

January 2018

## A Defense Of Epistemic Invariantism And Intellectualism

Lucas Lockard

Wayne State University, llockar0@gmail.com

Follow this and additional works at: [https://digitalcommons.wayne.edu/oa\\_dissertations](https://digitalcommons.wayne.edu/oa_dissertations)



Part of the [Epistemology Commons](#)

---

### Recommended Citation

Lockard, Lucas, "A Defense Of Epistemic Invariantism And Intellectualism" (2018). *Wayne State University Dissertations*. 2044.

[https://digitalcommons.wayne.edu/oa\\_dissertations/2044](https://digitalcommons.wayne.edu/oa_dissertations/2044)

This Open Access Dissertation is brought to you for free and open access by DigitalCommons@WayneState. It has been accepted for inclusion in Wayne State University Dissertations by an authorized administrator of DigitalCommons@WayneState.

**A DEFENSE OF EPISTEMIC INVARIANTISM AND  
INTELLECTUALISM**

by

**LUCAS RYAN LOCKARD**

**DISSERTATION**

Submitted to the Graduate School

of Wayne State University,

Detroit, Michigan

in partial fulfillment of the requirements

for the degree of

**DOCTOR OF PHILOSOPHY**

2018

MAJOR: PHILOSOPHY

Approved By:

---

Advisor

Date

---

---

---

© COPYRIGHT BY  
LUCAS RYAN LOCKARD  
2018  
All Rights Reserved

## ACKNOWLEDGMENTS

First, I would like to thank my committee chair Dr. Bruce Russell. Over the last four years, he has participated on three committees related to my PhD work reading several drafts of various manuscripts. His feedback has been invaluable to the development of the ideas in this dissertation. Throughout my time at Wayne, Dr. Russell has shown a willingness to thoughtfully listen as I worked out partially developed ideas, pointing out missteps when needed, and, perhaps, the occasional step in the right direction. His influence on my philosophical development is without question.

I would like to thank my committee members: Dr. Susan Vineberg, Dr. Eric Hiddleston, and Prof. Vincent Wellman. Both Dr. Vineberg and Dr. Hiddleston have provided helpful feedback on work related to this dissertation over the last few years. Whether through commenting on drafts or through direct conversation, their suggestions and advice have been influential in shaping some of the central ideas presented in this work. I'm also grateful for Prof. Wellman's helpful advice which led to some important changes in the early chapters.

I would also like to thank several colleagues including Graeme Cave, Joseph Dunne, Dr. Daniel Yaekel, and David Baxter. Each has been a source of support both personally and academically. In particular, Graeme Cave spent a significant amount of time reading most, if not all, of this dissertation. His advice has been invaluable in the process.

Thank you to friends and family for their support and polite interest at holidays. Thank you to my Mom and Dad, Jean and Michael Lockard, who are everything a son could ask for—their loving support has always provided a foundation for me to stand upon. Thank you to my boys, Graham and Whittaker, who, without knowing it, sacrificed much in their early years while Mom and Dad both pursued doctoral degrees.

Finally, I would like to thank my wife Megan, who willingly moved to Detroit with a 1-year-old, so I could join the department at Wayne. None of this would have happened without her support, of course. But it was specifically Megan’s relentless determination and self-discipline in her work that inspired me to forge ahead in my own. I’m pretty sure everything I do is to impress her, but don’t tell her that.

## TABLE OF CONTENTS

Acknowledgments	ii
List of Tables	v
Chapter 1: The Traditional View and Its Alternatives	1
Chapter 2: Knowledge Attributions and Mistakes	33
Chapter 3: Stakes and Intuition	72
Chapter 4: Intellectualism and Practicalism	95
Bibliography	123
Abstract	126
Autobiographical Statement	128

## List of Tables

Table 1: Comparing contextualist and invariantist explanations of first-person cases _____	19
Table 2: Comparison of Alien-Airport and Printer-Airport _____	60
Table 3: Comparing each of the three airport cases _____	66
Table 4: Reed's Caesar Case _____	96

## Chapter 1: The Traditional View and Its Alternatives

### *Introduction*

The late David Lewis's 1996 paper "Elusive Knowledge" begins with the ostensibly innocent claim, "We know a lot." This seems right if we consider a wide selection of common, everyday beliefs. If that is true, then it also seems right that we must be fallibilists about knowledge. On that view, a subject can know that  $p$  even if their justification for  $p$  is not sufficient to rule out every alternative possibility in which  $p$  is false. Yet, as soon as we start considering alternative possibilities, the appeal of fallibilism begins to wane—and along with it the idea we know a lot—for two reasons. First, most of our beliefs lack sufficient justification to rule out these possibilities and, second, it is difficult to say which possibilities must be ruled out and which can be properly ignored.

Attempts to resolve this 'problem of the alternatives' have multiplied in recent years with increased focus on the colloquial use of 'know'. Some argue that colloquial use reveals a direction relation between our knowledge claims and non-traditional factors. For instance, they argue that, in some cases, the difference between a person who knows and a person who doesn't know is only explainable by a difference in risk.

To elicit intuitions in support of this view, two cases are compared. The first case includes a subject who is in a low-risk context and for this case it seems right to conclude he does know. This is compared with a second case which is like the first except for two additional factors: there is a significant increase in the practical



risk and an alternative possibility is introduced. In this high-risk case it seems right to conclude the subject *does not* know.

Providing an acceptable interpretation of our responses to these cases has become a central requirement for any theory dealing with knowledge attributions. Most importantly, the theory must be able to explain our intuitions. This is especially important for the traditional view because it implies at least one of our intuitive responses to the original cases is wrong. On the traditional view, there is a standard amount of justification required to know that a proposition is true; in other words, the *truth-conditions* for 'know' do not vary. Since the same truth-conditions apply, one of the responses to the types of cases mentioned above must be false.

Some argue that accommodating our intuitions will require shedding the traditional invariantist position. On the *contextualist* view, the truth-conditions for 'know' *vary* according to context-dependent features (e.g. risk). This implies the invariantist feature of the traditional view is false. To others, the better explanation is that practical reasons factor directly into our epistemic justification. These *practicalists* argue that beliefs are not justified solely according to truth-conducive reasons but also the practical reasons relevant to a given situation.<sup>1</sup> Their view conflicts with a second feature of the traditional view that we will refer to as

---

<sup>1</sup> The name 'practicalist' or 'practicalism' comes from an unpublished paper by Matthias Steup. It is also sometimes referred to as 'pragmatic encroachment'.

*intellectualism* – the idea that only *truth-conducive* or *epistemic* reasons are relevant to the justification of belief.

This dissertation defends a traditional view of the truth-conditions for ‘know’. For the traditionalist, there is both an invariantist and an intellectualist component. In chapter 1, I begin by reviewing the contextualist and practicalist theories and the implications they have for invariantism and intellectualism. We’ll discuss the motivations for contextualism and practicalism as well as the central arguments offered in support of these views. Moving to the later chapters, I defend invariantism in chapters 2 and 3 and in chapter 4 turn my efforts to a defense of intellectualism.

In chapter 2, I argue that the contextualist view, most prominently defended by Keith DeRose and Stewart Cohen, fails in an important respect – it does not offer a plausible explanation of our intuitions as they have claimed. While contextualism purports to explain ordinary cases akin to those mentioned above, it fails to do so with a wider range of intuitions. In particular, contextualism is incompatible with the intuition that there is a contradiction between the knowledge attribution and denial. Since contextualism implies there is no contradiction, they must explain we why think there is in a way that is consistent with ordinary usage.

One the other hand, the traditional view implies that there *is* a contradiction between the two claims. So, as mentioned above, one of the knowledge claims is false. The task for the invariantist is to explain what mistake has occurred in the

original cases; that is, they must explain what independent reasons there are to conclude one of our intuitions from the original case is false. Although it is not important to my view *which* of these claims is false, I will argue that the knowledge denial is false because it rests on a mistaken belief. Furthermore, we'll find that understanding the contextualist cases according to the invariantist view is consistent with a broader range of intuitions.

Turning to chapters 3 and 4, the primary focus is on the idea that variations in practical issues can alter the justificatory requirements for knowledge. For chapter 3, I focus on the contextualist view that practical issues can affect the truth-conditions for knowledge claims. Borrowing a set of cases from Jonathan Schaffer, I argue that our intuitions about knowledge are not driven by practical issues, but instead by epistemic concerns. However, this is not a mere issue of whether the alternative possibility is salient to the subject. Only when that alternative is probable does it prompt an intuitive denial of knowledge. (And in that case, salience is irrelevant.) If this is correct, practical reasons don't affect our intuitions in the way that contextualists have assumed. Only when we have—or believe we have—good epistemic reason to think the person's belief might be false does it become intuitive to deny that they know.

Finally, in chapter 4 I argue against the practicalist view of Jeremy Fantl and Matthew McGrath. Beginning with their paper "Evidence, Pragmatics and Justification", Fantl and McGrath have argued that practical reasons are directly

relevant to one's epistemic justification.<sup>2</sup> To put it roughly, whether *S* knows that *p* is due to a combination of both his epistemic *and* practical reasons – an idea expressed in Fantl and McGrath's *KJ* principle. Several counterexamples to this principle have been offered and I begin by defending an example from Baron Reed. I argue that Reed's example does apply to *KJ* and, more importantly, to the principles that Fantl and McGrath use to support it. The chapter concludes by addressing an important assumption underlying their view: that knowing *p* is sufficient to rationally justify acting on *p*. If this assumption is correct, then practicalism quickly follows. Against this, I argue that knowledge alone isn't enough to rationally justify acting on *p* – the act must also be worth it.

So, you might say that this dissertation is an attempt to support the claim that we do know a lot. However, it is not an attempt to defend fallibilism directly – that will be assumed – but indirectly by defending a traditional view in which our (fallible) knowledge adheres to invariantist and intellectualist standards. In the end, we'll find a way to interpret and understand ordinary uses of 'know' that agrees with commonsense and does better on the whole than the non-traditional theories offered thus far.

### *§1.1 Skepticism and Ordinary Knowledge.*

In line with Lewis's sentiment above, we ordinarily take ourselves to know a fair number of things and we take it that we speak *truly* when making various

---

<sup>2</sup> Fantl and McGrath (2002)

knowledge claims from “Detroit is in Michigan” to “I have hands.” But there is also the ordinary thought that knowledge requires ruling out alternative possibilities—or, at least plausible alternatives—in which your belief is false. It seems that our evidence or reasons for belief should be good enough to rule out these possibilities.

For instance, if I always leave my keys in my coat pocket, then it seems reasonable enough to say that I know they are there even though I’m in the other room and my coat is in the closet. However, if I put them on the top of the fridge an equal amount and often forget where I put them, then it doesn’t seem that I know they are in my coat pocket. In the second case, there is an *alternative* outcome in which my keys are on the fridge rather than in my coat that my evidence can’t rule out. Rationally speaking, this is enough to question my belief they are in my coat and so, it would seem, enough to deny I know where they are *even if* they are still in my coat.

Examples like this indicate one way in which alternative outcomes factor into our ordinary knowledge assessments. The skeptical arguments we find in philosophical discussion play off from this intuitive notion of knowledge and what is required to know; specifically, it focuses on the idea of ruling out alternatives. Yet, there is an important difference between the way in which alternatives impact ordinary knowledge in the example above and the sense in which it impacts knowledge for the skeptical philosopher. In the keys example, I’m given a fair reason to at least question my belief: “I sometimes put them on the

fridge.” But what if my kid said, “I saw an alien eat your car keys!” No sane person would take this seriously, of course. Yet, it is *possible* that a key-eating alien snuck in and ate my keys. For the skeptic, this unlikely possibility is enough to destroy knowledge. It is simply another alternative we can’t rule out because my evidence—I remember putting them in my coat pocket—isn’t enough to distinguish between the case in which they are in my coat and I put them there, and the case in which I put them in the coat pocket and aliens soon after came and stole them. Saving ordinary knowledge will require, at the very least, showing that the keys-on-fridge case is relevantly different than the implausible key-eating-alien scenario. We need to explain why the key-eating alien doesn’t affect whether I know while the possibility I left my keys on the fridge does. Maybe not impossible, but certainly a difficult task.

Once we recognize how alternative outcomes can overturn claims to knowledge, it starts to look like we don’t know a lot of the ordinary things that we think we do. And yet, as Lewis points out, the intuition usually remains that we *do* know a lot. As it stands, we end up with two competing intuitions about what we know. On the one hand, it sure seems that we know a lot. But, on the other hand, it seems like there are a lot of alternatives we can’t rule out, implying that we *don’t* know much at all. So, what is it? In its two-thousand-year history, no response to philosophical skepticism has gained a foothold and this has led some to think it is an unresolvable issue. Yet, skepticism has not received broad acceptance either. Intuitively, it just seems false, but it is difficult to say why. Left

with this pseudo-paradoxical deadlock, it is at this point that the contextualist steps in.

### *§1.2 Skepticism and The Contextualist Solution.*

To take a closer look at these competing intuitions and how they factor into the contextualist view, let's take a brief look at the situation in argument form. The standard skeptical argument begins with an ordinary proposition like "I know I have hands." Focusing on the fact such claims are justified by perceptual experience, the argument then presents an alternative possibility in which I have all the same perceptual experiences, but I don't have hands. For instance, I might be completely ignorant that I am in the Matrix and so go on believing that I have hands among many other things, when I'm just a handless body floating in a vat of goo.

Intuitively, we *cannot* rule out the Matrix possibility, at least, with any certainty. Because of this it seems that I don't know I have hands. And still, it seems as if this quite ordinary matter of having hands is something that I do know. So, we have two propositions that are inconsistent with each other, but both seem correct:

(OK) I know that I have hands.

(SK) I don't know that I'm not a handless brain-in-a-vat.

The problem, of course, is that (OK) implies (SK) is false and (SK) implies that (OK) is false. For some, (SK) is the stronger claim despite the fact (OK) seems correct.

For the skeptic, once we see the implication between the two, while also realizing

we cannot disprove (SK), it then follows that I don't know that I have hands. And by further implication, we don't have much knowledge, in general.

- 1) If (OK) I know that I have hands, then (not-SK) I know that I'm not a handless brain-in-a-vat.
- 2) (SK) I don't know that I'm not a handless brain-in-a-vat.

Therefore, OK is false.

As mentioned earlier, responses to skepticism have utilized several different strategies, but whether these responses attack (OK) or (SK), none has captured broad appeal. Still, it isn't as if skeptical arguments have overrun the intuition we know a lot of ordinary things.

In response to this tension, as well as the lack of any direct solution to the problem, contextualists attempt to provide a way around the skeptical conclusion and validate the belief that we know a lot of ordinary things. Interestingly, they do this while avoiding a wholesale rejection of skepticism. To put it roughly, they argue that there are contexts in which SK turns out to be true and contexts in which OK turns out to be true. So, in effect, contextualism validates the intuitiveness of *both* OK and SK.

At first glance, this strategy might seem to land them in a relativist position because OK and SK are inconsistent with one another. However, the contextualist avoids this result by looking at how we use indexical terms such as 'tall' or 'large'. When using a term like 'tall', the *truth-conditions* of that term will be determined by specific features relevant to the context in which it is uttered. It follows from



this that the truth-conditions can *vary* from context to context. Since the truth-conditions vary, the meaning of ‘tall’ varies and so the proposition expressed also varies. Thus, two sentences such as “Steve is tall” and “Steve is not tall” might not be contradictory if uttered in distinct contexts. For example, in the context of an NBA basketball team, Steve, who is six-feet-tall, would not be considered tall and so “Steve is not tall” would be true. However, in an ordinary context he would be considered tall and so “Steve is tall” would be true. The explanation for this difference is that the *standards* for being tall are lower in the ordinary contexts than in the context of NBA players. Because the shift in standards results in a difference in the proposition being expressed, there is no contradiction between the two claims.

Just as we use and understand indexical terms like ‘tall’, the contextualist argues we should see ‘know’ in the same way. If correct, then what is expressed by the sentence “S knows that *p*” will vary based on the context similar to the varying expressions of ‘tall’ just described.<sup>3</sup> When we apply this reasoning to the tension between ordinary and skeptical intuitions about knowledge, we’ll find that the truth-conditions for ‘know’ are weaker in ordinary contexts and higher in skeptical contexts. And just like the example using ‘tall’ above, the sentence in

---

<sup>3</sup> Although contextualists make reference to terms like ‘tall’ in explaining their theory, explanation is the only purpose. The intent is not to make an argument by analogy and for good reason given important differences between ‘tall’ and ‘know’. See DeRose (2009) 67-68 and especially Stanley (2004) for further discussion.

question does not express the same proposition in the ordinary and skeptical contexts. Thus, there is no contradiction.

Varied truth-conditions also occur between different ordinary contexts. In a context with little at stake the truth-conditions are, again, weaker than a context with a lot at stake – we’ll look at some examples shortly. The result of this is that knowledge can be attributed or denied on the basis of, for example, a change in the pragmatic issues of a speaker’s context. If correct, this result would deal a swift blow to the traditional view that only epistemic factors – belief, evidence, reasons, etc. – are relevant to knowledge assessments.

Before moving on, let me add two points. First, contextualists do not see their theory as something which applies beyond the semantic level. So, even though the truth-conditions for ‘know’ might vary, it does *not* follow that the standard for *knowledge* varies. In other words, ordinary *use* of know may or may not accurately represent a correct understanding of the concept underlying these uses. This distinction has brought about a few important criticisms. For instance, it is often alleged that the contextualist treatment of skepticism is superficial because it doesn’t address the actual problem. While that is an interesting issue, it will not be addressed in this dissertation.

Second, contextualists *do not* take discussion of indexical terms like ‘tall’ as an analogical argument for their view that ‘know’ is an indexical term (See fn. 3). They do not, for instance, argue that ‘tall’ is an indexical term and since ‘know’ functions in a similar way, we should conclude that ‘know’ is also an indexical

term. Rather, the main evidence for their view comes from looking at ordinary cases in which knowledge attributions change and appear to do so without an epistemically relevant explanation. Contextualists argue that their view provides the best explanation of these ordinary knowledge attributing and denying practices. So, we now turn to these cases.

### *§1.3 First Person Cases*

In the last section, we discussed the long-standing tension between the intuition we have ordinary knowledge and the skeptical conclusion that we don't know much at all. This is the main problem contextualists hope to solve. However, instead of correcting the problem by showing that one or the other is wrong, contextualists propose an explanation that validates both. The most prominent support of this view comes from the work of Keith DeRose.<sup>4</sup> In this section we'll look at DeRose's contribution to the discussion which comes from his comparison of various ordinary uses of 'know'.

To support the contextualist position, DeRose developed sets of cases that reflect ordinary uses of 'know', but seem to draw out contradictory intuitions similar to the tension between OK and SK. The basic features of one set of cases—his now famous 'Bank Cases'—are as follows: in one version of the case, the subject(s) has almost nothing at stake and it seems right to say he/she knows. But

---

<sup>4</sup> See Gail Stine (1976) and David Lewis (1996).

when we compare this ‘low-stakes’ case to a case in which there is a lot at stake, it then seems incorrect to attribute knowledge to the subject.

One important aspect of this case is that the *subject* is both attributing and denying knowledge. So, the subject is both *attributor* and *subject* – this is an important distinction that will come into play later. Because the subject is making both the assertions, we’ll refer to these cases as *first-person*. Here is the original version of the bank case in which the subject has virtually no practical risk:

**Bank-Low:** Keith and his wife are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. As they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. They generally like to deposit their paychecks as soon as possible, but it is not especially important in this case that they be deposited right away, so Keith suggests they drive straight home and deposit their paychecks on Saturday morning. His wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” Keith replies, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”<sup>5</sup>

Looking at cases like Bank-Low we are then asked, “Is the subject’s attribution of knowledge true?” Following standard knowledge-attributing practices, it seems appropriate to say that Keith is correct when he says he *knows* the bank will be open the following day. His recent visit, in conjunction with the implicit idea that banks keep consistent hours and the fact he has no reason to think this bank won’t be open, seems sufficient evidence for believing it will be open the following day.

---

DeRose (2009), 2.

