Chapter Two

A Response to Radical Skepticism

Abstract: The radical skeptic provides a two-premise argument that is supposed to show that knowledge is impossible. In this chapter, the argument is scrutinized and one of the skeptic's premises is rejected: the principle of 'epistemic closure.' It is argued that epistemic closure (i.e., *modus ponens* across knowledge states) is not always true, so it cannot serve as the basis for the skeptic's conclusion that knowledge of ordinary propositions is impossible.

In chapter one it was described how knowledge is possible by stating the material conditions that are necessary and sufficient for knowledge. If knowledge exists, then these are its four material conditions. In describing how knowledge is possible, it was accepted that the 'skeptical hypothesis' is true. It is possible that I may have no knowledge at all. For example, I *may* be a brain-in-a-vat, because I can't know that I'm not. The proposition that it is possible that I may have no knowledge at all, is true. According to the skeptical hypothesis, I cannot know that I'm *not* being fed artificial perceptions by a demon giving me a false sense of independent consciousness. Every one of my sensory beliefs may be the product of a systematic elaborate deception. Unlike some theorists who consider it a desideratum of a theory of knowledge to show that this skeptical hypothesis is false, I affirm that the skeptical hypothesis is true, and accept that there *might* be no instances of human knowledge. It is maintained here that I don't know that I'm not a demon-controlled disembodied-brain-in-a-chemical-vat (BIV).

The Standard Skeptical Argument

With a two-premise argument, the skeptic doesn't just claim that we *may* be victims of a systematic deception, but the skeptic 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 BIV.' (Skeptical hypothesis)

(#2) If I do not know that 'I am not a BIV,' then I do not know **o**. (Closure)

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

The first premise states the skeptical hypothesis. The skeptical hypothesis maintains that I cannot know that I am not in a situation where all (or most) of my empirical beliefs are false. Everything that I believe that I perceive may not exist. This premise implies that I cannot verify the truth or falsity of whether or not I am a BIV because any means for testing this proposition are cognitively inaccessible.

In the second premise, it is asserted that if I cannot know that this skeptical hypothesis is false, then I cannot know any ordinary proposition **o**. The skeptic argues if I do not know that I am not a BIV, then I cannot consistently assert that I know **o**, because this unresolved premise that 'I don't know that I'm not a BIV' undermines any claims to knowledge. After all, if you don't even know that you're not a BIV, how is it possible that you can confidently know anything at all?[1]

The standard skeptical argument appears to be valid, and its premises are backed by very plausible intuitions. If the two premises are true, it seems necessary from the *modus ponens* structure of the argument, that the skeptical conclusion is true, and that one cannot know the truth of ordinary empirical propositions. In response, many philosophers believe that premise #1 is false, premise #2 is true, and that there must be a sound argument, theory, or explanation that shows that one can *know* that one is not a BIV, or alternatively that one can know in certain contextual situations. In order to have ordinary knowledge that 'I have two hands' it is believed by these philosophers that I must simultaneously know that I'm not a BIV. This majority position is disputed here.

In contrast, I believe that premise #2 is false, premise #1 is true, and the conclusion is false, so the argument is not sound. The reason premise #2 is false is that it assumes the truth of the following principle of 'epistemic closure':

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

---

[1] Intuitively, it may seem inconsistent to say that 'I know I have two hands,' but 'I don't know that I'm not a brain-in-chemical-vat.' The knowledge of an ordinary proposition **o** seems to imply (or presuppose) (1) the knowledge of one's independent existence, (2) the existence of an external world, and (3) that one isn't a BIV. This is the paradox.

Using several thought-experiment examples I will show that that this closure conditional is false by constructing deductive arguments where this principle is assumed true and where this leads to deductive arguments when all of the premises are assumed true, leads to a conclusion that is clearly false. This is the hallmark of an invalid argument form.

A Deductive Argument Proving the Skeptical Hypothesis is True

Unlike philosophers who react strongly against the skeptical hypothesis, I embrace premise #1 as a consequence of PE conditions 3 and 4a:

> (3) **S** believes **p** upon a set of reasons that are substantially *relevant* (i.e., truth-connecting) for why **p** should be believed, and
>
> (4a) In situations of critical doubt, **S** must possess 'adequate evidence' for believing **p**, and be able to resolve (i.e., rule-out, discard) any actual or logical possibilities that would undermine (or defeat) evidence possessed for **p**.

