Chapter Sixteen

Epistemology Since 1983

Abstract: A sketch of the recent history of epistemology is presented. Two modern approaches to epistemology have distinguished themselves: *epistemic contextualism* and *virtue epistemology*. David Lewis, Stewart Cohen, and Keith DeRose are advocates of a contextualist semantics for knowledge claims. Ernest Sosa, John Greco, Linda Zagzebski, and Duncan Prichard are all leading advocates of a virtue epistemology approach towards understanding knowledge. Two related principles, 'sensitivity' and 'safety,' have also been proposed and contrasted as a necessary condition for knowledge. These approaches followed from preceding 'reliabilist,' 'relevant alternatives,' and 'truth-tracking' theories. In this chapter, these approaches are critically assessed. Attention is centered on individual philosophers and their worldview beliefs, linguistic beliefs, use of definitions, use of analogies, and general methodology.

I. Introduction

It was Robert Shope's *The Analysis of Knowledge* (1983) that marked the unofficial end of the conceptual analysis of 'knowledge' in the form of necessary and sufficient conditions. Shope argued that since there were no successful conceptual analyses of knowledge over the past twenty years, that this shows that the approach is likely doomed to failure. It was time to try something new towards an understanding of knowledge that wasn't necessarily in the form of a definition.

There have simultaneously been two main alternative approaches. The first approach is that of *contextualism* as advocated by David Lewis (1996), Stewart Cohen (1999), and Keith DeRose (1992, 1995). The second major approach is the *virtue epistemology* of Ernest Sosa (2017), John Greco (2002, 2019), and Duncan Prichard (2007, 2009, 2019).[1] Preceding these two approaches were the relevant alternatives theories of Fred Dretske (1971), reliabilist theories of Alvin Goldman (1976), and Robert

[1] Two other major theories, that of Timothy Williamson who argues for a 'knowledge first analysis' and Hillary Kornblith's 'knowledge as a natural kind' are not discussed here. Kornblith (1993, 2002) argues that knowledge is a natural kind (to be analyzed the same way as other scientific kinds are). Kornblith's theory is outside the scope of this book. Williamson's views are discussed in chapter 18. For Williamson, knowledge is among the most fundamental psychological and epistemological (mental) states that there are. He argues that it is a mistake to analyze knowledge in terms of other, more fundamental epistemic notions, because knowledge itself, in at least many cases, is more fundamental. Williamson argues that there are other informative ways to characterize knowledge. He maintains that 1) Knowledge is the most general factive mental state, and that 2) S *knows* that p if and only if S's total evidence includes the proposition that p. Williamson endorses the safety principle but not as part of an analysis (2000, p. 126).

Nozick 's (1981) truth-tracking sensitivity theory as a necessary condition of knowledge. A 'safety' condition proposed by Ernest Sosa (1999) has been debated as an alternative necessary condition for knowledge (that is preferable to 'sensitivity'). Peter Unger's (1974, 1975) infallibilism was also very influential, especially for contextualists.

In this chapter these various approaches to 'knowledge' are critically assessed (mostly in footnotes). Otherwise, the background of the problems is described, and we cite the author's response using their own text. We will pay close attention to each author's worldview beliefs, linguistic beliefs, use of definitions, use of analogies, and general methodology.

II. Foundationalism and Empiricism

Contemporary theories of knowledge are mostly concerned with empirical knowledge.[2] Many epistemologists are foundationalists and to a large extent, *empiricists*. This is a position whereby **S**'s relevant reasons for having knowledge are ultimately traced back to 'basic' propositions of perception. Basic and non-basic propositions may be combined in inductive, abductive and deductive reasoning in order for an **S** to state a true proposition. It is assumed that knowledge *exists* either by 'foundation' or 'coherence.' It is also thought that a *constraint* of epistemological theorizing is to *prove* knowledge *exists* (rather than *explain* how it *can* occur). These platitudes are assumed:

> (1) Knowledge exists in concrete cases. We possess a great deal of knowledge about the world.
>
> (2) It is a Moorean fact that we *know* a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary. (David Lewis, 1996).

---

[2] In in earlier chapters, it was argued that besides perceptual (i.e., empirical) knowledge, that substantial knowledge also comes from (1) understanding natural and artificial languages, which is not directly obtained from sense experience. S's use of linguistic terms (written or spoken), syntax (mental), and concepts (mental) that are shared in natural language don't emanate from the senses. In other words, S's knowledge (or ignorance) about the nature of natural language, artificial languages, definitions, ethical assertions, speaker reference, etc., aren't from a sensual input. (2) there is knowledge from 'testimony' (oral, written), (3) there is knowledge by introspection**:** we can know something by examining our own psychological states, e.g., beliefs, values, moods (happy, angry, nervous, anxious, jealous). (4) There is an additional sense of 'know how' as knowing how to ride a bike, without propositional justification.

III. Formal Semantics

The contemporary discipline of epistemology intertwines with *semantic theory* mostly as a response to the radical skeptical argument. Semantic theory involves the notions of truth functional models, intensions, and extensions. Let's list elements from a (standard) philosophy of language that are important background assumptions:

> (1) A sentence *represents* the *state of affairs* that would make the sentence true; it represents its 'truth condition.' A sentence is true if a certain situation in the world obtains and is not true if the situation does not. This property of a sentence (its truth condition) is the core of a sentence's meaning.
>
> (2) To know the *meaning* of a sentence, is to know what the world has to be like, if the sentence were to be true. The meaning of a sentence is to be explained by relating it to the circumstances under which it would be true. A meaningful declarative sentence S represents the world as being a certain way and is either true or false.[3]
>
> (3) Formal semantic theories attempt to explain the *representational properties* of *linguistic symbols*. What needs to be explained is how certain linguistic forms contribute to a sentence's meaning, and ultimately its truth value. A truth-conditional approach to semantics attempts to 'interpret' the relationship which holds between *a sentence* and *the world*.
>
> (4) The semantic value of a sentence determines a set of truth conditions— a function from possible worlds to truth values. For sentences with indexicals and demonstratives in them, we don't speak of the semantic value of a sentence simpliciter: a sentence has a semantic value only relative to a context in which the designata of indexical expressions are fixed. The semantic value of a sentence is a function of the semantic values of its constituent expressions (and how those expressions are arranged), within a compositional theory of meaning.
>
> (5) The compositional theory of meaning states that to represent the meaning of a sentence (in precise structure) is to understand how its words contribute in a systematic way to the meaning of the sentence. Semanticists investigate in detail how the meaning of complex expressions are related to the meaning of the simpler expressions that they are constructed from.

---

[3] The characterization of *sentence meaningfulness* as *tied to its truth* effectively *excludes* the concept of 'prescriptive' utterances *as being meaningful*, since prescriptions are intended to be agreed-upon without being literally true or false. Critical comment: Do sentences *represent* truth conditions? Does a sentence (when asserted in context) have truth conditions, and thus meaning? The idea that language is essentially 'descriptive' in informational value is false, and thus can be challenged.

> (6) Semanticists normally use *models* to specify what sorts of things there are in the world, and with an assumed ontology, provide a formal symbolic interpretation of an object language (e.g., English). In such models, semanticists have an interest to investigate certain terms which remain invariant under changes in interpretation. It is generally assumed that certain basic expressions of English such as 'and,' 'or,' 'not,' 'every', etc. have a fixed invariant interpretation, i.e., they remain invariant from model to model, while indexicals such as 'I,' 'you,' 'now' and 'there,' have a variant contextual interpretation.[4]

The philosophy of language, as now practiced, is preoccupied to explain 'descriptive propositions' and their (necessary) consequences among other descriptive propositions. Truth evaluable assertions and their 'semantic composition' in *models* are the primary concern of explanation. Tim Button and Sean Walsh (2018) describe 'model theory' as follows:

> One of the most basic ideas in 'model theory' is that a structure assigns interpretation(s) to bits of vocabulary, and in such a way that we can make excellent sense of the idea that the structure makes each *sentence* (in that vocabulary) either *true* or *false*. Squint slightly, and model theory seems to be providing us with a perfectly precise, *formal way to understand* certain aspects of *linguistic representation*. It is no surprise at all, then, that almost any philosophical discussion of linguistic representation, or reference, or truth, ends up invoking notions which are recognizably model-theoretic (p. 3, italics added).

IV. Epistemic Necessities/Possibilities of 'Possible worlds'

Modern philosophers are greatly influenced by the German mathematician and philosopher G.W. Leibniz (1646-1716). With Leibniz's line of thought, it is claimed that there is knowledge of (1) *necessary truths*, e.g., '2 + 2 = 4,' which cannot fail to be true, and (2) *contingent truths*, e.g., 'Obama was the 44th President of the United States,' which, while true, might have been false. The truth of a sentence (or proposition) asserted in context is either contingent or necessary. This sort of necessity/possibility relation involves a commitment to the existence (by stipulation) of possible worlds in order to account for the epistemic status of modal assertions. The semantic properties of sentences

[4] Contextualists ask: *When* is the *sentence* "S knows p" *true* when stated by an attributor? Contextualists claim that the term 'knowledge,' like indexicals such as 'now' and 'there,' have variant truth conditions in relation to context. This is a semantic thesis. DeRose argues that the truth value of "S knows that the bank will be open on Saturday" is variant (i.e., either true or false) based upon (an attributor's) context.

(e.g., necessity, possibility, impossibility, contingency) are understood in terms of possible worlds and from the axioms of metaphysics, logic, and set theory.[5]

It is claimed that ordinary sentences are not used to merely talk about how things actually are, but they are also used to talk about how things (e.g., individual objects, ways of being) could have been different. The 'meaning' of a word or sentence (asserted in context) is understood as a referential concept and must be analyzed in terms of the notion of an extension. The referential force of a linguistic expression can extend beyond objects in the actual world to objects in other possible worlds.[6] These two facets of 'meaning' indicate that the meaning of a non-logical expression is a set theoretic entity, a function of possible worlds to extensions. It is a function that assigns to each possible world *the extension* the *expression has* when we *use it* in talking about that world. Thus, singular terms (e.g., proper names, definite descriptions) have as their meaning a function from worlds to objects; the meaning of n-place predicate is a function from worlds to sets of ordered n-tuples; and the meaning of a declarative sentence is a function from worlds to truth values.[7] A feature of propositions is that they can be said to be possible, necessary, impossible, or contingent.[8] The concept of 'possible world' is supposed to shed light on what is called the 'nature' of modality. To repeat, a proposition is 'necessarily true' if it is true in all possible worlds (i.e., a 'necessary proposition' is one

---

[5] Leibniz is classified as a "rationalist." This is a position where 'reason' has precedence over other (e.g., empirical) ways of acquiring knowledge. It is thought that we have a non-empirical and rational access to truth about the way the world is. Rationalists are attracted to mathematics as a model for knowledge.

[6] Linguistic expressions (e.g., words, sentences, numbers) *cannot* 'extend' to objects in other possible worlds. This is metaphorical.

[7] Michael Loux (1979) maintains that a variety of forms of discourse can be accommodated within the framework of possible worlds: discourse involving ascriptions of modality (both *de re* and *de dicto*), counterfactual discourse, discourse about meanings, and discourse about intentional abstract entities. Loux claims that possible worlds are part of the ontology of common sense.

[8] This is a major metaphysical *axiom*. It is neither true nor false. Axioms are the foundational principles that lie behind the exposition of the syntax and semantics of a formal system. In a deductive axiomatic theory, the set of axioms are the basis of a system, while the remaining definitions and propositions (e.g., theorems) are the logical consequence of the axioms. Axioms are sometimes considered 'first principles' known to be necessarily true, without need for justification. Axioms cannot be deduced from other sentences in a formal system. An axiom is typically (but not always) adopted if it helps map (or represent) the physical world (or linguistic discourse) in a fruitful way.

that could not be false). A proposition is 'contingent' if it is true in some possible worlds, and false in others. A 'necessary falsehood' is true in no possible world.

Possible-world modal realism postulates that there are ‘essential properties’ of some objects which are just those properties, that without them, the object could not be itself. A property possessed by a given object **x** is 'essential' only if **x** has this property in all possible worlds in which **x** exists. For example, the object Ludwig van Beethoven (1770-1824) was a person essentially, but only a piano player accidently. Beethoven might not have played piano in another world (a contingent property), but he is essentially a person (a necessary property) in worlds where he exists. According to orthodoxy, 'metaphysical possibility' cannot be reduced to linguistic rules and conventions: it constitutes a mind-independent subject matter for thought and talk.

We can summarize the propositions that make up the construction of a possible world semantics (loosely following Edwin Mares, 2011, pp. 14-22):

1) There is a set of possible worlds.

2) A 'possible world' is a way the universe might have been.

3) In different possible worlds, different propositions are true.

4) We say that a world W2 is 'nomically accessible' from another world W1, if and only if the laws of nature of W1 are true in W2.

5) In W, "it is physically necessary that S" is true if and only if S is true in every world nomically accessible from W.

6) In W, "it is *metaphysically necessary* that S is true" if and only if S is true in every possible world.

7) In W, "it is possible that S" is true if and only if in some world S is true.

8) In W, "S1 and S2" is true if and only if S1 is true in W and S2 is true in W.

9) A 'proposition' is a sentence that is made in a context, so that the sentence and a context are what allow the evaluation of whether a sentence is true or false in a possible world.

10) A proposition is 'contingently true' if it is possible, but not necessary. A proposition is contingent, if it is true in some possible worlds and false in others.

11) A 'necessary falsehood' is true in no possible world.

In essence, theorists assume that there is an ordering of worlds according to how much they resemble the actual world. Mares states that this 'possible worlds semantics' is the leading theory of 'necessity.'[9] From Mares' perspective, these propositions about 'necessity' are truths that are the result of *a priori* reasoning and are not subject to empirical observation.[10]

V. The Radical Skeptical Argument

In the 1970's, many epistemologists' attention turned to Descartes' 'evil demon' scenario and the ancient problem of a radical skeptical argument. With a simple two-premise argument, the skeptic doesn't just claim that we *may* be victims of a systematic deception, but deductively concludes that it is *impossible* to have knowledge.

The standard skeptical argument makes self-reference to 'I' assuming a single conscious thinker, and **o** 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**.

Certainly, we don't want to concede the conclusion to the skeptic that we cannot know ordinary propositions. In order for conclusion #3 to be false, either (or both) of the two intuitive premises must be false. Or as argued by contextualists, **S** can know **not-BIV** and **o**, in ordinary contexts but cannot know **not-BIV** nor **o** in skeptical contexts (where premise 2 is always true). Neomooreans, invariantists, and virtue epistemologists (not exclusive groups) typically argue that **S** can know **not-BIV** (at least in certain contexts) and thus the skeptical hypothesis #1 is false, and epistemic closure (premise #2) is always true.

---

[9] Ot these 11 propositions (above), we can identify sentences 1 and 3 as prescriptive propositions, premise 6 as an axiom, while the other sentences are all understandable as 'stipulative definitions' that form fixed definiens concepts.

[10] Many proponents of possible world semantics understand philosophical claims to *not* be just about what *is* the case (i.e., how things are), but that philosophy seeks to assess how things *mus*t be. Philosophy is presumed to be study of *necessary truths* and what could *possibly* be true, and thus, the study of modality.

VI. The Rise of Externalist Theories: A Causal Theory and Reliabilism

Alvin Goldman's "A Causal Theory of Knowing" (1967) started what became understood as an 'externalist theory' of knowledge, with a 'causal theory':

A Simple Causal Theory:
S knows that **p** if and only if:
(1) **p** is true.
(2) **S** believes that **p**.
(3) **S**'s belief that **p** is caused by the fact that **p**.

With the essay "Discrimination and Perceptual Knowledge" (1976), Goldman modifies the 'causal' theory with a 'reliabilist' theory. He eliminates the requirement that **S**'s belief that **p** must be connected to the fact that **p**. Instead, knowledge is defined as beliefs formed by reliable mechanisms, specifically to distinguish the truth of **p** from relevant alternatives:

A Simple Process Reliabilism Theory:
S knows that **p** if and only if:
(1) **p** is true.
(2) **S** believes that **p**.
(3) **S**'s belief that **p** was produced by reliable cognitive source.[11]

What an 'externalist' theory of knowledge does, is eliminate the internal justification condition and explains justification in terms of reliable methods (or processes) that support a belief. A method of 'belief formation' is 'reliable' when it tends to produce mostly true beliefs, rather than false ones, and a belief is justified when it is generated by a reliable method of belief formation. For externalists, knowledge is a natural relation between **S**'s beliefs (and belief-forming processes) and external states-of-affairs (or facts). Whether **S** knows **p** depends exclusively (or primarily) on factors such as how the belief was caused, and how reliable the faculty or mechanism was by which **S** came to believe **p**. Process reliabilists believe that a justified belief is formed by cognitive process(s) which tend to produce a high proportion of true beliefs relative to false ones.

Externalist theories are contrasted to 'internalist' theories. Internalists believe that whether a human belief is 'justified' depends on states (in some sense) *internal* to the

[11] With reliabilism, *methods* of *'belief-formation'* becomes a key subject of attention. Nozick's sensitivity theory assumes that S's method of belief formation is important in determining whether S has knowledge.

subject. Internal factors are primarily relevant to whether **S**'s true belief constitutes knowledge. Internalism is presupposed by the justified true belief (JTB) theory.

From the perspective here, a major problem with externalist theories (including Dretske's relevant alternatives) is that they can't satisfactorily respond to some case examples without using an artificial 'possible worlds' explanation (although metaphysicians don't view this as a weakness). For example, if **S** sees a real barn based upon causally produced perception, but with 99 nearby facades (unknown to **S**), we wouldn't attribute knowledge to **S**. In this case, **S**'s belief is a natural relation between **S**'s belief-forming processes (visual perception) and a states-of-affairs (a real barn), but still **S** doesn't know that he is seeing a barn. **S** is lucky, given his environment. Although **S**'s true belief that **p** was produced by reliable cognitive source **(S**'s visual capacity), **S** doesn't know **p**. With this barn façade example, process reliabilism as defined above is shown to be false.

A second major problem with reliabilist theories is how to determine exactly what constitutes a 'reliable process' in particular contexts (the 'generality problem'). A third major problem is that at times, even reliable processes (e.g., vision) are acknowledged to fail (e.g., wrongly perceiving an oasis in a desert) leaving **S** capable of possessing a (usually) reliably held, but false (perceptual) belief. In other words, what if the reliable source fails? Finally, it seems 'reliability' is something that persons must determine (i.e., interpret, evaluate, decide). Reliability is not an external objective commodity.[12]

VII. Dretske's Relevant Alternatives

The basic idea of Dretske's 'relevant alternatives' approach as a requirement for **S** to know **p**, is that for **S** to know **p**, **S** must be able to ‘rule out' competing hypotheses to **p**, but that *only some sub-set* of all **not-p** possibilities *are "relevant"* for knowledge attributions. In "The Pragmatic Dimension of Knowledge" (1981, 2000) Dreske proposes to use "relationally absolute concepts" as a model for understanding knowledge (p. 52):

[12] Reliabilism as an externalist theory diminishes S's decision process and personal justification from the equation of what knowledge is. It ignores S's subjectively picking up on 'red flags' (i.e., possible inconsistencies indicating p might be false) and adequate evidence evaluations. It also favors the analysis of perceptual empirical knowledge and the thought that with evidence e, S should rationally believe p.

> Knowledge is an evidential state in which *all relevant* alternatives (to what is known) *are eliminated*. This makes knowledge an absolute concept, but the restriction to *relevant* alternatives makes it, like 'empty' and 'flat' applicable to this epistemically bumpy world we live in.[13]

The 'relevant alternatives' approach described in various writings of Dretske is complex. Dretske's epistemic theory includes a counterfactual definition of knowledge (involving the concept of sensitivity) and a possible worlds relevant alternative semantics.[14]

The question of what is a 'relevant alternative' to a belief **p** is, is one that has always plagued proponents of "relative alternatives" theories. Dretske says:

> …the difference between a relevant and an irrelevant alternative *resides not* in what *we* happen to *regard* as a real possibility (whether reasonably or not), but in the *kind of possibilities* that *actually exist* in the objective situation. Whether or not a birdwatcher knows that the bird he sees is a Gadwall depends on whether or not, in some objective sense, it could be a look-alike Grebe (or any similar looking creature). (p. 63, italics are added).

But, as discussed below, some theorists, such as David Lewis disagree, directly contradicting Dretske, arguing that the difference between a 'relevant' alternative and an 'irrelevant' one *does reside* in what **S** happens to *regard* as a possibility.

