (i.e., epistemic closure) must be false, if the skeptic's conclusion is to remain false.

But with the explicit acceptance of the skeptical hypothesis, a major concern remains: How can I know that there is an external world, and that I am an independent conscious entity? This was Descartes' problem, and the question persists. The question of 'How can **S** know that there exists an external world?' has a concise answer. The PE definition states **S** can have knowledge that 'there exists an external world' *if* the following material conditions are satisfied:

- a) S believes that 'there exists an external world.'
- b) S has many premises with references to experiences of pain, happiness, interaction with other people, intentional action, touching external objects, etc. that are relevant (wide sense) for believing there exists an external world.
- c) S believes that the existence of an external world is the best explanation for her premises and dismisses the BIV possibility as being an unlikely defeater.
- d) It is true that there is an external world. There is no defeater (e.g., evil demon).

With the contingent satisfaction of these four conditions, it happens that S (in fact) knows that there exists an external world. Further, it is true that 'I am not a BIV' if these propositions are true, even if I don't *know* that I'm not a BIV. So, *it is possible to know that I am an independent living creature in an external world (if conditions a-d obtain), it is impossible for me to know that I am not a BIV.* Against DeRose (1995), it is *not* abominable that "I know that I have hands, but I don't know that I'm not a handless BIV."

The Distinction between Epistemic Closure and Propositional Closure

Why does premise #2 (that presumes epistemic closure) in the standard skeptical argument seem plausible to many philosophers? First, there is the intuition that in order to know any proposition, it seems that you should know of your own self-existence as an independent entity in a natural external world. Knowing that 'I have hands' implies that I'm not a BIV. If you don't know that you're not a BIV, how can you be sure you know anything? For Stewart Cohen, epistemic closure "expresses something deep about the nature of knowledge. How could you know **p** and know that **p** entails **q**, and yet fail to (at least be in a position to) know **q**?" (1999, p. 64). He further states that closure

"seems to me to be something like an axiom about knowledge" (2002, p. 312). In addition, for those who endorse epistemic closure, it seems intuitive that by validly deducing the consequences of previously known premises, we can extend our knowledge.

But let us distinguish between 'propositional' and 'epistemic' closure. The principle of *propositional closure* is found in mathematics and logic. Closure is a property of sets. The members of sets in propositional logic are that of 'true propositions.' A set of true propositions is closed under a valid deductive argument because true propositions will entail other truths. When reasoning by *modus ponens* across true propositions, it is accepted that if **p** is true, and if **p** implies **q**, then the truth of **q** is entailed by the truth of **p**. In contrast, in the case of *epistemic closure*, the members of the set are 'the items of knowledge.' With the principle of epistemic closure, it is maintained that if **S** knows **p**, and **p** implies **q**, and if **S** is aware of this implication, then **S** knows that the entailed proposition **q** is true. In opposition, I agree with a minority of well-known 'anti-closure' theorists who argue that the adaptation (or extension) of propositional closure to epistemic situations among known propositions using *modus ponens* is *not* always a truth-preserving form of inference. With an interest in accepting propositional closure, but denying epistemic closure, I contrast the following definitions:

Principle of Propositional Deductive Closure (*Modus ponens* entailment across *true propositions*): If **p**, and **p** entails **q**; then **q**.

Principle of Epistemic Deductive Closure (*Modus ponens* entailment across *knowledge states*): If **S** knows **p** and **S** knows **p** entails **q**; then **S** knows (or is in a position to know) **q**.

The first principle is uncontroversial. This mode of argument states that if proposition **p** is contingently true, and if **q** is true in every possible situation where **p** is true, then from these two propositions, it is implied (or entailed) **q** is true:

- 1) **p**.
- 2) If **p**, then **q**.
- 3) Therefore, **q**.

In contrast, the epistemic closure principle across knowledge states says that 'Necessarily, if **S** knows that **p**, and **S** knows that **p** entails **q**, then **S** is in a position to know that **q**.'

Modus ponens across knowledge states has this argument form:

- 1) **S** knows **p**.
- 2) If **S** knows **p** and that the truth of **p** entails the truth of **q**, then **S** knows **q**.
- 3) Therefore, **S** knows **q**.

The Argument against Epistemic Closure

In the parked car situation, it was shown that it seems intuitive (to many people) that *when* **S** knows **p**, **S** cannot always possess relevant reasons to know the falsity of every proposition that is inconsistent or undermining to belief **p**. Compare two arguments (A) and (B) where the validity of *propositional* closure is illustrated:

(A) An argument from propositional closure:

- 1) **S**'s car is parked on Nelson Street.
- 2) It is contradictory for **S**'s car to be parked, and simultaneously being stolen and driven down a street.

3) If S's car is parked on Nelson Street, then the car has not been stolen and driven down a street.

4) Therefore, S's car has not been stolen and driven down a street.

(B) Another argument from propositional closure:

1) S *knows* that his car is parked on Nelson Street.

2) S *knows* that it is contradictory for his car to be parked, and simultaneously being stolen and driven down a street.

3) If S *knows* that his car is parked on Nelson Street, then the car has not been stolen and driven down a street.