The reasoning for this requires some hard thinking. With condition 3, it is maintained that the *relevant reasons* for why an empirical **o** should be believed are *truth-connecting, objective* and *independent of persons*.[2] With condition 4a, it is maintained that **S** is justified in believing **p** if **p** is an immediate non-inferential (e.g., perceptual) belief or **S** has reasonably (non-recklessly) acquired strong evidence and used good inferential reasoning for believing **p**. With this conception of relevant evidence and how **S** is personally responsible for holding a belief, I argue that **S** cannot have relevant (truth-connecting) evidence to deny the truth of skeptical premise #1. If **S** cannot possess relevant evidence to know #1 is false, then **S** cannot know that she is not a BIV. The following deductive argument affirms the truth of the skeptical hypothesis #1:

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

---

[2] This sense of 'relevant' evidence is consistent with a dictionary definition: 'relevant' (an adjective) means 'bearing on the matter at hand, germane, material, applicable, pertinent.' For example, suppose that investigators seek the cause of a house fire. Intuitively, there are *relevant reasons* for why a house fire occurs. Investigators want to determine whether an arsonist was involved, or if there was an electrical problem or a lit candle fell, and so on. Investigators seek objectively *relevant evidence* for why the house caught fire and wish to discard any extraneous *irrelevant* facts and reasons not associated with the cause of the fire.

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

(3) As an assumed material condition of the BIV possibility, **S** cannot possess relevant (i.e., pertinent) reasons to believe that the BIV possibility is false, because any objective and independent evidence for testing the falsity of the BIV possibility is cognitively inaccessible. The phenomena of one's experiences are the same, no matter whether one is or isn't a BIV, so **S** cannot have adequate reasons or evidence to deny the physical possibility of being a BIV.

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

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

In this valid deductive argument, premise 1 is a restatement of knowledge condition 3. Premise 2 defines what a 'relevant' reason is.[3] Premise 3 states the nature of the BIV possibility. Premise 4 states that if **S** is unable to possess strong and relevant reasons for believing the BIV possibility is false, then 'I do not know that I am not a BIV.' The truth of the 'skeptical hypothesis' is entailed by the truth of the four premises. Given that premise #1 in the skeptical argument has been proven true, it is implied that skeptical argument premise #2 is false if the radical skeptical conclusion is to remain false (as intuition suggests).[4] We need to explain why the second premise is false.

---

[3] There is another sense of 'relevant' that is *not* being used in PE knowledge condition 3. It is sometimes said that what counts as 'relevant evidence' is any evidence that is 'somehow related' or 'might have some significance or probability' for why **p** is true. Fire investigators at the start of their inquiry will be interested in any *relevant evidence* that *may* have a bearing on why a fire started. Pragmatic interests guide investigators to what items are considered 'relevant.' In this sense, *any* and *all* evidence (e.g. the electrical system, smoking materials, etc.) that *might* have significance as to why a fire started, are items 'relevant' to determining the cause of a fire. This *wide sense* of 'relevant' is not the sense being used here.

[4] The skeptical hypothesis of premise #1 isn't as crazy as it initially seems. It has been noticed that persons as perceptual and information processing mechanisms, have the same self-measurement problem that is analogous to a gas gauge on an automobile. The reading of a gas gauge can accurately and reliably measure the amount of gas in a tank but has no mechanism for measuring its own accuracy. A person is at a similar disadvantage, relative to the possibility of an elaborate hoax created by an evil demon, because a person cannot determine whether one is systematically misreading one's environmental status. David Armstrong, Fred Dretske, and Alvin Goldman have drawn attention to this problem. With current science fiction movies such as the *Matrix*, the possibility of this perceptual problem is better appreciated.

The Distinction between Epistemic Closure and Propositional Closure

Why does premise #2 (that presumes epistemic closure) in the standard skeptical argument seem plausible? In order to know any proposition, it certainly seems that you should know of your own self-existence as an independent entity in a natural external world. Knowing that 'I have hands' implies that I'm not a BIV. If you don't know that you're not a BIV, how can you be sure you know anything? Stewart Cohen (1999), a strong proponent of closure, asks "how could you know **p** and know that **p** entails **q**, and yet fail (at least to be in a position to) know **p**?" (p. 64). He says that closure seems to be "something like an axiom about knowledge" (2002, p. 312). In addition, many philosophers believe that strong closure principles are needed to theoretically account for how deduction can be used to extend one's knowledge (John Hawthorne 2014, p. 43).

The principle of *propositional closure* is found in mathematics and propositional logic. Closure is a property of sets. The members of sets in propositional logic are that of 'true propositions.' With respect to propositional logic, a set of true propositions is closed under a valid deductive argument because true propositions will entail other truths. When reasoning by *modus ponens* across true propositions, it is accepted that if **p** is true, and if **p** implies **q**, then the truth of **q** is entailed by the truth of **p**.