In "Epistemic Operators" (1971) Dretske rejects the 'epistemic closure principle' in the 'zebra at a zoo' example. He claims that **S** can know that the animal **x** in the pen is a zebra on the basis of evidence that does *not* enable **S** to know that **x** isn't a cleverly disguised mule. He contends:

(1) **S** knows **x** is a zebra.

(2) **S** knows that **x** is a zebra is incompatible with **x** being a disguised mule.

(3) But **S** doesn't know that **x** isn't a disguised mule.

Dretske believes that knowledge is closed under some known entailments, but that it isn't closed in other situations, such as this one. Although **S** isn't in a perceptual position to

---

[13] Dretske's talk about 'absolute' concepts (like 'empty' and 'flat') is influenced by Unger (discussed below).

[14] We won't discuss the technnical details of Dretske's relevant alternatives semantics. Goldman (1976) and Stine (1976) also propose 'relevant alternatives' theories.

know that **x** *isn't* a disguised mule in an ordinary circumstance, **S** can still know that **x** is a zebra since the possibility of **x** being a disguised mule *isn't* a relevant alternative. But if **x** being a cleverly disguised mule were to become relevant (similar in the way barn facades became relevant for Henry's situation), then **S** wouldn't be able to rule out the mule possibility, and **S** would no longer know that **x** is a zebra. It should be emphasized that Dretzke is denying epistemic closure only in *some* situations. He says:

> To deny closure is *not* to say that you *never* know (find out, discover, learn) that **q** is true by inferring it from a **p** that you know to be true. It is merely to deny that this can be done for *any* **q** (2014, p. 31).

> Some skeptical alternatives are deliberately formulated so as to be immune from exclusion on the basis of evidence or reasons. One is often told that such skeptical scenarios are not plausible. And so they aren't. That, though, isn't the issue. The question is: do you know them to be false? If you do, *how* do you know it? (2008, p. 81).

The results of Dretske's analysis of the Henry and the Barn case study is disappointing. He says with respect to Henry and the Barn (paraphrased):

> To know **p**, **S** must have evidence that rules out all relevant alternatives to **p**. Irrelevant alternatives to **p** don't need to be ruled out. *In the land of fake barns*, in order to know from the road that **S** sees a barn, ***S** has to rule out the alternative* that **S** is *looking at a fake barn*. But in *ordinary farmland*, this *alternative is irrelevant* and *doesn't need to be ruled out*.[15]

VIII. 'Knowledge' described as Externalist Modal 'Sensitivity.'

Nozick's contribution to epistemology comes from the third chapter of his *Philosophical Explanations* (1981) that is about knowledge and skepticism. He says that typically, the arguments of the skeptic do not lead us to conclude that we don't have knowledge; but they do leave us wondering how we can know what we do (p. 10). The question is that, granted that we do know, *how* can we know? How is knowledge possible? He says that we want to formulate a hypothesis to show how knowledge can exist given the skeptic's possibilities.

---

[15] But the original force of the Henry and the Barn example is to assume that Henry is in barn façade country but doesn't know it. Comparing when S has knowledge in a *normal situation* and what S (normatively) has to 'rule out' (or is relevant) in contrast to what S needs to 'rule out' (or what is a relevant alternative) in *a barn façade neighborhood situation* seems to generate an uninteresting (either modal or normative) thesis of what needs to be 'ruled out' in different physical situations for S to know.

Nozick adds a third (subjunctive) condition for 'knowledge' that excludes justification (p. 172):

(1) **p** is true,
(2) **S** believes **p**, and
(3) If **p** weren't true, **S** wouldn't believe **p**.

He says (pp. 173-174):

> The subjunctive condition is powerful and intuitive, not so easy to satisfy, yet not so powerful as to rule out everything as an instance of knowledge. A subjunctive conditional "if **p** were true, **q** would be true," if **p** then **q**, does not say that **p** entails **q** or that it is logically impossible that **p** yet **not-q**. It says that the situation would obtain if **p** were true, **q** would also be true. This point is brought out especially clearly in recent 'possible worlds' accounts of subjunctives: the subjunctive is true when (roughly) in all those worlds in which **p** holds true that are closest to the actual world, **q** also is true. (Examine those worlds in which **p** holds true closest to the actual world, and see if **q** holds true in all these).
>
> Whether or not **q** is true in **p** worlds that are still farther way from the actual world is irrelevant to the truth of the subjunctive. I do not mean to endorse any particular possible-worlds account of subjunctives, nor am I committed to this type of account. I sometimes shall use it, though, when it illustrates points in an especially clear way.

A fourth condition is added (p. 176):

(4) If **p** were true, **S** would believe it.

Knowledge, he maintains, is connected to these 'sensitivity' principles as a 'necessary condition' for its attainment. At page 178 he says (paraphrased):

> '**S** knows that **p**' *when* **S** not only does truly believe it, but also would truly believe it and wouldn't falsely believe it. **S** not only actually has a true belief, **S** subjunctively, has one. It is true that **p** and **S** believes it; if it weren't true, **S** wouldn't believe it, and if it were true **S** would believe it. To know that **p**; is for **S** who would believe **p** if it were true, and who wouldn't believe **p** if it were false.
>
> It will be useful to have a term for this situation when **S**'s belief is thus subjunctively connected to the fact. Let us say of **S** who believes that **p**, which is true, that when conditions 3 and 4 hold, his belief *tracks* the truth that **p**. To know is to have a belief that tracks the truth. Knowledge is a particular way of being connected to the world, having a specific real factual connection to the world: tracking it.

The (externalist) "sensitivity condition" is as follows:

> **S**'s belief that **p** is sensitive iff, if **p** were false, **S** would not believe it.

Nozick suggests that **S** knows that **p** only if **S**'s belief that **p** is sensitive to the truth, that is, only if **S** would not believe that **p** if **p** were false. In evaluating sensitivity's counterfactual conditional, consider the *nearest possible world* in which ***p*** *is false*; the state of affairs, or the world, in which **p** is false, but is otherwise as similar to the actual state of affairs as it can be, and then determine whether, in that world, **S** believes that **p**. If **S** does believe that **p** in that world, then **S**'s belief is insensitive, and **S** doesn't possess knowledge. If **S** does not believe that **p** in that world, **S**'s belief is sensitive and **S** can therefore know **p**. This is consistent with Lewis's (1973) semantics which leaves the sensitivity condition as equivalent to the requirement that, in the nearest possible worlds in which **not-p**, **S** does not believe that **p**. A necessary condition for knowledge is suggested:

> **S**'s true belief that **p** is knowledge only if, if **p** were not the case, **S** would not believe that **p** via the method **S** actually uses to form this belief.

Nozick thought that this principle was consistent with process reliabilism.[16] Also, this condition responds to Gettier cases. Consider the case where **S** coincidentally has a *true belief* that it is exactly noon by looking at the town's clock tower (after the clock stopped functioning twelve hours earlier). In a nearby possible world, if it were 12:10PM (not exactly noon), **S** would nonetheless continue to believe **p**. In this given environment, **S** is not sensitive to the truth value of this belief. In a fortuitous situation where **S** reads the clock, it is coincidentally exactly 12PM, **S**'s true belief doesn't track its truth.[17]

For Nozick, there is an intuitive sense in which knowledge requires not merely being correct; but "tracking the truth" in other possible circumstances.[18] Also, for

---

[16] Nozick's central idea, likely influenced by Dretske, is that knowers are like thermometers. With a properly functioning thermometer, a true temperature reading is well-connected with reality. When S is well-connected to reality, S expresses an output of a true belief. Knowers track the truth of their propositions in a similar way to the way thermometers track the temperature in their environments.

[17] Nozick uses the Nogot example as a 'Gettier case.' The stopped clock case is easier to explain.

[18] Although Nozick's intuition that in order to know p, S's reasons must 'track' (or be narrowly relevant) for why p is true is reasonable, I don't share Nozick's intuition that truth-tracking involves a comparison with

example, **S**'s belief that there is distant water in a desert is insensitive as to whether water was there. If there was no water, then **S** would still believe that there is water in an environment of frequent mirages. Nozick (p. 179) defines knowing via a method:

> **S** knows, via method (or way of believing) M, that **p** iff:
>
> (1) **p** is true.
> (2) **S** believes, via method or way of coming to believe M, that **p**.
> (3) If **p** were false and **S** were to use M to arrive at a belief whether (or not) **p**, then **S** wouldn't believe, via M, that **p**. (Sensitivity).
> (4) If **p** were true and **S** were to use M to arrive at a belief whether (or not) **p**, then **S** would believe, via M, that **p**.

Nozick's (anti-skeptical, closure-denying) 'sensitivity condition,' and its comparison with the (anti-skeptical, closure-affirming) 'safety condition' has been a central subject of discussion in the following decades and to the present.[19]

IX. Peter Unger's Skepticism

Unlike most epistemologists (who are anti-skeptical), Peter Unger (1971, 1975) argued in favor of a skeptical thesis.[20] He argues that basically no one ever knows anything (about anything). He contends:

'counter-factual circumstances' in evaluating whether I or someone else knows. I don't believe that the fixed definiens concepts of 'sensitivity' and 'tracking' help explain the concept of knowledge, nor are they the necessary conditions that explain (in a positive fashion) what knowledge is.

[19] It is *questionable* whether the *worth (or value)* of the *analytic project* to adduce a (single) necessary (e.g., anti-luck, safety, or sensitivity) principle(s) which applies to all instances of 'knowledge' and responds to key epistemological issues (e.g., skepticism, Gettier situations, lottery questions). Outside of prolonged metaphysical discussion of possible worlds, such formulations are mostly idle otherwise. They are not intended to provide a complete theory of knowledge; but are largely intended to include (or exclude) certain cases of 'true belief' as being knowable or not knowable. This approach treats these epistemological issues as 'puzzle solving' with competing mapping systems of how sentences represent true or false propositions. Such single (or multi-) 'necessary principle' approaches don't ultimately contribute to the understanding of the theoretic natural kind concept of knowledge that is found in ordinary language.

[20] Besides Unger's concerns with the semantics of 'knowledge attributions,' his non-skeptical essay "An Analysis of Factual Knowledge" (1968) was influential. It is here that Unger presents an analysis of what it is for S to know that something is the case. His analysis that – "For any sentential value of p (at time t), S knows that p if and only if (at t) it is not at all accidental that S is right about its being the case that p" is a reaction to the Gettier problem attempting to present what a justified, true, non-Gettierized belief is. The essential idea that in order for S to have knowledge, S *cannot* be in a Gettier situation where p is true because of intervening coincidental (accidental) circumstances is ultimately proposed as a *negative* necessary condition. It *eliminates* cases of justified true belief which aren't knowledge. Like Nozick's sensitivity condition, (1) Unger's 'non-accidentally' *doesn't* explain what knowledge is, with a positive

1) If you know that **p**, then you have to be absolutely certain that **p**.
2) For most propositions **p** that you believe, you're not absolutely certain that **p**.
3) So, for most of the propositions that you believe, you don't know that **p**.

Unger's linguistic intuitions and assumptions are a key part of his skeptical philosophy. Unger's key linguistic intuition is that "certain" is an "absolute term" and that there is a big difference between (1) what's strictly speaking 'true' and (2) what's 'acceptable to say' or what's 'near enough to the truth' for practical purposes.

Unger argues that the meanings of key terms are formed by their use in the history of our ancestors. These meanings were formed and made to connect with those of other words, in order to accommodate developing thought. The task for philosophers is to *articulate* those *meanings*, thus articulating the inherited theory in whose terms we now verbally think (1975, p. 5). Unger writes his essay to articulate 'the meanings' of key terms which prove his skeptical theory.[21]

Unger suggests that the words 'reasonable' and 'justified' may be inapplicable terms. It may be impossible for **S** to be reasonable in believing things. He suggests that we may best regard his skeptical reasoning as indirect arguments against the suppositions embodied in our language, or against our commonsense beliefs. He conjectures that "every human knows at most, hardly anything to be so" and that it may be reasonable for *us to suppose this* without knowing it true (1975, p. 48).

Unger says that our language habits serve us in practical ways, even while they involve us saying what is (strictly) not true. And this often does occur when our positive assertions contain terms with special features of a certain kind, of which he calls *absolute (limit) terms*. We do not speak truly when saying that '**x** has a top which is flat.' Basic 'absolute terms' generally fail to apply to the world (e.g., flat, certain). He writes:

---

explanation, and (2) what (inaccessible) undermining propositions are, or (3) the status of the evaluation of what 'sufficient evidence' is. The idea that 'accidental truths' don't count as knowledge also conflates Gettier and Harman cases of 'accidental' true belief.

[21] Unger's theory of 'meaning' and 'meanings' is idiosyncratic, vague, and not obvious nor intuitive.

> The terms of knowledge, along with many other troublesome terms, belong to a class of terms that is quite pervasive in our language. I call these terms *absolute terms*. The term 'flat' in its central, literal meaning, is an absolute term. (With other meanings, however closely related, perhaps as in 'His voice is flat' and 'The beer is flat,' I have no direct interest). To say that something is flat is, so far as content goes, (is) no different from saying that it is absolutely, or perfectly flat. To say that a surface is flat is to say that some things or properties *which are matters of degree* are *not* instanced in the surface *to any degree at all*. Thus, something which is flat is not at all bumpy, and not at all curved. Bumpiness and curvature are matters of degree. When we say of a surface that it is bumpy, or that it is curved, we use the relative terms 'bumpy' and 'curved' to talk about the surface. Thus, many absolute and relative terms go together, in at least one important way, while other terms, like 'unmarried,' have only the most distant connections with terms of either of these two sorts. (pp. 54-55).[22]
>
> It is at least somewhat doubtful, then, that 'flat' ever applies to actual physical objects, or to their surfaces. And the thought must strike us that if 'flat' has no such application, this must be due in part to the fact that 'flat' is an absolute term.[23] We may then do well to be a bit doubtful about the applicability of any other given absolute term and, in particular, about the applicability of the term 'certain.' As in the case of 'flat' our paraphrase highlights the absolute character of 'certain': As a matter of logical necessity, if someone is certain of something then there never is anything of which he or anyone else is more certain… Thus, if it is logically possible that there be something of which any person might be more certain than he is now of a given thing, then he is not actually certain of that given thing (p. 67).

Unger argues that **S** knows something to be so, only if he is certain of it. But of **S**'s being certain; there is hardly anything of which **S** is certain. He suggests that with some terms we error systematically. If skepticism is right, we go around saying 'I know,' 'he

---

[22] This distinction between *absolute* terms (e.g., flat) and *relative* terms (e.g., bumpy, curvature) is arbitrary and unintuitive. There is no such thing as flatness *simpliciter*. The thesis that we do not speak truly saying that 'x has a top which is flat,' and that *absolute terms* fail to apply to the world (e.g., 'flat,' 'certain') is certainly extraordinary. 'Bumpy,' 'curved,' 'bald,' 'rich,' 'happy,' and 'sad' are group resemblance concepts.

[23] In opposition to Unger, disagreement (or different assessments) about whether x is flat, reflects the linguistic fact that 'flatness' is a group resemblance concept. 'Flat' objects have an appearance which are associated with paradigm prototypes or exemplars within the domain specified. The *standards* used by speakers *result from the context* of attribution (the intentions, presuppositions, etc., of speakers and listeners). Similarly, the term 'tall' is a group resemblance concept. In its normal use, the term involves a specified domain (e.g., fifth graders, college basketball players, buildings). There are no necessary or sufficient (contextual) conditions for determining whether 'x is tall,' there are only normative (or pragmatic) standards for judging whether x is tall (relative to a domain). The truth of sentence 'x is tall' is in part, judged by *S's standards* and *audience standards* (which are usually the same, but not always).

knows,' all while *what we say* (and believe) is actually false. The terms of knowledge lead us to error systematically. Unger's focus is on what words really do mean (p. 74).

In *Philosophical Relativity* (1984), Unger is interested in constructing a semantics to explain the use and truth-conditions of 'knows.' He shifts his position from being a skeptic to being a relativist. He thinks that the truth of skepticism is relative to semantics and since there is no true objective semantics; philosophical relativity must be the case.

The problem of knowledge, as he sees it, is that we know, or we don't (i.e., skepticism) (p. 4). He says that a crucial aspect of a philosophical problem may depend on *the meaning of*, or on the *semantic conditions* of, certain *linguistic expressions* in terms of which the problem is directly and standardly formulated.[24] For example, the problem of knowledge might thus turn upon *the meaning of 'know'* as it occurs in typical sentences of the form '**S** knows that **p** is the case.' He distinguishes two semantic approaches: 'contextualism' and 'invariantism' (p. 6). In understanding linguistic behavior, we distinguish what **S** says, and what is informally implied. Unger now understands that in 1975 he was an invariantist, denying that there were contextual standards. (He is the first philosopher to use these terms).[25]

Unger uses the sentence "That field is flat" as an example (p. 6). Does the semantic evaluation of this sentence involve the context for the demonstrative subject term 'that field' or the predicate 'flat' as well? The contextualist says it is both; the invariantist says it is just 'that field'. For the invariantist, what the uttered sentence means is that the field is absolutely flat (no matter about any purposes). An invariantist can *assign a semantics* to the philosophically important terms consistent with skepticism and contextualist does otherwise (p. 46).[26] If one really knows **p**, the skeptic says, **S** must be

---

[24] Again, a semantic theory about linguistic reference is *supposed* to help solve epistemological issues.

[25] That the problem of epistemic skepticism might be tied to the meaning of 'know' as it occurs in ordinary language sentences is a novel solution in part due to the rise of formal semantics in the twentieth century.

[26] The contextualist-invariantist debate centers upon semantic representation. Both models rest upon a view about knowledge ascriptions. These theories leave aside what knowledge is; and concentrate on a linguistic model of how 'know' (and its cognates) should be represented in a formal language. It centers upon case studies of S's material environment and the 'standards' that are 'salient' in a given case.

in a position to rule-out as untrue logically conflicting propositions (p. 46).[27] Unger distinguishes 'standards for ruling out' and 'ranges of relevant competitors' (p. 48). He says that a skeptic embraces an invariantist account of the semantics of know. The contextualist does not, which is more in line with common sense. Which semantics is correct? For Unger, as is the case of all practitioners of formal semantics, words have extensions, and it makes sense to ask, "what are our words true of?" (p. 85). In *Philosophical Relativity*, he argues that semantics is relative to a theory.

Unger has had considerable influence among epistemologists and practitioners of formal semantics. Dretske, Cohen, DeRose, and Prichard are all influenced by Unger's concept of an *'absolute term'* or *'absolute idea.'* Dretske in "The Pragmatic Dimension of Knowledge" (1981, 2000) recapitulates Unger's ideas (p. 50, slightly edited):

> *Flat* is an absolute term in the sense that a surface is flat only if it is *not at all bumpy or irregular*. Any bumps or irregularities, however small and insignificant they may be (from a practical point of view), mean that the surface on which they occur is not really flat. It may be *almost* flat, or *very nearly* flat, but (as both of these expressions imply) it is not really flat. We do it seems, compare surfaces with respect to their degree of flatness, but Unger argues that this must be understood as a comparison of the degree to which these surfaces approximate flatness. They cannot be both flat and yet one flatter than the other. Hence if A is flatter than B, then B (perhaps also A) is not really flat. Flatness does not admit of degrees, although a surface's nearness to flatness does, and it is this latter magnitude that we are comparing when we speak of one surface being flatter than another. Unger concludes that not many things are really flat. Almost any surface will exhibit *some* irregularities. Hence, contrary to what we commonly say (and presumably believe) these surfaces are not really flat. When we describe them as flat, what we say is literally false. Probably *nothing* is really flat. This according to Unger is the price we pay for having absolute concepts.

Dretske argues that Unger doesn't really show that nothing is really flat. For although nothing can be really flat if it has any bumps or irregularity, what counts as a bump or irregularity depends on the type of surface being described. Something is empty, if it has nothing in it, but this doesn't mean that an abandoned warehouse is not really empty because it has light bulbs or molecules in it. Light bulbs and molecules do *not* count as *things* when determining the emptiness of a warehouse. These are irrelevant.

---

[27] This skeptical ethical proposition (or demand) is both unreasonable and theoretically impossible.

X. David Lewis's Contextualism

David Lewis emerged as a contextualist with essays "Scorekeeping in a Language Game" (1979) and "Elusive Knowledge" (1996).[28] We will discuss this latter essay.