Moving to a second scenario, DeRose increases the stakes while also presenting an alternative outcome that was not discussed in the first scenario:

**Bank-High:** Keith's wife reminds him that they have just written a very large and very important check. If their paychecks are not deposited into their checking account before Monday morning, the important check they wrote will bounce, leaving them in a very bad situation. And, of course, the bank is not open on Sunday. She then points out that, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as he was before that the bank will be open, Keith replies, "Well, no, I don't know. I'd better go in and make sure."<sup>6</sup>

Following standard knowledge attributing practices, it seems appropriate to say that Keith doesn't know the bank will be open. Putting the two cases together, we get the following:

(KB) Keith knows the bank will be open.

(~KB) Keith doesn't know the bank will be open.

Explaining the shift from attributing to denying knowledge is the primary focus of contextualism. As discussed earlier, contextualists argue that the best explanation of this shift is a change in the truth-conditions for 'know'. As a result, each speaker has expressed a different proposition. When Keith claims to know in Bank-Low, the standards for attributing knowledge are lower and he does know. However, in Bank-High the change in circumstances—increase of risk and alternative possibility the bank changed its hours—results in a change in the

---

<sup>6</sup> *ibid*, 2. Again, this example has been slightly modified from the original. Note that the stipulation he "remains as confident as before" is original to DeRose's example. I find this claim to be suspect but will not address the matter here. See Bach (2005) for a critique of DeRose's claim.

standards for knowledge while Keith's evidence/reasons remain the same. So, Keith doesn't know in Bank-High.

#### *§1.4 First-person cases and Invariantism*

Recall that invariantists hold the truth-conditions for 'know' do *not* vary. On this view, any attribution or denial of knowledge must satisfy the context-*independent* truth-conditions for 'know' thus putting it at odds with the contextualist view. So, if the contextualist view discussed in 1.3 is correct, invariantism is false. Yet, due to the nature of first-person cases, there are 'traditionalist-friendly' explanations available which we will explore in this section.

One plausible explanation for the change is that Keith simply doesn't believe he knows the bank will be open after considering the alternative possibility raised by his wife. This explanation is further supported by the fact that he says "Well, no, I don't know." It is difficult to see what would justify attributing knowledge to someone when they (sincerely) deny that they know. My point here is not to rule out any possibility of knowledge despite sincerely believing otherwise.<sup>7</sup> However, in most cases, denying that you know  $p$  is a good indication that you do not, in fact, know. If that is correct, it would go a long way to supporting the explanation that Keith simply doesn't believe he knows in the context of Bank-High.

---

<sup>7</sup> Bruce Russell pointed out in conversation one example; a student in an epistemology class might find herself perplexed enough when discussing skepticism to deny that she knows despite still knowing. It seems plausible that such momentary 'lapses' don't constitute an actual loss of knowledge.

Now we have a plausible reason to explain the change between that context and Bank-Low. Is our problem then solved? We do have an explanation for the shift, but unfortunately, that won't settle the debate between contextualism and the traditional view. Even if Keith has lost belief, there is still the question of whether he would know if he claimed to know. In other words, if Keith replied instead "I *know* the bank will be open – I was just there!" would we be inclined to attribute knowledge? If the answer is yes, then the contextualist explanation is weakened. If no, then the traditionalist view is weakened. We'll address this point more in depth in chapter 2. For now, it is important to see that looking solely at whether Keith believes does not settle the issue in favor of either the contextualist or invariantist.

Also relevant is that Keith's wife raises the possibility the bank has changed its hours. In turn, Keith accepts her suggestion as *significant enough* to call into question his reliance on the recent visit as evidence the bank is open. This point is like the previous on changing belief. But whereas the change of belief aspect is somewhat controversial, it is obvious that Bank-High introduces new information to the subject. When new information is introduced into a situation, there is an inherent change in one's epistemic position – assuming, of course, that the person is aware of the new information and its connection to the belief in question. This doesn't immediately imply that Keith's epistemic position has been weakened, however. One might have good reason to disregard the new information or evidence to override it. Nevertheless, Keith exhibits the sort of response we would

expect from someone who has taken the alternative seriously and lacks – or at least believes that they lack – further reasons or evidence against it.

If someone believed they had good reason to reject an alternative suggestion, then we would expect that person to defend their position.<sup>8</sup> Looking back at Bank-Low, this is precisely what happens.<sup>9</sup> Keith’s wife asks how he knows the bank will be open and he responds by noting that he visited this exact bank on a Saturday no less than two weeks prior. But once she makes salient the fact that banks sometimes change their hours, he is without recourse (or so it seems). The moral of this is that a change in Keith’s epistemic position toward “The bank will be open” is a plausible explanation for why he denies that he knows. This seems right for two reasons. First, new information has been introduced to the context and, second, Keith denies that he knows. Since it is natural to assume that someone who denies they know also believes they lack evidence to justify belief, it looks as if Keith claims he doesn’t know because he no longer believes he has sufficient evidence to support his prior belief that the bank will be open. It is possible that new information doesn’t automatically imply a change in epistemic position, but Keith has taken it seriously enough to change his mind about what to do and whether he knows the bank will be open. Again, this would imply that

---

<sup>8</sup> Of course, this is only on the assumption that the person is not under threat, duress, or any other psychologically or physically limiting factor.

<sup>9</sup> Thanks to Bruce Russell for pointing this out.



he has either lost belief or his epistemic position has changed and perhaps it is both.

It would be unsurprising for a subject, upon concluding their evidence to be unreliable, to then withhold belief. And that seems to be exactly what is occurring the Bank Cases. Though DeRose doesn't agree a loss of belief or change in epistemic position is the correct explanation, he does concede that it is a possibility:

Perhaps both the [knowledge-attribution] and the [knowledge-denial] are true, not because the semantic standards for 'knowledge' vary from one case to the next, but because, though the same standards govern both cases, the speaker meets those standards in [Bank-Low], but fails to meet those same standards in [Bank-High].<sup>10</sup>

If the same standards govern both cases, as DeRose concedes might be correct, then an invariantist understanding of 'know' is the better explanation. Additionally, a further reason to prefer this account is that it doesn't resort to relying on the idea of varying truth-conditions. Of course, DeRose eventually rejects this explanation in favor of his view, but the cohesion with traditionalism and avoiding controversial claims about varying truth-conditions, gives further support to these 'traditionalist-friendly' interpretations.

As a final point, note that invariantism is consistent with our *intuitions* of each case; that is, the attribution in Bank-Low and denial in Bank-High seem correct, but we can account for that in a way that requires no extra work to defend

---

<sup>10</sup> DeRose (2009), 60.

invariantism. Later, when we look at *third-person* cases, we'll find that invariantism appears to conflict with our intuitions. That conflict has been a main point of contention between contextualists and invariantists. Whereas contextualists argue that the best theory should vindicate our intuitions, invariantists argue that there are better explanations which retain the 'intuitiveness of our intuitions', but are also favorable to their view.

### *§1.5 Summary of the competing explanations*

In the previous section, we discussed three different interpretations of the switch from attributing knowledge in Bank-Low to denying knowledge Bank-High. Each implies something about broader concerns such as whether there is a change in the person's epistemic position, denial of invariantism, and so on. Looking at the table below, the column on the far left indicates the explanation for the change from attributing knowledge in the low-stakes case to denying knowledge in the high stakes. The top row lists different implications for each of the explanations under consideration.

*Table 1: Comparing contextualist and invariantist explanations of first-person cases*

	Compatible with invariantism?	Compatible with change in epistemic position?	Compatible with change of belief?
Change of belief (COB)	Yes	Yes	Yes
Change in epistemic position (CEP)	Yes	Yes	Yes

Change in truth-conditions for 'know'. (CTC)	<i>No</i>	Perhaps, but brings into question why changing TCs is needed.	Perhaps, but brings into question why changing TCs is needed.
--	-----------	---	---

Note that each explanation seems to be compatible with one another. For example, it is quite possible that the subject has changed his belief as well as undergone a change in epistemic position. And both a change in belief and epistemic position are compatible with changing truth-conditions. Yet, COB and/or CEP suggest alternative explanations that are plausible in their own right. Since both plausibly occur in the bank cases discussed above, we have very little reason to push our explanation beyond one or both explanations.

DeRose does make a minor attempt to subvert COB by claiming that the subject "remains as confident as before", but this stipulation by DeRose is awkward and seems to go against what someone would normally claim in such a situation. Even if that does manage to show the subject's belief remains constant, it fails to explain why the new possibility introduced by Keith's wife (that sometimes banks change their hours) does not lead to a change in epistemic position.

So, the bank cases, and first-person cases more generally, can be explained quite easily by the traditionalist view. What the contextualist needs to provide is a case in which belief and epistemic position do not change—or, at least, it is plausible that they do not—and there is a shift from attributing to denying knowledge. Seeing this problem, DeRose moved to 'third-person' cases which

focus on the intuitive responses of a third-party who is separate from the speaker's context. In the next section, we'll discuss how these cases differ from first-person cases and why the contextualist believes they provide the best support for their view.

### **§1.6 Third Person Cases**

As we saw in the last section, our intuitive approval of both the subject's knowledge attribution and denial does not provide adequate support for contextualism. Even if those cases do exhibit standard knowledge attributing practices, they don't give us much reason on their own to think the truth-conditions for 'know' vary because they can be explained by a loss of belief or change in epistemic position. So contextualists need examples that are more favorable to their view and they believe such cases can be found in the form of *third-person* cases.

In the first-person cases, we are asked whether it seems correct that the *subject* has attributed/denied knowledge to himself. (Again, the subject in the first-person cases also acts as the *attributor*.) Since the shift from attributing to denying seems correct, we need to provide an explanation for it. Third-person cases follow a similar format but differ in that they separate the role of the *attributor* from the subject and focus on the intuitive correctness of knowledge attributions/denials of the attributor who is placed in a different conversational context. As a result, the subject is unaware of any new alternatives or increases in practical risk since he is now separated from the context in which these factors are introduced. But it

was just those factors that led to the traditional-friendly explanations discussed in the previous section (loss of belief, loss of evidence). Since the 'traditionalist-friendly' variables have been removed from the subject, the contextualist argues, the best explanation we have left is that the standards for attributing knowledge have changed.

For third-person cases, we are again asked whether the third-person's attributions/denials of the subject are correct. As with the bank cases, the assumption has been that it is intuitively correct for the speaker to attribute knowledge in the low-stakes scenario but deny knowledge in the high-stakes scenario. Here is such a case from Stewart Cohen that will be the basis for our discussion:

**Airport:** 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 if anyone knows whether the flight makes any stops. Another passenger, Smith, replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "How reliable is that itinerary, anyway? It could contain a misprint. They could have changed the schedule since it was printed, etc.," Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline.<sup>11</sup>

Considering the case from Smith's view, it seems appropriate to say that he does know the flight will stop in Chicago (Airport-Low). However, it also seems appropriate for Mary and John to conclude neither they nor Smith know the flight

---

<sup>11</sup> Stewart Cohen (2000), 95.

will make a stop in Chicago (Airport-High). So, again, we have two claims that are appropriate, but seem to conflict.

As with the Bank cases, contextualists argue that the change from attributing to denying knowledge is best explained in terms of a change in the truth-conditions for attributing knowledge. The contextualist explanation gains credibility, or so it is claimed, from at least three important differences between the Bank Cases and Airport. First, it is obvious that Smith's practical situation has remained constant and so arguing that he has lost confidence in his belief due to a change in circumstances doesn't make much sense. Second, Smith does not deny that he knows as Keith does in the bank case. So, explaining the change from attributing to denying knowledge as a loss of belief seems implausible as well (COB).

Finally, the subject is unaware of any new information. Though Mary and John consider the possibility the itinerary is not reliable, Smith is ignorant of these issues. Because of this, it looks like the traditionalist cannot argue Smith's epistemic position has changed in any way (CEP). Doing so would require making the somewhat odd claim that Smith's epistemic position is, "substantially affected by such factors as what kind of far-away conversation is taking place about the subject."<sup>12</sup> On this point, and the previous two, I agree with DeRose. Clearly, Smith's belief, confidence, and epistemic position have not changed—far-off

---

<sup>12</sup> DeRose (2009), 64.

conversations don't have that sort of effect on individuals. If that is the case, then invariantism looks to be false – truth-conditions are *not* invariant across contexts.

In chapter 2 I argue that this is not the whole story. Much of this debate has been driven by the assumption that any view which validates these intuitions is on better footing than a view that does not. To be sure, whether a view validates our intuitions is important, but the focus of these debates has been on the intuitions from contextualist cases with little discussion of related cases. In chapter 2, I argue that once we take these related intuitions into account, contextualism quickly loses any ground that it had over invariantism. Most importantly, it refocuses the discussion on what is of central importance – whether we have good reason to conclude one of the claims is false. Either Smith claims that he knows or Mary and John's claim that he doesn't know.

### §1.7 *Practicalism*

Since the growth of contextualism, several other non-traditional views have developed. One alternative approaches cases like Bank and Airport by arguing that there is a direct link between knowledge and action. In short, this type of view implies that, if it wouldn't be practically rational to act on  $p$ , then you don't know that  $p$ . An implication of this view is that practical reasons directly factor into one's justification. Thus, knowing is a product of one's epistemic or 'truth-conducive reasons' *and* their practical reasons as well.

*Practicalism* contrasts with *intellectualism* which holds that only truth-conducive reasons count towards one's justification for believing a proposition.

Since intellectualism is a central component of the traditionalist view, it is important that the threat it poses is addressed. Here I will briefly discuss the practicalist view of Jeremy Fantl and Matthew McGrath and later, in chapter 4, I will argue that it is not a sustainable position.

The central goal of Fantl and McGrath's view is to provide an answer to the threshold problem for a fallibilist view of knowledge. Although the fallibilist takes it that knowledge does not require certainty, it is difficult to say exactly how much justification one needs for knowledge. Fantl and McGrath think that if we just look at how practical reasons factor into one's justification for belief, we can get a firm answer to this problem. Their answer is summed up in the following principle which links knowledge to action (note that  $\Phi$  is a placeholder for both beliefs *and* actions):

**(KJ):** If you know  $p$ , then  $p$  is warranted enough to justify you in  $\Phi$ -ing, for any  $\Phi$ .

Applying KJ to contextualist cases, Fantl and McGrath (F&M) argue that the subject doesn't know in the high stakes case because their belief is not "warranted enough to justify" the action in question. Like the contextualists, F&M hold that the switch from attributing to denying knowledge in the contextualist cases cannot be explained, at least not fully, by a change in the epistemic position of the subject or a loss of confidence/belief. Unlike the contextualists, they explain this loss of knowledge as occurring from the direct effect that practical factors have on



knowledge. Exactly what this entails will be clarified shortly. First, a bit more background.

In response to KJ, several counterexamples have been offered. For instance, Baron Reed asks us to consider a subject who we'll refer to as Jack.<sup>13</sup> Jack is taking part of a social experiment and given the option of answering the question "When did Caesar die?"<sup>14</sup> There is no threat to his life if he chooses not to answer, but if he answers and does so incorrectly, they will give him a series of electric shocks. If he answers correctly, they will give him a jellybean. Clearly, the risk is not worth it, so he says, "I know the answer, but it isn't worth the risk." For most—including Fantl and McGrath—this seems like an acceptable response.<sup>15</sup> Yet, if KJ is right and if he knows when Caesar died, then he should be justified in answering. Nevertheless, it seems as if he should not answer from which it follows, according to KJ, that he does *not* know when Caesar died even though it seems quite appropriate for Jack to claim he does know.

Unfortunately, this sort of debate tends to get bogged down with arguing over which intuitions are correct. For the invariantist, the counter-example seems

---

<sup>13</sup> I've modified the case from the original where Reed has the subject state "I remember that Caesar was born in 100 BC, but I am not so sure of it that it is worth risking electrocution. Nevertheless, I quietly say to myself, 'I know it's 100 BC.'" I don't see why the hedging 'but I'm not so sure of it' is needed for reasons that will become clear in chapter 3. In short, the point I argue is that knowledge is not sufficient to justify action—it must also be 'worth it' as Reed slightly indicates here. The problem is that Reed seems to understand the 'worth of it' as being related to how sure he is when, in fact, it has to do with the ridiculously low payoff if he is successful.

<sup>14</sup> Reed (2012), 467.

<sup>15</sup> Fantl and McGrath (2012b), 485.

obvious enough, but Fantl and McGrath maintain that the fact that it would seem unjustified to act is best explained by a loss of knowledge because the practical implications have changed. In chapter 4, I will argue that Fantl and McGrath have not adequately addressed these counterexamples. For starters, while the traditional view implies that both are true, their view implies that Jack's claim is *false* and so that can't merely sidestep that issue by admitting that it seems correct. They need to explain why their view is a better explanation given that it implies one of the intuitions is false.

To avoid disputes over the right intuitions, F&M acknowledge that 'a principled argument' would go much further and so attempt to support KJ with more basic principles that they believe match up with ordinary ways in which we use knowledge and reasons in our deliberations about what to believe and how to act. The first regards the connection between knowledge and reasons—if you know  $p$ , then  $p$  can be a reason for further belief and action:

**Knowledge-Reasons (KR):** If you know that  $p$ , then  $p$  is warranted enough to be a reason you have [for  $\Phi$ -ing], for any  $\Phi$ .<sup>16</sup>

And the second states that if  $p$  is a reason you have for some further belief or action, then it can justify further belief and action:

**SafeReasons:** If  $p$  is a reason you have to  $\Phi$ , then  $p$  is warranted enough to justify you in  $\Phi$ -ing, for any  $\Phi$ .<sup>17</sup>

<sup>16</sup> Fantl and McGrath (2009), 69. In the original statement of KR, the principle ends with 'to  $\Phi$ , for any  $\Phi$ .' A change to "for  $\Phi$ -ing, for any  $\Phi$ " comes later in the discussion without mention. Perhaps, just an editorial mistake. I don't see any reason why this change would cause problems for their argument.

<sup>17</sup> It's easy to see that the consequent of KR is not the same as the antecedent of SafeReasons. Fantl and McGrath do address this and argue it is not an issue. I will assume they are correct as nothing seems to hinge on it. See Fantl and McGrath (2009) pages 82-83 for further discussion.

The first of these principles seems relatively uncontroversial and I won't address it any further. The main point of attack will be SafeReasons and, specifically, I will argue in chapter 4 that SafeReasons is susceptible to the counterexample Reed first used to attack KJ.

To conclude the chapter, we'll take a brief look at why KJ implies intellectualism is false, but first an important clarification. Each principle discussed employs the idea of being 'warranted' or 'warranted enough'. Fantl and McGrath explain this to mean that there is

“no weakness in your epistemic position with respect to  $p$  – no weaknesses, that is, in your standing on any truth-relevant dimensions with respect to  $p$  – stand in the way of  $p$  justifying you in having further beliefs. We will recruit the term 'warrant' to express this idea. Knowing that  $p$ , as we shall say, makes  $p$  warranted enough to justify you in believing any  $q$ .”<sup>18</sup>

In the paragraph following, the authors raise the question that most would ask in response to the last point: “Any  $q$ ?” While they agree that a 'relevant connection' would 'stand in the way' of  $p$  justifying believing that  $q$  in fact, a weakness in *epistemic position* does not. So, it is 'warranted enough' in that sense.

For instance, I know that I'm wearing a blue shirt ( $b$ ) and it would seem there is no 'truth-relevant weakness', ordinarily speaking, that stands in the way of it justifying further beliefs. Thus, on their account, we can say that  $b$  is 'warranted enough to justify' my believing that Paris is the capital of France ( $c$ )

---

<sup>18</sup> Fantl and McGrath (2009), 64.

although it does not justify  $c$  in fact. On Fantl and McGrath's view, their use of 'warrant enough to justify' implies that a known proposition *could* justify believing any other proposition—and later adding any *action—in principle* because it is a well-justified belief itself.<sup>19</sup> But whether it justifies believing any specific proposition is an open question and dependent upon other factors—preferences, relevance, and so on. With this understanding in mind, let's discuss the implications it has for intellectualism.

### §1.8 KJ and Intellectualism

KJ acts as a *necessary* condition on knowledge.<sup>20</sup> So, unlike contextualism, practicalism is a full-fledged theory about knowledge and not only a semantic-level theory. As noted above, KJ implies that knowledge of a proposition  $p$  requires one can put that knowledge to use towards *any* further belief or action  $\Phi$ , but this does not imply that  $p$  will justify  $\Phi$  all things considered. In many cases, that a further action is not practically justified can be explained by other non-epistemic factors; for example, the belief may be irrelevant to the act—as it is in the example above—or it might be immoral.