In contrast, in the case of *epistemic closure*, the members of the set are 'the items of knowledge.' With the principle of epistemic closure, it is maintained that if **S** knows **p**, and **p** implies **q**, and if **S** is aware of this implication, then **S** knows that the entailed proposition **q** is true. This adaptation (or extension) of the principle of 'propositional closure' to *epistemic situations* has been argued against by Fred Dretske (1970) and Robert Nozick (1981). I agree that the alleged entailment among 'known propositions' using m*odus ponens* is *not* always a truth-preserving form of inference.

With an interest in accepting propositional closure, but denying epistemic closure, I contrast the following definitions:

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

The content of the first definition of 'propositional deductive closure' is not controversial. *Modus ponens* across true propositions is a standard form of valid inference found in beginning logic texts. This mode of argument states that if proposition **p** is contingently true, and if **q** is true in every possible situation where **p** is true, then from these two propositions, it is implied (or entailed) **q** is true. *Modus ponens* across true propositions has this deductive argument form:

1) **p**.

2) If **p**, then **q**.

3) Therefore, **q**.

If **p** is true in premise 1, and if it is true in premise 2 that whenever **p** is true then **q** is true, then it is necessary that **q** is true.

While the principle of propositional closure is an accepted part of standard logic, the principle of closure across knowledge states is controversial. This principle asserts that if **S** knows that **p**, and comes to believe **q** by correctly deducing it from belief **p**, then **S** knows that **q**. Epistemic closure asserts that 'Necessarily, if **S** knows that **p**, and **S** knows that **p** entails **q**, then **S** is in a position to know that **q**.' *Modus ponens* across knowledge states has this argument form:

1) **S** knows **p**.

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

3) Therefore, **S** knows **q**.

For philosophers who endorse the epistemic closure principle, it seems intuitive that by validly deducing the consequences of previously known premises, we can expand our knowledge. We can learn new propositions from propositions that we already know.

The Argument against Epistemic Closure

Is the principle of epistemic deductive closure principle true? Should this principle be accepted as an epistemic axiom? Logic is the study of the principles of good reasoning. A good principle of reasoning (or rule of inference) should *always* be truth-preserving, so that whenever the principle is combined with other true premises in a deductive argument, the truth of the conclusion is always entailed. If a principle of

reasoning fails to be truth-preserving, it should be rejected. Below I show that the epistemic closure principle is an invalid mode of material reasoning because when it is adopted, it sometimes implies a false conclusion from otherwise true premises. There are conceivable instances where true premises are combined with epistemic closure that leads to the deduction of a false conclusion. Here is an example where the following three premises can be assumed true, including the epistemic closure principle (in premise 3), where the deduced conclusion is (intuitively) false:

The Parked Car Situation

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

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

3) If **S** knows that if his car is parked on Nelson Street, then **S** knows that the car has not been stolen. (This is assumed true by epistemic closure).

4) Therefore, **S** knows that his car is not stolen.

In chapter one, it was shown that the conclusion of this argument, depicting a normal circumstance, could plausibly be materially false, even if the first two premises are true.

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

To repeat, the car case, with its four contextual variants, illustrate our intuitions about *when it is true* that '**S** knows where his car is parked.' These cases support an intuitive principle:

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

If the conclusion (i.e., that **S** knows that his car is not stolen) is materially false in the above deductive argument, then this suggests that the conditional premise 3 (i.e., epistemic closure) is false. This indicates that in (at least some) deductive arguments that include an epistemic closure premise, that this will result in an invalid argument form.

In contrast, compare the following two arguments where the validity of *propositional* closure is illustrated. In both of these arguments, if the premises are true, then the form of the argument necessitates a true conclusion.

(A) An argument from propositional closure:

1) **S**'s car is parked on Nelson Street. (True).

2) It is contradictory for **S**'s car to be parked, and simultaneously being stolen and driven down a street. (True).

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

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

(B) Another argument from propositional closure:

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

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

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

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

The form of each of these two arguments is deductively valid by *modus ponens* across propositions. In both arguments, the truth of premises 1, 2, and 3 entails the truth of the conclusion. *If* propositions 1, 2, and 3 are assumed true, the truth of those premises makes conclusion 4 necessarily true. There is no debate about the validity of the above arguments. In both arguments, the conclusion makes no assertion about what **S** *knows.* The conclusion states a true proposition that is validly entailed by the truth of its three premises.

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

(C) An Argument from Epistemic Closure:

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

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

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

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

But, again, we have asked, is it true that if **S** knows where his car is parked, then **S** knows that it hasn't been stolen? Does **S**'s knowledge of ordinary proposition **o** entail that **S** is in an epistemic situation to know the falsity of any propositions that contradict **o**?