In the second paragraph of "Elusive Knowledge" Lewis says (p. 691):

> We have all sorts of everyday knowledge, and we have it in abundance. To doubt that would be absurd… it is a Moorean fact that we know a lot. It is one of those things that we know better than we know the premises of any philosophical argument to the contrary.

It is Lewis's worldview that knowledge plainly exists, so it must be explained how this occurs. With this in mind, Lewis grapples with linguistic intuitions about epistemic closure and fallibilism in conjunction with the skeptical argument (italics added):

> For no sooner do we engage in epistemology--- the systematic philosophical examination of knowledge—than we meet a compelling argument that we know next to nothing. The skeptical argument is nothing new or fancy. It is just this: it seems as if knowledge must be by definition, infallible. If your claim that **S** knows that **p**, and yet you grant that **S** cannot eliminate a certain possibility in which **not-p**, it *certainly seems* as if you have granted that **S** does not after all know that **p**. To speak of *fallible knowledge*, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory.

But Lewis concedes that uneliminated possibilities of error are everywhere. Lewis goes on to say that because there are always uneliminated possibilities of error, we have 'fallible knowledge' or none. Those possibilities are far-fetched, of course, but are possibilities all the same. We never have infallible knowledge.[29]

---

[28] "Elusive Knowledge" was originally published in *Australasian Journal of Philosophy* 74, 4 (1996), pp. 549-67. It is republished in *Epistemology: An Anthology*, 2nd edition (2008), Sosa et. al (eds), pp. 691-705. All page references are to the anthology.

As mentioned above, Unger (1984) first coins the terms 'contextualism' and 'invariantism.' It is suggested that we can contrast between what a sentence literally says and the social and material context in which it is asserted. Contextualists are interested in analyzing the truth conditions of 'knows' in sentential contexts in which it is asserted. Contextualism is the view that ascriptions of knowledge are context-sensitive—the *truth values* of *sentences* containing the words 'know,' and its cognates depend on *contextually determined standards.* The truth value of a sentence containing the knowledge predicate can vary depending on things like the purposes, intentions, expectations, presuppositions, etc., of the speakers who utter these sentences.

[29] According to the PE definition of knowledge, S can know p, despite not knowing it is impossible that not-p. Open logical possibilities do not circumvent the possibility of knowledge. (Lewis rightly concludes that we never have infallible knowledge).

But then conflicting intuitions about 'context-dependence' and 'justification' arise. Lewis doesn't agree that the mark of knowledge is 'justification' (as one of the three necessary conditions). His reasoning is summarized (p. 692):

> (1) Lewis claims that 'justification' isn't sufficient for knowledge: your true opinion that "you won't win the lottery" isn't knowledge, whatever the odds.
>
> (2) Lewis claims that 'justification' is *not always necessary*: (a) What (non-circular) argument supports our reliance on perception, on memory, and on testimony? And yet **S** gains knowledge by these means (This is a reference to Hume's problem of induction). (b) We can sometimes forget our justification for **p**, and yet retain knowledge. (c) 'Facial recognition' is an example of 'knowledge without justification': **S** knows the proper name of a face without knowing how.

He states that the link between 'knowledge' and 'justification' *must be* broken. It's clear that epistemology doesn't destroy knowledge simply by raising the standards of personal justification. What is the correct story?

Lewis refers to Unger and the infallibility of knowledge as a starting point. Lewis considers the following definition (p. 693):

> '**S** *knows* **p**' iff **p** holds in every possibility left uneliminated by **S**'s evidence; equivalently; iff **S**'s evidence eliminates every possibility in which **not-p**.
>
> In order for **S** to know **p** on this definition, **S** would need infallible evidence.

Lewis then 'clarifies' what he chooses to call 'propositions':

> …'propositions' are individuated coarsely, by necessary equivalence. For instance, there is only one necessary proposition. It holds in every possibility; hence in every possibility left uneliminated by **S**'s evidence, no matter who **S** may be and no matter what his evidence may be. So, the necessary proposition is known always and everywhere.[30]

He 'clarifies' what 'possibilities' are: A possibility will be specific enough if it cannot be split into sub-cases in such a way that anything we have said about possibilities, or anything we are going to say before we are done, applies to some cases, and not to others. For instance, it should never happen that proposition **p** holds in some but not all sub-cases; or that some but not all sub-cases are eliminated by **S**'s evidence (p. 693).[31]

---

[30] Lewis's Leibnizian metaphysical conception of what a 'proposition' is, is contrasted to the conceptual analysis of this concept in chapter 11.

[31] 'Clarifications' are oftentimes 'stipulations' (prescriptions). The metaphysics of when a possibility is 'specific enough' in relation to S's evidence, in terms of cases and sub-cases, is stipulated here.

Lewis later offers two ethical propositions (p. 694, paraphrased):

> (1) **S** is *not* entitled to ignore any **p** that he pleases.
>
> (2) **S** *may* properly ignore *some* uneliminated possibilities; **S** *may not* properly ignore others.

A new definition of knowledge is implied:

> '**S** *knows* that **p'** iff **S**'s evidence eliminates every possibility in which not-**p** (except for the possibilities that **S** is properly ignoring).

Lewis says that the rest of (modal) epistemology examines the question of what may we *properly* presuppose in our ascriptions of knowledge? Which of all the uneliminated alternative possibilities may *not* properly be ignored? Which ones are the 'relevant alternatives' –relevant, that is, to what **S** does and doesn't know? In reply to this question, Lewis continues with *normative rules* about what possibilities **S** may *not* properly ignore. He says that some of these rules (not all) are taken from Dretske (1970, 1981), Goldman (1976), Stine (1976) and Cohen (1988). A quick and incomplete summary of these rules are paraphrased as follows:

> (1*) Rule of actuality:* The possibility that actually obtains is never properly ignored; actuality is always a relevant alternative; nothing false may properly be presupposed. The rule is 'externalist,' **S** may not be able to tell what is properly ignored. In judging which of **S**'s ignorings are proper, hence what **S** knows, *we judge **S**'s success* in knowing, *not* how well **S** tried.
>
> (2) *Rule of belief:* A possibility that **S** believes to obtain is not properly ignored., whether or not **S** has the right to so believe… Since belief admits of degree, and since some possibilities are more specific than others, the rule should be reformulated in terms of degree of belief… A possibility may not be properly ignored if **S** gives it, or ought to give it, a degree of belief that is sufficiently high and high not just because the possibility in question is unspecific. How high is sufficiently high? Depends on what is at stake (and risks involved). This is the only place where Lewis allows 'belief' and 'justification' to enter into his explanation.
>
> (3) *Rule of Resemblance*: Suppose one possibility saliently resembles another. If one of them may not be properly ignored, neither may the other.

Lewis claims that these rules can explain lottery cases. It is the rule of resemblance that explains why you do not know that you will lose the lottery, no matter what the odds are against you, and no matter how sure you should be, that you will lose. For every ticket, there is the possibility that it will win. These possibilities are saliently similar to one another: so, either every one of them may be properly ignored, or else none may. But one of them may not be properly ignored: the one that actually obtains (p. 696). He continues with more rules, including the final 'rule of attention':

> (4) *Rule of Attention*: More of a triviality than a rule. When we say that a possibility is properly ignored, we mean exactly that; we do not mean that it *could have been* properly ignored. Accordingly, a possibility not ignored at all is *ipso facto* not properly ignored. What is and what is not being ignored is a feature of the particular conversational context. No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in *this* context we are *not* in fact *ignoring it* but attending to it, then for us now, *it is a relevant alternative*. It is in the contextually determined domain. If it is an uneliminated possibility in which **not-p**, then it will do as a counterexample to the claim that **p** holds in every possibility left uneliminated by **S**'s evidence. That is, it will do as a counterexample to the claim that **S** knows **p**.

Lewis asks us to do some epistemology and let our fantasies rip. We find uneliminated possibilities of error everywhere. In a context with an enormously rich domain of potential counterexamples, it can never happen (well hardly ever) that an ascription of knowledge is true. Epistemology destroys knowledge. But is does so only temporarily. The pastime of epistemology does not plunge us forevermore into its special context. We can still do a lot of proper ignoring, a lot of knowing, a lot of true ascribing of knowledge to ourselves and others, the rest of the time. Lewis states (p .698, italics added):

> What is epistemology all about? The epistemology we've just been doing, at any rate, soon became an investigation of ignoring possibilities.[32] But to investigate the ignoring of them was *ipso facto* not to ignore them. Unless this investigation of ours was an altogether atypical sample of epistemology, it will be inevitable

[32] The investigation of ignoring possibilities *isn't* what epistemology should be all about. More often when determining whether S has knowledge, attention is paid to the *adequacy of S's evidence* and whether there are any discrepancies or inconsistencies in S's evidence ("red flags") that should render doubt and further investigation. Also, of concern is whether S's evidence is obtained from reliable sources in a favorable environment, and whether S's rational and perceptual faculties and background knowledge are adequate. S's implicit or explicit elimination (discarding, ruling out, ignoring) of possibilities that undermine, or defeat S's belief is certainly a *part* of the *adequacy of S's evidence*, but this isn't what epistemology is all about.

> that epistemology must destroy knowledge. *That is how knowledge is elusive. Examine it, and straightway it vanishes.*

Lewis later discusses epistemic closure (p. 701). He states that if we analyze knowledge as a modality, as he has done, we cannot escape epistemic closure: *Knowledge is closed under (strict) implication*. He acknowledges Dretske's denial of closure. Dretske claims that having hands *does imply* not being handless and deceived; yet *knowing* that I have hands *does not* imply *knowing* that *I'm not* handless and deceived.[33] Lewis says that what Dretske says is close to being right. However, Lewis believes that in the skeptical argument there was a "context switch" midway, and the semantic value of the context-dependent word "know" switched with it. Lewis says:

> The premise 'I know that I have hands' was true in its everyday context, where the possibility of deceiving demons was properly ignored. The mention of that very possibility of switched the context midway. The conclusion 'I know that I'm not handless and deceived' was false in *its* context, because that was a context in which the possibility of deceiving demons was being mentioned, hence was not being ignored, hence was not being properly ignored. Dretske gets the phenomenon right, and I think he gets the diagnosis of skepticism right; it is just that he misclassifies what he sees. He thinks that it is a phenomenon of logic, when really it is a phenomenon of pragmatics. Closure, rightly understood, survives the test. If we evaluate the conclusion for truth not with respect to the context in which it was uttered, but instead with respect to the different context in which the premise was uttered, then truth is preserved. And if, *per impossible*, the conclusion could have been said in the same unchanged context as the premise, truth would have been preserved (p. 701).

In other words, in the argument below, Lewis believes that in *a non-skeptical context*, premise #1 is false, and premise #2 is true when there is no change of context, and thus conclusion #3 is false in a non-skeptical context:

> (#1) I do not know that 'I am not a BIV.' (Skeptical hypothesis)
>
> (#2) If I do not know that 'I am not a BIV,' then I do not know 'I have hands.'
>
> (#3) Therefore I do not know 'I have hands.'

In a *skeptical context*, since the evil demon hypothesis is *not* being ignored, premise #1 and premise #2 are both true, and the conclusion #3 is true (i.e., I don't know that I have

---

[33] On the PE theory of knowledge, Dretske's assessment against epistemic closure is deemed correct.

hands). Contextualist theories have the consequence of granting the skeptic the skeptical conclusion in skeptical contexts, but not in non-skeptical contexts where the demon (or BIV) possibility is properly ignored.

In essence, Lewis argues that conversational context and pragmatic considerations determine what possibilities are properly ignored, and when doing epistemology, such as when discussing the BIV possibility, where far-fetched possibilities are being discussed, they cannot be ignored. So, knowledge is 'elusive' when determining which possibilities are properly ignored (since they *aren't* being ignored when they are brought into consideration). In ordinary situations, **S** knows the ordinary proposition **p** that he has hands, but not in skeptical contexts where the skeptical hypothesis is assumed true, and from epistemic closure, **S** doesn't know he has hands.[34] Context and pragmatics play a role in whether **S** knows **p**.

XI. Stewart Cohen's Contextualism

In the late twentieth century, Stewart Cohen became a leading advocate of epistemic contextualism with essays including: "Knowledge, Context, and Social Standards" (1987), "How to be a Fallibilist" (1988), "Skepticism, Relevance, and Relativity" (1991), and "Contextualist Solutions to Epistemological Problems" (1998).

In his widely read essay, "Contextualism, Skepticism, and the Structure of Reasons" (1999) Cohen states (paraphrased, italics added, p. 57):

> Suppose speaker **S1** says about a subject **S** and proposition **p**, "**S** knows **p**." At the very same time, another speaker **S2** says of the very same subject and proposition, "**S** does not know **p**." Must one of the two be speaking falsely? According to the view I will call 'contextualism,' both speakers can be speaking the truth. Contextualism is the view that ascriptions of knowledge are context-sensitive—the *truth values* of *sentences* containing the words 'know,' and its cognates depend on *contextually determined standards*.[35] Because of this, *sentences* of the form '**S** knows **p**' can, at one time, have *different truth values* in different contexts. Now when I say 'contexts,' I mean 'contexts of ascription.' So,

[34] The contextualist concession to the skeptic that S sometimes (in context) does not know that he has hands is astounding and unintuitive. It is more plausible to deny closure and (putatively) know that I have hands.

[35] As it has been emphasized (and is well understood), Cohen's interest is about the 'truth values' of sentences containing the word 'knows.' Again, a semantic theory (about linguistic reference) is supposed to help solve an epistemological problem. It responds to the problem of radical skepticism.

> the truth value of a sentence containing the knowledge predicate can vary depending on things like the purposes, intentions, expectations, presuppositions, etc., of the speakers who utter these sentences.[36]

Cohen defends the view that ascriptions of knowledge are context sensitive by arguing that a contextualist account of knowledge ascriptions, when combined with a particular view about the structure of reasons can help provide a satisfactory response to radical skepticism. This essay (1999) is a revised version of "How to Be a Fallibilist" (1988).

Cohen considers the 'entailment principle' (p. 57):

> **S** knows **p** on the basis of (reason or evidence) **r** only if **r** entails **p**.

He says that the entailment principle leads to skepticism. Most philosophers reject the entailment principle thereby embracing fallibilism, which he defines as follows (p. 58):

> **Fallibilism: S** can know **p** on the basis of **r** even if there is some alternative to **p**, compatible with **r**. (Fallibilism allows us to know on the basis of non-entailing reasons).

Cohen then raises the normative question: But how *good* do the reasons have to be? He presents a case study (p. 58):

> Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks over the flight itinerary he got from the travel agent and responds, "Yes, I know it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." John and Mary *agree* that Smith doesn't really *know* that the plane will stop in Chicago. They decide to check with the airline agent. (p. 58).

---

[36] The form of the argument:

> (1) Hypothesis: Contextualism is the view that ascriptions of knowledge are context-sensitive; the truth values of sentences containing the words 'know' and its cognates depend on contextually *determined* standards"
> (2) Consequence*:* *Sentences* of the form 'S knows p' can, at one time, have *different truth values* in different contexts.
> (3) Stipulative Definition (3c, 3b): 'contexts' = contexts of ascription.
> (4) Hypothesis: Truth value of a sentence containing a knowledge predicate can vary depending on the purposes, intentions, expectations, presuppositions, etc., of the speakers who utter these sentences.
> (5) Conclusion: Thus, descriptions of knowledge are context sensitive. When this conclusion is combined with a view about the structure of reasons, this helps to provide a response to skepticism.

> What should we say about this case?[37] Smith claims to know that the flight stops in Chicago. Mary and John deny that Smith knows this. Mary and John seem to be using a stricter standard than Smith for how good one's reasons have to be in order to know. *Whose standard is correct?*

Cohen then considers three normative responses to this question: (1) Smith's is correct and so John and Mary's standard is too strong. (2) John's and Mary's standard is correct and so Smith's standard is too weak. (3) Neither Smith's nor John and Mary's standard is correct—both are too weak. He says that none of these answers seem satisfactory, and proposes this as the best answer (p. 59):

> Neither standard is simply correct or simply incorrect. Rather, context determines which standard is correct. Since the standards for knowledge ascriptions can vary across contexts, each claim, Smith's as well as Mary and John's can be correct in the context in which it was made. When Smith says, "I know **p**" what he says is true given the weaker standard operating in that context. When Mary and John say "Smith does not know **p**" what they say is true given the stricter standard operating in their context. *And there is no context independent correct standard*.

Cohen says that this case (and others like it), suggests that ascriptions of knowledge are context sensitive. The standards that determine how good one's reasons have to be in order to 'know' are determined by the context of ascription.

In a section entitled "semantical considerations," Cohen presents his linguistic intuitions (p. 60):

> Many, if not most, predicates in natural language are such that the truth-value of sentences containing them depends on contextually determined standards, e.g., 'flat,' 'bald,' 'rich,' 'happy,' 'sad,' and so on. These are all predicates that can be satisfied to varying degrees and that can also be satisfied *simpliciter*. So, e.g., we can talk about one surface being flatter than another and we can talk about a surface being flat *simpliciter*. For predicates of this kind, context will determine the degree to which the predicate must be satisfied in order for the predicate to apply *simpliciter*. So, context will determine how flat a surface must be in order to be flat.[38]

---

[37] The PE definition suggests that condition 4a (strength and adequacy of evidence) is satisfied for Smith, but not for Mary and John. Smith knows p, but Mary and John don't know because of the 4a violation.

[38] Does a speaker's context determine how flat a surface must be in order to be flat? Yes, to some extent. 'Flatness' is a group resemblance concept. 'Flat' objects have an appearance associated with paradigm prototypes or exemplars within the domain specified. The *standards* used by speakers *result from the context* of attribution (the intentions, presuppositions, etc., of speakers and listeners). The sentence 'x is flat' is in part, judged by *S's standards* and *audience standards* (which are usually the same, but not always).

> Does knowledge come in degrees? Most people say no (although David Lewis, 1996, says yes). But it doesn't really matter. For on my view, justification, or having good reasons, is a component of knowledge, and justification comes in degrees. So, context will determine how justified a belief must be in order to be justified *simpliciter*.

Cohen argues that since justification is a component of knowledge, an ascription of knowledge involves an ascription of justification, and ascriptions of justification are context sensitive. He continues with his linguistic and worldview intuitions about 'semantic models' (p. 61):

> - How from the viewpoint of formal semantics should we think of this context-sensitivity of knowledge ascriptions? We could think of it as a kind of indexicality. On this way of construing the semantics, ascriptions of knowledge involve an indexical reference to standards. So, the knowledge predicate will express different relations (corresponding to different standards) in different contexts.
>
> - But we could instead view the knowledge predicate as expressing the same relation in every context. On this model, we view the context as determining a standard at which the proposition involving the knowledge relation gets evaluated. So, we could think of knowledge as a three-place relation between a person, a proposition, and a standard.
>
> These semantic issues, as near as I can tell, are irrelevant to the epistemological issues. As long as we allow for contextually determined standards, it doesn't matter how formally we construe the context-sensitivity.[39]

In a section entitled 'closure and tracking,' Cohen considers Nozick (1981) and Dretske's (1970) claim that knowledge is subject to what Nozick calls a "tracking" condition (pp. 63-64, paraphrased):

> **S** knows **p** *only if*: **(T)** If **p** were false, **S** would not believe **p**.
>
> Where **p** is a proposition that we ordinarily claim to know, this condition can be satisfied. But where **p** is the denial of the skeptical hypothesis, **(T)** fails to be satisfied. Thus, **(T)** falsifies premise (1) of the skeptical paradox, thereby falsifying the closure principle.

---

[39] Critical comment: Context doesn't determine standards. Persons in context determine individual standards (e.g., what undermining possibilities are relevant (in the wide sense), how much positive evidence is needed, is the evidence from a reliable source, how pragmatically important is it that p is known, and not just believed).