In those instances,  $p$  would be 'warranted enough' to justify  $\Phi$ , but the act would be still practically irrational. For example, if I am sitting at the bar and see

---

<sup>19</sup> Looking at this idea from a different angle, the point of KJ seems to be something like the following: Consider any further belief or action of which  $p$  is at least a partial reason and justifiably so. If  $S$  *knows* that  $p$ , then it seems quite odd to criticize that further belief or action because  $S$  lacked justification for  $p$ .  $S$  may lack justification for believing that  $p$  *supported* the belief or action or there may be some other reason (practical or moral) for criticism, but it wouldn't be because their belief that  $p$  lacked sufficient justification.

<sup>20</sup> Put another way, KJ implies that knowing  $p$  is *sufficient* to justify acting on  $p$ . In chapter 3 I argue that this is false and, in many cases,, it isn't even necessary to know  $p$  to justify acting on  $p$ .

a beer sitting on the table before me, I know there is a beer next me. But let's say I am the designated driver for the evening. Clearly, I should not take a drink. Is the right explanation for this that there is a weakness in my justification for believing there is a beer sitting before me? Of course not—it is the fact I shouldn't drink because I am the designated driver. So, it isn't as if KJ implies that  $p$  must be warranted to justify  $\Phi$ -ing *no matter what*; that is, no matter what the non-epistemic factors might be (e.g. the moral permissibility of the action). As with any assessment of practical rationality, those factors must be satisfied as well.

The problematic case for intellectualism occurs when the epistemic factors between two agents are the same, and yet it only seems practically rational for one of the subjects to act upon  $p$ . If KJ is true, then one knows  $p$  and the other does not, even though they have the same evidence for  $p$ . To make their case against intellectualism, Fantl and McGrath offer an example that meets these conditions. The example begins with two people who've both planned trips to Foxboro. The first, Matt, is sight-seeing and prefers Foxboro, but wouldn't mind Providence either.<sup>21</sup> The express train would take him straight to Providence, but the local will stop in Foxboro first and then head on to Providence. So, either train really isn't an issue, but Matt is looking at the 201 and we are to assume that he knows it is the local (it is) and so he knows it will take him to his preferred destination of Foxboro. It seems perfectly reasonable for Matt to board the train without further

---

<sup>21</sup> Fantl and McGrath (2009), 84.

inquiry. To put it in the language of KJ, there is no weakness in Matt's justification for "the 201 is the local" to prevent it from justifying him to board the train.

On the other hand, Jeremy needs to attend a meeting in Foxboro. If he accidentally got on the express, then he would be late and most likely lose out on a huge and important sale. He is also looking at the 201 and has the same evidence as Matt, but it seems that he should be more careful –he ought to double check to make sure that it is the local before boarding. Again, in the language of KJ, it seems there is a weakness in Jeremy's justification which 'stands in the way' and so would not justify his boarding the train; i.e. the consequent of KJ is not satisfied. Since it would not justify him boarding the train, KJ implies that Jeremy doesn't know the 201 is the local.

The 'Train Case' attempts to provide an example much like the Bank Cases and Airport. As before, we have two subjects with the same evidence, two choices –one that maximizes their interests against one that does not, and only one subject is justified to act. For Fantl and McGrath, knowledge is restricted by what it can practically justify which implies that practical reasons will factor into our epistemic evaluations in certain cases. This also suggests that practical reasons can be cited to justify or deny knowledge. For example, in the Train Case, one could say to Jeremy "You don't know it is the local –there is simply too much at stake!" To my ears, this sounds quite odd. Perhaps it is true that Jeremy doesn't know. Even so, explaining his lack of knowledge by referencing practical reasons

seems to go against both a theoretical understanding of knowledge as well as common usage. Looking ahead, I will take up this point more fully in chapter 4.

## Chapter 2: Knowledge Attributions and Mistakes

### *Introduction*

In the current debate over knowledge attributions, both invariantists and epistemic contextualists agree that ordinary speakers tend to use the term 'know' appropriately. The disagreement begins when discussing the truth-values of these sorts of cases in which there is a shift from knowledge being attributed to a subject and then denied. Contextualists argue these changes occur due to a shift in the standards for uses of 'know'. On their view, both the attribution *and* denial are true despite an apparent conflict between them. Invariantists, however, deny that epistemic standards will vary, and they deny that context is relevant to the standards for knowledge *and* knowledge attributing sentences.

For the invariantist, explaining the responses to cases like Bank and Airport will require that at least one of the claims is false. Some argue this sort of explanations should be avoided because they imply we intuitively judge a false claim to be appropriate and even true. But invariantists are not alone as contextualism leads to similar issues. For instance, when looking at cases like Cohen's Airport (see 1.4), it certainly appears as if the knowledge denial and attribution are at odds since Smith's claim to know the plane will land in Chicago seems to contradict Mary's claim that he does not. Attempts to explain away this intuition also imply that we don't know how terms like 'know' function even if we happen to use them correctly. To support their view, the contextualist must



provide a plausible explanation for why the denial and attribution appear contradictory but are in fact not.

Putting this all together, we end up with are three intuitive claims: that Smith knows, that Mary is correct to deny he does, and that those two claims are contradictory. If they are in fact contradictory, then a case like Airport provides no evidence for contextualism since one of the two claims is false. To regain support from these cases, the contextualist will need to show that Smith and Mary's claims are not contradictory. That is easier said than done, and I argue that the prospects are dim.

Turning to defend the invariantist view, I argue that we have good reason to conclude either Smith or Mary's claim is false. To motivate this position, I discuss modified versions of the Airport case in which Mary either has or lacks support for her denial. If Mary's denial is well-supported, then we have reason to conclude that Mary's denial is true, and Smith's claim is false despite his initial claim and our intuitive response. Smith's mistake is obvious in this example since he is unaware of some key facts relevant to his situation. If Mary's denial not well-supported, then we have good reason to conclude that Smith's initial claim is true. In this case, the mistake occurs because Mary's denial rests on overestimating the probability of the alternative outcome.

If I am right, we don't have a plausible basis for rejecting the claim that Smith and Mary are contradicting each other. However, we do have a plausible basis for rejecting either Smith or Mary's claim depending on the support that

Mary has for her denial. Thus, we have good reason to hold onto the invariantist view.

### §2.1 Knowledge Attributions and Mistakes

Epistemic contextualism is the view that the truth-conditions for sentences which attribute or deny knowledge vary according to context-relative features.<sup>22</sup> In the first chapter we discussed some of the problems a contextualist faces when supporting their view with subject-based examples. For instance, it is difficult to devise a case that doesn't seem to imply a loss of belief on the part of the subject. If the truth of 'S does not know that *p*' can be explained by loss of belief, then the case provides no evidence for contextualism. As a result, the most widely held version of contextualism is *attributor* contextualism whereby the context of a third-party, i.e. the attributor, determines the relevant truth-conditions. A prominent example comes from Stewart Cohen's Airport case discussed in section 1.4. Here it is again for review:

**Airport:** 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 if anyone knows whether the flight makes any stops. Another passenger, Smith, replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "How reliable is that itinerary, anyway? It could contain a misprint. They could have changed the schedule since it was printed, etc.," Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline."<sup>23</sup>

<sup>22</sup> See DeRose (1999) for a discussion of the varieties of contextualism.

<sup>23</sup> Cohen (2000), 95.

Considering the case from Smith's view it seems appropriate (i.e. 'correct usage') to say he does know the flight will stop in Chicago. He has little at stake, itineraries are generally reliable and there is no reason to question this fact. For ease, we will take LOW to refer Smith's low-risk claim to know. However, it also seems appropriate for Mary and John to conclude that Smith does *not* know the flight will make a stop in Chicago. They have a lot at stake and they have also raised the possibility that the itinerary is not reliable. We will refer to Mary and John's denial as HIGH.

So, we have two claims that both seem appropriate, but that also seem to conflict. As with the Bank cases (see 1.3), contextualists argue that the change from attributing to denying knowledge is best explained in terms of a change in the truth-conditions for attributing knowledge. Because Mary and John have more at stake, the truth-conditions for knowledge are higher as compared to Smith.<sup>24</sup> This suggestion implies that, not only are these claims appropriate, they are also *true*.<sup>25</sup>

DeRose explains:

As we can see [from cases such as Bank and Airport] ... speakers do in fact use 'know(s)' in the way described [that is, in the way that contextualists contend are the standard responses to their cases] and appropriately so - they will in fact, and with apparent propriety, ascribe 'knowledge' in situations like [Bank-Low], yet will deny 'knowledge' when they find

<sup>24</sup> It may also be the result of the additional alternative considered by Mary and John. This raises the question of how that alternative would affect Smith. We'll discuss these issues more directly in chapter 3.

<sup>25</sup> For a fuller defense of this claim see, see chapter 5 of DeRose (2009) and also his (1999). Cohen discusses his reasons for concluding both are true in his (2000).

themselves in conversational circumstances like [Bank-High]. This supports the premises that both of the imagined claims are true...<sup>26</sup>

And elsewhere:

...if a serious assertion is appropriate, that strongly suggests the intuition that it is true, at least from the speaker's point of view. (And if the speaker isn't mistaken about any underlying facts, that the claim is simply true).<sup>27</sup>

I agree with DeRose that appropriateness is a good indication of truth. If it wasn't, we humans would be in a bad position, as far as language use is concerned, since appropriateness is something of a base-line indication that we understand the 'give-and-take' of ordinary conversation. To put that differently, if our intuitions about appropriate usage fail to indicate language use that is generally correct and meaningful, it is difficult to see what in fact would indicate correct usage. Thus, when usage seems appropriate, we should give it the benefit of the doubt and assume it is true unless we have good reason to say otherwise.

However, granting truth from appropriateness comes with a caveat—the claim cannot rest on a mistake—a point which DeRose recognizes in the parenthetical “if the speaker isn't *mistaken* about any underlying facts.”<sup>28</sup> If it turns out that one of the intuitive responses to the Airport cases does rest on a mistake, that would suggest one of the claims is false and it would vindicate invariantism. On the other hand, if there isn't a mistake, that would suggest both claims are true

---

<sup>26</sup> DeRose (2005), 173-174.

<sup>27</sup> Ibid, 175.

<sup>28</sup> Emphasis added.

and so vindicate the contextualist's claim that these cases provide good evidence for their view.

To be sure, the claims made in examples like Bank and Airport seem appropriate and invariantism does require that one is false. But DeRose assumes too much when he claims, "it's a bad strike" simply if it rules LOW or HIGH to be false.<sup>29</sup> Claiming one or the other to be false isn't a problem, in itself, but doing so *without identifying where a mistake has occurred* that raises an issue. So, to avoid a 'bad strike', the invariantist must show that a mistake has occurred and do so in a way that is charitable to our intuitions.

### ***§2.2 A Contradiction between the Cases?***

Bank and Airport provide some *prima facie* support for contextualism since both the attribution and denial seem correct. However, if we expand our discussion to one case further, we'll find an ordinary and intuitive example that provides support for invariantism. The ordinary case I have in mind is one in which an additional subject, Carl, suggests that Smith and Mary/John are contradicting each other. Upon overhearing all involved, Carl says, "Mary and John are contradicting Smith. Both can't be right." Carl's statement in this example – which I'll refer to as CON – seems appropriate; that is, it does seem as if the two parties are contradicting each other even though it appears both

---

<sup>29</sup> DeRose (2005), 173.

statements are appropriate and perhaps even seem to be true when considered separately.

Following DeRose's point above that appropriateness "strongly suggests the intuition... is true" and given the fact that Carl's statement is appropriate, we ought to accept that Carl's statement is, in fact, true. At least, we ought to accept it as true unless we can identify a mistake Carl has made. If we can't, then the cases are not immediately of help to the contextualist because the truth of Carl's statement implies that either LOW or HIGH is false. Denying that CON is true without good reason would result in the same 'bad strike' criticism for contextualists that was previously brought against invariantists. (To be sure, invariantists don't win by default as they will need to show that either LOW or HIGH rests on a mistake.) So, the task for contextualists is to show that Carl has made a mistake because their view implies that what he said is false.<sup>30</sup>

Granted that the appropriateness and possible truth of LOW and HIGH does raise the possibility that contextualism is the correct theory. However, once we take CON into consideration, it is clear we cannot isolate LOW and HIGH to attack the invariantist position since that would require assuming LOW and HIGH are true, and CON implies that is not the case. In short, some further support is

---

<sup>30</sup> What sort of mistake would result in an attribution being false? Of course, if the fact in question was false, it would rest upon a mistake (though not necessarily a mistake for which the subject or attributor is epistemically responsible). The subject or attributor may have a mistaken belief regarding, as DeRose puts it, "a relevant underlying matter of fact." (2005, 173) For example, they might possess some misleading evidence, wrongly believe some fact to support their claim that does not in fact support it, or be unaware of some important evidence against their claim.

needed that does not come from assuming LOW and HIGH are true. Moreover, since contextualism implies that CON is false, that suggests we fail to recognize an important aspect of the way we employ 'know'. Claims like LOW and HIGH ordinarily seem to represent competing claims, but if contextualists are right, that is not the case. So, in addition to showing that there is a mistake, an explanation as to why we think CON is appropriate and even true is required.

Contextualists have attempted to explain away this problem by suggesting we make similar mistakes in other aspects of our language use. Stewart Cohen points out that we often express sentences which seem to contradict at first glance, but if the issues are clarified we will recognize they do not and so the contradiction was only apparent.<sup>31</sup> For example, consider that two people might debate over whether a 6'1 male is tall and thus think there is an actual contradiction between "S is tall" and "S is not tall". However, once it is clarified that the first speaker means 'tall generally' and the later speaker means 'tall for an NBA player', what seemed to be a contradiction initially is quickly recognized to have been a simple misunderstanding. Cohen suggests that we make the same sort of mistake when thinking the knowledge denial and attribution in cases like Airport contradict. There only *seems* to be a conflict until we clarify that LOW is true relative to a 'low-risk standard' for knowledge and HIGH is true relative to a 'high-risk standard'.

---

<sup>31</sup> Cohen (2005), 'Contextualism Defended' and 'Contextualism Defended Some More'.

There are three problems with this response. The first, is that we *do not* find similarities in ordinary use between the apparent contradictions of 'tall' mentioned by Cohen and the sort of contradiction pointed out by Carl. Quite unlike the clarifications which occur in the 'tall' example, a contradiction between uses of 'know' *does not* seem to dissipate after explaining the apparent contradiction as simply a misunderstanding about what standards are at play. For example, it would be unordinary if Smith clarified to Mary and John, "Oh, I meant that I know relative to a low-standards context" and Mary and John replied with a similar remark about the high-standard context when denying that Smith knows. To add to that, it would also seem odd if they agreed that Smith truly knows *C* relative to his context, but not Mary and John's.

Finally, arguing that 'know' is analogous to 'tall' requires an entirely new indirect argument to show that there *is* in fact an analogy between 'know' and 'tall'. Recall that contextualists initially compared 'tall' and 'know' simply to clarify how 'know' is to be understood on their view. 'Know' seems to have some similarities to certain terms already acknowledged to be context-sensitive and that helps to clarify the contextualist position. Surely, that was helpful in explaining their view. However, the contextualist was not claiming 'tall' has properties *x*, *y*, and *z* and 'know' has properties *x* and *y*, therefore 'know' has property *z* as well. The point was simply that 'tall' exhibits certain properties which are context-dependent and *that is what they mean* by saying that 'know' is context-dependent; i.e. that 'know' is also an indexical term like 'tall'.



Using 'tall' as an example isn't an issue, of course, because it is only to clarify their position. However, once analogies to 'tall' are offered to explain away apparent contradictory uses of 'know', *we must assume that 'know' is context-dependent* for the analogy to work. At the outset, the only reason we have to conclude 'know' is an indexical term comes from cases like Bank and Airport. Since the truth-values of these claims are up for debate, they provide little assistance in showing that 'know' is an indexical term. So, the contextualist first needs support for the claim that 'know' is indexical independently. Only *then* could an analogy to tall be made to explain away the problem with CON such that the explanation applied to apparent contradictory uses of 'tall' will also apply to 'know'. Absent this step, the contextualist can, at best, offer the analogy to 'tall' as an example of how we *could* explain apparently contradictory uses of 'know' if 'know' is an indexical term like 'tall'.

This leads us to the second problem with Cohen's argument, and the contextualist view in general. Because CON raises the possibility that either LOW or HIGH is false, the contextualist cannot offer the original cases in support without addressing the problem raised by CON. So, Cohen must show that 'know' is context-sensitive *without* the support of cases like Airport since. Given any attempt to use LOW and HIGH to support contextualism, the invariantist can simply point to the equally intuitive CON in response. (To be sure, LOW and HIGH together raise the possibility that CON is false. So, invariantists have their

own concerns.) So, the cases don't provide uncontested support for contextualism until CON is addressed; that is, until a mistake in CON has been identified.<sup>32</sup>

Can contextualists respond by putting forward an indirect argument? That is an option, of course, but it is important to note this move would open the discussion to further debates over the significant differences between uncontested context-sensitive terms and 'know' – in fact, DeRose specifically avoided offering an indirect argument for this very reason:<sup>33</sup>

Note that this argument from ordinary usage [i.e. contextualist cases] is not an indirect argument that takes as its premise that some other term, like 'tall', is context-sensitive, and then argues that because 'know(s)' is so similar to 'tall', 'know(s)' too is context-sensitive. Such an indirect argument would be very insecure, in my opinion, because, while there are very important similarities between the behaviors of 'tall' and 'know', there are also many important differences.<sup>34</sup>

The main difference between 'tall' and 'know' – which DeRose points out shortly following this passage – is that the former is an adjective and the latter a verb. In addition to this point, Jason Stanley discusses a number of further differences in his paper "On the Case for Contextualism". For instance, he points out that 'know' doesn't allow for modifiers in the same way as standard context-sensitive terms.

---

<sup>32</sup> One might respond that invariantists have the same issue; that is, they cannot use CON to argue against LOW or HIGH. That thought is entirely correct, and it must be clear that I am not arguing CON rules out the possibility that both LOW and HIGH are true. At this point, CON has merely raised that possibility and, as a result, shown that contextualists have a difficult issue of their own to address that is on par with the invariantists concern with LOW and HIGH. Further argument, independent of relying on the cases, is needed to support *either* view. This point will be taken up in section 2.3 and those following.

<sup>33</sup> DeRose points out one significant difference: 'know' is a verb and 'tall' is an adjective.

<sup>34</sup> DeRose (2009), 68.

Phrases such as “very tall” are quite natural whereas “very knows” is not. ‘Know’ also resists ‘comparative modifications’. Whereas ‘Sam is taller than Stan’ is an ordinary formulation, neither “Sam is knower than Stan” nor “Sam knower than Stan” are found in ordinary use.<sup>35</sup> If the contextualist is to make an indirect argument, then these dissimilarities between ‘know’ and words like ‘tall’ need to be explained. Without first explaining why these dissimilarities are not an issue, we have little reason to think the analogy between contradictory uses of ‘tall’ and ‘know’ will be successful.<sup>36</sup>

Certainly, there is room left to discuss an indirect argument and my point isn’t that an indirect argument cannot be successful. Rather, the point is that there are good reasons to think its success is unlikely because of the significant dissimilarities between context-sensitive terms and ‘know’. Even if an indirect argument does work in some capacity, the fact is that applying the same explanation of apparent contradictory uses of ‘know’ is simply not intuitive as it is with apparently contradictory uses of ‘tall’. So, the contextualist still needs to explain why we fail to understand the context sensitivity of ‘know’ or recognize

---

<sup>35</sup> We do use some phrases such as “Sam knows more than Stan”, but this sentence commonly refers to having more knowledge of independent facts rather than ‘knowing more’ than another about a particular fact. Even so, there isn’t anything odd about saying that someone knows something better than another. The invariantists is not at all committed to holding that every knower knows to the same degree. The position is simply that the baseline for knowledge is invariant for every possible (and actual) knower.

<sup>36</sup> While ‘know’ doesn’t seem open to these modifications, ‘confidence’ is compatible with them. This comparison might be an interesting route to pursue by those invariantists who argue that there is a loss of confidence in Bank-High.

apparent contradictory uses of the term between the sort of uses found in cases like Bank and Airport.