Premise 3 is false, and the conclusion is false. In a case of knowing where my car is parked, I can know where my car is parked, but at the same time I do *not* know whether 1) a car thief has just spotted my car, or 2) another auto has just smashed into my car moving it several feet away, or 3) a car thief has driven my car a mile away from its former location. Whether I know where my car is parked is contingent upon the belief-forming mechanisms that lead to my belief, and the contingent material conditions that surround me. When **S** knows **o**, **S** does *not know the falsity* of every proposition that is undermining to belief **o**.

In the following sections, we will use additional case examples to explore the issues and intuitions that motivate a misplaced acceptance of epistemic closure. The overall effort will be to explain our ordinary (and theoretical) linguistic intuitions and beliefs when dealing with skeptical possibilities. If we are successful with the following arguments, we will be better placed to appreciate the four PE conditions that allow us knowledge about the external world, and why it is impossible to know that the radical skeptical hypothesis is false.

(#1) Does Samantha Know that She Sees a Goldfinch?

Suppose that Samantha is an ordinary observer looking at a prairie field near her home in North America. Samantha says to her young daughter that she sees a goldfinch, and points to it from a distance of fifteen yards away. A 'goldfinch' is a small yellow bird that is native and common to the area. Samantha asserts that she knows that she sees a goldfinch in the field.

But, when asked about how assured she is of her belief, Samantha candidly concedes that from her vantage point, she is not able to distinguish the sighted bird as being different from what could be a 'canary,' which is non-native to North America, and could appear in the prairie as a lost pet. With this admission, does Samantha still know that she is seeing a goldfinch?

One scenario is that Samantha cannot rule-out the unlikely possibility that the bird is a canary, and according to her evaluation, Samantha admits that she does *not* know that she sees a goldfinch. If Samantha acknowledges that with her visual and statistical evidence that she is unable to discard this potentially defeating possibility, then PE conditions 4a and 4b are violated. Samantha would concede that she does not know that she is seeing a goldfinch and would continue to have a strong belief that she is seeing a goldfinch, without knowing it.

But alternatively, what if Samantha dismisses (i.e., rules-out, discards) the remote possibility that she is seeing a canary, and steadfastly claims that she *knows* that she is seeing a goldfinch? For example, if Samantha believes that the possibility of seeing something other than a goldfinch is so small, and that her visual vantage point is 'adequate enough' to strongly believe that she is seeing a goldfinch does she now *know* that she is seeing a goldfinch?

In this situation of 'critical doubt' it might seem that the normative dimension of 4a and the subjunctive conditional 4b leave us at an impasse as to whether Samantha knows that she sees a goldfinch. But according to the PE definition, whether **S** knows that she is seeing a goldfinch is partially contingent upon her material surroundings. If **S** claims to *know* that she is seeing a goldfinch, the question of whether she *really* knows, depends partially upon the fact of whether there exists something nearby that is not a

goldfinch, but resembles a goldfinch. Whether the first sentence of condition 4b (i.e., that there is no undermining evidence that would weaken **S**'s belief) is satisfied contingent upon whether there is an entity, in direct view or in the vicinity that would undermine Samantha's strong belief that she sees a goldfinch. For instance, if someone's pet canary had escaped and flown into the air space only a half-mile from where **S** now stands, this fact would be undermining to **S**'s belief, no matter whether **S** was aware of it or not. **S** would likely concede that she does *not know* that she was seeing a goldfinch *if there actually exists* a canary in the vicinity of where **S** now gazes (and *if* **S** was to learn of this fact). (The half-mile parameter is arbitrary and reflects **S**'s tacit evaluation of the importance of **p** being true).

In sum, if **S** believes she sees a goldfinch, and if it is true that she is seeing a goldfinch, and if **S**'s belief that she is seeing a goldfinch is from evidence that is relevant for why it should be believed that she sees a goldfinch, and if **S**'s visual faculties and background beliefs are 'adequate' for her to rule-out the canary and other conflicting hypotheses, and if there are no nearby yellow birds (or goldfinch resembling objects) that could be mistakenly identified as a goldfinch; then **S** *knows* that she sees a goldfinch. The PE definition states that if these material conditions are obtained, then '**S** knows **p**' is true, where **S** designates Samantha and **p** designates 'I see a goldfinch in the field.' Although spatial considerations and characteristics of resembling objects are left very imprecise, the PE definition brings into account all of the factors for the possibility that Samantha knows that she sees a goldfinch, as well as the alternative possibility that she does not know that the is seeing a goldfinch.[5]

With respect to this goldfinch example, a staunch advocate of epistemic closure will maintain that *when Samantha knows* that she is seeing a goldfinch, *she also knows* that she is not seeing a canary. But is this correct? Is it true that if **S** knows **p**, then she must *know* that any and all undermining evidence (and potential defeaters) are false (or irrelevant) to the truth of **p**? We answer 'no' to this question and appeal to the principle of fallibilism and condition 4a. With 4a, it is stated that **S** must subjectively 'rule-out'

---

[5] J.L. Austin (1979) states that in order to know that one is seeing a goldfinch in a prairie, one doesn't need to prove that it isn't a *stuffed* goldfinch (p. 84). This is consistent with the position adopted here.

possibilities (in a pragmatic and normative context) that imply ~**p**, as a necessary condition for knowing.