> If the closure principle is false, then what explains the appeal of skeptical arguments? I suppose Dretske and Nozick (D&N) could say that it is our failure to appreciate the deep truth about the nature of knowledge as revealed by their theory. Is it that we mistakenly believe in the closure principle because we fail to see how the tracking condition falsifies it?
>
> The problem with the D&N response is that many think the closure principle 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**?[40]
>
> The very fact that the tracking condition is inconsistent with the closure principle gives us *reason* to *reject* the tracking view of knowledge. This point of view is best expressed by Richard Fumerton:
>
> > In his discussion of empirical knowledge, probably the most startling, original, and dialectically ingenious move that Nozick makes is to take the most devastating objection to his view (the failure of closure) and embrace it as one of its advantages (Fumerton 1987).
>
> Because I, and many others, share Fumerton's view, let's move on to consider other responses to the skeptical argument.[41]

Cohen thinks that the Nozick-Dretske approaches simply deny closure without explaining the appeal of skeptical arguments:

> For if I can know on the basis of my evidence that I see a zebra, and know on the basis of my seeing a zebra that I'm not seeing a cleverly disguised mule, then what explains the intuition that *I fail to know* that I am *not* seeing a cleverly disguised mule? (p. 65, italics added).[42]

---

[40] Cohen states that closure "seems to be something like an axiom about knowledge" (2002, p. 312).

[41] Cohen's argument for epistemic closure is largely based upon strong linguistic intuitions, shared among a number of philosophers. But there is an alternative answer (contrasting propositional and epistemic closure) that affirms that **S** could know **p**, and know that **p** propositionally entails **q**, and yet fail to be in a position to know **q**. It affirms that: If **S** *knows* that his car is parked on Nelson Street, then (actually) *the car has not been stolen*. It denies, however, that if **S** *knows* that his car is parked on Nelson Street, then ***S*** *knows* (epistemically) that the car has not been stolen. See chapters 1 and 2.

[42] The PE explanation: In order to know that 'I see a zebra' I don't need to *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 based upon strong evidence and background beliefs, and *assume* them false, *without knowing* that defeating possibilities (e.g., a disguised mule) are false. The mule possibility is discarded as irrelevant (wide sense).

In the airport case, the context of ascription determines how good **S**'s reasons have to be in order for **S** to know. The truth value of an ascription can vary with either the strength of the subject's reasons or the strictness of the standards.

> On the contextualist view, we explain our confidence in the truth of our everyday knowledge ascriptions by supposing that our reasons are sufficient for us to know, relative to the standards of everyday contexts. When confronted with skeptical arguments however, the chance of error becomes salient, and the standards can shift. Skeptical arguments are forceful precisely because they can have this effect on us. In this new context, the standards are stricter and knowledge ascriptions true in everyday contexts are false. So, while the strength of our reasons remains fixed, the strictness of the standards for how strong those reasons have to be, varies across contexts. By supposing that knowledge ascriptions are context sensitive in this way, we can do justice both to our strong inclination to say we know and to the undeniable appeal of skeptical arguments (pp. 65-66).[43]

For Cohen rejecting closure is completely unacceptable. Contextualism defends the closure principle while explaining why there is an appearance of closure failure.

> We can illustrate this through Dretske's zebra case: My reasons for believing I see a zebra consist of the animals looking like zebras and being in pens marked 'Zebra.' My reason for believing I do not see a cleverly disguised mule consists of the inductive evidence that I have against the possibility of such a deception. It looks as if I know that I see a zebra, but I fail to know I do not see a cleverly disguised mule.
>
> We might be tempted to think my reasons for believing I see a zebra are stronger than my reasons for thinking that I do not see a disguised mule. But surely if we accept the closure principle, we accept that where **p** entails **q**, the strength of my reasons for believing **p** can be no greater than the strength of my reasons for believing **q**.
>
> So, my reasons for believing *I see a zebra* can be no stronger than my reasons for believing *I do not see a cleverly disguised mule*. According to contextualism, however, the standards for how strong my reasons have to be in order for me to know can vary across contexts.
>
> In contexts where we consider whether I know I do not see a cleverly disguised mule, the chance of error is salient, unlike everyday contexts where we consider whether I know I see a zebra. And when the chance of error is salient in a

[43] Skeptical arguments have much less 'intuitive appeal' than philosophers attribute to them. The salience of error may affect some assessments of 'adequate evidence' (affecting PE 4a), but this scrutiny only tends to *sometimes* diminish claims of knowledge. In reality, persons typically treat 'skeptical hypotheses' as similar to 'regress of reasons' skepticism. They are properly ignored. See chapter 3.

> context, the standards tend to rise. Thus, we evaluate whether I know I do not see a cleverly disguised mule at a stricter standard than at which we evaluate whether I know I see a zebra. This gives rise to the appearance of closure failure. But if we hold the context, and so the standards fixed, we see that the closure principle is not threatened (p. 66).
>
> So in everyday contexts the standards are such that my reasons are good enough for me to know that I see a zebra. And since my reasons for denying that I see a cleverly disguised mule can be no worse, my reasons are sufficient for me to know that proposition as well, given the standards of those contexts. Thus in everyday contexts, I can know that I don't see a cleverly disguised mule, on the basis of inductive evidence I have against such a scenario. In skeptical contexts, where the standards are higher, I fail to know, on the basis of the inductive evidence, that I do not see a cleverly disguised mule. But since my reasons for believing that I see a zebra can be no better, I fail to know that proposition as well, given the standards of that context. The appearance of closure failure results from the shift in standards that occurs when we move from considering whether I know that I see a Zebra to considering whether I know that I do not see a cleverly disguised mule (p. 66).
>
> So on a contextualist view, the appearance of closure failure results from our evaluating the antecedent and the consequent of the principle, relative to different standards. This happens in general when we consider instances of the closure principle where the consequent concerns knowing the falsity of a skeptical alternative. Again, this is because thinking about skeptical alternatives can cause the standards to rise. But if we evaluate the closure principle relative to a fixed context, thereby fixing the standards, it comes out true. So the paradox arises because of our failure to be sensitive to contextual shifts (p. 67).

Cohen concludes this section by stating that it will depend on context on which proposition of the three inconsistent propositions that a contextualist will deny. He affirms that epistemic closure is true in every context. In everyday contexts, **S** knows **p**, and the skeptical hypothesis is false. In skeptical contexts, the skeptical hypothesis is true, and **S** doesn't know that **not-BIV**.[44]

---

[44] Summary: Cohen's contextualist view proceeds in terms of *internalist* notions of evidence and rationality and is explained as a matter of 'personal justification.' In everyday contexts, S sees a zebra. In skeptical contexts, S doesn't know that she sees a zebra. The apparent failure of closure results from shifting standards. According to Cohen, the paradox arises because of our failure to be sensitive to contextual shifts. The contextualist argues that in everyday contexts, S can know not-BIV, but not in skeptical contexts.

XII. Keith DeRose's Contextualism

Keith DeRose is a leading advocate of epistemic contextualism with many published essays including "Contextualism and Knowledge Attributions" (1992), "Solving the Skeptical Problem" (1995), as well as a two-volume book set, *The Case for Contextualism* (2009) and *The Appearance of Ignorance* (2017). We will survey some of the contents of his two earlier essays for a good overview of his thought.

"Contextualism and Knowledge Attributions" (1992)[45]

DeRose begins with a case example involving the question of whether a certain bank location will be open on a Saturday. (This was pre-internet; financial transactions needed to be done in-person at a bank). This case presents linguistic intuitions that are a basis for addressing the skeptical paradox. We construct this case study in its essential form (just eliminating original reference to DeRose (**S**) and his wife (**S1**)).

The Bank Context A:

> **S1** and **S** are driving home on Friday afternoon. They plan to stop at the bank to deposit their paychecks. When driving to the bank, they observe typical Friday crowds on the inside of the bank.
>
> **S** suggests to **S1** that they drive straight home and deposit their checks on Saturday morning (since there is nothing financially urgent as to when the checks are deposited). **S1** responds "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." **S** replies, "No, I know it will be open. I was there just two weeks ago on Saturday. It is open until noon.

The Bank Context B:

> **S1** and **S** are driving home on Friday afternoon. They plan to stop at the bank to deposit their paychecks. When driving to the bank, they observe typical Friday crowds on the inside of the bank.
>
> **S** again suggests to **S1** that they drive straight home and deposit their checks on Saturday morning explaining that "I know it will be open. I was there just two weeks ago on Saturday. It is open until noon." But in this case, **S** and **S1** have written a very large and very important check. If their paychecks are not deposited before Monday morning, the important check will bounce, and of

[45] "Contextualism and Knowledge Attributions" *Philosophy and Phenomenological Research*, Vol. LII, No. 4, December 1992.

> course the bank is closed Sunday. **S1** reminds **S** of this urgent financial situation. **S1** then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as before (that the bank will be open), still **S** replies, "Well, no. I'd better go in and make sure."[46]

DeRose continues with interpreting the example (again edited):

> Assume that in both cases, the bank will be open on Saturday, and there is nothing unusual about either case that has not been included in my description of it. It seems to me that:
>
> **(1)** when ***S*** *claims to know* that the bank will be open on Saturday in Case A, **S** is saying something true.
>
> But it also seems that **(2) S** is saying something true in Case B when ***S*** *concedes not knowing* that the bank will be open on Saturday. Yet **S** seems to be in no better position to know in Case A than in Case B.
>
> It is quite natural to say that **(3)** If **S** knows that the bank will be open on Saturday in Case A, then **S** also knows that it will be open in Case B.[47]

DeRose wants to investigate and defend a view that *all three* of these propositions are *true*.[48] He says it would be *inconsistent* to claim that (1) and (2) are true, and also hold:

> **(4)** If what **S** says in Case A claiming to know p is true, then what **S** says in Case B in conceding not knowing that **p** is false.

But DeRose says that there is a big difference between (3) and (4) and this difference is crucial to the view that he wants to investigate and defend.

---

[46] Is S 'as confident as before' when going into the bank? This shouldn't be stipulated if not really true.

[47] I suggest that we might rephrase DeRose's example as follows (*= false):

> 1) 'S knows that p' (The bank will be open Saturday morning) is true in A.
> 2) 'S knows that p' (The bank will be open Saturday morning) is false in B.
> *3) If S knows p in Case A, then S knows p in Case B.
> *4) If what S says in Case A claiming to know p, then what S says in Case B in claiming not to know p is false.

[48] I maintain that premise 3 is false (i.e., that if S knows p in case A, then S knows p in case B). Also, the term epistemic "position" suggests interpreting S's perceptual (and inferential) knowledge as a function of S interacting with states-of-affairs found in the external world. What exactly is an 'epistemic position'? It is a view that there are facts and that epistemologists should describe how S acquires knowledge of these (external) facts. To ask whether S is in *position* to know an empirical p, is a metaphysical question. Also, shouldn't S's epistemic "position" include internal states (beliefs, knowledge, situational pragmatics)? Is the concept of 'epistemic position' specified by semanticists-metaphysicians needed for explaining knowledge?

De Rose follows Unger's (1984) terminology of 'contextualism' and 'invariantism.'

The view that DeRose is investigating is a "contextual theory of knowledge attributions":

> It is a theory according to which the truth conditions of sentences of the form '**S** knows that **p**' or '**S** does not know that **p**' vary in certain ways according to the contexts in which the sentence is uttered. The contextualist can deny (4) even while admitting that **S** is in no better position to know in Case A than in Case B. The contexts of **S**'s utterances in the two cases, make it easier for a knowledge attribution in Case A than in Case B. [49]

DeRose says that there are important contextual differences between Case A and Case B which one might think relevant:

> 1) The *importance* of being right. The requirements for making a knowledge attribution goes up as the stakes go up. In Case B there is more importance for p to be true.
>
> 2) The *mentioning* of a possibility. In Case B, **S1** raised the possibility that the bank's hours might have changed over the course of two weeks.
>
> One might think that if this possibility has been mentioned, **S** cannot truly claim to know p (on the evidence that 'p was true two weeks ago'), unless **S** can rule-out the possibility that the bank's hours have changed since then.[50]
>
> On the other hand, perhaps **S** doesn't need to be able to rule-out this possibility in order to truly say **S** knows, as in Case A, since no possibility has been suggested.
>
> In considering the bank cases, the invariantist will assert (4), which seems very plausible, and will therefore deny either (1) or (2). Typically, the invariantist will deny (1). In fact, Unger uses the term "invariantism" to denote the position that the standards for true knowledge attributions remain constant and very high—as high as they can possibly be. This position I will call "skeptical invariantism," leaving the more general term "invariantism" to denote any position according to which the truth conditions for knowledge attributions do not vary in the way the contextualist claims they do, whether or not the standards are said to be very high.

---

[49] A review of formal semantics: The "truth condition" of a sentence is the core of a sentence's meaning. A sentence represents the situation (or 'state of affairs') that would make the sentence true. To know the meaning of a sentence, is to know what the world has to be like, if the sentence were to be true. A meaningful declarative sentence represents the world as being a certain way and is either true or false

[50] The term 'rule-out' is ambiguous. Some philosophers (curiously) take the phrase 'rule-out the possibility' as synonymous with 'knowing the defeating possibility is false.' In another ordinary sense, S may internally 'rule-out' a potential defeater possibility, deeming it *improbable*, but still acknowledge that a *defeater possibility* nevertheless *exists*. In the parked car case, S may rule-out the possibility that the car is stolen (in a context) as part of PE #4a to be *personally* justified in believing where the car is parked (while admitting that the *possibility* of the car having been stolen *always exist*s).

> I will then use "non-skeptical invariantism" to refer to a position according to which the standards are held to be constant but relatively low (pp. 915-16).[51]

In defending contextualism, DeRose admits that contextualists can disagree about what features of the context of utterance really do affect the truth conditions of knowledge attributions and to what extent that they do. But in this essay, he will address general issues that confront the contextualist, and distinguish it from the 'relevant alternatives' theory. DeRose states (p. 917, italics added):

> Contextualist theories of knowledge attributions have almost invariably been developed with an eye towards providing some kind of answer to philosophical skepticism.[52] For some skeptical arguments threaten to show, not only that we fail to meet very high requirements for knowledge of interest to philosophers seeking absolute certainty, but also that we don't meet the truth conditions of ordinary, out-on-the-street claims to know. They thus threaten to establish the startling result that we never, or almost never, truly ascribe knowledge to ourselves or other human beings.
>
> According to contextual analysis, when the skeptic presents her arguments, she manipulates various conversational mechanisms that raise the semantic standards for knowledge and thereby creates a context in which *she can truly say* that *we know nothing or very little*. But the fact that the skeptic can thus install very high standards which we don't live up to has *no tendency* to show that we *don't satisfy* the more *relaxed standards* that are in place in ordinary conversations. Thus, it is hoped, our ordinary claims to know will be safeguarded from the apparently powerful attacks of the skeptic, while at the same time, the persuasiveness of the skeptical argument is explained.[53]
>
> Many find such contextualist resolutions of skeptical arguments very attractive since their main competition is the skeptical invariantist resolutions according to which the persuasiveness of various skeptical arguments is explained in a way that is alarming as it is simple: They seem persuasive because they are indeed sound and successfully establish the startling conclusion that we never or almost never truly ascribe knowledge.

---

[51] DeRose's four propositions under consideration are very artificial. "Invariantism" defined as "any position according to which the truth conditions for knowledge attributions *do not vary* in the way the contextualist claims they do" is a fixed definiens concept that strays far from ordinary language.

[52] It isn't good that an epistemic theory is centered on a single issue. A comprehensive view is needed.

[53] With a contextualist theory of personal justification (chapter 3), I have implicitly argued that there is no manipulation of conversational mechanisms. Typically, any normal S will dismiss (i.e., exclude or rule-out) 'skeptical' and 'regress-of-reasons' challenges. The radical skeptical argument isn't that persuasive.

> But many while finding the contextualist resolutions a preferrable alternative to an unacceptably radical form of skepticism, at the same time fell an initial resistance, closely related to the appeal of (4), to the thought that contextual factors of the types I've mentioned can really affect whether or not a subject knows.

In part II of the essay, DeRose wants to carefully distinguish contextualism from 'relevant alternatives' theories. DeRose says that relevant alternatives theories are often thought of as a popular form of contextualism, but that it is tricky to say how 'relevant alternative theories' are contextualist.

> According to RA, a claim to know that **p** is made within a certain framework of relevant alternatives which are incompatible with **p**. To know that **p** is to be able to distinguish **p** from these relevant alternatives, to be able to rule out these relevant alternatives to **p**. But not every contrary or alternative to **p** is a *relevant* alternative (p. 918).

DeRose discusses Dretske's Zebra case, where **S** claims to know that '**x** is a zebra,' the proposition that '**x** is a disguised mule' is not an alternative, and **S** need not be able to rule it out in order to *truly claim* that *x is a zebra*. Only in a very extraordinary case, might that alternative become relevant.

But DeRose asks concerning Dretske's analysis, how can the "**x** is a disguised mule" hypothesis become relevant? In order to consider this question, he examines a standard presentation of relevant alternatives, that of Alvin Goldman (1976). DeRose considers Goldman's analysis of the Henry and the Barn case and suggests his own analysis. A number of technical concepts are used in the very dense discussion: 'subject factors,' 'attributor factors,' 'contextualist,' 'invariantist,' 'lucky,' 'meaning,' 'character,' 'content,' and 'epistemic position.' In the final portion of the essay, DeRose returns to a discussion of the zebra case. DeRose continues with concerns about 'knowledge attribution' within a framework of formal semantics. He says that insofar a RA theorist allows attributor factors to influence which alternatives are relevant, this theorist is a contextualist. An invariantist can be a RA theorist, if he allows only subject factors to influence which alternatives are relevant. An RA theorist need not be a contextualist.[54]

---

[54] DeRose's use of technical terms and concepts leads to a technical conclusion.

"Solving the Skeptical Problem" (1995)[55]

This DeRose essay was published in *The Philosophical Review* and is reprinted in *The Appearance of Ignorance* (2017). The page references are to the book. We will follow DeRose using the section titles as found in the essay.

(1) The Puzzle of Skeptical Hypotheses

The Argument from Ignorance (AI):

1) I don't know that not-H. (Skeptical Hypothesis)
2) If I don't know that not-H, then I don't know that o. (Epistemic Closure)
3) So, I don't know that o. (Skeptical conclusion, I don't know o).

DeRose reacts to this argument by saying that it is a paradox, or a puzzle, and DeRose sets out to prove (or explain) that 'I have hands' is a proposition that is *known* by **S** in ordinary context. Both premises seem plausible, but the conclusion doesn't. He is looking for a correct solution to this puzzle (p. 1).

(2) Contextualist solutions: Basic Strategy

Suppose speaker **A** (attributor) says "**S** knows **p**" of **S**'s true belief that **p**.

> According to contextualist theories of knowledge attributions, how strong an epistemic position **S** must be in with respect to **p** for **A**'s assertion to be true can vary according to features of A's conversational context.[56]
>
> Contextualist theories of knowledge attributions have almost invariably developed with an eye toward providing some kind of answer to philosophical skepticism. The skeptic's conclusion that we almost never truthfully ascribe knowledge to ourselves, or to others, is startling.

In response to skeptical arguments, contextualists maintain that the skeptic manipulates the semantic standards for knowledge, thereby creating a context in which it can *truthfully* be said S knows nothing or little. Once *standards are raised*, *we* correctly sense

---

[55] "Solving the Skeptical Problem" *The Philosophical Review*, Vol. 104, No.1 (January 1995).

[56] The worldview that we should be concerned with the *truth* of an *attributor's assertion* (e.g., the skeptic) about whether S knows p or not, as part of a semantic-epistemic analysis seems mistaken. The PE definition of 'knowledge' (assumed as a natural kind with necessary and sufficient truth conditions) resolves the skeptical argument in a simple, more direct way, than does the contextualist semantic analysis.

that we only could *falsely claim* to *know* such things as 'we have hands.'[57] But DeRose observes, that as soon as we find ourselves in more ordinary conversational contexts, it will be not only true for us to claim to know **p** but it would also be wrong for us to deny that we know these things. DeRose says:

> For the fact that the skeptic can invoke very high standards that we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in more ordinary conversations and debates.

Part of the contextualist solution is to identify the mechanism by which the skeptic at least threatens to raise the standards for knowledge. He says:

> The contextualist's ultimate point will be this: To the extent that the skeptic does succeed, she does so only by raising the standards for knowledge, and so the success of her argument has no tendency to show that our ordinary claims to know are in any way defective.

DeRose says that the contextualist explanation involves the *standards* for knowledge being *changed* by what is said in *conversation*. He continues:

> For the most part, I will frame the contextualist explanation in terms of such conversational rules, largely because that's what has been done by my contextualist predecessors, with whom I want to make contact.[58]

DeRose wants to explain why the skeptical argument is so appealing. DeRose's solution is to employ a contextualist strategy to explain the conversational rule or mechanism by which the skeptic raises (or threatens to raise) the standards for knowledge.