4) Therefore, S's car has not been stolen and driven down a street.

Each of these two arguments is valid by *modus ponens* across propositions. If propositions 1-3 are assumed true, the truth of those premises entails that its conclusion is necessarily true. In both arguments, the conclusion makes no assertion about what S *knows*. That S *knows* that his car is parked on Nelson Street *entails* that the car has not been stolen.

The defender of epistemic closure extends from argument (B), the assumption that the following argument is also valid where the first two premises remain the same, but where '*knows*' is inserted in the consequent of the third premise and in the conclusion:

(C) An argument from epistemic closure:

1) S *knows* that his car is parked on Nelson Street.

2) S *knows* that it is contradictory for his car to be parked, and simultaneously being stolen and driven down a street.

3) If **S** *knows* that his car is parked on Nelson Street, then **S** *knows* the car has not been stolen and driven down a street.

4) Therefore, **S** *knows* that his car has not been stolen and driven down a street.

But as we have asked above, is this conclusion always the case? Does **S**'s knowledge of proposition **o** entail that **S** is in an epistemic situation to always know the *falsity* of *any* proposition that contradicts **o**? Premise 3 is false, and the conclusion is false. Whether **S** knows where his car is parked is contingent upon the belief-forming mechanisms that lead to the belief, and the contingent material conditions that surround **S**.²

Fred Dretske (2014) offers a similar example. This is where **S** can know that she is seeing oranges (a kind of fruit) in a grocery store, with visual evidence; but still not have visual evidence (from a distance) to know that the objects aren't wax imitations of oranges. In this case, it is possible and plausible for **S** to know that 'I'm seeing oranges' but not know that 'I'm not seeing wax imitations of oranges.'

In order to for **S** to know 'where his car is parked' or 'the bag is in the basement,' or 'this is five pounds of grapes,' or 'I see oranges' **S** may not *know* that all possible defeaters are false. The PE condition 4a allows **S** to fallibly resolve possibilities that would otherwise undermine (or defeat) **S**'s premises for believing and knowing **p**. **S** may discard improbable possibilities (e.g., stolen car, snags, unreliable scales, wax oranges) and *assume* them false, *without knowing* these undermining possibilities are false.

² In experimental studies, John Turri (2015) has found strong support for anti-closure, concluding that rejecting epistemic closure "doesn't conflict with ordinary practice or common sense" (p. 14).

The Safety Principle

The 'grape weight' case study, above, poses a serious problem for 'safety':

(1) Safety Principle: If **S** knows **p**, then **S**'s true belief that **p** must be 'safe' in the sense that **p** couldn't have easily been false.

(2) Safety Condition: **S**'s true belief that **p** counts as knowledge only if in all (or most) nearby possible worlds in which **S** believes that **p** via **S**'s actual method, **p** is also true.

With the safety condition, **S** knows a true **p** only if **S** could not *easily* have falsely believed **p**. Safety explains why **S** doesn't know that he sees a sheep in a field (hidden from view, the actual sheep could have *easily* moved) and why Henry doesn't know that he sees a barn (because of a risky situation, the true belief could have *easily* been false).

But it seems that **S**'s environment when purchasing grapes is risky, and the safety principle seems to suggest that **S** *cannot* (ever) have knowledge (in the store) that she has five pounds of grapes. The principle that **S** knows **p** only if **p** is safe from error, in that **S** couldn't have easily believed falsely, significantly limits what is knowable. The status of conditions in 'nearby worlds' where **S** weighs a quantity of grapes is probabilistic. This skeptical (or otherwise indeterminate) result is intuitively unacceptable.

(#8) Does Professor S know that Professor Brown is in Room 222 at 4PM?

Professor **S** is a professor at a major university. He works with his colleague Professor Brown. Professor **S** believes that Professor **B** will present a lecture on his recently completed book at a meeting of ten graduate students and visiting professor at 4PM on a Friday afternoon in room 222. At 10AM on the same Friday morning, **S**

briefly chats with **B**, and **B** is prepared and eager to present his paper. **S** knows that **B** is dedicated, reliable, enthusiastic, punctual, and is a great speaker. Because of prior commitments, **S** cannot attend **B**'s reading, but strongly believes (upon strong evidence) that '**B** will be in room 222 at 4PM on Friday.'

However, at 11AM, the anticipated paper reading runs into problems. Several graduate students have unexpected circumstances that will prevent their attendance. The visiting professor cannot attend because of an airline delay. These parties communicate this to Professor **B** by phone or e-mail by 1PM on Friday. Around the same time, Professor **B** begins to feel slightly ill with a sore throat. He decides to postpone the presentation and phones the remaining students announcing the cancellation with the promise to reschedule in the future. The department staff is notified, the website is updated, and a sign is placed on the door of room 222 announcing the cancellation.

As it so happens, Professor **B** lives near campus, and it is a fifteen-minute walk to campus. At 3:45PM on that Friday, **B** remembers that he left his reading glasses in room 222 on campus, while inspecting the room earlier in the day. He walks to campus to pick up his eyeglasses and is in room 222 at 4PM. At this same time, Professor **S** thinks about his colleague, Brown, and says to himself, I know that '**B** is now in room 222.'