The reason why knowledge condition 4a is endorsed, and epistemic closure is rejected is that condition 4a allows **S** to *fallibly rule-out and resolve* actual and logical possibilities that would undermine (or defeat) **S**'s premises for believing **p**. With condition 4a, **S** can tacitly or explicitly discard improbable undermining possibilities, and *assume* them false, *without knowing* them false. The acceptance of fallibilism suggests '**S** can know **p** upon strong reasons, but **S**'s strong reasons for believing **p** do not guarantee the truth of **p.**' Fallibilism has strong support among many epistemologists and is endorsed here. Two common statements of fallibilism in terms of 'justification' and 'evidence' are as follows:

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

By accepting condition 4a and the principle(s) of fallibilism, it is affirmed that **S** can know **p**, based in part upon ruling-out potential defeating propositions that cannot be known to be false. Sherman and Harman (2011) similarly endorse this principle.

(#2) Dretske's 'Zebra at a Zoo' Example

To continue the argument against epistemic closure-*modus ponens*, and to advocate fallibilism, let's consider how Dretske's (1971) well-known 'zebra at a zoo' argument lends support to the thesis that the closure principle is false. His example:

> You take your son to the zoo, see several zebras, and when questioned by your son, tell him they are zebras. Do you know that they are zebras? Well, most of us would have little hesitation in saying that we did know this. We know what zebras look like, and, besides, this is a city zoo and the animals are in a pen clearly marked 'Zebras.' Yet something being a zebra implies that it is not a mule and, in particular, and not a mule cleverly disguised by zoo authorities to look like a zebra. Do you know that these animals are not mules cleverly disguised by the zoo authorities to look like zebras?

When epistemic closure is assumed in the third premise, this example can be stated in the same form as the 'parked car' example:

1. **S** knows that 'I see a zebra.'

2. **S** knows that a zebra is not the same thing as a disguised mule.

3. If **S** knows that 'I see a zebra' then **S** knows that 'I am not seeing a disguised mule.'

4. Therefore, **S** knows that 'I am not seeing a disguised mule.'

Dretske asks what are the circumstances that make the first premise true?[6] He says that in an ordinary zoo, **S** would spontaneously believe (from perceptual evidence) that she is seeing a zebra, without having any conscious inference about why she sees a zebra. Similarly, **S** wouldn't consciously rule-out the proposition 'I am not seeing a disguised mule.' This possibility is so bizarre that it wouldn't enter **S**'s conscious thought process.

Dretske then asks what if someone walked up to **S** and suggested that she might be seeing a *mule* that is cleverly disguised as a zebra. Dretske suggests that **S**, from her distance and perspective, would likely concede that she would not be able to distinguish a zebra from a cleverly disguised mule. She would admit that she couldn't rule-out, and know that this bizarre possibility is false, based upon her perceptual reasons alone. Given that **S** does not know that this error possibility is false, does it follow that **S** no longer knows that she sees a zebra in the public zoo?

Dretske contends that **S** can still know that she sees a zebra, even if she doesn't know that she isn't seeing a disguised mule. Premise 3 in the above argument is false. It is false that if **S** knows that 'I see a zebra' then **S** knows 'I am not seeing a disguised mule.' The falsity of this premise mirrors the falsity of the closure principle, 'If **S** knows **z**, and **S** knows **z** entails **~m**, then **S** knows **~m**.

---

[6] Dretske's response to what makes this premise true will not be discussed here. Dretske's theory involves acceptance of a counterfactual definition of knowledge, and a possible worlds 'relevant alternative' semantics and neither of these is endorsed here. With comparison to the PE account, Dretske's concept of 'relevant alternative' is eliminated. It is loosely replaced by 'undermining evidence' where the existence of undermining evidence is contingent upon **S**'s background beliefs and material conditions.

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

We conclude: In order to know that 'I see a zebra' or that 'I'm seeing oranges' 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**. One can discard improbable possibilities (e.g., disguised mules, wax imitation oranges) based upon one's background beliefs, and *assume* them false, *without knowing* that every undermining possibility is false.