(4) Some Old Contextualist Solutions: The "Relevant Alternatives" Approach and the Rule of Relevance

DeRose begins by discussing the "relevant alternatives" (RA) approach as a popular solution to the skeptical puzzle.

Suppose that **S1** says '**S** knows that **p**."

---

[57] In contrast, I contend that I know that 'I have hands' in all contexts. I never infer from the radical skeptical argument that there are contextual occasions where 'I know that I have hands' is false.

[58] The primary target audience of the contextualist explanation is a narrow field of specialists who are familiar with (and adopt) each other's inter-connected terminological jargon. Whether philosophers *should* be exclusively debating with their predecessors is questionable. Also, the idea of "conversational rules" is murky. (A less-technical, less-stipulative vocabulary, to engage the average person is better).

> According to RA, such an assertion is made within and must be evaluated against a certain framework of *relevant alternatives* to **p**. To know that **p** is to have a true belief that **p** and to be able to rule out these relevant alternatives. But not every contrary of or alternative to **p** is a *relevant* alternative.
>
> In an ordinary case of claiming to know that some animals in the zoo are zebras, the alternative that they're cleverly disguised mules is not relevant. Thus, S can truthfully claim to know that **x** is a zebra despite the *inability* to *rule out* this fanciful alternative.[59] Only in extraordinary cases would the disguised mule hypothesis become relevant (e.g., at a zoo that fairly consistently uses painted mules to fool the zoo-going public).

DeRose says the *Relevant Alternatives Solution (RAS)* is as follows (paraphrased, p. 11):

> The AI skeptic's *mentioning* of the BIV hypothesis of the first premise of AI *makes* that hypothesis relevant. Once the skeptical hypothesis has been made relevant, we correctly sense that we cannot truthfully claim to know anything contrary to it unless we can rule it out. Since we are unable to rule it out, and since it is an alternative to both *I am not a BIV* and to *I have hands*, we correctly sense that we could only falsely claim to know these things. So, the skeptic truthfully asserts that we don't know that the hypothesis doesn't obtain, and then truthfully concludes that we don't know that we have hands.
>
> We realize that in most of our conversational circumstances in which we find ourselves, our inability to rule-out the skeptic's far-fetched hypothesis is no bar to our truthfully claiming to know such things as that we have hands. Thus, even as we find the skeptic's denials of knowledge persuasive, we realize that when we again find ourselves in more ordinary contexts, it will not only be correct for us to claim to know such things, it would be wrong to deny that we know them merely because we can't rule out the BIV hypothesis.[60]

S's claim to *know* 'I have hands' is *compatible* with the skeptic's *denial* that **S** knows.

> RAS, then, is an instance of the general contextualist strategy—one according to which the raising of the standards consists in enlarging the range of alternatives that are relevant and that **S** must therefore be in a position to rule out in order to count as knowing.[61]

---

[59] According to the PE theory, 'S can know z despite not knowing not-m' is true, but S must rule-out m (despite not knowing it) in order to have a consistent belief (i.e., not an inconsistent belief that z and maybe m, which contradicts the belief that z). In other words, S may (fallibly) rule-out m in order to believe z.

[60] Is the skeptical argument persuasive? Not if epistemic closure is denied.

[61] Is knowledge largely a 'ruling out' process? This is a distorted view of knowing. Knowing is more concerned with paying attention to inconsistencies (so-called 'red flags') that could be evidence that one's present belief is false. In order to know, we don't typically engage in counter-possibilities to our belief.

> The conversational rule or mechanism that RAS posits for enlarging that range (raising the standards), then, is that *mentioning* a proposition **q**—ceteris paribus with within certain limits, no doubt—tends to make **q** a contextually relevant alternative to any **p** that is contrary to **q** (p. 11).[62]

DeRose names this conversational rule (or mechanism) the *Rule of Relevance*. He says that this rule can explain how a positive claim to know the skeptical hypothesis doesn't obtain has the same effect on the meaning of sentences containing 'know' as would a denial of such knowledge. But DeRose is critical of the relevant alternatives solution (RAS):

> But to explain the persuasiveness of AI (particularly of its first premise) and thereby to solve our puzzle, a treatment of AI must tell us what it is about skeptical hypotheses that makes it difficult to claim to know that they don't obtain. The key feature of the skeptical hypotheses that RAS seizes on is clearly this: we can't rule them out (p. 12).[63]

Following two additional deliberative paragraphs, he continues:

> What accounts for the plausibility of saying that I don't know that I'm not a BIV? The fact that I can't discern that I'm not one? This is no explanation. It seems just as good (in fact, to me, better) to reverse things and claim that the fact that I don't know that I'm not a BIV accounts for the plausibility of saying that I can't discern that I'm not one.
>
> Likewise, for ruling out. It is indeed plausible to suppose that we can't rule out skeptical hypotheses. And it's plausible that we don't know that they don't obtain. But it doesn't advance our understanding much to explain the plausibility of either by the other (p. 12).[64]

---

[62] That the conversational mention of a defeating possibility to S about p might make it a (wide sense) relevant possibility is plausible.

[63] With the PE theory, we personally have to rule-out skeptical hypotheses (even if not known false). Although we don't know that a well-proposed skeptical hypothesis is false, we still (materially) know p.

[64] These are two contentious paragraphs and, it seems to me, are subject to counterargument, even if one isn't advocating the RAS.

9) Nozick's Own Solution and the Abominable Conjunction

DeRose embraces Nozick's 'subjunctive conditionals account' (SCA) and the concept of 'sensitivity" as an account of the plausibility of the skeptical hypothesis. From above, DeRose has stated:

> SCA's generalization can be restated as follows: We tend to judge that **S** doesn't know that **p** when we think that **S**'s belief that **p** is insensitive. **S**'s belief that **p** is "insensitive" if **S** would believe that **p** if **p** were false.

He says that Nozick's own treatment of the AI paradox *fails*, where Nozick accounts for knowledge as true *sensitive* belief, where roughly, **S**'s true belief that **p** is sensitive to the truth value of **p** if **S** would not have believed that **p** if **p** had been false. He continues:

> Thus, Nozick's treatment of AI involves accepting the skeptic's first premise. But at the same time, and much more unfortunately, it also involves denying the second. You don't know that you're not a BIV, Nozick claims, because any belief that you might have to this effect is insensitive: You would have held this belief even if it were false (even if you were a BIV). By contrast, Nozick claims your belief that your belief that you have hands *is* a sensitive belief: If *it* were false—you didn't have hands—you would not hold it. So, you do know you have hands even though you don't know that you're not a BIV. The skeptic's mistake—the second premise—is supposing that you can have hands only if you also know that you're not a BIV (p. 21).

DeRose strongly criticizes Nozick's rejection of the epistemic closure principle:

> Nozick takes his implausible stand on the issue of the second premise, denying it in the face of its evident intuitive appeal. Accepting this treatment involves embracing the abominable conjunction that while you don't know you're not a bodiless (and handless) BIV, still you know you have hands. Thus, while the account does quite well on the relevant particular intuitions regarding what is and isn't known, it yields an intuitively bizarre result on the comparative judgment the second premise embodies (pp. 21-22).

(10) Strength of Epistemic Position and AI's Second Premise

Given that DeRose wants to preserve the "evident intuitive appeal" of the closure principle which maintains that knowledge is closed under known propositional implication, he proposes incorporating SCA (subjunctive conditionals account) into a

contextualist solution of the puzzle. He proposes a very strong endorsement of the AI second premise. He states:

> Recall that according to contextualist theories of knowledge attributions, how strong a subject's epistemic position must be to make true a speaker's attribution of knowledge to that subject is a flexible matter that can vary according to the features of the speaker's conversational context. Central to contextualism, then, is the notion of *(relative) strength of epistemic position* (p. 22).

DeRose says that most people understand (as an intuition):

> …that sometimes the standards for knowledge are higher than usual, or that in some conversational situations one's epistemic position must be stronger than in others to count as knowing.[65]

> But it would be good to *clarify* this important notion of *strength of epistemic position* as best we can by, for instance, supplying an intuitive test for when one epistemic position is stronger than another. The best such device is that of *comparative conditionals*.[66]

In a single paragraph, DeRose presents a short and complex theory of comparative grounds (involving 'tallness') for assenting to conditionals and comparative epistemic positions.[67] He thinks that these logical intuitions are a good intuitive test for understanding 'strength of epistemic position.' He follows with another single paragraph analysis of the Henry and the Barn case with his complex and questionable analysis using comparative conditionals.[68]

---

[65] It should be understood that when the standards for justification are higher than usual, the standards for knowledge (the four PE conditions) remain the same. Thus, we *agree* that sometimes the evidential standards for knowledge are higher than usual, or that in some situations one's epistemic position (quality and quantity of reasons) *must* be stronger than in others to count as knowing.

[66] Do ordinary persons have these linguistic and worldview intuitions leading to the formulation of a concept of 'strength of epistemic position'? Or is the initial naming of this formal technical stipulative definition a mechanism (or tool) for defending DeRose's worldview about radical skepticism and the radical skeptical argument?

[67] DeRose say that comparative conditionals are a test: If S knows that p in A, then S knows that p in B. S is in at least as strong epistemic position with respect to p in situation B as he is in situation A, and this comparative conditional serves as a good intuitive test for that comparative fact: It brings that fact to light.

[68] DeRose's analysis: A normal situation is exactly like façade situation, except that there are no fakes in the area—the things Henry has taken to be barns have all been actually barns. With regard to these examples, the conditional *If Henry knows in façade situation, then he knows in normal situation* seems to get the comparison right, indicating that Henry's in at least as strong an epistemic situation in the normal situation

(14) The Contextualist Solution Clarified

In this section, DeRose presents the 'puzzle' of the 'skeptical hypotheses' involving two premises and a conclusion. He says (p. 30):

> The puzzle of the skeptical hypotheses, recall, concerns the premises of AI together with the negation of its conclusion:
>
> (1) I don't know that not-H.
> (2) If I don't know that not-H, then I don't know that O.
> (3) I do know that O.
>
> A solution to the puzzle must, of course, issue a verdict as to the truth of each of these three, but it must also explain why we find all of them plausible.
>
> Let's be clear about what our present contextualist solution has to say about each of these. Our verdict regarding (2) is that it's true regardless of what epistemic standards it's evaluated at, so its plausibility is easily accounted for. But this, combined with a similarly enthusiastic endorsement of (1), would land us in bold skepticism. We avoid that fate by endorsing (1) as true, not at all standards, but only at the unusually inflated standards, conducive to skepticism. Thus, on our solution, we do know for instance, that we're not BIVs, according to ordinary low standards for knowledge.

In short, DeRose's verdict is that premise (2) is true regardless of epistemic standard it's evaluated at. He endorses premise (1) as (sometimes) true, but not on all standards, only at the inflated standards conducive to skepticism. On DeRose's solution: We do *know* that we are *not* BIVs, when practicing the ordinary low standards for knowledge.[69]

---

as the façade situation. The evident failure of *If Henry knows in normal situation, then he knows in façade situation* to get the comparison right shows that Henry's not in as strong a position to know in *façade situation* as in *normal situation.* Together, these results indicate that Henry's in a stronger epistemic situation in *normal situation* than in *façade situation.* Comparative conditionals can similarly be used to test the relative strength of epistemic position of a single subject with respect to different propositions that subject believes in the same situation: Thus, the intuitive correctness of *If S knows that p, then S knows that q* and *If S doesn't know that q* then *S doesn't know that p* can indicate that S is in at least as strong an epistemic position with respect to q as she's in with respect to p.

[69] This conclusion that in ordinary low standard situation, we *can* know not-BIV, clashes with the PE theory which maintains that we *cannot* (ever) know not-BIV.

XIII. Ernest Sosa

Ernest Sosa is an influential scholar. He is credited as being the founder of the "virtue-theoretic" approach to epistemology. The virtue approach treats knowledge as a particularly 'successful or valuable form of belief.' A virtue-theoretic theory rejects the 'JTB + X' form of account of the definition of 'knowledge' and replaces it with a truth functional combination of independent (metaphysical) epistemic properties. Besides his virtue epistemology, Sosa has contributed work to the following question: If a belief is to be an instance of knowledge, what modal link must there exist between the belief and its truth? According to Nozick, knowledge requires 'sensitivity' but for Sosa it requires 'safety.' This is the topic of "How to Defeat Opposition to Moore."

"How to Defeat Opposition to Moore" (1999)[70]

In following other epistemologists, in "How to Defeat Opposition to Moore" (1999), Sosa proposes a "safety condition" to improve upon the "sensitivity condition":

Sensitivity: If p were false, S would not believe that p.[71]

Safety: If S were to believe that p, p would not be false.[72]

Sosa asks: what modal relation must a fact bear to a belief in order for this belief to constitute knowledge of that fact? Sosa thinks 'sensitivity' is open to serious objections, with 'safety' being preferable. He says that with safety, it is possible to defend plain "Moorean common sense" against skepticism.[73]

We follow Sosa's thinking (p. 141):

---

[70] Sosa, Ernest 1999. 'How to Defeat Opposition to Moore' in *Philosophical Perspectives* 13, Epistemology. James E. Tomberlin (ed.). Malden, MA: Blackwell Publishers.

[71] Various forms of externalism share a modal conception of 'sensitivity.' This concept is important for the theories of Dretske, Nozick, and DeRose.

[72] Characterizing 'safety' in counterfactual terms depends on substantive assumption about the semantics of counterfactual conditionals: In all nearby worlds in which S believes that p, p is not false.

[73] Moorean "common sense" is that (1) the skeptical hypothesis is false (and S knows that not-BIV), (2) epistemic closure is true, and so (3) S knows ordinary propositions. This "common sense" is false.

(1) The concept of sensitivity: A belief by **S** that **p** is 'sensitive' iff were it not so that **p**, **S** would not believe that **p**. Sensitivity is a necessary condition for knowledge.

Sensitivity: In order to constitute knowledge a belief must be sensitive.

(2) An "alternative" to a proposition is any incompatible possibility. (Among the truths, only the contingent truths have alternatives, since no "possibility" can be incompatible with necessary truth. To "rule out" such an alternative is know that it is not the case. The following principle of exclusion now seems plausible:

Principle of Exclusion (PE): In order to know a fact **p** *one must rule out* (i.e., know to be false) every alternative that one knows to be incompatible with it.[74]

That creates a problem for sensitivity requirement:
My belief that: (**o**) here is a hand

might constitute knowledge even though my belief that

(-h) I am not now fooled by a demon into believing that here is hand.

is not sensitive, despite my knowing that **o** entails **~h**. But if my belief of **~h** is not sensitive, then the sensitivity requirement precludes my knowing **~h**, and precludes thereby my ruling out **h**, which, in combination with PE, precludes in turn my knowing **o**. Advocates of the 'relevant alternatives' approach reject PE in its full generality. Instead, they propose this:

PE-relevant alternatives: In order to know a fact **p** one must rule out every *relevant* alternative that one knows to be incompatible with it.

Sosa continues (p.142, paraphrased):

Thus, **S** might know that **o**: here is a hand, despite being unable to rule-out the hypothetical possibility that **h**: one is a BIV. Replacing PE with PE-relevant enables them to exclude alternative **h**, *if they can marginalize that alternative as irrelevant*.

What is the difference between relevant and irrelevant alternatives? What makes an alternative irrelevant? No answer is generally accepted, even among relevantists, and the notion of relevance remains obscure, no published account having yet much relieved this darkness.[75]

---

[74] The 'principle of exclusion' seems false. S isn't required to rule-out possibilities by knowing them false.

[75] The concept of 'relevance' is demystified in chapter 1.

Sosa proposes an alternative approach (pp. 142-43, paraphrased):

> "Safe" (stipulative definition): Call a belief by **S** that **p** "safe" iff: **S** would believe that **p** only if it were so that **p**.[76]
>
> (Alternatively, a belief by **S** that **p** is "safe" iff: **S** would not believe that **p** without it being the case that **p**; or better, iff: as a matter of fact, though perhaps not a matter of strict necessity, not *easily* would **S** believe that **p** without it being the case that **p**).[77]
>
> *Safety*: In order to (be said correctly to) constitute knowledge a belief must be safe (rather than sensitive).
>
> While akin to sensitivity, safety has important advantages. Principle PE, for example, does not give *safety* the problem, we saw it give to *sensitivity*.
>
> Suppose the belief above to be a *safe* belief, and consider the paired skeptical proposition **~h** that one knows to be entailed by **o**. Although one's belief of **~h** is clearly not sensitive, it does seem quite safe. In other words, unlike sensitivity, safety is preserved in this known entailment. No belief constitutes knowledge unless safe, we may now say, while leaving ourselves free to exclude such skeptical scenarios that we know to be incompatible with something we know. If you know that **o**, and you know that some such scenario **h** is necessarily incompatible with **o**, you are not precluded by the safety requirement from knowledgably excluding that scenario.
>
> Replacing the sensitivity requirement with the safety requirement may thus unable a unary conditionals-theoretic account of knowledge in need of no distinction between relevant and irrelevant alternatives. (This is counter to some of the rationale for the relevant alternatives tack).[78]

---

[76] The core claim of the safety condition on knowledge is that knowledge requires belief that is safe from error. Roughly what this means is that one's belief must not only be true, but it must also be the case that one might not easily have believed falsely. Or, given a standard possible world semantics for 'might,' to say that one believes safely is to say that one not only believes the truth at the actual world, but one also avoids errors at nearby possible worlds. In this way, safety is a modal condition on knowledge.

[77] 'Easily' is a group resemblance concept.

[78] In sum, Sosa is claiming that his 'safety' condition is better than 'sensitivity' and 'relevant alternatives' theories, which are his metaphysical competitors.

Sosa follows with an analysis of the skeptical argument, comparing the views of Moore, Nozick, and DeRose. Let us note what Sosa takes to be as a counterexample to 'sensitivity' (pp. 145-46):

> 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 snagged somehow in the chute on the way down (an incredibly rare occurrence), or some such happenstance.
>
> But none such could affect my predictive belief as I release it, so I would still predict that the bag would soon arrive in the basement. My belief seems not to be sensitive, therefore, but constitutes knowledge anyhow, and can correctly be said to do so.

In sum, Sosa takes it as intuitive (and granted) that in this circumstance that he *knows* that the bag is in the basement. Since this belief isn't sensitive, this shows that the sensitivity principle is false.[79] He completes his essay by defending safety and the Moorean stance.[80]

---

[79] 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 when releasing the bag, the chute normally functions in a way for the bag to reach the basement. But what about the possibility of a rare snag? For **S** to know **p**, all four conditions of the PE definition must be satisfied and so typically **S** knows **p**. This is the normal situation for a well-maintained high-rise (that Sosa assumes). But, if conditions 1, 3, and 4b are false (e.g., the bag is snagged in the chute, after normal wear, or bad installation, or squirrel nest), then **S** *fails* to know **p**. If **S** is (somehow) aware of the live possibility of a malfunction, and for some reason cannot dismiss this counter-possibility **q**, then **S** doesn't know **p**. In a case of **S**'s knowledge (i.e., the trash *is* in the basement), **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 reasons for believing **p**. In determining whether '**S** knows **p**,' the contingencies of **S**'s external environment take precedence.

[80] The necessary (modal) condition proposed by Nozick is 'sensitivity':

> Sensitivity: S knows that p only if were p false, S wouldn't believe that p.
> (1) S's belief that p (I have hands) is sensitive: If S didn't have hands (e.g., lost in an accident), S wouldn't believe p (i.e., I have hands)
> (2) S's belief that p (I am not a BIV) is not sensitive: If S was a BIV, everything would seem the same, and S would continue to believe p (i.e., I am not a BIV).
> (3) On this account, premise #2 is false (epistemic closure).

The necessary (modal) condition proposed by Sosa is 'safety':

> Safety: as a matter of fact, though perhaps not a matter of strict necessity, not *easily* would S believe that p without it being the case that p
> (1) S's belief that 'I have hands' is *safe* because there are no nearby worlds in which S falsely believes that he doesn't have hands. S knows 'I have hands.'

'Safety' unlike 'relevant alternatives' and 'sensitivity conditions,' permits S to *know* that the skeptical hypothesis is false. Although, 'safe' is a gradable adjective, the safety condition is not presented within a contextualist semantics, where the truth conditions of the proposition "S knows p" depends on the context of the attributor. Since safety isn't a part of contextual semantics, Sosa is (by definition) an invariantist.[81]

***Epistemology* (2017)**

We turn to Sosa's *Epistemology* (2017). It is a textbook that expounds a virtue epistemology that Sosa has developed over several decades. He states that the main objective of this book is to make the view maximally accessible to the uninitiated.