Moving forward, we have two favorable reasons to explore an explanation of the cases which is consistent with the invariantist view. First, there is the fact that LOW and HIGH are intuitively contradictory. Contextualists have taken cases like Bank and Airport to support their view, but the natural addition of CON shows that the dialectical situation is not so much in favor of contextualism as they have taken it to be. I don't mean to imply that invariantism is better off in light of this issue, but, at the very least, contextualists and invariantists are on even footing. Second, contextualists have yet to offer a plausible reason to conclude CON is false. These two points are enough to quell the thought that invariantists are in a worse position with cases like Airport. The fact of the matter is that CON is just as, if not more, intuitive than LOW and HIGH. The task here is, of course, to show that invariantism is the better view. With that in mind, the following sections will focus on the possibility that either LOW or HIGH rests on a mistake.

### ***§2.3: The General Strategy and A Contextualist Rebuttal.***

Consider again the events in Airport-High. Mary raises the alternative possibility that Smith's itinerary is not as reliable as he takes it to be. Given the high stakes and the fact the alternative could be true, the denial looks to be appropriate. So, what could be the mistake in these cases? The point I wish to argue is that, even though the alternative raised by Mary is *possible*, she and John

*overestimate* the probability of this possibility. If correct, we have identified a mistake in HIGH which indicates that HIGH is false, and LOW is true.

Once the alternate possibility *is* supported with good reasons, we then have good reason to deny that Smith knows and so conclude that LOW is false. So, based on the reasons supporting the alternate possibility we have one of two options: if the alternative possibility lacks support, HIGH is false, but if it has good support, LOW is false. If correct, this result would vindicate the invariantist position and it would do so intuitively because we would have good reason to conclude one of the claims rests on a mistake.

To support this conclusion, the strategy is as follows: first, we'll look at two modified versions of the Airport case. In one case, Mary offers an unlikely possibility – the itinerary is not reliable because aliens may have attacked Chicago and now the plane has been rerouted. Even though Mary raises an alternative possibility *and* the stakes are high, the right conclusion is that Smith knows the plane will land in Chicago despite these concerns and Mary and John's claim that he does not. In the second case, Mary suggests that the itinerary is not reliable because the Airline has been experiencing issues with its printer software the last few weeks. In this instance, it seems right to say Mary and John's denial that Smith knows is true because they have offered sufficient reason to question the itinerary. However, once we become aware of the printer issue, it also seems right to say that we made a mistake by first attributing knowledge to Smith. Looking at all three cases together, I will argue that the original Airport case is more like the first

case in that Mary and John lack sufficient support for the possibility that they raise. So, it is of no threat to Smith's claim to know *in either context*.

Before diving into the main argument, I must briefly address one lingering issue. Contextualists often argue that we should keep the cases separate and, in some sense, isolate our intuitions for the individual cases. In a footnote, DeRose writes:

Of course, we may begin to doubt the intuitions above [regarding the Bank case] when we consider them together, wondering whether the claim to know in the first case and the admission that I don't know in the second can really both be true. But when the cases are considered individually, the intuitions are quite strong, and, in any case, the linguistic behavior displayed in the cases quite clearly does accurately reflect how 'I know'/'I don't know' is in fact used.<sup>37</sup>

DeRose's point seems to be that, upon first considering each case individually, we have good reason to accept our intuitions because of their strength in conjunction with ordinary use. So, I might be accused of simply engaging with such comparisons and thus ignoring these concerns and ruling out contextualism from the start by assuming that there is an unresolvable clash between LOW and HIGH.

In response to this charge, I will make two points. First, we must keep in mind the intuitiveness of CON. Though the contextualist argues we shouldn't draw comparisons, the intuitive appropriateness of CON implies that is exactly what we should do, and it is something which is a part of our ordinary use of 'know'. Without showing that CON is false, the contextualist has given us no

---

<sup>37</sup> DeRose (2009), p. 55 fn. 7.

reason to ignore the intuitiveness of this claim and so conclude that comparing LOW and HIGH is legitimate.

Second, it would be a misunderstanding to think I am comparing various claims in a way that implicitly rules out contextualism. The reason is that my conclusion is not based on comparing the cases, but rather, by simply evaluating the reasons that support LOW and HIGH *individually*. To be sure, our assessment of one claim will have implications for the other, but the argument does not begin and proceed by comparing these cases and assuming one of them must be wrong. Looking just at high, for instance, the assessment will be entirely in terms of Mary and John's belief concerning the alternative possibility and *their reasons supporting that belief*.

Finally, as much as the contextualist cases accurately reflect linguistic behavior of 'know', surely *reasoning about and assessing* such claims *individually* is equally reflective of this behavior. In fact, this is precisely what occurs in the high stakes cases since the attributors raise and consider further possibilities. Since the contextualist cases already engage with this sort of ordinary reasoning, it would be inconsistent for them to argue that additional cases are not admissible into the discussion.<sup>38</sup> Even more so if the additional data ends up providing a more

---

<sup>38</sup> This last thought also addresses a point DeRose often raises about the 'best cases' — cases in which there is little to no debate and the separate contexts don't converge in some sense. Surely, those are the best cases *for contextualists*, but I see no independent reason to accept that the debate between contextualists and invariantists ought to be limited to only those cases.

intuitive understanding of the original cases and ordinary uses of 'know', in general.

#### §2.4: *The Mistake in Contextualist Cases.*

Let's turn to the first variation of the Airport example in which Mary and John raise an unlikely possibility:

**Alien-Airport:** Two alien enthusiasts, 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 if anyone knows whether the flight makes any stops. A passenger, Smith, replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "It is possible that aliens have attacked Chicago – if that is the case, the plane will need to reroute to somewhere else. How do we know that hasn't happened? We should go check to make sure." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary.<sup>39</sup>

This example is a bit extreme, but it is important that we have a case in which the knowledge denial is obviously false. (At least, in ordinary circumstances it would be obviously false.) Mary and John think that the possibility of the alien attack is enough to conclude Smith doesn't know *C* (land in Chicago) because *A* (alien attack) undermines Smith's evidence *I* (the itinerary) for *C*. But *A* is wholly ineffective as a defeater of Smith's evidence.<sup>40</sup> So, despite raising the alternative possibility *A*, Smith knows that *C*.

<sup>39</sup> As noted, this is a modified version of Cohen's original case from his article "Contextualism and Skepticism" (2000).

<sup>40</sup> The skeptic may want to push back here and argue that *A* is enough to undermine *I*. Something should be said about hyperbolic skeptical worries, but we must remember here that the focus is on ordinary, everyday knowledge attributions. In such situations, hyperbolic concerns are cast aside as irrelevant.



This seems right, but why isn't the alternative enough to undermine *I*? To put that differently, if Mary and John's knowledge denial is false in Alien-Airport, upon what mistake does their false claim rest? Two issues come to mind. First, she has offered no reason to think an alien attack will occur or has occurred in this instance. Instead, she simply mentions the possibility. Second, given our background reasons, including what we know about our little part of the universe, that we have little reason to think aliens exist or that they have visited earth, we have good reason to reject the idea that such an event is likely to occur.<sup>41</sup> As a result, it stands to reason that Smith's itinerary retains its status as being reliable and Smith knows that *C* even in Mary and John's context. So, *A* fails to undermine Smith's evidence *I* because *A* itself lacks sufficient epistemic support and, furthermore, we have sufficient reason to reject *A*.

I take it that this explanation of Alien-Airport will satisfy most readers, but there is an underlying point that will become more important once we turn to the original Airport example. Clearly, Mary and John's belief that *A* is mistaken because *A* is very unlikely. However, their belief that *A* is possible, in the strictest sense, is surely true. So, it might look as if denying that Smith knows *C* isn't based

---

<sup>41</sup> Throughout this chapter and chapter 3, I've assumed an objective view of the relevant probabilities. I take these probabilities to be based on the background evidence relevant to the issue at hand (e.g. the probability of an alien attack) and any relevant evidence directly stated in the examples. To be sure, there will be disagreements over how probable an event might be. However, in such cases, that merely implies a disagreement over the subject's epistemic position and so it is not directly relevant to this discussion. An important goal of these examples is to invoke probabilities which are relatively uncontroversial and, for a case like Alien-Airport, the significant lack of any evidence supports the claim that the probability is very low. Since it is low, a belief that it is probable will be false. This sort of reasoning seems to be on par with our everyday assessments of evidence and so relevant to this debate.

on a false belief. However, the knowledge denial isn't simply based on the idea that the alien attack is *possible*, but it seems that they also believe the alien attack has a sufficiently high probability. In other words, Mary and John are using 'possible' in the sense that is often used in ordinary conversation. So, they believe it is a likely alternative to C.

If we consider the alternative – that is, if we assume they believe A has a very low probability – it would be odd for them to give it any attention and then deny that Smith knows. For instance, assume that there is a one in a million chance of the alien attack and Mary states something like the following:

There is a one in a million chance that Aliens have attacked Chicago. How can we be sure that hasn't happened or won't happen? Smith's itinerary doesn't rule that out. So, he doesn't know the plane will land in Chicago.

Though she acknowledges the low probability of the alien attack, Mary's claim in this passage is disingenuous. In everyday situations – even risky ones – such incredibly poor odds are not something we sincerely consider. So, it isn't just *possible* alternatives that are of concern, but alternatives for which we have reason to think the probability of the alternative is sufficient to raise concern that the itinerary is not reliable.

The other option is that Mary and John have greatly overestimated the probability of the alien attack and so believe that it is likely enough to warrant sincere consideration. If that is the case, then Mary and John do hold a false belief. This explanation is much more plausible compared to believing that it is probable.

If correct, then the knowledge denial in Alien-Airport rests upon the mistaken belief that *A* is sufficiently likely to present a legitimate alternative to *C*.

So, we have one case in which the attributor, Mary, denies that Smith knows *C*, but the denial is false, and it rests upon a mistake (i.e. a false belief). Now we turn to a case in which the alternative comes with good epistemic support and the attributor does not make a mistake:

**Printer-Airport:** 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 if anyone knows whether the flight makes any stops. A passenger named Smith replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "You know, I just read in the newspaper yesterday that this airline has recently made a lot of mistakes on their itineraries. Something about their new printing software that is causing problems and they are trying to work out the bugs. So, I'm not so sure we can take that itinerary as reliable." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline.

The first thing to notice about Airline Printer is that Mary and John provide some reason to support the possibility that the itinerary is not reliable. Of course, they do so in Airline Alien as well. However, in this instance Mary and John's reasons support rejecting *I* since the airline has had trouble with their printing software lately. Thus, their worry that the itinerary contains an error is reasonable and, given the strong reasons cited against it, denying that Smith knows is both appropriate and true even though it is in fact true the plane will land in Chicago. Yet, if Smith's claim to know is false, that implies he has made a mistake which needs to be identified. Of course, the mistake is not direct in the sense that he has

reasoned incorrectly. Rather, he is simply ignorant of the fact that his itinerary is not reliable.

This leads to a concern about the effect that *unknown defeaters* can have on knowledge. We'll notice that in every case discussed so far – the original, alien, and printer cases – Smith is ignorant of the further alternatives that are raised. In some cases, those 'defeaters' undermine his knowledge (e.g. the printer case), but in others not (the original and alien cases). Yet, in each case the evidence that Smith has for believing the plane will land in Chicago is the same. Smith is unaware of any differences, of course, which raises two concerns: First, why is his evidence not sufficient for knowledge in Printer-Airport and, second, doesn't this imply that he is just lucky in Original-Airport?

The crucial difference between these cases has to do with Smith's background evidence. For Alien-Airport, Smith already has background information that will override the proposed alternative, but for Printer-Airport he does not. So, although he is unaware of the alternative possibilities, it is the information that he possesses that determines whether he does in fact know. In the next section, we'll explore extensions of each case to look at this issue more in depth and to provide support for concluding that the knowledge-denial (HIGH) in the original cases rests on a mistake.

### ***§2.5: Rebutting and Conceding.***

The two cases discussed in the previous section will act as the basis for arguing the knowledge denial in the original airport case rests on a mistake. To

strengthen that conclusion, I propose two further tests that will help determine whether a mistake has occurred. The main idea is that we consider an extended case in which Smith offers a rebuttal to Mary and John's knowledge denial that does not rest upon Smith acquiring any more evidence than he already has. Smith can defend his knowledge claim against Mary and John by turning to his background evidence. Importantly, this test is simply a reflection of the way that we ordinarily reason when an assertion is made, and reasons are offered to counter that claim. Beginning with Alien-Airport, something like the following dialogue might occur:

**Alien-Rebuttal:**

**Mary:** I heard you say that you know the plane will land in Chicago, but did you consider that aliens might have attacked Chicago and so prevent the plane from landing there. Your itinerary can't rule that out.

**Smith:** That seems far-fetched. What gives you the idea it might happen?

**Mary:** Alien attacks are possible and that is something we should be careful not to ignore.

**Smith:** If you don't have any reason to think an alien attack will happen, we don't have any reason to conclude the itinerary is not reliable. So, I have no reason to question the reliability of the itinerary and will confidently restate I know the plane will land in Chicago.

Smith's response to Mary seems appropriate. Simply by pointing out the mistake in her reasoning, and by explaining why it's still reasonable for him to believe the itinerary is reliable, he has upset her attempt to show that he lacks knowledge. In particular, notice that Smith offered no further evidence to support his own

position. This is significant because further evidence would alter the case in a way that unfairly biases it towards invariantism by adding to Smith's evidence.

For instance, if Smith checked a few news outlets or called a friend in Chicago to confirm there was no attack, his epistemic position would be stronger than in the original case. But since the original cases are importantly based on Smith having the same epistemic position, the contextualist could simply point out that there was a change which implies the rebuttal case is irrelevant. So, it is important that further evidence – at least, on the part of Smith – is kept out of the discussion.

For Alien-Airport, Smith's background information is the primary reason that he avoids losing knowledge. However, a subject need not always possess background evidence to prevent a defeater from undermining knowledge.<sup>42</sup> Keith Lehrer and Thomas Paxson's Grabit-Case is one example.<sup>43</sup> In that case, Mrs. Black witnessed her student, Tom Grabit, while he was stealing a book. Elsewhere, Tom's mentally ill mother heard about this and stated that Tom's identical twin brother, who doesn't exist, committed the crime. Of course, Mrs. Black is unaware of this and she lacks any background evidence to address the defeater; for instance, that Tom's mother has mental health issues or that Tom has mentioned being an only child on several occasions. Intuitively, however, she knows it was Tom who stole the book. But it seems that she would believe Mrs. Grabit's claim if she became aware of it and didn't know about her mental health. In effect, Mrs. Black is *lucky* to have avoided this misleading defeater.

---

<sup>42</sup> Thanks to Bruce Russell for making this point and bringing the Grabit-Case to my attention.

<sup>43</sup> Lehrer & Paxson (1969)

For this sort of case, Mylan Engel Jr's distinction between evidential and veritic luck is helpful.<sup>44</sup> Evidential luck occurs when there is some amount of 'luck' that occurred in one's acquisition of good evidence and perhaps, as we see with Mrs. Black, their avoidance of bad evidence. This luck does not undermine knowledge, however, because, lucky or not, that evidence reliably distinguishes between  $p$  and  $not-p$ .<sup>45</sup> In the case of Mrs. Black, her first-hand experience of seeing him steal the book, in conjunction with any background evidence, is reliable.

On the other hand, veritic luck occurs when one's evidence doesn't reliably distinguish between something being true and its being false. Turning to an extended, rebuttal-version of Printer-Airport, we can see that the same sort of response given in Alien-Rebuttal would not be successful. Unlike Smith's response in Alien-Rebuttal, he doesn't have any evidence to defeat the alternative possibility discussed by Mary. Since he lacks further supporting evidence, he has little recourse but to straightforwardly deny that Mary's evidence for the printer issue is relevant or that it is enough to suggest the itinerary is not reliable:

---

<sup>44</sup> Engel Jr. (1992)

<sup>45</sup> Perhaps, someone can get *too* lucky in terms of having the right evidence. Even so, some amount of evidential luck is permissible. And, more to the point, whatever issue this raises for my view, it would seem to raise the same issue for contextualism as well. So, I don't see a need to fully address this here. Nevertheless, it does seem to me that this essentially comes down to the reliability of one's environment. Tom's mother isn't enough to upset that reliability because it is a single misleading defeater that is not even salient to Mrs. Black. However, if Mrs. Black lived in an environment with an extraordinarily high number of twins, it greatly raises the possibility that Tom does have a twin. In that case, it might be enough to upset her knowing it was Tom who stole the book.



**Printer-Rebuttal:**

**Mary:** I heard you say that you know the plane will land in Chicago, but did you know that the airline has had issues with their printer software over the past few weeks? Are you sure that your itinerary doesn't have an error?

**Smith:** That seems far-fetched. What gives you the idea that might happen?

**Mary:** I heard it on the news this morning and there is also an article in the New York Times. As far as I know, the airline has yet to say it has been fixed.

**Smith:** I don't think that gives us any reason to worry the itinerary is not reliable. So, I don't see any reason to conclude that I don't know the plane will land in Chicago.

In this instance, Smith's response is false. The issue with the printers is a good reason to conclude that Smith's itinerary is not reliable. So, it does not provide good evidence that the plane will land in Chicago because the probability that the itinerary contains a mistake is too high.<sup>46</sup> As a result, the itinerary doesn't stand as a reliable piece of evidence that will support Smith's claim to know the plane will land in Chicago.

Smith's true belief is *veritically* lucky in this case because his true belief is not the result of having good evidence, but because he just happened to 'get it right' – he's just lucky that *his* itinerary is correct despite the fact that itineraries *in general* are unreliable. Based on this bad evidence, it is also likely that his belief could have turned out false. Of course, Smith could offer further evidence to reestablish his claim to know. But that supports the general point – the issue with

---

<sup>46</sup> The use of 'good evidence' here is admittedly vague, but as we will see, my argument does not rest on a precise definition. Rather, it is driven by our intuitive assessments of these cases and those we will discuss moving forward.

the printer software undermines the reliability of the itinerary and so further evidence is required. And this is true whether we are discussing Smith's claim to know or Mary and John's denial.

This last point underlies a critical aspect of the printer case – whether the issue is salient to Mary or Smith is *irrelevant* to whether he knows the plane will land in Chicago. The fact of the matter is that he doesn't know because his evidence is not reliable, and it would still be unreliable *even if he and Mary were entirely unaware of the printer issues*. The reason is that the issue with the printer creates an environment in which itineraries are not a reliable source of evidence and it doesn't get to stand as good evidence simply because the problem is not salient. I'll return to this point in chapter 3.

However, salience can affect a subject in other ways. For example, let's say that Mary reads a newspaper headline about printer issues with the airline, but the newspaper has incorrectly reported the problem to be with JetBlue when the issue is affecting Southwest. Mary believes the report and she has good, although misleading, evidence to justify her belief. Smith is unaware of this. However, were Mary to make this issue salient to him, he would lack sufficient evidence to override the misleading evidence. In short, this false news report would *defeat* his

original belief because he lacks any evidence to suggest that the newspaper has made such an error.<sup>47</sup>

To be sure, there is evidence that would defeat the false news report – i.e. ‘defeat the defeater’ – but Smith is not aware of this either. So, despite the fact the news report is false, he would no longer know that the plane will land Chicago if the news report were salient to him. However, despite his lack of evidence to counter Mary’s misleading evidence, he knows because his evidence is in fact reliable.<sup>48</sup> This seems correct even though he is in some sense lucky to have avoided Mary’s misleading evidence; i.e. he is *evidentially* lucky.

So, when rebutting the alien possibility, Smith’s knowledge claim is true, but when rebutting the printer issue, it is false. What happens if Mary concedes to Smith? For instance, what if she responds: “I see what you’re saying Smith. I guess you do know the plane will land in Chicago.” Applied to both cases, the obvious answer is that her claim is true in the alien case, but false in the printer case. Here is a ‘concession’ version of the latter:

**Printer-Concession:**

**Mary:** I heard you say that you know the plane will land in Chicago, but did you know that the airline has had issues with their printer software over the past few weeks? Are you sure that your itinerary doesn’t have an error?