It is again emphasized that the above argument is only directed against the principle of *epistemic* deductive closure, and *not* against the principle of *propositional* deductive closure. The principle of 'propositional closure' under *modus ponens* is true: If **z**, and **z** entails **~m**, then **~m**. This argument is valid:

1. **S** knows that she sees a zebra. (True).
2. **S** sees a zebra. (Implied from the truth of the first premise). (True).
3. **S** knows that a zebra is not the same as a disguised mule. (True).
4. If **S** knows that she sees a zebra, then **S** does not see a disguised mule. (True).
5. Therefore, **S** does not see a disguised mule. (*Modus ponens* from 1, 4) (True).

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

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

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

(1) **S** knows that she has five pounds of grapes.

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

(3) Therefore, **S** knows that the scale isn't malfunctioning.

But **S** *doesn't* outwardly *know* that the scale isn't malfunctioning! On the PE definition, *if* **S** knows that she has five pounds of grapes, then **S** has dismissed the possibility of malfunction (as part of her background evidence) and if there is no material malfunction (satisfying PE 4b), then **S**'s reading a reliable scale satisfies PE 3, and **S**'s true belief ('I have five pounds of grapes') is knowledge. All conditions are satisfied. But does **S**'s presumed knowledge that she has five pounds of grapes *entail* that she knows the scale isn't malfunctioning? Intuitively, it *doesn't* seem that **S** knows that the scale *isn't* malfunctioning; it is just *contingent* that it isn't. Of course, **S** might verify the weight of the grapes with a reliable scale at her home, and then know that the store's scale was accurate, and if this is the case, then this just confirms that **S** already knew **p**.

Similar to cases where **S** knows where his car is parked (when it might be stolen), **S**'s knowledge of purchasing five pounds of grapes doesn't seem to allow **S** to know that the scale isn't malfunctioning. Instead, **S** must fallibly rule-out this counter-possibility as false as part of condition 4a, and condition 4b must be contingently satisfied. With this example, where **S** may know the weight of the grapes without knowing that *the scale isn't malfunctioning*, the principle of epistemic deductive closure is shown false:

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

In comparison, however, the principle of propositional deductive closure remains true:

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

With propositional deductive closure, (where **p** = '**x** is 5 pounds') if '**S** knows **p**,' and '**S** knows **p**' implies **q** (i.e., the scale is reliable or accurate), then **q** (the scale is reliable, or at least was accurate on this occasion). This deduced propositional entailment that the scale was reliable (or accurate) is true. Propositional closure is a valid rule of logic.

Is Epistemic Deductive Closure Sometimes True?

By restricting closure to transparent (i.e., immediately verifiable) implications, there are examples where epistemic closure is plausible. For example:

1) I know that there is a cat in front of me.

2) A cat is a type of animal.

3) If I know that there is a cat in front of me, then I know that there is an animal in front of me.

4) I know that there is an animal in front of me.

But as discussed, when assuming epistemic closure with opaque (i.e., immediately unverifiable) implications, the principle sometimes fails to transmit a true conclusion.

Supporters of the principle of epistemic closure (*modus ponens* across knowledge states) sometimes offer the following examples where epistemic closure is true. These four arguments are found in the literature to support epistemic deductive closure:

> 1) **S** knows today is Friday. 2) **S** knows that Saturday follows Friday. 3) If **S** knows today is Friday and that Saturday follows Friday, then **S** knows that tomorrow is Saturday. 4) Therefore, **S** knows that tomorrow is Saturday.
>
> 1) **S** knows that **L** was born in Kentucky. 2) **S** knows that Kentucky and Ohio are two of the fifty exclusive states in the United States. 3) If **S** knows **L** was born in Kentucky and that Kentucky and Ohio are separate states, then **S** knows that **L** was not born in Ohio. 4) Therefore, **S** knows that **L** was not born in Ohio.
>
> 1) **S** knows that Andy is taller than Bradley. 2) **S** knows that Bradley is taller than Charlie. 3) If **S** knows that Andy is taller than Bradley and that Bradley is taller than Charlie, then **S** knows Andy is taller than Charlie. 4) Therefore, **S** knows Andy is taller than Charlie.
>
> 1) **S** knows that his friend is wearing a yellow tie. 2) **S** knows that ties come in different colors. 3) If **S** knows that his friend is wearing a yellow tie, then **S** knows that his friend is wearing a tie. 4) Therefore, **S** knows that his friend is wearing a tie.