In chapter one, entitled "Descartes' Pyrrhonian Virtue Epistemology," Sosa finds affinities in the writings of Descartes and the Pyrrhonists, with his virtue analysis (p. 19):

> In all important structural respects, Cartesian virtue epistemology takes the same view as a virtue epistemology on the contemporary scene, *virtue perspectivism* (found in *Knowledge in Perspective,* 1991). The structure of this view does not require the theological content that Descartes gives to his own version. The role of theology can be played instead by science, by common sense, or by the two combined. … it should become increasingly clear as the contemporary view is developed to include more explicitly the epistemic agency that played so central a role in Cartesian epistemology. This is a development of the contemporary view that is now underway (in *Judgment and Agency*, 2015).

---

(2) S's belief that 'I am not a BIV' is *safe* because there are no nearby worlds in which S falsely believes that he is not a BIV. S knows 'I am not a BIV.'
(3) Premise #2 (epistemic closure) is retained. Premise #1 that S cannot know not-BIV is rejected.

[81] A safety-based account *explains* the absence of knowledge in the 'stopped clock' and 'fake barn' cases:
(1) Safety requires avoidance of error not only in the actual world but also at nearby possible worlds.
(2) In the 'stopped clock' and 'fake barn' case, S obtains a true belief in the actual world. However, S falls into error at the vast majority of nearby possible worlds.
(2a) In stopped clock, there are many nearby worlds in which S looks at the clock at least a minute earlier or later. (S would have a *false* belief that it is noon, violating that p must be true, as a necessary condition for knowledge).
(2b) In fake barns, there are many nearby worlds in which S looks at x, but is driving further down the road, and sees façade rather than a barn (since there is a 99% chance of façade).
(3) S's belief isn't safe in either 'stopped clock' or 'barn façade.'
(4) Since S's belief that p isn't safe, S doesn't know p. This conclusion answers (i.e., explains) why S doesn't have knowledge in either 'stopped clock' or 'barn façade' cases. In S's circumstance and environment, S's true belief isn't knowledge, because p could easily have been false. In these situations, S would (or might well) retain the (true) belief, even if it was coincidental, and this isn't knowledge.

Chapters two and three are entitled "Dream Skepticism" and "Regress Skepticism." We will proceed straight to chapter four, "Knowledge: What It Is and How We Might Have It." We begin with the initial paragraphs of that chapter (p. 55):

> According to the foregoing chapters, we can and do enjoy knowledge despite the protestations of skeptics. But what is this knowledge? At a minimum, how might we determine the conditions necessary and sufficient for its attainment?
>
> (1) A gambler rolling dice might be persuaded that the dice will come up 7 or 11 in his next toss, and he might be right. But he still would not know.[82]
>
> What more would he need in order to know, besides a belief that is true? A reason, perhaps, a good reason?

At this point, Sosa recites two Gettier cases (the stopped clock case, the Nogot case) as well as the Henry and the Barn case, where in these cases, even a true belief based on excellent reasons may fall short of knowledge.

> (2) Perhaps we need not a true belief that has a good rational basis but rather one that is sensitively true.

He suggests the sensitivity condition:

> *Sensitivity Condition*: In order for a true belief to be a case of knowledge, it must be *sensitive*: it must be one that the believer *would not* hold if its content were false rather than true.
>
> This requirement is violated by the true believers in our four examples above. All four of those beliefs are insensitive. The gambler would (or anyhow might well) still believe that the dice would come up 7 or 11 even if it were false that they would do so. The clock reader would (or might well) still believe that it was noon (from reading the stopped clock) even if he were reading the clock at a time other than noon (say at 11:50). You would (or might well) still believe that someone owned a Ford, even if no one did (say because not even Havit owned one or because Havit was not here after all). And you would (or might well) still believe that the structure you saw was a real barn even when that was not so (and you saw a *mere* barn façade).

---

[82] According to the PE theory, S has a personally unjustified true belief. Conditions 3 and 4a are violated.

In these situations, **S** would continue to believe **p** even if **p** was false. For example, when **S** reads the (stopped) clock which generates a true belief (it is noon), this belief isn't knowledge since **S** would have this belief even if it was 11:50.[83]

With respect to lottery cases, where **S** expects correctly and justifiably to lose a lottery, on a rational basis; this isn't knowledge because **S**'s belief that his ticket will lose is insensitive: it would still be held by **S** on the same probabilistic basis, even if against all odds, as a matter of contingent fact the ticket would win. He continues:

> (3) We shall find that our sensitivity condition is open to serious objections. Fortunately, its undeniable intuitive attractiveness can be explained through an easily confused but far preferable notion of "safety." The denouement of our reflections will be to show how replacing sensitivity with safety makes it possible to defend plain common sense against the spurious advantages over it claimed by alternative accounts. What follows supports that replacement by showing how the sensitivity requirement runs against serious problems (p. 57).

Sosa then asks that we suppose that sensitivity is required for knowledge and that we *agree* to *epistemic closure*: Any conclusion that is reached through competent deduction from something known to be true is also known to be true.[84]

He then asks us to suppose that we are in a BIV case. In such a material case, **S** would continue to have ordinary beliefs even if they were false (and so **S**'s beliefs are insensitive and not knowledge). He then asks us to suppose that we are not in a BIV case. But the problem with this is that any such belief would still be insensitive. He argues that sensitivity is false by *reductio*. He says:

> (4) Better Safe Than Sensitive: The problems we have found to affect sensitivity do not affect a closely related "safety." A belief is sensitive iff had it been false, **S** would not have held it, whereas a belief is *safe* iff **S** would not hold it without its being true (p. 59).

In chapter five, entitled "Knowledge as Action," Sosa starts with the question: What constitutes a fully apt performance? The initial five paragraphs of the chapter are presented below (pp. 71-73):

---

[83] As an aside, in Sosa's footnote #2, Sosa *correctly* describes what is here recognized as a distinction between 'Gettier cases' (intervening luck) and 'Harman cases' (environmental luck). He references his *A Virtue Epistemology* (2007, Lectures 2 & 5) as providing details.

[84] Sosa asks the reader to *agree* to the epistemic closure principle.

Judgement and knowledge itself are forms of intentional action—that is the thesis to be argued in this chapter. Such action falls under a certain normative structure of success, competence, and aptness or success that manifests competence. Judgment is a special case falling under that structure.[85]

A concept of the fully apt performance figures importantly in my account. In order to introduce this concept, we turn first to practical performances, to how they are constituted and to a special normativity that pertains to them. An example yields initial insight, whereupon we turn to performances that are epistemic rather than practical, on which a further example will also shine its light. These examples show the way to a better view of human knowledge, which is our main objective.[86]

(1) Actions and Performances

(a) Intentional actions come in two sorts. An attempt is an intentional action, an endeavor to attain a certain objective. An attempt can fail and remain a mere attempt, whereas as an achievement is a certain sort of *successful* attempt.[87] Thus, if someone intentionally flips a switch in the endeavor to turn on a light, that is an *attempt*. If the attempt succeeds, then the agent does turn on the light by flipping that switch, which if all goes well, is an *achievement*.

(b) Intentional actions (both mere attempts and also achievements) are one sort of performance. Some performances are also aimings, however, without being intentional. The heart, for example, aims by pumping to circulate an organism's blood, which it does through functional teleology without doing so intentionally.

Such functional aimings also come in two sorts. Those, too, can fail and remain mere aimings, to be distinguished from those that are successful. Footnote: Thus, the expression "intentional attempt" is redundant. Not so "intentional aiming," since an aiming could be just functional or teleological without being intentional. The heart beats, for example, so as to circulate the blood; that is its (teleological) end in so doing.

(2) The AAA Structure

(a) In what follows we focus mainly on (intentional) attempts. These fall under the sort of normative structure that pertains, for example, to an archer's aim to hit

---

[85] This thesis that “judgement, and knowledge itself, are forms of intentional action” seems flat out wrong.

[86] I don't believe that the adoption of new concepts in the form of technical stipulated definitions provides a better view of knowledge.

[87] 'Successful' is a group resemblance concept.

a target. The shot might hit its target, in which case it was accurate. Second, it might be competent: as the arrow leaves the bow, its speed and orientation derived from the archer's skill would, in normal conditions, take it to be a bullseye. A shot can be accurate without being competent and competent without being accurate. But even a shot that is both accurate and competent might still fall short in a further respect. Thus, a gust of wind might take the arrow off course, whereupon a second gust might come along and put it back on course.

That shot is then accurate and competent, but the accuracy is attributable not so much to the archer's competence as to the luck of the second gust. Such a shot then falls short in a further respect. A shot is *apt* iff its success manifests the competence then exercised by the archer. Its *success* must manifest that competence seated in the agent and exercised on that occasion. Real achievement would seem to require not just success but apt success.

(b) Here is the AAA structure, which applies to that shot, and then more generally to attempts. In first approximation:

A shot is *accurate* iff it hits the target.
It is *adroit* iff it is an exercise of competence.
It is *apt* iff it is accurate because adroit.

Generalizing to attempts generally, because these all have constitutive aims, we can also distinguish as follows:

An attempt is *successful* iff it attains its aim.
It is *competent* iff it is an exercise of competence.
It is *apt* iff it is successful because competent.

(3) Epistemology as a Special Case:

Epistemology concerns a domain of epistemic performance, such as a belief and judgment (and also inference, learning, teaching, inquiry, and so on).

He later suggests this 'AAA" model of evaluation is applicable quite generally for the evaluation of any action or object with a characteristic 'aim.' It is especially applicable to 'belief 'with respect to 'truth.'

1) A belief is *accurate* iff it is *true*.
2) A belief is *adroit* iff it is produced *skillfully*.
3) A belief is *apt* iff it is *true* in a way manifesting, or attributable to, the believer's *skill*.

According to Sosa, “knowledge” is apt belief.[88] Knowledge entails both truth (accuracy) and justification (adroitness), but they are not merely independent components out of which knowledge is truth-functionally composed.[89] It requires that the skill explain the success.[90]

In the final chapter (thirteen), entitled "Knowledge and Justification" Sosa concludes with a sketch of the main components of the earlier chapters while providing further historical context (p. 207):

> 1) Knowledge is a form of action, to know is to act, and knowledge is hence subject to a normativity distinctive of action, including intentional action.[91]
>
> 2) We begin with a notion central to Aristotle's virtue ethics, a notion of "apt performance" (for a handy label).

A quotation from Aristotle’s *Nicomachean Ethics* is presented (p. 207):

---

[88] With Sosa's theory, and its expansive vocabulary, a limited number of philosophers are capable of understanding what knowledge is. In his most recent book, *Epistemic Explanations: A Theory of Telic Normativity and What it Explains* (2021) Sosa provides “a substantially improved telic virtue epistemology” and “refigures earlier virtue epistemology” which “now seems a first approximation” (xii).

[89] To summarize, in *Epistemology* (2017) and *A Virtue Epistemology* (2007) Sosa makes use of the skilled archer analogy. He says that there are two ways in which an archer's shot might be evaluated:

> 1) Was the shot successful? Did it hit its target?
>
> 2) Did the shot's execution *manifest* the archer's skill? Was it produced in way that makes it likely to succeed?

The kind of success at issue in (1), Sosa calls *accuracy*. The kind of skill discussed in (2), Sosa calls *adroitness*. A shot is adroit if it is produced skillfully. Adroit shots needn't be accurate, as not all skilled shots succeed. And accurate shots needn't be adroit, as some unskilled shots are lucky. In addition to accuracy and adroitness, Sosa suggests that there is another respect in which a shot may be evaluated, relating the two. This, Sosa calls *aptness*.

> 3) Did the shot's success manifest the archer's skill?

A shot is apt if it is accurate because adroit. Aptness entails, but requires more than, the conjunction of accuracy and adroitness, for a shot might be both successful and skillful without being apt. For example, if a skillful shot is diverted by an unexpected gust of wind, then redirected towards the target by a second lucky gust, its ultimate accuracy does not manifest the skill, but rather reflects the lucky coincidence of the wind.

[90] Understanding knowledge as apt belief 'explains' Gettier examples by this analogy. But Sosa has problems with a 'Harman case' (e.g., Henry & Barn) where S is deprived of his ability in his current environment. His competence does not manifest itself in his belief, since its truth is attributable more to luck than to his skill in recognizing barns. Since S's belief is apt, Sosa (counter-intuitively) accepts that Henry knows that he sees a barn.

[91] This proposition is unintuitive. Knowledge is better viewed as a natural kind concept.

> It is possible to do something that is in accordance with the laws of grammar, either by chance or at the suggestion of another. A man will be a grammarian, then, only when he has both done something grammatical and done it grammatically; and this means doing it in accordance with the grammatical knowledge in himself. (Aristotle, EN II 4, 1105a22-26).

He continues (pp. 207-208):

> In order to be "in accordance with" the grammatical know-how in a speaker, an utterance must attain its grammatical success under the guidance of the speaker' own competence, so that it does not succeed just "by chance."
>
> That introduces the notion of aptness, of success attributable to the competence of the agent, so that it is not just "by chance" (in a distinctive way). Whether grammatical or otherwise, a success will be apt not just because it happens to satisfy some proper standard of quality but rather because of the agent's relevant competence to bring about the satisfaction of that standard. This success is, then, more deeply "in accordance with" the relevant know-how or competence in the agent.
>
> So far aptness has been attributed no normative import. But that changes dramatically in the fuller development of Aristotle's account.
>
> 3) How such a concept of aptness is in Aristotle's ethics may be seen in the following passage: "Human good turns out to be activity of the soul in accordance with virtue, and if there is more than one virtue, in accordance with the best and most complete" (Aristotle EN I 7, 1098a16-17).
>
> Just as the grammaticality of an utterance can be in accordance with the grammatical knowledge or competence of the agent, so the good quality of an action or an activity can be in accordance with the virtue seated in the agent. We are thus told by Aristotle what in his view human good *is*, how it is constituted.[92]

XIV. John Greco

We now turn to a second virtue epistemologist, John Greco. We will examine his essays "Virtues in Epistemology" (2002) and "Virtue, Knowledge, and Achievement" (2019). Greco is also the author of *Achieving Knowledge* (2010).

---

[92] Why defer to Aristotle's prescriptive normative value of what he takes to be "human good," as if there is such a thing? As mentioned, 'good' is a group resemblance concept that is prescriptively attributed to x by S. Assertions of what is good (from a normative perspective) are not truth-apt.

We will outline some of Greco's discussion (2002), with the author's frequent references to Sosa, starting with several of the initial paragraphs (quoted):[93]

> What is a virtue in epistemology? In the broadest sense, a virtue is an excellence of some sort. In epistemology, the relevant kind of excellence will be "intellectual." But then what is an intellectual virtue? Some philosophers have understood intellectual virtues to be broad cognitive abilities or powers. On this view, intellectual virtues are innate faculties or acquired habits that enable a person to arrive at the truth and avoid error in some relevant field. For example, Aristotle defined "intuitive reason" as the ability to grasp first principles and "science" as the ability to demonstrate further truths from these.[94] Some contemporary authors add accurate perception, reliable memory, and various kinds of good reasoning to the list of intellectual virtues. These authors follow Aristotle in the notion that intellectual virtues are cognitive abilities or powers, but they loosen the requirements for what count as such (p. 287).[95]
>
> Other authors have understood intellectual virtues quite differently, however. On their view intellectual virtues are more like personality traits than cognitive abilities or powers. For example, intellectual courage is a trait of mind that allows one to persevere in one's ideas. Intellectual open-mindedness is a trait of mind that allows one to be receptive to the ideas of others. Among these authors, however, there is disagreement about why such personality traits count as virtues. Some think it is because they are truth-conducive, increasing one's chances of arriving at true beliefs while avoiding false beliefs (e.g., Zagzebski, 1996). Others think that such traits are virtues independently of their connection to the truth—they would be virtues even if they were not truth-conducive at all (e.g., Montmarquet, 1993) (p. 287).[96]

---

[93] Greco, John 2002. 'Virtues in Epistemology' in Paul K. Moser (ed.). *The Oxford Handbook of Epistemology*. New York: Oxford University Press.

[94] In a footnote, Greco references Aristotle's *Nicomachean Ethics*, book VI. While Aristotle is an admirable philosopher, his definitions of 'virtue,' 'intuitive reason' and 'science' in terms of first principles and demonstration seems to be outdated. If philosophy is going to progress as a social science, some of its history must be studied just as such. That there are intellectual activities thought of as 'virtuous' by Aristotle isn't evidence (or a reason) to support a 'virtue epistemology.'

[95] The whole ontology of the existence of 'virtues' is misguided. Defining a virtue as an excellence of some sort and then trying to make the notion more precise is nothing more than an *ad hoc* artificial stipulation. What are considered as 'virtues' in ordinary language are *group resemblance* items (generosity, truthfulness, loyalty). That 'moral virtue' is primary to 'epistemic justification' is questionable.

[96] This debate about the nature of intellectual virtues is artificial and 'academic' in the highest sense. There is no 'nature' of intellectual values. Intellectual virtues are group resemblance and are prescribed as values.

Who is right about the nature of the intellectual virtues? One might think that this is a matter of semantics— that different authors have simply decided to use the term "intellectual virtue" in different ways. In what follows I will argue that there is some truth to this analysis (p. 288).[97]

However, it is not the whole truth. This is because epistemologists invoke the notion of an intellectual virtue for specific reasons, in the context of addressing specific problems in epistemology. In effect, they make claims that understanding the intellectual virtues in a certain way allows us to solve those problems (p. 288).

The intellectual virtues made their contemporary debut in a series of papers that were reprinted by Ernest Sosa (1991). In these papers Sosa is primarily concerned with two problems in the theory of knowledge. The first is the debate between 'foundationalism' and 'coherentism.' The second is a series of objections that have been raised against reliabilism (p. 288).

Foundationalism and coherentism are positions regarding the structure of knowledge. According to foundationalism, knowledge is like a pyramid: a solid foundation of knowledge grounds the entire structure, providing the support required by knowledge at the higher levels. According to coherentism, knowledge is like a raft: different parts of the structure are tied together via relations of mutual support, with no part of the whole playing a more foundational role than do others (p. 288).[98]

Let us use the term "epistemic justification" to name whatever property it is that turns mere true belief into knowledge.[99] We may then define "pure coherentism" as holding that only coherence contributes to epistemic justification, and we may define "pure foundationalism" as holding that coherence does not contribute to epistemic justification at all. In the papers that introduce the notion of intellectual virtue, Sosa argues that neither pure coherentism nor pure foundationalism can be right… (pp. 288-289). This is where the notion of an intellectual virtue is useful, according to Sosa.[100]

---

[97] In deciding to use the term "intellectual virtue" in different ways, these authors are prescribing different stipulative technical formalizations (3c definitions) in order to fit their overall (purportedly explanatory) theories (or models).

[98] In chapter 3 it is argued that both foundationalism and coherentism are false and that a contextualist theory of personal justification is true. Sosa is debating the merits of two false positions. Also, there will be objections to reliabilism because it seems that any *purely externalist* epistemic theory is *false*.

[99] Like most philosophers, Greco is cavalier in introducing a new stipulative definition with terms that have some prior meaning (or use). Using Greco's stipulative definition of "epistemic justification," the satisfaction of PE conditions 3, 4a, and 4b are the properties (or conditions) added to true belief to constitute knowledge.

[100] Greco turns his attention to the writings of Sosa in the following several paragraphs.

Sosa argues that virtues in general are excellences of some kind: more specifically, they are innate or acquired dispositions to achieve some end. Intellectual virtues, Sosa argues, will be dispositions to achieve the intellectual ends of grasping truths and avoiding falsehoods. This notion of an intellectual virtue can be used to give a general account of epistemic justification as follows:

Definition of 'epistemic justification': A belief **B(p)** is epistemically justified for **S** (i.e., justified in the sense required for knowledge) if and only if **B(p)** is produced by one or more intellectual virtues of **S**.