---

<sup>47</sup> I use the term ‘defeater’ here with some qualification. It seems a defeater in the sense that Smith would be justified to believe that the news report is true and unjustified to believe otherwise were the news report salient to him. So, the sense in which it is a defeater is *subjective* because this requirement ‘kicks in’ only when the issue is salient to the subject. However, it is not a defeater in an objective sense. In the objective sense, salience is irrelevant to whether the defeater must be addressed to have knowledge. An example of this occurs in the printer case.

<sup>48</sup> Thanks to Bruce Russell who pointed out these issues.

**Smith:** That seems far-fetched. What gives you the idea that might happen?

**Mary:** I heard it on the news this morning and there is an article in the New York Times. The airline has yet to say it has been fixed.

**Smith:** I don't think that gives us any reason to worry that the itinerary is not reliable, and I will confidently restate I know the plane will land in Chicago.

**Mary:** That's a good point -- maybe the printer issue has been fixed. I guess you do know that the plane will land in Chicago.

The evidence is against Smith in this case, but Mary has conceded to him anyway and that seems incorrect. Notice that she attempts to justify the response with an unsubstantiated claim—*maybe* the printer issue has been fixed. Interestingly, this is the same issue that arises in the alien case since she suggests an alternative possibility for which she has little to no evidence. On the other hand, she has very good reason to think that the printer issue *has not* been fixed. So, again, if Mary concedes in the Alien case, her concession seems appropriate and true, but when doing so in the printer case it is false.<sup>49</sup>

### ***§2.6 Applying Rebuttal and Concession to the Original Case***

In the previous section, we discussed two types of cases—one in which Smith offers a rebuttal to Mary and John and a second in which Mary concedes to Smith's rebuttal. Below is a table to organize our thinking on these cases so far. Each 'T' and 'F' refers to whether the knowledge claim/denial in the given case is intuitively true or false. For example, there is a 'T' in the Alien-Airport/Low box

---

<sup>49</sup> This example points to an underlying issue raised by Printer-Concession—whether Mary raises the possibility is irrelevant to whether Smith knows. In other words, if there is a significant possibility that the itinerary contains an error, then the itinerary is simply not reliable. I'll return to this point in chapter 3.

below and an 'F' in Alien-Airport/High. In the first case, Smith claimed to know the plane will land in Chicago and that claim is intuitively true. In Alien-Airport/High, Mary and John's denial that Smith knows is intuitively false.

(Left Blank)

Table 2: Comparison of Alien-Airport and Printer-Airport

	LOW	REBUTTAL	CONCESSION	HIGH
<b>Alien-Airport</b>	T	T	T	F
<b>Printer-Airport</b>	F	F	F	T

Table 2 covers the four cases we have discussed so far in this chapter. LOW and HIGH are Smith's original claim to know *C* and Mary and John's denial, respectively. For Alien-Airport, it seems correct to say that Smith knows *C* even though Mary and John raise the possibility of an alien attack and Smith hasn't yet addressed this issue. However, in Printer-Airport it seems correct to say that he doesn't know *C*. The second column, Rebuttal, indicates whether it would be correct to say Smith knows that *C* after he offers a rebuttal to Mary and John's denial that he knows. For Alien-Airport, it seems correct to say that he knows there is no alien attack. In that instance, simply pointing out that Mary and John don't

have any good reason to believe *A* is sufficient to undercut their claim. That is not the case in Printer-Airport. If Smith stubbornly claims to know, that claim is false.

Putting this altogether, we get the following general point: when an attributor lacks good reason to deny that subject knows, the subject's rebuttal successfully defends against the attributor's denial – assuming he has background reasons to support the rebuttal – and the attributor's concession that the subject does know (or 'probably knows') is correct. I take these points to be a commonplace since we typically reject a claim when the person is unable to provide sufficient reason to accept it.

On the other hand, when the attributor *has* good reason for the denial, the type of responses given in Rebuttal and Concession don't work. Again, this seems to follow the ordinary way in which we reason on a daily basis. When good reasons are offered to suggest that a subject's evidence is unreliable, we typically reject the subject's claim to know unless better reasons are offered in support of their belief. Smith's response does not work in Printer-Rebuttal because the issues with the printer, which Mary points out, imply that his evidence is not reliable. To regain knowledge, Smith needs further evidence to support his belief that the plane will land in Chicago or that the itinerary is reliable.

Turning to the original case, the same ideas apply. If we find Smith's rebuttal and Mary's concession for the original case to be appropriate, we have good reason to think that this is because Mary's denial rests on a mistake. (And if we do not, then we have reason to conclude that Smith's original claim is false,

and it rests on a mistake.) These results alone won't get us all the way to concluding a mistake occurs, of course. That is an issue we'll need to address later. But to begin, we will analyze the original Airport case in the same way that we have analyzed these modified versions. Here is the original case once more:

**Airport:** 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 if anyone knows whether the flight makes any stops. Another passenger, Smith, replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "How reliable is that itinerary, anyway? It could contain a misprint. They could have changed the schedule since it was printed, etc.," Mary and John agree that Smith doesn't really know that the plane will stop in Chicago on the basis of the itinerary. They decide to check with the airline.<sup>50</sup>

As we see in the original case, the issue that Mary and John raise is the possibility the flight changed or the itinerary itself contains an error, either of which implies the itinerary is unreliable. Based on these considerations, they deny that Smith knows the plane will land in Chicago and that seems like a reasonable and appropriate claim for them to make. As discussed earlier and in chapter 1, I accept the denial is appropriate. The main issue is whether the denial – i.e. HIGH – is *true*.

To assess whether HIGH in the original case is true, it will be instructive to see how the alternative possibility that Mary and John raise in the original case turns out when looking at the rebuttal and concession cases. Here is a version of the original case with a dialogue similar to the other rebuttal examples:

---

<sup>50</sup> Cohen (2000), 95.

**Original-Rebuttal:**

**Mary:** I heard you say that you know the plane will land in Chicago, but did you consider the possibility that the itinerary contains an error, or the flight plan was changed?

**Smith:** No, I didn't – did you hear something?

**Mary:** Not specifically, but it's possible that something has happened. Flight plans do change from time to time, among other mistakes, and so I don't think you can conclude that you know without ruling out that possibility.

**Smith:** Of course, it is a possibility and I agree that such things happen, but as far as I know they are quite rare. And since you don't have any reason to think that it has actually happened to this flight, why do I need to prove that it won't or hasn't? It's not typical and aside from the mere possibility, I have good reason to think the itinerary is reliable. So, it seems safe to say that I know the plane will land in Chicago.<sup>51</sup>

Smith's rebuttal here seems on point. He's offered good reason to reject Mary's claim by showing that she lacks support for the possibility that an error of some kind has occurred. I'll say a bit more about this shortly, but for the moment we'll assume Smith successfully defends his knowledge claim in this dialogue.

Moving to the second case, Mary concludes the discussion by conceding the Smith does know:

**Original-Concession:**

**Mary:** That's a good point—I don't have a reason to think something has happened. I guess you do know it will land in Chicago.

---

<sup>51</sup> If we took Smith's response a bit further, he might point out that there are an infinite number of possibilities of which he has no reason to believe actually obtain. Without some reason to think that one (or more) of them do, it doesn't seem like he has to rule them out.



For my part, Mary's concession seems appropriate, but I can see how someone might think her statement, "I guess you do know" is forced. Perhaps, that is because attributing knowledge in this instance isn't typical or because Mary is hedging somewhat. Even so, this is no problem for the invariantist, in part, because Mary's 'hedging attribution' isn't necessary. What matters is that it is reasonable for her to *not* deny Smith knows even if she doesn't attribute knowledge to him. Consider two other responses that seem reasonable for Mary:

- You're probably right the plane will land in Chicago, but I'm a bit anxious and so I just want to double-check.
- I see your point and I guess I don't have much reason to think some error or problem has occurred. However, I'd like to go ask the attendant anyway just to make sure.

Again, these responses seem appropriate and the implicit knowledge claims ("You're probably right") seem true. Mary might refrain from attributing knowledge for any number of reasons and certainly we can understand that she and John are being cautious given their situation. Yet, being cautious doesn't say anything about whether LOW is true or false.

If attributing knowledge to Smith or, at least, not denying he knows seems appropriate, then the idea that her original denial is true is called into question. It is also reasonable to think we are just following Mary's reasoning rather than tracking variable truth-conditions for 'know'. When Mary reasons that Smith doesn't know, the claim seems appropriate, but when she reasons that he does (or

probably does) that also seems appropriate.<sup>52</sup> Whether or not those claims are *true* is a separate question and we should be wary of answering based on appropriateness alone.

So, we have one case in which Mary and John deny Smith knows and another in which they do not. Unless the contextualist can show that Mary's concession is mistaken, he will have to account for the difference between HIGH and her concession by arguing there has been a change in context. Such a move should raise an eyebrow since there has been no change in the practical issues and no further evidence or alternatives have been introduced. Smith has only pointed out that they lack reason(s) to believe the flight plan has changed and that was sufficient to override it. In short, arguing that the context has changed looks like an *ad hoc* move to avoid the problem at hand. So, I think it's safe to say that Mary's concession is reasonable and appropriate.

### ***§2.7 The Error in Airport-Original***

Assuming that our discussion of Original-Rebuttal and Concession is correct, we can start filling out the table below with Airport-Original added into the mix. In this case, Smith's original attribution (LOW) seems correct. His claim to know also seems true when offering a rebuttal to Mary and John. Finally, it

---

<sup>52</sup> The fact both responses seem appropriate raises a further issue for contextualism—people tend to *defer* to the speaker in these cases which is sometimes referred to as the 'agreement bias'. In a series of experiments, John Turri (2017) showed that people tend to agree with the speaker's claims. So, whether Mary and John attribute or deny knowledge, people tend to agree with whatever they claim. This makes sense because most are not actively ready to reject a claim unless something 'sticks out' as it does in the alien case.

seems true when Mary concedes to Smith—even if she hedges somewhat in doing so. So, we have the following table, with the most important box still missing an answer:

*Table 3: Comparing each of the three airport cases*

	LOW	REBUTTAL	CONCESSION	HIGH
Airport-Alien	T	T	T	F
Airport-Printer	F	F	F	T
Airport-Original	T	T	T	?

Immediately, we can see that there is more similarity between Airport-Alien and Airport-Original. In both cases, Smith knows when the possibility is not salient, but also when he stands-in to rebut the possibility suggested by Mary and John. To be sure, some important similarities are shown in the chart above, but these similarities are not the crucial. Rather, my argument is that the best explanation of all cases considered is that HIGH in the original case is false because Mary and John's claim rests on the same type of mistake that occurred in the alien case.

Notice that Mary lacks good reason to support the suggested alternative in both the alien case and the original case. In both cases, her claim is based on it being possible, but she provided no evidence to suggest it was at all likely. Because she lacks good reason to support her claim, we have good reason to think the rebuttal will work and the concession will seem appropriate which is what we find when applying these cases to the original airport case as we discussed in the

previous section. The best explanation of these results is that HIGH in the original case is false.

How might the contextualist respond? First, one might balk at concluding the original case is like the alien case because the possibility discussed in the latter is somewhat ridiculous. Alien attacks don't happen, but flight plans *do* sometimes change, and so it is safe to say that flight plan errors are much more likely than alien attacks. If that is the case, how can it be helpful to compare them?

Focusing on the difference in probability misses the point of the argument. I have not argued, for instance, that the probabilities in the cases are similar, therefore we should conclude they are alike in other respects. Rather, I have argued that there is a similarity in the *type* of mistake(s) that occurs in alien case which also appears to have occurred in the original case; namely, the false belief that an alternative possibility is probable. This similarity in type is shown via the rebuttal and concession cases. Because of those similarities and because the high-stakes knowledge denial is false in the alien case, we have reason to conclude HIGH is false in the original case. So, if one wishes to attack this claim, they must address the rebuttal and concession cases.

Even if it is argued the probabilities are too different, contextualists must explain why this difference is enough to suggest that HIGH is not false and do so by offering some account of what sort of probabilities are high enough to matter. This leads us to our second point since it could also be argued that Mary and John do have good reasons to support their claim that Smith doesn't know. Of course,

those reasons need to be identified since merely asserting they have good enough reason to deny Smith knows is not sufficient.<sup>53</sup> For instance, imagine that Mary replies with something like the following:

**Mary:** I'm not sure that this flight plan was changed, but I did hear recently that this airline is notorious for changing flight plans at the last moment.

If Mary replied in this way, then Smith's rebuttal would be far less convincing, and HIGH would seem to have the upper hand. Given that she has offered some reason to believe the itinerary has an error, we also have some reason to conclude Smith's initial claim to know (LOW) is false. So, when Mary and John *have* good reason to support their claim, then we've just created another scenario similar to Airport-Printer. When the attributor *lacks* good reason, and we are careful to analyze a bit beyond our initial reaction, we end up with an outcome similar to Airport-Alien.

But perhaps this is too hasty – do Mary and John actually need to support their claim with further reason? Isn't the mere suggestion of the possibility enough? For example, let's consider a version of the case from Jessica Brown:

**Mary:** ...The itinerary might be mistaken; perhaps there's been a late change of schedule. We better check.

**John:** Well, it's not very likely that the itinerary's incorrect. We've certainly no reason to think there's been a late change. How often have you come across an incorrect itinerary? Of course, the plane's going to Chicago.

---

<sup>53</sup> I qualify this claim to distinguish it from their having good reasons to seek further evidence. As discussed earlier, this is perfectly acceptable response for practical purposes. However, the implication of this argument is that it isn't good enough for epistemic (or semantic) purposes.

**Mary:** OK, I know it's unlikely that the itinerary's wrong. I believe the plane's going to Chicago too. But that's not the point. Just imagine what would happen if we took the flight and it doesn't stop at Chicago. We'd miss the meeting. [Smith] can't rule out the possibility of an itinerary error [and so he doesn't] know the plane will stop. [We better go and check.]<sup>54</sup>

In this version of the case Mary offers no further reason(s) to support the possibility of an error. However, she maintains that the *mere possibility* alone is sufficient to support the claim that Smith doesn't know. If Mary's claim is true, then my view is called into question because Mary doesn't seem to require any further reason to support the possibility the itinerary has an error.

Is this a convincing case? I think not for at least two reasons. First, we should compare this case to the case in which Mary concedes ("You're probably right"), but states that she would like to double-check anyway. For my part, that is a much more intuitive response than stubbornly holding that Smith doesn't know. It is quite appropriate for Mary to remain concerned and desire to double-check. However, an individual's desire to double-check and be more secure in their thinking does not have broader implications about the truth of attributing or denying knowledge of another individual.

Second, in ordinary circumstances it is disingenuous for someone to acknowledge that an alternative possibility is unlikely and yet take this unlikely alternative to provide sufficient reason to reject a possibility for which we have good reason to believe that it is true. In ordinary contexts, at least, mere

---

<sup>54</sup> Brown (2005), 147. I've slightly modified the end, so it is applicable here. Also, note that Brown offers this case as a criticism of Kent Bach's 'belief removal' interpretation of these cases. See Bach (2006).

possibilities don't destroy knowledge even if they might grab our attention at first as Mary's suggestion may have done in the original Airport case. Knowledge is truly affected when there are possibilities that we have good reason to think are likely to occur. This seems to be the way that people ordinarily reason and draw conclusions about what they or someone else knows. If this is correct, then we have good reason to think that HIGH is false and so reject the contextualist view.

## *2.8 Conclusion*

In this chapter, I have argued that the debate between contextualism and invariantism comes down to three different intuitions: an intuition that *S* knows *p*, an intuition that *S* doesn't know *p*, and an intuition that the first two claims contradict. To defend their view, contextualists must show that the third intuition is false because it rests on a mistake. At this point, the prospects for this are weak. On the other hand, the invariantist must show that one of the first two claims rests on a mistake. I have argued that the high-stakes claims in the original Bank and Airport cases do rest on a mistake because the alternative possibility offered lacks sufficient support to suggest that the subject's original evidence is unreliable.

Nevertheless, the contextualist still has some resources to defend their view. For example, they might argue that changes in what is at stake is the primary issue 'motivating' our intuition to deny in the high-stakes cases. This would suggest that, when stakes are high, the truth-conditions to 'know' increase. In the

next chapter, I will address the issue of changing stakes in addition to returning to questions about the role of salience.



## Chapter 3: Stakes and Intuition

### *Introduction*

In Chapter 2, I argued that Mary and John's knowledge denial in Cohen's airport case rests on a mistaken belief. Since the denial rests on a mistake, we have good reason to conclude that Mary and John's denial is false. Yet, even if the denial rests on a mistaken belief, the contextualist can reply that an increase in what is at stake *directly* motivates Mary and John's denial (and our intuition to accept the denial).<sup>55</sup> If a mere change in stakes motivates the denial, then we might have some evidence for contextualism. In this chapter, I examine the possibility that stakes motivate our intuitions. My aim is to show that there is little reason to conclude an increase in what is at stake affects the truth-value of a knowledge claim.

To support this conclusion, I first look at a set of cases from Jonathan Schaffer who offers a modified set of the original bank cases, but without the bias of having a subject or attributor claim to know/not know. Schaffer's cases suggest that the introduction of a salient alternative (e.g., that the itinerary is mistaken), rather than a change in what is at stake, is what prompts a change in intuition. This might seem to imply that *mere* salience is enough to undermine a knowledge claim and to be sure, salience can affect whether a subject has knowledge. However, the

---

<sup>55</sup> An increase in stakes may affect claims to know and intuitions about such claims in ways that are not controversial. For example, an increase in stakes may result in losing confidence that one's belief is correct. (This is why DeRose argues for the controversial claim that the subjects in his cases *do not* lose confidence. For more on this see Bach, 2005.) Of interest for us, is whether stakes may *directly* affect our intuitions by which I mean that no other factor acts as an intermediary (e.g. loss of confidence/loss of belief that *p* is true) between the increase in stakes and change in intuition.

mere fact that an alternative is salient is not enough. As I argued in chapter 2, the alternative in question must provide good reasons to call into the question the reliability of one's evidence.<sup>56</sup> On this view salience is irrelevant and only the probability of *not-p* determines whether it must be ruled out to know that *p*.<sup>57</sup> So, while Schaffer's general point about the irrelevance of stakes is correct, he did not see the further implication about the type of alternative that is required to undermine knowledge.

To conclude this chapter, I will discuss two possible responses by the contextualist. The first is that salience will sometimes determine if an alternative must be ruled out and the second is that a sufficiently high change in stakes will alter the truth-conditions for 'know'. I argue that neither response is sufficient.

### *§3.1 Schaffer's Unbiased Minimal Pairs: It's Not About the Stakes*

In his article "The Irrelevance of the Subject", Jonathan Schaffer argues that stakes are not the motivating factor behind our intuitions. His argument begins by criticizing the original bank cases for being 'biased pairs' because they contain multiple important differences (change in belief, change in evidence, etc.).<sup>58</sup> As a

---

<sup>56</sup> This is, of course, just one way the justification for some proposition can be defeated. It can also be undermined if there is more evidence for *not-p*. However, we won't discuss that style of defeater here. See Casullo (2003) Chapter 3 for discussion on these types of defeaters.

<sup>57</sup> There are instances in which the subject *believes* that an alternative represents a good defeater when it in fact does not. But these cases result from a mistake in reasoning by the subject, and so they are not relevant to our discussion.

<sup>58</sup> Schaffer is specifically arguing against the Subject-Sensitive Invariantism of Jason Stanley and John Hawthorne. Even so, the same point—that an increase in stakes is not the relevant issue—also applies here.

result, it is not clear which difference motivates the shift from intuiting that Smith knows to intuiting that he does not.<sup>59</sup> Some of these differences lead to problems for first-person cases that we discussed in Chapter 1 since a change in belief, for example, can account for the difference in our intuition without any need for the further contextualist explanation. This difference isn't present in third-person cases. However, there are two additional differences that become problematic even for third-person cases.

In the original Airport cases, for instance, the high stakes context includes at least two major differences from the low stakes context. The first difference is that Mary offers an alternative possibility in which the itinerary is not reliable. Since this alternative is not salient to Smith (or us) when looking solely at Smith's context, it is quite possible that the salient alternative is the motivating factor behind accepting Mary's denial.<sup>60</sup> This suggests that, if the alternative is also salient in Smith's context, we might be prone to denying that he knows. To test this, we will have to look at a set of cases in which the alternative is present in both and a set of cases in which it is salient in neither. The second difference is the explicit change in what is at stake. For Smith there is very little at stake, but for Mary and John there is a lot at stake, and it is possible that a change in what is at

---

<sup>59</sup> As I understand the contextualist view, this isn't so much a 'change in intuition', but two entirely distinct intuitions given that we are actually considering two different propositions.