In these cases, epistemic deductive closure holds. Competent deductions across knowledge states are achieved. In the first example, the conclusion is known necessarily based upon **S**'s knowledge of the fixed definiens of the sequential days of the week. In the second and third examples, the conclusion is known necessarily from a set of

stipulations in set-theoretic reasoning. In the fourth example, **S**'s knowing that his friend is wearing a yellow tie implies that **S** knows his friend is wearing a tie. This illustrates a logical rule of conjunction elimination. The principle of epistemic deductive closure is *sometimes true* when the truth of the conclusion is based upon premises that crucially involve fixed definiens concepts[7], stipulated states of affairs, or implied (or deductive) inferential consequence. That epistemic closure holds in these limited cases doesn't imply that epistemic deductive closure holds in all cases.

The Validity of Conjunction Decomposition across Epistemic States

Besides the above deductions based upon immediate verification and prior stipulations, it is alleged that there are empirical cases where it is inconceivable that closure using *modus ponens* across knowledge states doesn't hold. For instance, this example:

1. Sally knows that she is holding one animal, her cat, Fluffy.

2. Sally knows that her cat Fluffy is not identical to her dog, Rover.

3. *If* Sally knows that she is holding one animal, her cat Fluffy, *then* Sally knows that she is not holding dog Rover. (A material conditional).

4. Therefore, Sally knows that she is not holding dog Rover.

But despite the validity of the argument, the *interpretation* of the structure of the reasoning in this epistemic situation can be questioned. *The fact that Sally knows* that she is holding Fluffy and not Rover in premise 3 is better interpreted as a *conjunction*, rather than a material conditional. It isn't natural for Sally to conclude that she is not holding Rover as a consequence of the three premises. It is more likely that Sally has relevant (truth-connecting) reasons for knowing *the conjunction* that she is holding a cat and not holding a dog. The conjunction is represented in premise 3a:

1. Sally knows that she is holding one animal, her cat, Fluffy.

2. Sally knows that her cat Fluffy is not identical to her dog, Rover.

[7] A 'fixed definiens concept' (i.e., 'closed concept,' 'formal concept') is defined to *unequivocally identify* any item(s) that fall under its definition. The definiens is precise enough to distinctly exclude entities that don't fall under the definition. It is stable and not subject to alteration (without creating a new concept).

3a. Sally knows that she is holding one animal her cat Fluffy, *and* Sally knows that she is not holding dog Rover. (A conjunction).

4. Therefore, Sally knows that she is not holding dog Rover.

The deduced conclusion, that Sally knows that she is not holding Rover is a *conjunction decomposition* of the conjunction 'Sally knows that she is holding cat Fluffy, *and* Sally knows that she is not holding dog Rover' in premise 3a. The valid inferential form of conjunction decomposition across knowledge states can be stated as a rule of inference as follows:

> **Principle of Epistemic Closure- Conjunction Elimination** (Conjunction Decomposition across knowledge states): If **S** simultaneously knows **p** and **q**, and **S** deduces **p**, then **S** knows **p**.

In semi-formal notation, this principle is: if **S** knows **holding cat** *and* knows **not-holding dog**, then it is entailed that **S** knows **not-holding dog**.[8] Conjunction elimination is recognized as a valid form of inference for propositions in standard deductive logic. Likewise, epistemic closure by conjunction elimination is a valid mode of inference across knowledge states.[9]

The Denial of the KK Principle

The KK principle is short for the 'knowing that you know' principle. It asserts that 'For any **p**, if one knows that **p**, then one knows that one knows it.' This principle is false, and most philosophers regard it as false. Knowing that one knows would require **S** to be omniscient of all pertinent facts related to a proposition, including an externalistic perspective of one's own material situation. There is no meta-knowledge. Having

---

[8] Normally Sally can know that she is not holding a dog by having direct perceptual reasons for why the item in her lap is a cat (and not a dog), and not from any deductive reasoning based upon knowing that she is holds a cat. *Neither* epistemic closure by *modus ponens* nor conjunction decomposition across knowledge states explains how Sally typically knows that she isn't holding a dog.

[9] Hawthorne (2004, p. 31) worries that if epistemic closure under *modus ponens* should be false, does this imply that other closure principles involving knowledge states will also fail? It is illustrated here that the falsity of epistemic closure using *modus ponens* doesn't imply the falsity of all other epistemic closure principles. (In fact, he is concerned with this same conjunction decomposition across knowledge states that he labels 'distribution').

knowledge is simpler. Knowledge is the material situation where the four PE conditions are contingently true. It is false to assert the conditional that 'if **S** knows **p**, then ***S*** *knows* that the four material conditions of knowledge are satisfied.' Assuming an external (third person) position for assessing whether **S** knows **p** is a heuristic for defining knowledge, but we cannot assert that when knowledge is attained, that we know that we know.[10]

Why Contextualist Theories of Knowledge Fail

At this point we should compare the PE theory of knowledge with contextualist theories of knowledge as offered by DeRose (1995), Cohen (1999), and David Lewis (1996). Contextualist theories of knowledge state that there are shifting standards of personal justification (high, ordinary, low) that are relative to different epistemic and social contexts. According to contextualism, differences in context, and contextual shifts are crucial to whether we grant knowledge to ourselves or others, in given situations.