This is a Gettier situation where despite strong evidence and true belief, **S** *doesn't* know **B** is room 222. It is also clear that this is a case where **S**'s true belief couldn't have easily been false. It took an unlikely sequence of coincidences (graduate cancellations, airplane delay, sore throat, forgotten eyeglasses, timing of walk) for the lecture to be cancelled, and for **S**'s (true) belief not to be a case of knowledge.

But the safety principle indicates that this necessary condition for knowledge is satisfied because **S's** belief that '**B** is in now in room 222' is true and couldn't have easily been false. But it seems otherwise: although the condition is satisfied, **S** does *not* know **p** despite the fact that **p** is true. The satisfaction of safety incorrectly predicts **S** knows **p**.

Safety similarly suggests that in nearby worlds, **S** would *know* **p** ('**B** is in Room 222' at 4PM) because **p** would not easily be false. But the reasons for the truth of **p** in close worlds is that there *wouldn't* be this *number* of individual non-sufficient *possible events* that would lead **B** to cancel the lecture. **S** would know **p** in these close worlds. The problem with this, is that even if **S's** belief **p** is 'safe' in the actual world and **S** knows **p** in nearby worlds (by using the same method for acquiring reasons for a true belief in the actual world), this doesn't respond to the problem that in the actual world, **S's** (true) safe belief isn't known. Safety fails as an anti-luck condition.

Conclusion

In this essay it has been argued that a successful and explanatory definition of knowledge doesn't guarantee that knowledge exists. Given the accepted truth of the skeptical hypothesis, the definition of 'knowledge' hypothesized here, doesn't assure that **S** knows *any* **p**. The PE definition just states the necessary and sufficient conditions for the *possibility* of knowledge. Against the skeptical conclusion that denies the possibility of knowledge, it is shown that knowledge is *possible* and very obtainable, and the definition describes when it occurs. Of course, this short essay doesn't establish the truth of the PE definition since there are numerous other interesting case studies, and conceptual questions that need to be addressed, but this is a start in the right direction.

References

- Annis, David. (1978). 'A Contextualist Theory of Epistemic Justification.' *American Philosophical Quarterly* 15: 213-219.
- Baker-Hytch, Max and Benton, Matthew A. (2015). 'Defeatism Defeated.' *Philosophical Perspectives* 29: 40-66.
- Brendel, Elke and Jäger, Christoph (2004). 'Contextualist Approaches to Epistemology: Problems and Prospects.' *Erkenntnis* 61: 143-172.
- Cohen, Stewart. (1988). 'How to be a Fallibilist.' *Philosophical Perspectives* 2: 91-123.
- Cohen, Stewart (1999). 'Contextualism, Skepticism, and the Structure of Reasons.' *Philosophical Perspectives* 13: 57-89.
- Cohen, Stewart (2002). 'Basic Knowledge and the Problem of Easy Knowledge.' *Philosophy and Phenomenological Research* Vol. LXV, N. 2.
- DeRose, Keith (1992). 'Contextualism and Knowledge Attributions.' *Philosophy and Phenomenological Research* 52: 913-929.
- DeRose, Keith (1995). 'Solving the Skeptical Problem.' *The Philosophical Review* 104: 1-52.
- Dretske, Fred (2014). 'The Case Against Closure' in Steup, Turri, and Sosa (eds.). *Contemporary Debates in Epistemology*, (2nd ed). Malden, MA: John Wiley & Sons, Inc.
- Engel, Mylan Jr. (1992). 'Personal and Doxastic Justification in Epistemology.' *Philosophical Studies* 67: 133-150.
- Gettier, Edmund. (1963). 'Is Justified True Belief Knowledge?' *Analysis* 23: 121-123.
- Goldman, Alvin. (1976). 'Discrimination and Perceptual Knowledge.' *The Journal of Philosophy* 73: 771- 791.
- Harman, Gilbert. (1973). *Thought*. Princeton: Princeton University Press.
- Hawthorne, John. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Leite, Adam. (2004). 'On Justifying and Being Justified.' *Philosophical Issues* 14: 219-53.

Leite, Adam. (2005). 'A Localist Solution to the Regress of Epistemic Justification.' *Australasian Journal of Philosophy* 83: 395-421.

Prichard, Duncan (2009). *Knowledge*. New York: Palgrave Macmillan.

Sosa, Ernest (1999). 'How to Defeat Opposition to Moore.' *Philosophical Perspectives* 13: 141-153.

Turri, John. (2015). 'An Open and Shut Case: Epistemic Closure in the Manifest Image.' *Philosophers' Imprint* 15-2: 1-18.

Stine, Gail. (1976). 'Skepticism, Relevant Alternatives, and Deductive Closure.' *Philosophical Studies* 29: 249-261.

Williams, Michael. (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*. New York: Oxford University Press.

Zagzebski, Linda (1994). 'The Inescapability of Gettier Problems,' *The Philosophical Quarterly* 44: pp. 65-73.