<sup>60</sup> That is not to say they are *good* reasons. As I argued in the previous chapter, Mary and John's denial that Smith knows rests on a mistake.

stake prompts a shift in our intuitions because we inherently take it that better evidence is required when more is at stake.

The support that contextualism receives from these cases depends greatly on which difference motivates the change in intuition. If the change is the result of salient alternatives, the invariantist can plausibly argue that the epistemic position of Smith—or, at least, our evaluation of it—has changed. However, contextualism is favored if the change is the result of an increase in what is at stake.

To avoid these issues and get a helpful assessment of what is motivating our intuitions, Jonathan Schaffer has modified the original cases to produce sets of “unbiased minimal pairs.” They are unbiased in the sense that the subject and/or attributor do not make claims about knowing or not knowing. Instead, whether the subject seems to know is left entirely up to the reader.<sup>61</sup> And they are minimal in the sense that there is only a single difference between each case.<sup>62</sup> They only differ in terms of what is at stake. So, as mentioned above, there will *not* be a difference in terms of salient alternatives. In each set, the alternative will be salient in both or salient in neither.

---

<sup>61</sup> Because Schaffer takes out any aspect of the subject attributing or denying knowledge to himself, these cases are also not first-person cases. So, they do not run into the same issues that we discussed in regard to DeRose’s original bank cases in Chapter 1.

<sup>62</sup> Schaffer (88-89).

Schaffer first offers a set of cases in which no additional possibility has been made salient in either case (e.g. the bank has changed its hours), but there is a difference in what is at stake:

**NOT-SALIENT-LOW:** On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. So, does Sam know that the bank will be open this Saturday?

**NOT-SALIENT-HIGH:** On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. So, does Sam know that the bank will be open this Saturday?<sup>63</sup>

Schaffer intuitively thinks that Sam knows in both cases, and I would say the same. There is the further question of whether Sam would be rational to act on this knowledge and those issues will be addressed in Chapter 4.<sup>64</sup> Setting that aside for now, the

---

<sup>63</sup> Ibid. Though my intuition is that Sam knows in both cases, posing the question “Does he know?” has something of an inherent skeptical effect. I would imagine that is because questioning whether someone knows is often asked with an element of skepticism and especially in cases in which the person has already stated their reasons for belief. We don’t usually follow up a claim with “Do you know that such and such is the case?” unless we are challenging the claim in some sense (e.g. seeking for the evidence/reasons to be established, reiterated or supported, etc.). It seems to me that this might raise some methodological issues regarding the way in which surveys are conducted which purport to test our intuitions about knowledge attributions. See Turri 2016 for another discussion of possible methodological issues with these examples.

<sup>64</sup> In short, the argument is as follows. Some object that, if you know  $p$ , then it is rational to act on  $p$ . I argue in Chapter 4 that this claim is false. For one, it fails to recognize that you also *know* the bank is open now and the task will be completed if you don’t wait. What you *don’t know* is whether waiting till the next day will be successful—there are, in fact, a vast number of other factors than just the bank being open which take part in determining whether you’ll be successful if you wait till the next day. Since the added benefit of waiting till the next day is quite trivial, why take the risk?

main question here is whether a change in stakes – specifically, an increase in what is at stake – motivates our intuition to conclude that Sam doesn't know. If a mere difference in stakes directly motivates a change in intuition, we would expect to intuit that Sam doesn't know the bank will be open. At least as far as Schaffer and I are concerned, that intuition does not follow from a mere increase in stakes which suggests that stakes don't have the proposed effect. Without the possibility Sam is wrong, we don't think it is reasonable to conclude he doesn't know.

On the other hand, what if we compare the intuitive responses to a set of cases in which the possibility of error *is* salient? Here again Schaffer provides two minimal and unbiased cases in which the only difference is in terms of what is at stake for each subject. In contrast to the original bank cases and the two cases just discussed, *both* include a salient alternative:

**SALIENT-LOW:** On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, but he has no pressing need to deposit the check. He has little at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. But banks do change their hours, and Sam has not looked into this. So, does Sam know that the bank will be open this Saturday?

**SALIENT-HIGH:** On Friday afternoon, Sam is driving past the bank with his paycheck in his pocket. The lines are long. Sam would prefer to deposit his check before Monday, and indeed he has pressing financial obligations that require a deposit before Monday. His entire financial future is at stake. Sam remembers that the bank was open last Saturday, so he figures that the bank will be open this Saturday. He is right – the bank will be open. But banks do change their hours, and Sam has not looked into this. So, does Sam know that the bank will be open this Saturday?<sup>65</sup>

---

<sup>65</sup> Schaffer (89-90).

The second set of cases are word-for-word the same as the first set except for the brief addition of “But banks do change their hours, and Sam has not looked into this.” Again, you’ll also notice that Sam doesn’t make any claims about whether he knows since that would only serve to confuse what is motivating our intuition. By leaving out such claims, we avoid any ‘bias’ between the cases.

Schaffer intuits that Sam does *not* know in either of the salient cases. For these cases, my intuition differs from Schaffer as it seems to me that Sam still knows in both cases. The reason is that the salient alternative offered rests on the same sort of mistake as the alternative Mary proposes in the Airport case. For that case, I argued the alternatives discussed in the high-stakes scenario fail to stand up to scrutiny and this is why, rationally speaking, they can be ignored. Mary offers an alternative possibility that she takes to undermine Smith’s evidence (i.e. the itinerary). However, her alternative comes with no supporting evidence to suggest that the issues she has mentioned are common or that such an error has occurred.

In the same way, the alternative suggested in the Bank case—that banks sometimes change their hours—also rests on a mistake. To be sure, banks most likely change their hours from time to time, but we’ve been given no reason to think the bank has changed its hours in this instance or that this bank (or banks in general) changes its hours often enough to raise concern.

Even though I intuit differently than Schaffer, we agree in the most relevant and important sense in that *we do not intuit a difference between the low and high-*

*stakes cases*. For these cases to support contextualism, we should intuit differently in at least one set of the cases; that is, we should intuit Sam knows when the stakes are low, but not when the stakes are high. The fact we don't suggests that what is at stake is not the motivating force behind our intuition to attribute or deny knowledge or agree that an attribution or denial is true. So, the contextualist's prediction that we will is questionable at best.

The contextualist might argue that Schaffer's intuitions as well as my own are something of an anomaly but looking at some recent studies we don't seem to be alone. In these studies, a change in what is at stake doesn't appear to lead to a change in people's intuitions regarding what a subject knows. Schaffer offers support for this in his own study with Wesley Buckwalter.<sup>66</sup> Additional work from Adam Feltz and Chris Zarpentine indicates that individuals are not "sensitive to the practical facts of a subject's situation."<sup>67</sup> To be fair, a few studies seem to counter these conclusions.<sup>68</sup> So, I won't hang my hat on the studies that support

---

<sup>66</sup> Buckwalter and Schaffer (2015). Also see Buckwalter (2010) and Turri (2016). In the 2015, Buckwalter and Schaffer conclude that stakes do not motivate denials, but salience does. I agree that salience can have an effect. However, as I will discuss later, these examples (and the corresponding empirical studies) fail to recognize that salience is often *irrelevant*. It isn't *mere* salience that motivates a knowledge denial. Rather, it is a salient possibility that is accepted by the 'intuiter' as sufficiently likely. In example cases (e.g. the Alien case from Chapter 2), in which the salient possibility is sufficiently *unlikely*, the intuition is to maintain that the person knows.

<sup>67</sup> Feltz and Zarpentine (2010).

<sup>68</sup> See Pinillos (2011) and Sripada and Stanley (2012). Buckwalter and Schaffer (2015) offers a successful counter to both, in my opinion. One particular argument from Buckwalter and Schaffer is of interest. In one of their studies, they offered participants cases that switched out uses of 'know' for related terms such as 'belief' or 'confidence' and so on. Such changes seemed to have no effect on the results which suggests that knowledge is not of concern for the participants, but instead the more practical issue of what the subject should do. The term in use is at best a secondary issue.



the invariantist position. The point, however, is that neither can the contextualists hang their hat on the claim that the tendency to attribute or deny knowledge in ordinary situations is influenced by changes in what is at stake since there is good evidence to the contrary.

To sum up, the contextualist's initial reports are unreliable because they rest on cases that contain multiple differences as well as claims which are biased toward the contextualist position. Once the issues in these cases have been removed, the intuitive responses don't appear to support the contextualist's view.

### §3.2 *Salient Alternatives and Attributors*

Schaffer's revised cases suggest that only when there is a salient alternative to the original claim does a change in intuition result. (Also, recall that some empirical studies provide support to this view.) Underlying this argument, we'll find a further important detail – whether the alternative is salient to the attributor is largely *irrelevant* to what we intuit.

One point in favor of this is that the attributor's claim is sometimes counter-intuitive even though an alternative has been raised. We know this, for example, from the alien example in Chapter 2. Though Mary and John suggest the possibility of an alien attack, it doesn't motivate us to agree with Mary and John. Rather, it is correct to say Smith knows the plane will land in Chicago. So, it would be a mistake to conclude that mere salience is sufficient to motivate a denial; in other words, to merely suggest the *possibility* one is wrong in everyday use of 'know' isn't thought to be sufficient to undermine knowledge.

To be sure, a salient alternative can have that affect if we don't recognize it is unlikely. But the issue is easily explained since it is quite possible that a subject/*we believe* it is likely even though it is not. It seems plausible that this is what occurs with the Airport and Bank examples for those who intuit that high-stakes denials are correct.<sup>69</sup> If they believe it is likely  $p$  is false (as does the attributor), it is only rational to question the evidence for the original belief and so withhold believing that  $p$  is true. If that is the case, then "S doesn't know that  $p$ " is the rational conclusion.

We also know from the printer example in Chapter 2 that some alternatives must be ruled out whether they are salient to the attributor or not. The subject doesn't have easy access to knowledge simply because an alternative is *not* salient. Alternatives to  $p$  that are sufficiently likely *must* be addressed in order to truly say "S knows that  $p$ ." In Printer-Airport, Smith doesn't know about the issue with the printer software, but that is irrelevant to whether he knows. Given the high probability that the itinerary he's holding contains a mistake, it isn't a reliable piece of evidence. So, Smith doesn't know the plane will land in Chicago.

Applying this point to our intuitions about the software issue, the same results follow. Even though we are unaware of the issue when first considering his

---

<sup>69</sup> See section 2.4. The essential point is that, in ordinary situations, if we take some alternative  $q$  to undermine knowing that  $p$ , we don't concurrently believe that  $q$  is unlikely. That aside, it is of course possible that the alternatives *are* very likely. If that is the case, our intuition is correct, but, as I argue in Chapter 2, Smith's original claim to know is false. Not simply false from the attributor's context, but false independently of any context.

case, our original intuition was incorrect. Once we are informed of the problem with printer software, it is clear the itinerary is not reliable.<sup>70</sup> If we extended this to include the attributor, the same would be true. Were the attributor unaware of the printer issue and so attributes knowledge to Smith, the attributor would be wrong. So, being salient to the attributor is not a determining factor of whether something needs to be ruled out to correctly attribute knowledge.

Salience is relevant to our intuitions, but to motivate a change in intuition from “S knows that  $p$ ” to “S doesn’t know that  $p$ ”, the alternative must at least *appear* to be a sufficiently likely possibility that  $p$  is false. When it is quite low as in the Alien case, denying knowledge is not intuitive, but when it is sufficiently likely in the Printer case, then it is intuitive to deny knowledge. But if it is the probability of the alternative – whether real or only apparent – which drives our intuition, the use of the attributor in third person cases only serves to distract us from the central issue of whether there is sufficient reason to deny the subject’s claim to know.<sup>71</sup>

---

<sup>70</sup> Attributing knowledge in the alien case and denying it in the printer case seems correct *no matter what the stakes may be*. In the printer case someone wouldn’t, for instance, know the plane will land in Chicago if there were absolutely nothing at stake for them. The sufficiently high possibility the itinerary is in error seems to prevent knowledge no matter one’s circumstances and assuming their evidence has not changed. On the opposite end, no matter how high the stakes, it doesn’t seem at all correct to deny that someone knows because of the Alien possibility. The possibility of an alien attack is just too low to matter at all.

<sup>71</sup> At this point, it is natural to ask what factors determine whether an alternative to  $p$  is ‘sufficiently likely’ so that it poses a threat knowing that  $p$ . A full-blown account of ‘relevance’ goes beyond the purposes of this project. However, it does seem to me that certain ‘intuitive tests’ will help us determine whether an alternative must be considered and when it can be ignored which I discuss in the next section and elsewhere. For instance, one test I discuss is whether it holds up to the sort of scrutiny offered in the rebuttal case from Chapter 2. See section 2.4-2.6.


So, although the contextualist moved to third-person cases to avoid problems discussed in Chapter 1, they fare no better by doing so. The inclusion of an attributor is irrelevant to what we intuit and so those cases can be ignored.<sup>72</sup> Rather, we should focus on cases like Schaffer's which remove any bias or problematic differences. For the invariantist, this is good news because those cases appear to support their position.

### §3.3 Alternatives that 'Must' be considered.

As discussed in the previous section, it seems that some alternatives can be ignored outright – whether a subject knows is *not* dependent upon ruling out these alternatives. On the other hand, some alternatives must be addressed whether salient or not. Whether a subject knows *does* depend on ruling out these alternatives whether salient or not. I take it that these are uncontroversial ideas, at least, from an ordinary, everyday understanding of knowledge attributions. From the invariantist, *every* alternative will fall into one of these two categories. To illustrate, each box below represents one of the two categories, the arrow indicating an increasing probability and the space in-between, the 'probability threshold' that separates the two:

---

<sup>72</sup> John Turri (2016) points to a further issue—in many instances people simply defer to the speaker. Turri reasons that “they might assume that others are well positioned to report on their own mental states.” (142) I would further add that this seems a likely explanation in cases in which one is not acquainted with the relevant facts. For example, in Cohen's Airport, it seems perfectly reasonable that airlines do have late changes to flight plans from time to time. However, few people will specifically know how often that happens. So, without reasoning further—as Smith does in the rebuttal case—perhaps, we are prone to defer to Mary and John.




If this view is correct, the alternatives discussed in the original Bank and Airport cases (i.e. the bank may have changed its hours; the flight itinerary may contain an error) would fall into one of these two categories. Moreover, if correct, the contextualist would lose significant support for their view since the only determiner of whether an alternative must be ruled out is its probability – context is irrelevant.<sup>73</sup>

As a rebuttal to this view and the foregoing arguments, the contextualist might reply that the alternatives mentioned in the original Bank and Airport cases must be ruled specifically *because they are salient*. To be sure, if they were *not* salient, they wouldn't need to be ruled out. Yet, their probability is enough so that *if salient*, they must be ruled. This response follows from a central aspect of contextualism; that is, the idea that salience plays a primary role in what alternatives must be ruled out. To be sure, any alternative that must be ruled out will have to reach a minimal threshold and there is a higher-threshold that, if met, the possibility must be ruled out *even if not salient*. However, there is also an intermediate threshold in which the probability is high enough to require being ruled out *if salient* (e.g. changing bank hours), but not so high that it must be ruled out independently of

<sup>73</sup> In chapter 2, I argued that they fall into the first category because they rest on a mistake—the probabilities are low, but mistakenly thought to be high enough to warrant consideration.

being salient (e.g. printer issues – Printer-Airport). I'll refer to these as *intermediate alternatives*.

Below I've modified the illustration from above to coincide with the contextualist interpretation. Again, the arrows indicate the increasing probability of the alternative possibility and each threshold is represented by a separate box. As you can see, an additional box has been included in the middle to represent the proposed intermediate alternatives:



The most critical issue for the contextualist is whether or not there is an intermediate threshold which represents alternatives that must be ruled out if salient, but don't need to be ruled out if not salient. If so, the contextualist could argue that context—and specifically salience—is a determining factor of whether a knowledge attribution is true. The reason being that we must look at the context to see if the alternative in question is salient or not. For example, in the airport case the alternative suggested by Mary and John is not salient in Smith's context. So, his knowledge claim is true. But it is salient in Mary and John's context. So, their knowledge denial is true.

This interpretation also provides a contextualist-friendly response to Schaffer's unbiased minimal pairs. The contextualist can simply explain that

Schaffer's intuitions are correct because the alternatives raised are of the intermediate variety. So, Schaffer correctly intuits that both subjects *know* when the alternative possibility is *not* salient, and he correctly intuits that both subjects *do not know* when the alternative *is* salient.

For this view to be plausible, we should expect to find some instances of intermediate possibilities, and we can assume that the alternatives from the original Bank and Airport cases are paradigmatic. So, we would expect these alternatives to be of a kind which *must* be considered and ruled out—if they are *mentioned*—in order to appropriately and truly attribute knowledge to a subject. Again, these alternatives must be ruled out specifically because they raise a sufficiently high possibility that the original belief could be wrong. This interpretation implies that the evidence used in the case is not itself unreliable, but instead it is the salient possibility that is unreliable.

Take, for instance, a version of the original Airport case in which the itinerary is *not* questioned and instead is used to support Mary and John's belief that the plane will land in Chicago:

**Airport-Posted:** "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 if anyone knows whether the flight makes any stops. Another passenger, Smith, replies, "I know - I remember when I ordered the tickets that we first go to Chicago, then New York." It turns out that Mary and John have a very important business contact they have to meet at the Chicago airport. Mary says [to John], "That guy said he remembers seeing it on the flight plan when he bought his tickets, but what if he remembers it incorrectly and we are actually stopping over in Cincinnati? Or what if the flight plan has changed since then? We should find a flight itinerary to make sure he is right." Mary

and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to find an itinerary, which they do, and then conclude they know the plane will land in Chicago."

It seems reasonable for Mary and John to take the itinerary as good evidence in this version prompting us to wonder why it isn't good evidence in the original example. The difference cannot be a change in the possibility that there is an error – objectively, both cases are the same in that regard.

The contextualist might suggest that the itinerary stands as good evidence in one case because there is no salient reason that suggests the itinerary has an error, but it is not good evidence in the other because there the risk is higher and there is a salient reason. In contextualist terms, the truth-conditions change when there is a salient reason, and this is why the itinerary is good (i.e. sufficient) evidence in one, but not the other.

This view doesn't hold up if we look at further variations of the Airport case in which the possibility of error is salient but doesn't intuitively require being ruled out if the attributor ignores them for some reason. One such instance was already hinted at in section 2.5 when discussing the original airport case. There I pointed out that it would seem quite reasonable for Mary to concede Smith is probably correct even though she raised the possibility the itinerary could be wrong. She then suggests they double-check anyway and that seems like a prudent option for them. Even more important is the fact she *does not deny* that Smith knows, and this is reasonable despite the salient possibility of error. So, mere



salience does not explain why the itinerary is good enough evidence in one case, but not the other.

I think this case is enough to seriously question the idea that there are intermediate alternatives. Yet, the fact Mary doesn't *attribute* knowledge to Smith might leave an opening for the contextualist since they could argue that, at the very least, it would be false for Mary to attribute knowledge to Smith. In short, withholding a denial *is* compatible with the contextualist view, but attributing knowledge is not. Since we've only seen an example of the former, there is no evidence against intermediate alternatives.