This account of justification, Sosa argues, allows us to explain the unifying ground of the foundationalist's epistemic principles regarding perceptual beliefs. Specifically, such principles describe various intellectually virtuous dispositions. Thus, human beings are gifted with perceptual powers or abilities; that is, dispositions to reliably form beliefs about the environment on the basis of sensory inputs of various modalities. Such abilities are relative to circumstances and environment, but they are abilities, nonetheless. The foundationalist's epistemic principles relating perceptual beliefs to their experiential grounds can now be understood as describing or explicating these various abilities… It is possible to give similar accounts of other sources of justification traditionally recognized by foundationalism. Because they are reliable, such faculties as memory, introspection, and logical intuition account as intellectual virtues and therefore give rise to epistemic justification to their respective products. In similar fashion, various kinds of deductive and inductive reasoning reliably take one from true belief, and hence count as virtues in their own right. By defining epistemic justification in terms of intellectual virtue, Sosa argues, we get a unified account of all of the sources of justification traditionally recognized by foundationalism.

Once the foundationalist makes this move, however, pure foundationalism becomes untenable… In our world, in normal circumstances, coherence-seeking reason is also a reliable source of true belief and hence a source of epistemic justification (p. 290).

Finally, Sosa argues, we are now in a position to recognize two kinds of knowledge. First, there is "animal knowledge," enjoyed by any being whose true beliefs are produced by intellectual virtue. But second there is "reflective knowledge" which further requires a coherent perspective on one's beliefs and their sources in intellectual virtue. We may also label the latter kind of knowledge "human knowledge," recognizing that the relevant sort of reflective coherence is a distinctively human virtue. More exactly:

> S has **"animal knowledge"** regarding **p** only if
> (1) **p** is true, and
>
> (2) **S**'s belief **B(p)** is produced by one or more intellectual virtues of **S**.
>
> S has **"reflective knowledge"** regarding **p** only if
> (1) **p** is true,
>
> (2) **S**'s belief **B(p)** is produced by one or more intellectual virtues of **S**, and
>
> (3) **S** has a true perspective on **B(p)** as being produced by one or more intellectual virtues, where such perspective is itself produced by an intellectual virtue of **S**.

Greco cites Sosa's (1991) "Knowledge and Intellectual Virtue" and "Intellectual Virtue in Perspective" (chapters 13 and 16) as a source of the above definitions (p. 291).[101]

At pp. 310-311, Greco says that in Gettier cases, **S** believes a true **p**, but it is only by accident. This is opposed to cases of knowledge, where it is to **S**'s credit that she believes the truth, because she does so as the result of her own cognitive abilities:

> However, this diagnosis led to the following question: Why is it appropriate to credit **S** with true belief in cases of knowledge, but not in the two Gettier cases (Lehrer's Nogot, Chisholm's Sheep) given that in these two cases **S**'s abilities are part of the causal story regarding how **S** came to have a true belief?

Greco's response:

> In cases of knowledge, but not in Gettier cases, **S**'s abilities are a salient part of the causal story of how **S** came to have a true belief. It is plausible, in fact, that our cognitive abilities have a kind of "default" salience, owing to our interests and purposes in being information-sharing beings. In Gettier cases, however, this default salience is trumped by something abnormal in the case. For example, someone in the office owns a Ford, but it is not the person who **S** thinks it is. There is a sheep in the field, but it is not in the place that **S** is looking. In these cases, it is only good luck that **S** ends up with a true belief which is to say **S**'s believing the truth cannot be attributed to her abilities.

Greco says that these considerations suggest the following account of knowledge.

---

[101] Greco and Sosa's worldview intuitions about knowledge strike me as too metaphysical. Again, the methodology of analogy, metaphor, and stipulated existents isn't persuasive. However, this methodology isn't unique. Much of philosophy since the ancient Greeks has incorporated metaphysical theories with these features. Philosophy is part of the humanities as conceived by Sosa. I favor moving the discipline into the social sciences with conceptual analysis.

> **S** has '**knowledge**' regarding **p** if and only if
>
> 1) **S**'s belief **B(p)** is *subjectively* justified in the following sense: **B(p)** is produced by cognitive dispositions that **S** manifests when **S** is motivated to believe what is true.
>
> 2) **S**'s belief **B(p)** is *objectively* justified in the following sense: **B(p)** is produced by one or more intellectual virtues of **S**, i.e., by one or more of **S**'s cognitive abilities or powers, and
>
> 3) **S** believes the truth regarding **p** *because* **S** believes **p** out of intellectual virtue. Alternatively: The intellectual virtues that result in **S**'s believing the truth regarding **p** are an important necessary part of the total set of causal factors that give rise to **S**'s believing the truth regarding **p**.
>
> If we stipulate that 'intellectual virtues' involve a motivation to believe the truth, we may collapse the above account as follows.
>
> **S** has '**knowledge**' regarding **p** if and only if **S** believes the truth regarding **p** *because* **S** believes **p** out of intellectual virtue.

Greco notes that a number of philosophers have defended the idea that in cases of knowledge **S** believes the truth *because* **S** believes out of intellectual virtue, including the aforementioned Sosa and Linda Zagzebski (1996).

"Virtue, Knowledge, and Achievement" (2019)[102]

In "Virtue, Knowledge, and Achievement," Greco provides a concise overview of his virtue epistemology (p. 273):

> A number of philosophers have defended the claim that knowledge is a kind of achievement. The central idea is that, in cases of knowledge, the knower's getting things right can be attributed to her own doing.
>
> More exactly, her getting things right can be attributed to her own *competent* doing. Another way to put the general idea, then, is that knowledge is a kind of success from competence, or success from ability. Suppose that we think that intellectual virtues are a kind of intellectual excellence, ability or competence.

---

[102] Greco, John (2019) 'Virtue, Knowledge, and Achievement' in Heather Battaly (ed.). *The Routledge Handbook of Virtue Epistemology*. New York: Routledge.

> Then another way to put the general idea is that knowledge is a kind of success from virtue.[103]

In the first footnote, Greco states that:

> Intellectual abilities or competencies are here understood as abilities to reliably get things right, relative to some field or subject matter, and under appropriate conditions. For example, visual perception is an *ability* to form true beliefs about various features of middle-sized objects (e.g., color, size), under appropriate lighting conditions, with an unobstructed view, etc.[104] A number of authors have argued that knowledge is a kind of achievement including Sosa (2007), Greco (2010), and Zagzebski (1996).[105]

He continues:

> The idea that knowledge is a kind of achievement has considerable explanatory power regarding the nature, value and scope of knowledge. Regarding the nature of knowledge, the achievement account yields the following diagnosis of standard Gettier cases. In cases of knowledge, **S**'s arriving at the truth is due to her own competent cognition (p. 273). More exactly:
>
> (1) **S** has a true belief,
> (2) **S**'s belief is formed by cognitive ability and
> (3) **S** has a true belief *because* her belief is formed by cognitive ability.
>
> In Gettier cases, (1) **S** has a true belief, (2) **S**'s belief is formed by cognitive ability, but (3) **S** does not have a true belief *because* her belief is formed by cognitive ability. Rather, **S**'s forming a true belief is merely lucky (Greco 2003, 2010). (p. 273).

Greco distinguishes between (1) a case of the normal perception of a sheep in the field, as contrasted to (2) a misperception where **S** sees a dog in the field but mistakes it as a sheep

---

[103] The concepts of competence, success, ability, excellence, and virtue are all group resemblance concepts. The idea that "knowledge is a kind of success from virtue" is a metaphysical thesis. It is not based upon a linguistic conceptual analysis written in terms of a social science. Metaphysical theories based on interpretations of group resemblance concepts or their stipulations, can be rejected based upon their unreliable speculative methodology.

[104] Is visual perception an ability? This worldview about perception as a 'kind of ability' (and not merely as one of the five human sensory mechanisms) fits into Greco's worldview that embraces an epistemic externalist (virtue) reliabilism.

[105] This central worldview can be disputed. Arguments involving group resemblance concepts, analogies, metaphor, and stipulations aren't convincing.

(by an unusual trick of light), but there is a sheep in the field obstructed from **S**'s view. His view is paraphrased:

> In cases of successful perception, **S**'s true perceptual belief is attributable to his excellent perception; that is, **S** has a true belief because **S** exercised perception.
>
> In the Gettier case, **S** exercises excellent perception and **S** ends up with a true belief, but **S** does not end up with a true belief because **S** has exercised excellent perception. On the contrary, it is just good luck that there is a sheep in another part of the field, unseen and unknown to **S**. Suppose we think of *achievements* as successes that are attributable to competent agency, as opposed to mere lucky successes that are not so attributable.[106] Normal perceptual knowledge describes a cognitive achievement, whereas a Gettier case describes a mere lucky success.[107]

At page 274 Greco says this (abridged):

> The idea that knowledge is a kind of achievement also yields an *elegant* explanation of the *value* of knowledge.[108] That is because, in general, we think that achievements are both intrinsically and finally valuable. That is, we think achievements are both valuable "in themselves" and "for their own sake."[109] By understanding knowledge as a kind of achievement, we can explain the value of knowledge in terms of the value of achievements more generally. In the same way, the account elegantly explains the superior value of knowledge over mere true belief, in terms of the superior value of achievements over mere lucky successes (Greco 2010, chapter 6, italics added).

---

[106] We have the right to reject Greco's stipulative definitions based upon vague group resemblance definiendums and a definiens containing additional group resemblance terms. We are under no obligation to participate in this metaphysical model.

[107] Normal perceptual knowledge describes a cognitive achievement, whereas a Gettier case describes a mere lucky success? Again, we are well aware that in Gettier cases, that S's true belief was formed as a matter of luck. To describe knowledge as an achievement and non-knowledge is due to luck seems to state the vague and obvious.

[108] As an aside, some examples of 'reportive definitions' from a standard dictionary:

> "Elegant" (adjective)-
> 1) pleasingly graceful and stylish in appearance or manner.
> 2) (of a scientific theory or solution to a problem) pleasingly ingenious and simple.
>
> "Ingenious" (adjective)-
> 1) (of a person) clever, original, and inventive.
> 2) (of a machine or idea) cleverly and originally devised and well suited for its purpose. "Ingenious devices."

[109] I have never had the thought that, "achievements are both intrinsically and finally valuable." This is a moral realist metaethical (metaphysical) theory.

The forgoing considerations regarding the nature and value of knowledge also have implications for the *scope* of knowledge. In particular, the idea that knowledge is a kind of achievement provides traction for responding to a variety of skeptical arguments.

For example, a number of skeptical arguments trade on the idea that our perceptual abilities would not be able to discriminate between real-world scenarios and various skeptical "dream" scenarios. An implicit assumption of such skeptical reasoning is this: my perceptual abilities yield knowledge only if they could discriminate between real-world scenarios and skeptical dream scenarios.

But the idea that knowledge is a kind of achievement—a kind of success from ability—gives us traction against this sort of reasoning. For reflection on the *nature* of achievements in general reveals that achievement-grounding abilities need not be infallible, and need not be reliable in unusual or atypical environments.[110]

Knowledge requires success from ability, and perceptual knowledge requires success from perceptual ability.[111] Presumably, we are reliable perceivers in normal perceptual circumstances, and therefore often have perceptual knowledge.[112] This does not imply that we never make perceptual mistakes. Neither does it imply that we would be reliably successful if we were in very different perceptual circumstances—if we were disembodied victims of a Cartesian demon, or prisoners in the Matrix, for example. Therefore, the implicit assumption identified in the skeptical reasoning above—that my perceptual abilities yield knowledge only if they can discriminate between real-world scenarios and skeptical dream scenarios—is false, and reflection on the nature of achievement more generally gives us grounds for saying so (Greco, 2010).

These anti-skeptical considerations are directly related to the idea that knowledge cannot be merely lucky. On one hand, skeptical arguments exploit the idea that knowledge is inconsistent with lucky true belief. One way to understand the skeptical reasoning above is exactly in these terms: If our perceptual abilities cannot discriminate between real-world perceptual scenarios and mere dreams, then any true perceptual beliefs we have are merely lucky, as opposed to due to our own perceptual abilities. The idea that knowledge is a kind of achievement challenges that skeptical reasoning by deepening our thinking about the

---

[110] 'Nature' is italicized. Does 'achievement' have an objective nature (as a natural kind concept) or is this a group resemblance concept?

[111] This metaphysical assertion isn't truth-apt. Or if it is, it is only within the metaphysical system that it is a part of.

[112] This sentence is the presumption that the radical skeptic (with argument) outwardly disputes.

> relationship between knowledge and luck, and once again, it does so by considering the relationships among luck, ability, and achievement more generally (p. 275).
>
> Finally, the idea that knowledge is a kind of achievement yields insight into the nature of epistemic normativity and epistemic evaluation, but understanding these within the context of a broader normative domain. That is, in any domain of human performance that allows for success and failure, we make a distinction between success due to competent agency and success that is merely lucky. The present account exploits this familiar distinction to understand epistemic normativity as simply a species of performance normativity more generally (Greco 2010, chapter 1, and Sosa, 2015). In fact, this final point is not unrelated to the considerations above regarding the value of knowledge, the superiority of knowledge over mere true belief, and the relationship between knowledge and luck. In all of these contexts, we get insight into the nature of epistemic normativity and epistemic evaluation by considering the contours of performance normativity more generally.[113]

XV. Linda Zagzebski

A third advocate of 'virtue epistemology' is Linda Zagzebski. We will follow her "The Inescapability of Gettier Problems" (1994) that is republished in *Epistemology: An Anthology* (Sosa, et al 2008, 2nd ed). We will then examine "What is Knowledge?" (1999) published in *The Blackwell Guide to Epistemology* (Greco & Sosa, eds.). Finally, we will briefly comment on her monograph *Virtues of the Mind* (1996).

"The Inescapability of Gettier Problems" (1994)

Zagzebski says that Gettier problems arise in the theory of knowledge when it is only by *chance* that a justified true belief is true. Since the belief might *easily* have been false in these cases, it is normally concluded that they are not instances of knowledge.[114] The moral drawn in the thirty years since Gettier published his paper is that either:

---

[113] The last three paragraphs are certainly quite metaphysical in nature. I have no interest in considering the relationships among luck, ability, and achievement.

[114] Note that 'easily' is a group resemblance concept. With the Professor Brown case it was argued that there can be Gettier cases where S's true belief (Brown is room 222) could not have easily been false (e.g., only with a 'perfect storm' of factors; cancellations, airport delay, sickness); but is true upon another unlikely fact (Brown's walk to 222). The 'safety principle' (which includes 'easily') falls to this counterexample.

1) JTB is not sufficient for knowledge, in which an extra component is needed[115]
or
2) 'Justification' must be reconceived as to *make it* sufficient for knowledge.

Assuming that "the relation between justification and truth is close but not inviolable," Zagzebski argues that *neither* of these options will avoid Gettier counterexamples. She concludes that Gettier problems are inescapable for every analysis of knowledge.

She says that Gettier problems arise for both internalist and externalist notions of justification (p. 207, paraphrased):

> On internalist notions, **S**'s reasons for believing **p** are accessible to **S** and there is nothing wrong with the cognitively accessible aspects, but there is a mishap in something inaccessible to **S**. Since (personal) justification does not guarantee truth, it is possible for there to be a break in the connection between justification and truth, and for that connection to be regained by chance.[116]

Zagzebski recites an original Gettier case:

> **Brown in Barcelona**
> **S** has strong evidence for a proposition that he doesn't realize is false, namely **q**: 'Jones owns a Ford.' **S** selects a place name, 'Barcelona,' at random and constructs proposition **p**: 'Either Jones owns a Ford or **b**: Brown is in Barcelona.' Not having any idea of Brown's actual whereabouts, **S** accepts **p** on the basis of **q**. As it happens, Brown is coincidentally in Barcelona (**b** is true), so **p** is true.[117]

This can be summarized as follows:

q = Jones owns a Ford.
b = Brown is in Barcelona.
p = **q** or **b**.

(1) **S** believes that **q**.

(2) **S** believes the disjunction **p**: '**q** or (personally unjustified belief) **b**.'
(3**) b** is true, but **q** is false.

---

[115] The PE definition maintains that knowledge is JTB plus two components: PE 3 & PE 4b.

[116] This is a good paragraph.

[117] This original 'Gettier example' is very strange. It isn't typical for S to reason by disjunctive conjunction by adding an arbitrary conjunct to the proposition believed. But with this valid reasoning, S's reasons for believing a true p are irrelevant for why p is true. The truth of p is coincidental to S's evidence.

(4) **S** believes **p** on belief of **q**.

(5) **p** is true; but having a true belief ('Either Jones owns a Ford or **b**: Brown is in Barcelona') is the result of the good luck because **b** is coincidentally true. A true belief is generated by a coincidental (accidental) feature of the situation. The belief that **p** isn't based upon **S**'s cognitive ability.

Zagzebski says that what generates the problem for JTB is that an accident of bad luck (Jones false bragging about his Ford ownership) is cancelled out by an accident of good luck (Brown being in Barcelona).[118] The right goal is reached but only by chance. This is the Gettier problem for internalists.[119]

The problem also arises for reliabilism. **S** is justified when the belief is formed in a reliable truth-conducive manner. But there is no guarantee that justified beliefs are true, and a breakdown in the connection between a reliable belief-forming process and the truth is possible. *When* that happens, even if **p** is true, you don't have knowledge.

She says that the 'fake barn' case is an example of this sort (p. 208):

Henry & Barn: It is only an accident that visual facilities normally reliable in this sort of situation are not reliable in this particular situation and it is another accident that you are looking at a barn. Accident of bad luck is cancelled out by good luck.[120]

Zagzebski criticizes Alvin Plantinga. Plantinga (1993) asserts that the 'property' that in 'sufficient quantity' that converts true belief into knowledge is "warrant," rather than "justification." For Plantinga "warrant" is the property a belief has for **S** when the belief is produced in **S** by **S**'s faculties working appropriately in the appropriate environment, according to a design plan successfully aimed at the truth. Plantinga considers warrant

---

[118] Whether the concepts/terms of 'luck,' 'cognitive ability,' and 'chance' should be introduced into an explanation of knowledge is questionable.

[119] The PE theory simply explains a Gettier case as a PE condition 3 and 4b failure. The PE theory doesn't try to construct a theory of knowledge that is immune to skepticism. It is agreed that no necessary and sufficient conditions approach will present a set of conditions *where it is guaranteed* for when knowledge is attained. Instead, although knowledge can never be guaranteed (since one may be a BIV), its conditions of attainment can be described (as contingent) by the PE definition.

[120] On the contrary to Zagzebski, S's visual facilities *are reliable* in this situation. A better explanation is that it is the hostility of the environment (99 visually indistinguishable barn facades) that leads to the coincidental 'luck' of seeing a barn. With the PE definition, S does not know p, because if had S known that he was in the land of barn facades, he would acknowledge that he doesn't/didn't know that x is a barn.

and reliability as properties that admit of degree, but the degree of warrant sufficient for knowledge does not require faculties to be working perfectly in an environment perfectly matched to them.[121]

Zagzebski argues that as long as the property (e.g., warrant) that putatively converts true belief into knowledge is analyzed in such a way that it is strongly linked with the truth but does not guarantee it, it will always be possible to devise cases in which the link between such a property and the truth is broken but regained by accident. This is the nature of Gettier cases and JTB will never be sufficient for knowledge. Zagzebski goes on to say that we can construct Gettier cases with the following procedure:

> 1) Start with a case of justified (or warranted) false belief.
>
> 2) Make the personal justification (or warrant) strong enough for knowledge; but make the belief false.
>
> 3) The falsity of the belief is due to some element of luck.
>
> 4) Now emend the case by adding another element of luck. The situation might be described as one element of luck counteracting another. JTB will never be sufficient for knowledge.
>
> 5) No account of knowledge as true belief plus other condition(s) can *withstand* Gettier counterexamples. Gettier cases will never go away.[122]

---

[121] Alvin Plantinga's (1993) two-volume set *Warrant: The Current Debate* and *Warrant and Proper Function* is concerned with the concept of "warrant." Plantinga is a philosopher at the University of Notre Dame and explicitly is a theist. He argues for his worldview in *Knowledge and Christian Belief* (2015). Plantinga's metaphysical views are shared in *The Nature of Necessity* (1974). My thoughts about the concept of warrant: Intuitively, if we were asked, what is 'warrant,' we might think that it is about whether people are 'warranted' in believing p or acting in a certain way. This is the sense that Plantiga has in mind. From the PE perspective, the concept of *'personal justification'* pretty much covers the concept of *'warrant'* (and is nearly synonymous). Typically, whenever we say that '**S** was *warranted* in believing p' (or performing action a), we are apparently indicating that **S** was *personally justified* in believing **p** *given* his stated evidence (or grounds), even if some of that evidence was false or misleading. Also, warrant has a normative component, because in certain situations where a belief (or action) is pragmatically important, **S** must have *strong warrant* (or personal justification) for the belief or action. In short, when philosophers talk about 'warrant' (as opposed to simply 'personal justification'), they are giving a nod to a dubious concept that is both stipulative and metaphysically murky.