With respect to the BIV hypothesis, the contextualist proposes that in ordinary situations, I can *know* that I am not a BIV. As I conduct my daily activities, I am personally justified in believing that I am not a BIV and the possibility of being a BIV is remote. The skeptical conclusion is denied by rejecting premise #1 in the standard skeptical argument:

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

(#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**.

Contextualists maintain that *in ordinary contexts*, I can know that 'I am not a BIV,' so premise #1 is false. Thus, the assumption of epistemic deductive closure, implicit in premise #2, is retained.

---

[10] It seems probable that I *know* (based on history books and memory) that 'Napoleon was defeated at the Battle of Waterloo in 1815' by presuming this is a true belief with personally justified strong evidence and if there is no undermining evidence (e.g., there is recent uncovered evidence by a credible revisionist historian that renders the date questionable). But if I was to say that 'I *know* that I *know* that Napoleon was defeated at the Battle of Waterloo in 1815,' this apparently obligates me to further investigate my reasons (evidence) satisfying 4a, and to do some research to make sure that my memory is correct (and that there is no credible historian disputing this claim). Also, it seems implied that for me to *know* that I *know* p, I should *know* that my reasons are substantially relevant (i.e., truth-connecting) for why I know that I know p. This seems to require unobtainable omniscience. The KK principle is clearly false.

In contrast, when a contextualist considers a shift in epistemic context, such as moving to the more thoughtful skeptical context discussing the BIV possibility, one no longer knows that one is not a BIV. The skeptical hypothesis of premise #1 undermines my belief of having an independent conscious experience. In a philosophy classroom, I no longer know that 'I have two hands' because 'I may be a BIV' and 'I don't know that I'm not being deceived by a demon.' In a context where justification standards are higher, contextualists consider the above argument to be valid and sound. The skeptical hypothesis and the skeptical conclusion are true. One can only hope to have true ordinary beliefs (but not knowledge) in this less-permissive philosophical context.

This contextualist diagnosis of the skeptical argument is unsatisfying and false. It has been argued here that premise #1 is clearly true and #2 is clearly false. But if premise #1 is true, and I cannot know that I am not a BIV, then this leaves the unsettling question: *How can I know that there exists an external world?*

Can I Know That There Exists an External World?

A major concern with accepting the skeptical hypothesis (i.e., I do not know that 'I am not a BIV') is the related question of how can I know that there is an external world, and that I am an independent conscious entity? In chapter one, we have described examples of how knowledge is possible. The question about how **S** can know that there exists an external world has a concise answer. The PE definition states that **S** can know that 'there exists an external world' *if* the following four material conditions are satisfied:

a) **S** believes that 'there exists an external world.'

b) **S** has premises with references to experiences of pain, happiness, interaction with other people, intentional action, touching external objects, distinguishing waking and sleeping, etc. that are relevant (wide sense) for believing there exists an external world.

c) **S** believes that the existence of an external world is the best explanation for his premises and dismisses the BIV possibility as being an unlikely defeater.

d) It is true that there is an external world. There is no defeater (e.g., evil demon).

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

Summary and Conclusion

We have denied the soundness of the standard skeptical argument because the second premise and conclusion are false. The first premise has been proven true with a four-premise deductive argument.

The second premise is false because it assumes the truth of epistemic closure: If **S** knows **p** and **S** knows **p** implies **q**; then **S** is in a position to know **q**. This form of inference has been shown invalid by counterexamples (involving parked cars, gold finches, seeing zebras and oranges, and knowing the weight of grapes*)*. Contrary to the conclusion of the radical skeptical argument it is indeed *possible* that we can have *knowledge* of ordinary empirical propositions and an external world.[11]

In summary, the PE definition and its implied denial of epistemic closure responds to the skeptic's challenge by showing how knowledge is possible. It helps explain how and when we *may* have knowledge. But the PE definition never guarantees that we (ever) have knowledge. As the BIV situation indicates, every one of my beliefs *may* be the product of a systematic elaborate deception. In describing how knowledge is possible, it is accepted that I may have *no* knowledge at all.

---

[11] If one can know that 'there exists an external world,' we can assume that it is possible to know that 'there are other minds' and that 'the past is real' using the same a-d formula as above.