Such a response wouldn't free them completely of the problem raised by Mary conceding that Smith "is probably right", but it would alleviate some of the sting. Nevertheless, we can push this issue further and look at a case in which Mary *does* attribute knowledge to Smith:

**Airport-NoDeny:** 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 if anyone knows whether the flight makes any stops. Another passenger, Smith, replies, "I do. I just looked at my flight itinerary and there is a stop in Chicago." It turns out that Mary and John have a very important business contact they need to meet at the Chicago airport. Mary says [to John], "How reliable is that itinerary, anyway? It could contain a misprint. They could have changed the schedule since it was printed or something else." Mary pauses and takes a deep breath. "I guess I'm just worried about the meeting. I'm

sure he is right—He knows the plane is going to Chicago." John agrees. Both take a seat and wait to board.<sup>74</sup>

The same error possibilities from the original case are included here, yet Mary's reaction is appropriate. I don't find her claim that "he is right" odd or unreasonable nor do I find their decision to take a seat and wait to board imprudent.<sup>75</sup> However, if the 'intermediate alternative' view is correct, Mary's response should be intuitively false or, at least, imprudent. Once the alternative is mentioned, it must be ruled out, but that is not the case when she concedes he is "probably right" nor when she states he knows the plane will land in Chicago. Mary's response in these cases make further sense if we take into account the conclusions of the previous chapter.<sup>76</sup>

If I am right that Mary's denial in the original case rests on a mistake, there is nothing odd about the fact she doesn't follow through with the denial in the concession case. And the same can be said about her decision not to deny in the

---

<sup>74</sup> In section 2, I argued that the inclusion of an attributor is essentially irrelevant to these cases. So, the use of an attributor here might appear disingenuous, and I would be inclined to agree if that were not the contextualist strategy. However, since contextualists have typically employed an attributor to support their view and the claim is that a salient alternative of sufficient probability must be ruled out, we should expect the standard intuition to disagree with Mary and John. Moreover, if we arrange the case to focus solely on the subject, i.e. Smith, so that he raises the same concerns as Mary, it again seems fine to me if he points out that he is simply worried and just takes a seat.

<sup>75</sup> I slightly waver about the 'imprudence' issue since it would be very easy for them to double-check. In such cases it may seem imprudent if you *don't* check given the opportunity. Even so, it is not obvious that 'prudentially justified double-checking' is somehow indicative of not knowing. We'll touch on these issues in chapter 4.

<sup>76</sup> This same point can be applied to the case in which she concedes "he is probably right, but I want to check anyway."

case just described.<sup>77</sup> So, as it stands, arguing that there are certain ‘intermediate alternatives’ is not a plausible route.

### §3.4 What About Very High Stakes?

In this last section we will look at one more attempt to find a case that provides good support for contextualists. In general, the best support for contextualism would come from a case in which the alternative is salient *in both the high and low stakes*, but it is only intuitive that the high stakes subject does not know. Schaffer’s unbiased minimal pairs suggest that such a result isn’t supported by standard intuitions. So, it looks like mere differences in what is at stake does not motivate our intuitions about knowledge attributions and denials.<sup>78</sup>

Yet, this doesn’t entirely close off the idea that stakes can independently motivate a knowledge denial. One way for the contextualist to respond is to suggest that we look at an extreme case in which the stakes are incredibly high and there is no salient alternative. If our intuition is to deny knowledge in such cases, then perhaps the issue with the original cases is simply that the stakes weren’t high enough. So, we need a case that includes a sufficiently high increase

---

<sup>77</sup> On the other hand, when the denial *does not* rest on a mistake, then an unwillingness to check further and/or deny that Smith knows *would* be troubling. For instance, in the printer case from 2.4, she has very good reason to question the reliability of the itinerary. So, it would be epistemically (and practically) irresponsible to ignore that issue.

<sup>78</sup> This same point can be applied to the case in which she concedes “he is probably right, but I want to check anyway.”

<sup>78</sup> On the other hand, when the denial *does not* rest on a mistake, then an unwillingness to check further and/or deny that Smith knows *would* be troubling. For instance, in the printer case from 2.4, she has very good reason to question the reliability of the itinerary. So, it would be epistemically (and practically) irresponsible to ignore that

in stakes. In this section, we will look at a case with very high stakes to assess the plausibility of this suggestion. Whether we intuit that the person knows or not, I will argue that, in either circumstance, the least plausible explanation is that the high stakes have motivated our intuitions.

To begin, consider the following bank case in which the stakes have been dramatically increased:

**Bank-Ransom:** While waiting in line at the bank, Keith receives a phone call that his child has been kidnapped. (Let's assume he has good evidence that this is true.) The kidnappers demand that he transfers money to an offshore account as ransom for the child. If he doesn't get the money deposited by 12:00pm on Saturday, then his child will be murdered. To have enough money to meet their demands, Keith needs to deposit the check he is holding. Again, his evidence that the bank will be open on Saturday is his previous visit a few weeks prior. Does Keith know if the bank will be open the following day?<sup>79</sup>

As with the previous bank examples, assume that Keith retains his belief the bank will be open on Saturday and that this belief is true—the bank will be open. For my part, Keith knows that the bank will be open on Saturday, but that does not imply he would be justified to wait for the following day. For the same reasons discussed above, it would be practically irrational—not to mention morally irresponsible—for Keith to wait and deposit the money at a later time.<sup>80</sup>

---

issue.

example is a modified version of a case from Bruce Russell (2005, 35).

<sup>80</sup> Consider if Keith decided to stop in the Dunkin' Donuts next door to the bank before heading over to the deposit the check. There is virtually no risk that the bank will close in the 5 minutes time it takes him to buy a coffee and donut. However, doing so is still irresponsible for several reasons. For my part, this suggests that the cases have very little to do with banks changing their hours. As mentioned, we'll discuss this more in depth in Chapter 4.

For the sake of discussion, however, let's assume that the standard response is to intuit Keith *doesn't* know. Would the incredibly high stakes explain this? There are two good reasons to conclude that high stakes are not the motivating factor.

First, there are countless everyday examples in which there is a lot at stake, but intuitively they do not affect a person's knowledge. Consider the claim, "I know that my furnace is functioning correctly." On rare occasions home furnaces do explode. So, it is possible that my furnace has a malfunction which will eventually lead to a large explosion and the death or serious injury of anyone currently in the home. A furnace that is malfunctioning and at risk of exploding represents an incredible cost, but would the risk alone motivate denying I know my furnace is functioning correctly *without* any further evidence or reason to do so? It seems fair to say that it does not—if it did, then I should be on the phone immediately to get someone over to my house and look at it. Even though the stakes are very high, it doesn't motivate denying that I know.

We can apply these same thoughts to many common, every day issues in which there is a lot at stake, but it is incredibly unlikely that a bad outcome will occur: driving a car, walking along a street, standing under a tree, swimming in a lake, etc. In each of these cases, it seems incorrect to deny that a person knows certain relevant propositions even though the stakes are high. If high stakes

independently motivated denials, however, we should expect the opposite when the higher stakes are salient.<sup>81</sup>

Even if we assume that it is intuitive to deny Keith knows in the ransom case the contextualist must also explain why it is *not* intuitive to deny that I know my furnace is functioning correctly. If high stakes aren't enough to deny in one case, but are high enough in another, what accounts for the difference? Since both the ransom and furnace cases have incredibly high stakes, it would be arbitrary to try and explain the difference of intuition in terms of a difference in what it as stake.

This leads to the second issue: the idea Keith should check further seems to assume that waiting till the following day to make the deposit is a reasonable option *as long as* Keith knows the bank will be open. In effect, the question "Does Keith know the bank will be open?" is presented in such a way so that one's response will either support or reject his waiting till the next day. Otherwise, what is the purpose of checking further? Yet, as the case is written, *in no way* does waiting till the next day seem at all reasonable. (I would argue that it is also unreasonable to do so in the original case—whatever the evidence for the bank being open the following day.) Even if Keith checks further and it is confirmed that the bank will be open at 9:00am, it would still be morally and practically irresponsible for him to wait for the following day.

---

<sup>81</sup> By 'independently', I mean that what is at stake cannot motivate a change in attributing or denying knowledge without the support of something else; namely, further evidence or reasons.

So, a case like Bank-Ransom is perplexing because it implies that our intuition about whether he *knows* will determine what he should *do*. But *no amount of evidence* the bank will be open on Saturday would justify waiting for the next day. The reason being that deciding to wait till the following day is determined by *multiple* factors and not simply whether the bank will be open. For example, waiting to go to the bank on Saturday should be *worth it* and, for this example, there is no obvious sense in which that is the case. I will take up this issue further in the next chapter where I argue that knowing *p* is not sufficient to justify acting on *p*. Rather, the act must also be worth it, and it is clear in Bank-Ransom that waiting for the next day is not.

## Chapter 4: Intellectualism and Practicalism

### *Introduction*

In the previous chapter I addressed the problem that contextualism raises for the invariantist component of the traditional view. Although contextualists and traditionalist differ in this regard, they do agree that only epistemic reasons—i.e. truth-conducive reasons—are relevant to a subject's justification for belief. This 'intellectualist' view has been opposed by those who argue that practical reasons are also relevant. Beginning with their 2002 article "Evidence, Pragmatics, and Justification," Jeremy Fantl and Matthew McGrath suggested that practical issues (e.g. the level of risk) impose a necessary condition on knowledge.<sup>82</sup> In low stakes cases, one has, in a sense, more practical justification for acting on  $p$  and so less epistemic justification is needed to know that  $p$ . But when the stakes are higher, one has less practical justification and so greater epistemic justification is required.

The authors continued to support and refine this idea in their book *Knowledge in an Uncertain World*. There they argue for the KJ principle discussed in chapter 1:

**KJ:** If you know that  $p$ , then  $p$  is warranted enough to justify you in  $\Phi$ -ing, for any  $\Phi$ .

---

<sup>82</sup> The list of opponents to intellectualism also includes Jason Stanley (2005) and John Hawthorne (2004). Fantl and McGrath have continued support of their original article in their book *Knowledge in an Uncertain World*.



If KJ is true, then invariantism fails – as it also does if contextualism is true – but the intellectualist thesis likewise fails because practical issues directly factor into whether one’s epistemic justification is sufficient for knowledge.

In this chapter, I will defend the intellectualist thesis in two ways. First, I will argue that the argument for KJ fails as a result of a variation on Reed’s counterexample that we discussed in 1.7. Fantl and McGrath argue that Baron Reed’s counterexample isn’t successful, but I will show that they have failed to recognize the force of the example. In particular, we will see that intellectualism is consistent with the intuitive response given by both Reed’s original example and Fantl and McGrath’s proposed response, whereas KJ is only consistent with the latter. Following this, I will argue that Reed’s counterexample also applies to the Safe Reasons principle at use in their argument. As a result, the principled argument does in fact fail.

Second, once the argument for KJ is out of the way, I will defend the subject’s claim to know in Reed’s counterexample by showing that his response and refusal to act is plausibly explained by the tendency of individuals to be risk averse. We’ll look at some cases to support this conclusion in 4.3. Moreover, once we better understand the subject’s responses in light of these cases, we’ll also find that interpreting these cases in terms of KJ has some very odd results. For example, KJ implies that subjects can acquire and lose knowledge solely because of an increase or decrease of one’s payoff from successfully acting on that knowledge.

These issues will be discussed in section 4.4. Given these principled and explanatory problems, we have good reasons to reject KJ.

#### *§4.1 The Explanatory Benefits of KJ*

As mentioned in the first chapter, one of the motivations for adopting the KJ principle is that it solves the threshold problem for fallibilism about knowledge. On the fallibilist view, one can know that  $p$  even if their justification for  $p$  does not guarantee that  $p$  is true. Fallibilism gets a lot of intuitive support from the fact that it seems right we know a lot. But despite that intuitive support, it is not at all clear what amount of justification is needed for knowledge. Clearly, ambiguous phrases such as “the weight of reasons” won’t settle the issue and it would be arbitrary, for instance, to declare a specific probability which must be met.

KJ solves this problem by setting the threshold to a level that warrants “putting knowledge to work.” In other words, if your justification for  $p$  is enough to warrant using that knowledge to justify a further action or belief, then you’ve hit the threshold. If it doesn’t put you into a position to justify further action or belief, then you haven’t hit the threshold.

In addition to solving the threshold problem, the authors mention several other factors they take to favor this principle. In particular, they suggest that KJ

Makes smooth sense of our habits of citing knowledge to criticize or justify action. We tend to think that, once we’ve settled a matter theoretically – once we can say, “Oh, o.k., at least we know *that*”; we can then count on it in action as well as a belief.”<sup>83</sup>

---

<sup>83</sup> Fantl & McGrath (2012), 444.

For example, it's natural to say something like "I know the dog will escape so I shut the gate" or "I stopped looking because I know I left my keys on the fridge." Without question these sorts of examples are commonplace in ordinary situations. In the same way, the lack of knowledge often explains why someone shouldn't do something, rationally speaking. For example, if someone claims, "I know the dog won't escape", but it doesn't seem rational for them to leave the gate open, a plausible explanation is that they really *don't* know the dog won't escape. To be sure, such examples don't get us all the way to KJ. However, they do force us to consider the fact that knowledge and action are closely related in ordinary use.

In addition to solving the threshold problem and capturing how we ordinarily use knowledge, KJ also provides a straightforward explanation of cases like Airport and Fantl and McGrath's own train case (See 1.8). In each case, the important difference is that it is intuitively reasonable for the low-stakes subject to ignore the chance their belief is in error, but for the high-stakes subject it is not.<sup>84</sup> In the train case, both believe that the 201 is the local and both have the same evidence for that belief. Since they both have the same evidence—and their evidence doesn't guarantee the 201 is the local—they also have the same weakness in their epistemic position. The difference is that, while that weakness isn't too much for the low-stakes subject, it is too much for the high-stakes subject. Since

---

<sup>84</sup> Fantl & McGrath (2009), 25ff.

the high-stakes subject should inquire further, it is reasonable to conclude that he doesn't know. Further inquiry is typically a good indication that more evidence is needed to know. That seems especially true in cases when inquiring further seems to be the only rational option.

The problem with KJ is that it implies practical factors directly affect whether you know. This point is not lost on Fantl and McGrath as even they recognize that denying intellectualism looks 'mad' on the face of it. However, they suggest that we need to contrast this apparent madness with the explanatory benefits of such a view. But even more, we should also be honest about the madness of fallibilism as David Lewis suggested in the following:

"If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possibilities of error. Even if you've numbed your ears, doesn't this overt, explicit fallibilism still sound wrong?'"<sup>85</sup>

The underlying point is that, if we are to be fallibilist, we should really have something to say about what amount of error in our justification is permissible. An explanation that makes it sound *not* so wrong and KJ offers a very clear way of doing this even if it does fly in the face of the traditional view. While that may be true, in the following sections I will argue that these explanatory benefits are not sufficient to override the problems with the view.

---

<sup>85</sup> Lewis (1996), 550.

#### §4.2 JB-Know and JB-Deny.

Near the end of chapter 1, we looked at Baron Reed's counterexample to KJ. In the example, Jack is involved in a contest and asked when Julius Caesar is born. If he answers correctly he gets a jelly bean and if not, he is given a painful electric shock. There is no penalty (or reward) if he doesn't answer. The table below lays out the possible outcomes:

Table 4: Reed's Caesar Case

	<b>Correct Response</b>	<b>Incorrect Response</b>
<b>Answer</b>	Jelly Bean	Electric Shocks
<b>No Answer</b>	None	None

Having considered his options and noting the severe consequences if he answers incorrectly or something else goes awry, he says to himself "I know he was born in 100 B.C., but [answering] isn't worth the risk."<sup>86</sup> The point of this case is to provide an example in which acting on  $p$  would not be justified, but the subject still knows that  $p$ . If it is, then we have a counterexample to KJ because Jack knows that  $p$ , but that knowledge would not justify acting on  $p$ . Let's refer to this version of the case as JB-Know.

Fantl and McGrath acknowledge that Jack's response doesn't "raise any eyebrows", but they also suggest that it wouldn't be out of the ordinary to respond with the following:

<sup>86</sup> As noted in chapter 1, I've changed Reed's original example.

“Do I really know that Caesar was born in 100 BC or am I just pretty confident about it? Well, I thought I knew this before, but after thinking about the risk, I guess I really don’t know after all. I better not answer.”<sup>87</sup>

Facing the same circumstances, this subject isn’t so sure that he knows when Caesar was born. There is simply too much at risk, so he decides not to answer. (Note that the outcomes are the same as listed in Table 4.) According to KJ, if the subject isn’t justified acting on  $p$ , the subject doesn’t know that  $p$ . Since that seems to be what happened in this version—Jack isn’t justified in answering, so he doesn’t know—it looks like we have an alternative to JB-Know that is consistent with KJ. Let’s refer to this version of the case as JB-Deny.

#### ***§4.2 Counterbalancing Intuitions?***

Fantl and McGrath take it that their version of the case (JB-Deny) neutralizes Reed’s original case and so frees them from the counterexample:

“The point is not that Reed’s case doesn’t provide evidence against KJ. It’s that we have two *counterbalancing* intuitions about the case: one is that an insistence on knowledge would be natural, and another that a denial of knowledge would be natural.”<sup>88</sup>

Looking at this response, the most important claim is that there are ‘counterbalancing intuitions’. What could they mean by this? What Fantl and

---

<sup>87</sup> Fantl and McGrath (2009), 62.

<sup>88</sup> Fantl and McGrath (2012b), 485. Italics added.

McGrath seem to be thinking is that Jack's different responses contradict each other and so one of the responses must be wrong. In other words, even though we think that Jack's claim to know in JB-Know and that his denial in JB-Deny are natural, they are incompatible with each other because the claim to know and the claim to deny can't both be true. Both seem right, but one of these intuitions must be false.

Interpreting the cases in this way suggests they are like other cases in which we appear to have contradictory intuitions such as the standard trolley and transplant scenarios. The point of these cases is, in part, to answer the question "Are consequences or deontological considerations more foundational to moral justification?" In Trolley, we are asked if it would be morally acceptable (and perhaps obligatory) to switch the track so that a train bearing down on five people will change to a track in which only 1 person would be killed. Most think this is the right thing to do, or at least that it is morally permissible. However, when considering a case in which someone could save five people by harvesting the organs of one innocent person – and thereby killing him – most intuit that it would be the *wrong* thing to do and so not morally permissible. For the Trolley case, it looks as if the consequences justify the action, but for Transplant some deontological consideration puts a stop to the consequentialist justification.<sup>89</sup> On

---

<sup>89</sup> Granted, these two scenarios are usually given as a starting point for more complex variations and contradictory intuitions may not arise for some until those further variations are considered. Yet, the overall point remains.

both examples, we'll save five lives while sacrificing the life of one person and it is hard to say why sacrificing the one is acceptable in Trolley, but not Transplant.

The principle issue is that neither view seems capable of accommodating both intuitions – consequentialism can only accommodate the intuition in Trolley and a pure deontological view can only accommodate the intuition in Transplant. So, in terms of the competing views, our intuitions are 'counter-balancing' because both views receive equal support and opposition from the cases. In short, we are at an impasse between consequentialism and the deontological position.

Do we find something similar when comparing the invariantist and practicalist interpretation of JB-Know and JB-Deny? As with trolley and transplant, we have two cases that result in two intuitions that, at first glance, appear to oppose one another. And, as Fantl and McGrath seem to think, perhaps neither the invariantist nor the practicalist view can explain both. This is correct regarding the practicalist view because it implies that JB-Deny is true while JB-Know is false. Since we do not intuit that JB-Know is incorrect – as even Fantl and McGrath agree – their view is not consistent with one of our intuitions. If at least one intuition is also not compatible with intellectualism, that would suggest our intuitions are contradictory, and, like the consequentialist v. deontology debate, we are at an impasse.

The problem is that the intellectualist view *does not* imply that either intuition is incorrect. To explain the difference between the cases, the intellectualist can simply point to Jack's own denial in JB-Deny as the relevant difference. In



effect, Jack's denial implies that he has either had a change of evidence (CEP) and/or change of belief (COB) that results in his denying that he knows.<sup>90</sup> Since intellectualism provides a plausible and validating interpretation of *both* intuitions, there is no impasse between the views and, thus, talk of 'counterbalancing' intuitions is incorrect.

So, rather than two 'counterbalancing intuitions', the fact of the matter is that we have one view, intellectualism, which can accommodate both intuitions and another view, practicalism, which must deny Jack's claim to know in JB-Know even though it seems natural. On the intellectualist view, *both* responses found in JB-Know and JB-Deny are true. More to the point, there are no contradictory or 'counterbalancing intuitions' since we can offer CEP or COB as an explanation. When someone claims to know and has good reasons in support of that claim, then it is natural to think their claim is true. And when someone denies they know because they lack sufficient reason – and we have no reason to suspect that they are being insincere, etc. – then it is quite natural to think that claim is true. This is precisely what we see in these examples above, and it is consistent with the intellectualist view that only epistemic reasons are relevant to knowledge.