[122] The historical attempt to specify the conditions required for **S** to know **p** and *guarantee* that **S** knows **p** was a mistake. A necessary and sufficient conditions approach to knowledge doesn't need to be an attempt to formulate a Gettier-proof analysis. The PE definition describes the (natural kind) material conditions for *when* **S** knows **p** (and when **S** doesn't know **p**) and explains *why*. A definition shouldn't be an attempt at evading (or resisting) Gettier situations. Zagzebski rightly concludes that Gettier cases will never go away.

"What is Knowledge?" (1999)

The first paragraph of Zagzebski's "What is Knowledge?" is paraphrased:

> Knowledge is a highly valued state in which **S** is *in cognitive contact* with reality. It is, therefore, a relation.[123] On one side of the relation is a conscious **S**, and on the other side is a portion of reality to which **S** the knower is directly or indirectly related. While directness is a matter of degree, it is convenient to think of *knowledge of things* as a *direct form of knowledge* in comparison to which *knowledge about things* is *indirect*. The former has often been called *knowledge by acquaintance* since the **S** is in experiential contact with the portion of the reality known, whereas the latter is *propositional knowledge* since what the **S** knows is a true proposition about the world.
>
> Knowing **S1** (i.e., another person) is an example of knowledge by acquaintance, while knowing that **S1** is a philosopher is an example of propositional knowledge. 'Knowledge by acquaintance' includes not only knowledge of persons and things, but also **S**'s knowledge of his own mental states. In fact, the knower's own mental states are often thought to be the most directly knowable portion of reality.[124]

She later writes (p. 94):

> From what has been said so far it follows that knowledge is a form of believing a true proposition. At this point the process of defining knowledge it becomes much more difficult and more open to debate. All parties agree that knowledge is a good state, good as in the sense of desirable, and perhaps also in the sense of praiseworthy. But there are different kinds of praiseworthiness. Good looks, wit, and physical strength are desirable qualities, and we praise others for having them, but we typically do not blame them when they lack such qualities. In contrast, we praise persons for having qualities like courage, kindness, or fairness, and we also blame them for their absence. This suggests that it is a requirement of the moral sense of the praiseworthy that it is a quality whose presence is praised and whose absence is blamed. But this is only roughly right since blame for absence is also missing at the high end of moral praiseworthiness. We praise persons for being noble or saintly, but we do not blame them when they are not.

---

[123] This metaphysical assertion can be disputed. Instead of thinking of knowledge as a *relation*, why not think of 'knowledge' as a natural kind, where knowledge is possessed by (i.e., a *property* of) **S** (i.e., that of being in an *epistemic position* where all four PE conditions satisfied).

[124] This worldview is distorted. A distinction between 'knowledge by acquaintance' to 'propositional knowledge' should be abandoned. The foundationalist sensory evolution of propositional knowledge as explaining the structure of knowledge is dubious.

Further textual excerpts:

> It is significant that knowledge has not traditionally been treated as a moral concept, yet it has had many of the trappings of the moral—for example, the connection with epistemic duty and responsibility, as when we criticize a person by saying that she *ought* to know better, a criticism that is often accompanied by the type of distaste characteristic of the moral (p. 94).

> Major disputes over the definition of knowledge may turn on contrasting senses in which knowledge is good. According to the contemporary theory of reliabilism knowledge is true belief arising from a reliable truth-producing mechanism. This proposal makes the good of knowledge a natural good like that of beauty, wit, or strength. The traditional proposal that knowledge is true belief based upon good reasons is associated with the ethical concepts of responsibility, praise, and blame. One is praised for believing the truth upon good reasons and blamed for not doing so. The idea that knowledge is noble comes from Plato (*Protagora*s, 352d, 345b). In my judgment no definition of knowledge can succeed if it does not incorporate or at least adjudicate the senses of good used in these opposing types of theory (p. 95). It is common for philosophers to aim for a real definition of knowledge….it is an aim that deserves attention because it presupposes some disputable semantical and metaphysical views. Perhaps knowledge is not an ontological category for which real definition is possible. For example, no one would attempt a real definition of *rich*, *candy*, or *large plant*, and only some theorists would attempt a real definition of *food*, *intelligence*, or *virtue* (p. 96).

She notes the recent rebellion against the 'necessary and sufficient conditions' approach to knowledge, notably from contextualist theories, and authors such as Edward Craig and Hilary Kornblith. She says (p. 97):

> It is particularly desirable to question whether we should aim for a real definition since it is hard to determine whether knowledge is a single kind of thing for which a real definition is possible. Epistemologists almost always have the aim Plato had in the *Theaetetus*, where he says he is setting out to "bring the many sorts of knowledge under one definition" (148e). But do we know this aim is attainable?

> I believe we should begin by assuming that there is a single concept of knowledge about which philosophers have been debating for millennia and that *we should aim* for a *necessary truth* in our definition until forced to give up by continual failure in reaching the goal. I am less confident that knowledge is a single natural kind on par with *water* or *gold*, but it is tempting to hope that it is the case. In any event, if knowledge is not a natural kind it is unlikely that we will discover that unless we attempt to treat it as one. I will therefore tentatively accept the traditional aim of aspiring to a *real definition* of knowledge (p. 98, italics added).

Zagzebski recites some common criteria for a good definition (p. 98):

> 1) It shouldn't be *ad hoc*.
> 2) It shouldn't be negative when it can be positive.
> 3) It should be brief.
> 4) It shouldn't be circular.
> 5) It should use concepts that are less obscure than the concept (term) to be defined.
> 6) A definition is supposed to tell us something we didn't know.

In an earlier footnote (#8) she mentions Richard Robinson's *Definition* (1950).[125] These six criteria are found there. These criteria are often recited in other brief philosophical discussions of 'definition.' She later says (p. 99, italics added):

> If we want a definition to connect the concept to be defined with other key concepts in well-developed philosophical theories, the concepts that have a central place in such theories might turn out to be normative ones. And since *knowledge* is a *normative concept*, that is just what we would expect.[126]

Let us develop an enhanced (i.e., slightly edited) textual summary of Zagzebski's writing (below) where it is *integrated* with *elements* of the *PE theory*:

> So far, we have concluded that knowledge is *good true belief*. But this serves neither as a theoretical definition nor has a practical purpose. The concept of 'good' is at least in need of conceptual analysis as knowledge.[127]
>
> Since believing is something a person does, beliefs have customarily treated as analogous to acts, so beliefs are good in the sense in which acts are right. Right believing has traditionally been identified as JTB and knowledge is defined as thus. Sometimes, but not always, this has been understood as *true belief for the right reasons* (p. 100).[128]

---

[125] The scholarship on the nature of definition is surprisingly thin. Robinson's monograph is the only single volume dedicated to the topic that I'm aware of.

[126] I deny that knowledge should be called a 'normative concept.' Because there is a normative condition (4a) in the PE definition does this make 'knowledge' a normative concept? What exactly is a 'normative concept'? Is a 'normative concept' identified by group resemblance (based upon use) or is it fixed by stipulation? I don't understand what a 'normative concept' is.

[127] It is argued in Chapter 5 that the concept of 'good' (or 'virtuous') is group resemblance and that it is used in sentences to assert a prescriptive proposition.

[128] Sometimes 'right believing' is called a personally justified belief, i.e., a belief that is personally justified given S's available evidence. It is a prescriptive assessment that S isn't at fault for believing p on the basis of evidence e. Other times, 'right believing' is about S's possessing relevant truth-connecting reasons.

Gettier's examples are cases in which a belief is true and justified; but is not an instance of knowledge because it is only by *chance* that the belief is true… The structure of the Brown in Barcelona case study shows that this is a case of double luck. By case of double luck, you end up with a true belief (p. 100)

Gettier problems arise for any definition in which knowledge is true belief plus something else that is closely connected to the truth but does not entail it. It does not matter if the something else is a matter of believing for the right reasons or even if it is captured by the concept of justification. It need not even be anything accessible to the consciousness of the believer; for example, it may simply specify that the belief is produced by a reliable process or properly functioning faculties. All that is necessary is that there be a small gap between truth and the component of knowledge in addition to true belief in the definition (p. 101).

Defeasibility is discussed: Consider that there is a fact (unknown to **S**) which would defeat **S**'s personal justification, if **S** were to be aware of it. To be indefeasible, there are no truths when added to **S**'s evidential premises, that would lead (or cause) **S**'s strong belief to be doubted or diminished (PE #2 violation) or no longer possess evidence that is 'adequate,' (PE #4a violation). *But in the nature of induction*: The strongest of inductive evidence does not entail the conclusion. It is possible to make a mistake. Zagzebski provides an example where a physician correctly diagnoses a patient for a certain virus, but where the physician's accessible evidence isn't the reason for why the patient has the virus. **S** can believe **p** (which is true) and be personally justified and undefeated by any accessible evidence and not know **p**. Gettier cases are where **S** has strong evidence for **p**, but **p** is true for reasons unrelated to **S**'s evidence (pp. 101-102).

We may conclude that the prevalent method of defining knowledge as true belief plus something else cannot withstand counterexample as long as there is a small degree of independence between truth and that something else. It follows that there must be a necessary connection between truth and the other conditions of knowledge in addition to truth, whatever they may be. In the first section we saw that those other conditions can loosely be defined as believing in a good way. So the sense in which knowledge is believing in a good way must entail truth (p. 102). [129]

---

[129] Zagzebski says that knowledge as true belief plus something else (conditions 3, 4a, 4b) cannot withstand counterexample as long as there is a small degree of independence between truth and that something else. But (again) are we seeking (1) a Gettier-proof definition of knowledge or (2) seeking to explain what is a case of knowledge (and what isn't)? I'm interested in the latter question, not the former.

Since Gettier cases are those in which accidentality or luck is involved, it has often been suggested that knowledge is non-accidentally true belief. This definition is vague and negative, with little practical import (p. 103).

Since we *know* that the *concept of knowledge* is *normative*, it is a theoretical advantage if it can be related to central concepts in ethics since ethicists already have proposed theoretical structures in which these concepts have been analyzed. If it turns out that normative concepts are reducible to or supervene on nonnormative concepts, the demonstration that that is the case would be an independent project. Meanwhile, one of my purposes will be to integrate the concept of knowledge into a background ethical theory (p. 104, italics added).

The moral of Gettier examples: The conclusion is that the normative component of knowledge, the component that makes knowledge good, must entail the truth. Success in reaching the truth must be an intrinsic part of the sense in which instance of knowledge is good (p. 105).

Should we try to embed the concept of knowledge in a background normative theory because it is a normative concept? Should we instead embed it in a background metaphysical theory on the assumption that metaphysics is more basic than epistemology? Or should we embed it in a scientific theory on the grounds that knowledge is a natural phenomenon? I have already said that I will take the first of these alternatives, but I have not argued for it and I can see many advantages in defining knowledge in terms of very different concepts from the one I have chosen. In fact, even if the purpose is to embed the concept of knowledge in a background ethical theory, the choice of theory will obviously depend upon one's position regarding the kind of ethical theory most likely to serve our theoretical and practical purposes (p. 105).

The definition that I will propose arises from a virtue theory of ethics (1996). The complete theory includes intellectual as well as moral virtues within the same theory and aims to give a unified account of the morality of believing as well as acting. The concept of an '*intellectual virtue'* is central and it has a number of theoretical and practical advantages (p. 105).

The definition of knowledge must make success in reaching the truth an intrinsic aspect of that which makes knowledge good. The traditional concept of justification cannot serve this purpose, nor can any concept of a property of a belief. That is because no normative property of a belief guarantees its truth… In Aristotle, the concept of a virtue combines with an admirable internal state with external success… I suggest that it is beneficial to move back a step from properties of beliefs to *properties of persons* in our search for a concept that attaches the good of knowledge to its truth (p. 105).

> Virtues are properties of persons. Intellectual virtues are properties of persons that aim at intellectual goods, most especially truth. Moral virtues are properties of persons that aim at distinctively moral goods such as the well-being of others. Since the concept of virtue already has a rich history; if we can connect knowledge to virtue, that would be a theoretical advantage. In addition, the concept of virtue has practical uses. Ordinary people speak of such individual virtues as kindness, fairness, courage, open-mindedness, perseverance, generosity, discretion, and trust, and sometimes the names are used for both moral and intellectual virtues. Furthermore, the evaluation of acts is often made in terms of the virtues or vices they express…*Virtue* is not a technical concept, although it can be technically refined (p. 106).
>
> A virtue has two components. The first is a motivational component and the second is a component of success in reaching the end of the motivational component. The motivational component of a virtue is a disposition to have an emotion that directs action towards an end. Each virtue has a distinctive motivational component with a distinctive end, but groups of virtues can be categorized by their ultimate ends. Intellectual virtues have truth as their ultimate end (sometimes understanding).
>
> Moral virtues have other ultimate ends. The success component of a virtue is a component of reliability in bringing about the end of a virtuous motivation. The virtues of compassion, trust, and open-mindedness can be roughly defined:
>
> (1) The virtue of compassion is a trait that includes the emotion-disposition to alleviate the suffering of others and reliable success in doing so.
>
> (2) The virtue of trust is the trait that includes the emotion-disposition to trust those and only those who are trustworthy, and reliable success in doing so.
>
> (3) The virtue of open-mindedness is the trait that includes the emotion-disposition to be open to the views of others even when they conflict with one's own and reliable success in doing so.
>
> I suggest that the structure of all, or least, most of the virtues can be defined by this pattern (p. 106).

A technical definition of "an act of virtue" (p. 108):

> An act is *an act of virtue A* if and only if it arises from the motivational component of A, is an act that persons with virtue A characteristically do in the circumstances and is successful in bringing about the end of virtue A because of these features of the act.

A brief definition of 'knowledge' is proposed as well as a more comprehensive one that includes knowledge by acquaintance as well as propositional knowledge (p. 109):

> ‘Knowledge’ is a belief arising out of acts of intellectual virtue.
>
> ‘Knowledge’ is cognitive contact with reality arising out of acts of intellectual virtue.
>
> Knowledge is generally not reached through a single act, but through a combination of acts of one’s own, as well as through the acts of others and cooperating circumstances. We tend to think of knowledge as our own accomplishment, but this is rarely the case. The fact that our knowledge depends upon the knowledge and intellectual virtue of other persons in our intellectual community, as well as a cooperating universe, makes it clear that we cannot expect to isolate the conditions for knowledge in some set of independent properties of the knower, much less a set of properties over which the knower has control. Epistemic luck permeates the human condition for good or for ill.
>
> Gettier problems result from any definition in which the sense in which knowledge is good does not entail truth. The concept of an act of intellectual virtue does entail truth, and so my definition is not guaranteed to fail in the way I have outlined for those theories susceptible to the double luck strategy (p. 111).
>
> In the Brown in Barcelona case study (1994) and in the case of a doctor giving a diagnosis that **S1** has a certain virus, the believer **S** reaches the truth because of a feature of double luck. Despite reaching these respective beliefs because of intellectually virtuous motivations, **S** does not reach the truth because of the features of the situation. This means that reaching A *because of* B is a key element of the definition. We all have intuitions about what it means for something to happen because of something else, but this concept is in need of further analysis, and I do not know of one that is adequate (p. 111).
>
> The definition as I have proposed it here meets many of the criteria for a good definition, but it is vague, and it clearly needs more extensive analysis. We have already seen the need for an account of the *because of* relation in the third component of the definition of an act of virtue. It also needs an account of motivation, as well as an account of acting in a way that is characteristic of a virtue, the first and second components of the definition… There is also the matter of identifying and individuating the intellectual virtues. This is important not only because lists of the individual virtues and their analyses can result in account of knowledge that differ greatly in plausibility, but because it is possible that some of the virtues conflict. Virtue theories of ethics have this same problem… (p. 112).

***Virtues of the Mind* (1996)**

> On pages 9-10, in *Virtues of the Mind* (1996), Zagzebski says this:
>
> Ernest Sosa's (1991) insight that it would be fruitful for epistemology to make the primary object of evaluation intellectual virtues and vices and attach secondary justification to individual beliefs because of their source in intellectual virtues is a significant contribution to the field. In addition, that "we need to consider more carefully the concept of virtue and the distinction between moral and intellectual virtues" (p. 190) strikes me as exactly right and is advice that I am following in this book.

At pages 134-137, a definition of 'virtue' is stipulated by Zagzebski (with comments stated in the italicized parentheses):

> (1) A virtue is an acquired excellence of the soul, or to use more modern terminology, it is an acquired excellence of the person in a deep and lasting sense. A vice is the contrary quality; it is an acquired defect of the soul.
>
> (2) A virtue is acquired by a process that involves a certain amount of time and work on the part of the agent.
>
> (3) A virtue is not simply a skill… virtuous persons are expected to have the correlative skills in order to be effective at action, but skills don't have the intrinsic value of virtues.
>
> (4) A virtue has a component of motivation. A motivation is a disposition to have a certain motive, and a motive is an emotion that initiates and directs action to produce an end with certain desired features.
>
> (5) "Virtue" is a success term. The motivational component of virtue means that it has an end, whether internal or external. A person does not have a virtue unless she is reliable at bringing about the end that is the aim of the motivational component of the virtue. (*A pragmatically formalized definition or a personally formalized definition).*
>
> (6) A virtue therefore has two main elements: a motivational element, and an element of reliable success in bringing about the end (internal or external) of the motivational element.
>
> (7) A virtue can be defined as a deep and enduring acquired excellence of a person, involving a characteristic motivation to produce a certain desired end and reliable success in bringing about that end. *(A pragmatically formalized definition or a personally formalized definition).*

XVI. Conclusion

Contextualism is the thesis that the truth values of knowledge ascriptions (or ascriptions of epistemic justification) are context dependent. We have surveyed the attributor contextualism of Lewis, Cohen, and DeRose. This view contends that it is the context of the speaker (**S1**, the attributor) that determines the truth conditions for a given assertion that "**S** knows that **p**." It is thought that what induces a context change is when speaker **S1** raises the standards for knowledge, by introducing various error-possibilities.

But in critical response to contextualism, is it attributors (or the skeptic) that create the context and 'standards' for knowledge on occasions? The contextualist appeal to shifting conversational contexts that create various 'standards' isn't intuitive. This really doesn't happen. The 'standards' for how much evidence is needed in a context seems to be better understood as self-prescribed by **S** (e.g., how much evidence is adequate, what undermining possibilities are relevant, how pragmatically important is it that **p** is known and not just believed, etc.). The idea that there are 'standards' that can 'govern' contexts of assertion seems completely unintuitive. Moreover, the contextualist belief that epistemic closure (as a rule of inference) preserves the validity of every deductive argument is mistaken. Finally, the intuition that in the context of discussing the BIV possibility we *cannot know* that we have hands is bizarre. This claim that our ordinary knowledge claims are false (in these skeptical contextual situations) isn't intuitive, nor true. In sum, the idea that a solution to the skeptical argument is best explained by contextualized sentences of the form "**S** knows **p**" fails to answer the problem. It requires an elaborate model of (stipulated) possible worlds and the supposition of literal linguistic reference, both of which are questionable. The idea that formal semantics can help resolve epistemological questions clearly fails.

Virtue epistemology focuses on the epistemic evaluation of persons, and their intellectual abilities and character traits. It contends that persons are the primary objects of epistemic evaluation, and that "intellectual virtue" is theoretically more important than justification and knowledge, which are types of belief evaluation. The structure of virtue epistemology is analogous to that of virtue ethics and distinct from that of traditional

belief-based epistemology, which takes knowledge and justification, as types of belief evaluation, to be theoretically fundamental. Whether 'virtue epistemology' (or even 'virtue ethics') has interesting theoretical value is open to serious doubt. A single critical response to virtue epistemology is to question whether virtue epistemology is really informative, or whether it just expresses a systematic moral realist metaphysical view about the value of knowledge and the 'appropriate' dispositions of persons. The worldview intuition that knowledge is a kind of achievement, competence, or success from cognitive ability isn't standard among ordinary folks. This approach to epistemology cannot provide satisfying solutions to problematic case studies that help to reveal the nature of knowledge. Belief evaluation (and *not* personal evaluations) should remain central to a social scientific analysis of the concept of knowledge.