---

<sup>90</sup> This suggests that JB-Know and JB-Deny are another set of cases with problems similar to the Bank Cases discussed in chapter 1. As we discussed, those cases provide little to no support for the contextualist because the subject specifically claims that he knows when the stakes are low, but not when the stakes are high (See sections 1.3-1.5). Because of this difference, no further explanation beyond the change in belief is required.

On the other hand, Fantl and McGrath's view implies that Jack's claim in JB-Know is false because his belief that Caesar was born in 100 BC isn't warranted enough to justify answering the question. According to KJ, this means that he doesn't know Caesar's birth date which conflicts with our intuition that his claim to know is natural and appropriate. To be sure, intuiting that the claim is natural and appropriate doesn't mean it is in fact true. So, Fantl and McGrath can avoid the conflict with our intuitions if they can show that Jack's claim is false. One way to accomplish this is by showing that KJ is true via the 'principled argument' they offer in support of KJ. If their principled argument is sound, then KJ is true and Reed's counterexample fails because the subject's claim to know when Caesar was born is false—even if we take it to be true. In the next section, I will argue that the principled argument is not sound because the second premise, SafeReasons, is also susceptible to Reed's counterexample.

### *§4.3 Against the Principled Argument*

In the previous sections we have discussed Reed's counterexample to KJ which implies that KJ is false. Since we intuit JB-Know to be correct, we have some reason to think KJ is false. However, the problem with Reed's approach is that it has left the principled argument for KJ untouched, and Fantl and McGrath can use it to support KJ. So, to convincingly show that KJ is false we need to show where the argument fails, and that is what I intend to do in this section.

Let's begin with a review of the principled argument for KJ which includes the two premises, KR<sup>91</sup> and SafeReasons:

(1) **KR**: If you know that  $p$ , then  $p$  is warranted enough to be a reason you have [for  $\Phi$ -ing], for any  $\Phi$ .

(2) **SafeReasons**: If  $p$  is a reason you have to  $\Phi$ , then  $p$  is warranted enough to justify you in  $\Phi$ -ing, for any  $\Phi$ .

**KJ**: If you know that  $p$ , then  $p$  is warranted enough to justify you in  $\Phi$ -ing, for any  $\Phi$ .

The problem I want to raise for the principled argument is that JB-Know also seems to work as a counterexample to SafeReasons. In JB-Know, Jack says "I know Caesar was born in 100 BC, but it isn't worth the risk." As Fantl and McGrath accept, this response seems appropriate and won't "raise any eyebrows" even though KJ implies that it is false. Jack's claim to know is false *if* the principled argument is sound. However, if we consider a version of Reed's counterexample in terms of SafeReasons, we also seem to get a counterexample to that principle from which it follows that the principled argument is not sound.

So, instead of Jack claiming to know, let's consider that Jack instead claimed the following: "I have very good support for believing that Caesar was born in 100BC *and that is a reason for me to answer*, but it isn't worth the risk." As with Reed's

---

<sup>91</sup> Though I won't make the argument here, it seems to me that KR is susceptible to a very simple counterexample since  $\Phi$  could be: *believe not-p*. But that doesn't make much sense – how could  $p$  be 'warranted enough to be a reason' you have to believe *not-p*? Any warrant for  $p$  will be warrant *against not-p* and *vice versa*. Only warrant *against p* could be warrant for *not-p*, but that would imply you don't know that  $p$ . So, you can't *know* that  $p$  and also have  $p$  be 'warranted enough to be a reason' to *believe that not-p*.

original case, there is nothing odd or ‘eyebrow raising’ about someone saying they have a reason to act, but to further say they would not be justified to act on that reason.<sup>92</sup> So, it is natural to conclude that Jack’s claim is true. If correct, then we have a counterexample to a premise in the principled argument and so the argument fails. Since Fantl and McGrath need the principled argument to show Jack’s claim in JB-Know is false, we are without a good reason for concluding that it is.

Fantl and McGrath could respond to this counterexample in the same way they responded to Reed’s original case. They might point out that Jack could also say “I thought I had a reason, but after thinking about the risk I’m not so sure I do have that as a reason and so I’m not going to answer.” This response, however, wouldn’t provide much support to their view for two reasons. First, the response doesn’t align with ordinary use of reasons since Jack goes so far as to *deny* that his belief Caesar was born in 100AD is even *a reason* to answer. But in everyday life we have reasons to do any number of things that would not be practically justified and that doesn’t seem to discount them as being a reason *at all*. Someone with a heart condition, for example, might love sausage and pepperoni and so have that as a reason to eat an entire large pizza, but that wouldn’t justify doing so. Neither would it mean that it isn’t a reason to do so.

---

<sup>92</sup> Bruce Russell pointed out, and I agree, that this claim is even more intuitive than the claim in JB-Know.

Second, the same points made in the previous section also apply here – doubting one has reason and so refraining from action is a legitimate response, but it is something that poses no issue for intellectualism. A subject who sincerely doubts that they have sufficient epistemic support (i.e. warrant) to use  $p$  as a reason would lose  $p$  as a reason. That would result in a change to their epistemic position regarding  $p$ . So, a case in which the subject expresses doubt about  $p$  is not a strike against intellectualism.<sup>93</sup>

Now that we see how a version of JB-Know also applies to SafeReasons, we have good reason to conclude that the principled argument is not successful. Without the principled argument to stand on, Fantl and McGrath have little recourse to argue that JB-Know is false and, moreover, they also need explain why his claim to know is intuitive and natural even though, on their view, it must be false.

However, the intellectualist has some work of their own because they need to say why Jack's decision not to answer is reasonable even though, according to their view, he *does* know. So, we are left with a debate over which of two views, intellectualism and practicalism, provide the best explanation of Jack's claim and whether it is true or false – a situation they hoped to avoid with the principled argument. In the next few sections, I will argue that there is a plausible way to

---

<sup>93</sup> This is true even if their doubt is the result of some practical issue.

understand these ‘practical intuitions’ which is compatible with intellectualism. ‘We’ll also find that this leads to further problems for KJ.

#### *§4.4 Two Practicalist Assumptions*

If the conclusion of the previous section is correct, then the principled argument for KJ is not sound. A central goal of the principled argument was to avoid debates over competing intuitions, but with the failure of the argument we are again left with these debates. So, the failure of the principled argument doesn’t entirely rule out KJ since one could argue that KJ offers the best explanation of why we think answering the question in the Caesar case would be practically irrational. To be sure, it does seem as if there is an important relation between the amount of risk one is undertaking and the care that they should take in figuring out if it is the right course of action. When the risk is greater, it seems we should be more careful by checking our facts and evidence to make sure we got it right. But when the risk is lower, it seems that we don’t need to worry so much about being right.

For the practicalist, the best explanation of this relation is that being practically justified is a necessary condition on knowledge. Looking at the cases discussed throughout— Airport, The Train Case (1.7), and Caesar, we might lean towards this sort of explanation because it seems that knowledge should be good enough to justify action. And if the action isn’t justified, then asking whether the subject has knowledge seems quite reasonable. This way of understanding the relation between action and knowledge suggests the following question: “If

knowledge is good enough when the stakes are lower, why isn't knowledge good enough when the stakes are higher?" To offer a plausible explanation of their own, intellectualists need to provide an answer to this question.

The main problem for intellectualism comes about as a result of two implicit assumptions. The first is that knowing  $p$  is sufficient for rationally acting on  $p$  in low-stakes cases. In other words, the assumption that, in cases like Bank-Low, what rationally justifies the action is the fact the subject knows that  $p$ ; i.e. assuming that 'knowledge is good enough' in at least some cases. If we assume the subject's action is justified in Bank-Low because they know the bank is open the following day, then we've already accepted that knowing  $p$  is, at least in some instances, sufficient to justify acting on  $p$ . Let's refer to this as the Sufficiency Principle:

**Sufficiency Principle (SP):** If  $S$  knows that  $p$ , then it is practically rational for  $S$  to act on  $p$ .

My primary concern is to vindicate Jack's response in JB-Know. If SP is true, that claim would appear to be false since it does not seem rational for Jack to respond. I will argue that this principle is false – knowledge alone is not sufficient to justify action. So, the implicit idea that knowledge is 'good enough' even in the low-stakes cases is false.

This leads us to the second assumption that acting on  $p$  in the high-stakes cases is *not* rationally justified *because* the subject lacks sufficient justification to know that  $p$ . On Fantl and McGrath's view, an increase in the negative outcome leads to an increase in the justification required to know that  $p$ . So, the practical

rationality of acting on  $p$  is a direct relationship between one's justification for  $p$  and the risk of acting on  $p$ . When the risk goes up, so does the justification required to know and when it goes down, less is required. I will also argue that this assumption is false. Rather than a lack of epistemic justification, acting on  $p$  is not rationally justified because the action is not, practically speaking, *worth it*. What I mean by an action being 'worth it' will be discussed in the next section.

The central argument is hinted at by Reed in his paper "Resisting Encroachment." There he writes:

"It's not that knowledge gives you safe reasons in any situation, but rather that knowledge can be used as a reason *only when it is safe to do so*. But that much weaker conclusion does not underwrite a pragmatic condition on knowledge."<sup>94</sup>

I think Reed's idea of something being 'safe to do' is essentially correct, but we need some clarification. Specifically, we need a better understanding of how it can be that someone knows that  $p$  but wouldn't be justified to act on that knowledge. In what follows, I will argue that it isn't the increase of risk, but the fact that acting on  $p$  is no longer *worth it* that makes acting on  $p$  practically irrational. This is true, even in cases in which the possibility of being unsuccessful is exceptionally low – an action must still be worth it. We'll also find that this intuitive understanding of practical action is more compatible with intellectualism than KJ.

---

<sup>94</sup> Reed (2012), 468. Emphasis added.



#### §4.5 Actions that are Worth It

When we act, we consider the various payoffs and penalties that might occur and our decision to act is, in part, based on whether the payoff is good enough to risk the penalty. How they compare will often determine what course of action we take (or don't take). In general, looking at cases in which there is a non-zero chance of an unsuccessful outcome, the payoff must be sufficiently high in relation to the penalty. For instance, if the *payoff* of a successful act vastly outweighs the *penalty* if one is unsuccessful, then the action will likely be worth it and so practically rational.<sup>95</sup> On the other hand, if the penalty sufficiently outweighs the good outcome, the action is most likely not worth it and so not practically rational to engage. The important implication is that, when an act appears to be practically irrational, it is not simply a matter of the risk that the outcome will be unsuccessful, but whether the risk of being unsuccessful is *worth it* in comparison to the possible gain.

Before we examine some cases, there is one important point to discuss which I hinted at in the previous paragraph: there is no action of which the outcome is certain even though there are many actions for which the probability of a successful outcome is *highly likely*. Even in those cases the subject will incur some minimal amount of risk. Surely, the probabilities of each outcome will be an important aspect of decision making, but it is also a comparison of the outcomes

---

<sup>95</sup> Moving forward, I will take 'payoff' to mean the result of successfully acting on  $p$  and 'penalty' to mean the result of unsuccessfully acting on  $p$ .

that matter. The fact that no outcome is guaranteed (i.e. there is a non-zero chance of being unsuccessful) weighs heavily on our reasoning and especially so when the consequences of a failed action would be far beyond what we can bear. It is important to keep this in mind as we discuss the following examples.

Turning now to some modified versions of JB-Know, we'll find that our practical intuitions—i.e. our intuitions about what it is rational to do—tend to track the ratio between the penalty *and* the payoff as opposed to tracking a relation between one's justification for *p* and the risk of the penalty *alone* as KJ suggests. I'll refer to this as the *Worth It Thesis*:

**Worth It Thesis (WIT):** Acting on *p* is rationally justified for *S* only if the ratio between the outcome for successfully acting on *p* (i.e. the payoff) as compared to the outcome of an unsuccessful outcome (i.e. the penalty) is sufficiently high.<sup>96</sup>

Immediately, we might wonder what sort of ratio between the payoff and penalty counts as 'sufficiently high'. That is an important question but providing an answer to that question is beyond our purposes here. At the very least, there is an intuitive sense in which some actions 'make sense' because the possible gain or payoff is high enough to risk the possible loss or penalty. When the payoff is modified to achieve a favorable or unfavorable ratio between the payoff and penalty, we'll find that our intuitions regarding the practical rationality of acting

---

<sup>96</sup> There is a hint of expected utility theory present in WIT as it would seem to provide a useful way to understand the 'ratio'. I've set these issues aside to focus on the crucial point: it is not simply the increase in risk that alters our intuition, but the disparity that results in the payoff vs. the penalty of the separate outcomes.

on that knowledge also change. Some examples will be discussed in the following sections.

To put it another way, a poor ratio between the payoff and penalty results in what we ordinarily call an *unnecessary risk*; that is, a risk such that the benefit of performing an act is significantly outweighed by the possibility, even if unlikely, of a high negative outcome. In many cases, avoiding unnecessary risk has little to do with the probability that the negative outcome will occur. For example, the probability that I will be in a fatal car crash is low, but the negative outcome is clearly significant. Thus, it would be an unnecessary risk to forgo wearing a seatbelt since the benefits are very minimal—at best, you are marginally more comfortable in your seat. Looking to the examples ahead, these sorts of considerations appear to be a leading factor in our intuitions about whether an action is practically rational.

One more concern to discuss is that WIT might appear to suggest that changes in epistemic justification—even changes when one already has knowledge—do not alter whether an action is practically rational. Commonsense, however, suggests that epistemic justification is unquestionably relevant to practical rationality, and it is important to note that WIT does not conflict with this. Rather, WIT only suggests that there are at least two components that determine whether something is practically rational and one of these has to do with epistemic justification. In the cases ahead, we'll find that changes in whether the act is worth it affect whether the act is practically rational. However,

differences in epistemic justification might have the same effect as well. For example, if Jack was asked “What is 1+1?”, perhaps it is more reasonable to answer in that case than in the Caesar case.

Yet, we must be careful here to distinguish between a case in which the person should have more epistemic justification vs. those cases in which the act just isn’t worth it. Certainly, there are cases like former, but my concern is with cases that have to do with the latter. If Jack’s problem is a lack of justification, then my point is mute, but the examples discussed below suggest that justification is not the issue in cases such as this, but the disparity in outcomes. This points to an important difference between the Sufficiency Principle and WIT. WIT allows that one can know that  $p$  even if it is not practically rational to act on that knowledge because it isn’t worth it. On the other hand, SP (and so KJ) imply that, for any instance in which it is not practically rational to act on  $p$ , one doesn’t know that  $p$ . But that implies there will be cases in which it is nearly impossible to have knowledge—even of ordinary propositions—because the negative outcome is simply too high. For example, if the penalty for answering incorrectly is the death of innocent bystanders, that would rule out the possibility of knowing the relevant proposition at all. In such a case, it doesn’t seem reasonable to answer if the payoff is only a jellybean. WIT, however, avoids these awkward results because it recognizes that some acts just aren’t worth it.

#### §4.6 The First Case: 10K

Two cases will be discussed in support of WIT. First, we will look at a version of JB-Know in which the payoff increases for a successful outcome and a second in which it decreases. Aside from the change in payoff, everything else remains the same with specific regard to the following:

- (1) Jack's epistemic justification for believing  $p$ .
- (2) The probabilities (either subjective or objective) of either outcome.
- (3) The penalty if unsuccessful.

Jack has not received any additional information that would alter his current justification for  $p$  and the possibility of successfully acting on  $p$  has also not changed. So, Jack is just as likely to have a successful outcome as in JB-Know.

When the outcomes include a jelly bean as a payoff and electric shock as a penalty, it seems that it would *not* be rational for Jack to answer the question. According to KJ, this implies that Jack lacks sufficient justification to know when Caesar was born. But let's consider a further scenario in which the payoff is increased significantly:

**10K:** After initially refusing to answer, Jack is then offered a payout of \$10,000 if he answers correctly. While still risking an electric shock, Jack decides that \$10,000 is enough to offset that risk and so he decides to answer.

In this version, it *does* seem rational for Jack to act on  $p$ . At least, it seems rational if you think that \$10,000 is enough to risk an electric shock, but perhaps you think \$10,000 is more than needed or maybe even too little. If so, you can modify the

number to your liking because it isn't the specific number that matters. Rather, the essential point is that a high enough payoff will make it so that the action is practically rational without changing any other factors of the case. If merely changing the payoff to be more in line with the penalty is enough to alter our intuitions regarding the practicality of acting on  $p$ , we have little reason to conclude that the issue for Jack was a lack of justification for  $p$  and even less reason to think that a lack of knowledge was at issue. Rather, the better explanation is simply that *the disparity between the penalty and the payoff* is the crucial factor which determined that not answering is the rational thing to do.

This response runs counter to KJ which implies that the increase to the penalty increased the justification required to know  $p$ . On my view, the increased penalty is relevant to the practicality of acting on  $p$ , but its relevance concerns its relation to the payoff rather than to Jack's justification for  $p$ . Why does this matter? On Fantl and McGrath's view Jack loses knowledge because the risk of being wrong 'stands in the way' of his acting on  $p$ . In other words, when the stakes are low, the possibility he is wrong doesn't matter. But when the stakes increase, that possibility becomes more of a concern. Thus, he needs better justification to know that  $p$ . So, the increased risk implies an increase in the justification required to know  $p$  and his current justification can't account for the difference. This interpretation of the cases, as we discussed earlier, assumes that knowing  $p$  is sufficient to justify acting on  $p$ ; i.e. knowledge is 'good enough'.

However, if it is not the risk, but the poor ratio between the payoff and penalty that drives our practical intuitions in these cases, then the better explanation is that Jack should just avoid acting on  $p$  altogether because it isn't worth it—it's an unnecessary risk. This implies that Jack hasn't lost knowledge and his response "I know he was born in 100 B.C., but [answering] isn't worth the risk" is not only reasonable, but what we should expect in this situation. So, the justification he has for  $p$  and whether he knows that  $p$  is not the relevant issue.

#### *§4.7 The Second Case: Nickel*

One might think that these issues only arise at more extreme payoffs and penalties, but in fact the same seems to be true at lower amounts as well. Let's assume again that Jack's justification for  $p$  and the probabilities for each outcome stay the same. Again, there will be no change to his epistemic justification for  $p$ , to the penalty, or the probability of either outcome.

Beginning with a payoff and penalty of \$100, it seems quite reasonable for Jack to answer. However, what if the payoff is decreased to only a nickel while the penalty remains at \$100? Again, all that has changed is a decrease in the payoff.<sup>97</sup> Under these conditions, it doesn't seem practically rational to answer. As with the previous example, KJ implies that it isn't rational for Jack to answer when the payoff drops because he doesn't know.

---

<sup>97</sup> As with the first case, if you don't find the disparity to be sufficient, then simply alter the amounts. The issue is not whether this specific difference is enough, but whether a drop in the payoff that results in a high enough disparity leads to intuiting that the acting is not practically rational.

According to WIT, Jack should refrain because it simply isn't worth it. He *knows* the answer, but he lacks good reason to respond. Again, it would be quite reasonable for him to say, "I know he was born in 100 B.C., but [answering] isn't worth the risk." Even if he is extremely confident he knows the answer, it is unnecessary to take on the risk of losing \$100 when he would only win a nickel. The essential point is, that our intuitions vindicate Jack's claim because they seem to track the ratio between the payoff and penalty in these cases rather than any supposed issue with Jack's epistemic justification. When the ratio is favorable, the action is rational, but when it is unfavorable, it is not. As with 10k there is no need to explain the change in terms of a weakness in his justification.

#### *§4.8 Two Replies*

In this section, I will consider two replies. First, it might be argued that, even though an action must be 'worth it', that just supports the idea that practical interests factor into knowledge in the way that KJ suggests. One reason this might appear to be the better explanation is that it is rational to act when the payoff and penalty are an even \$10, but if we increase the amount so that both are \$100,000, it doesn't seem rational even though they are the same ratio. So, we have two sets of cases with the same ratio and so both are 'worth it', but we are only justified to act in one of the cases. One could then argue that the best explanation of this is that the subject's justification for  $p$  is sufficient in the \$10 case, but not the \$100,000 case. In the \$10 case, the subject knows  $p$  and is justified to act on  $p$ , but in the \$100K



case the subject neither knows  $p$  nor is he justified to act on  $p$ . Just as KJ would predict.

Though the explicit monetary values are even, it does not follow that there is an even ratio. If I lose \$10 that is a bit disappointing and I might miss out on a new book or a six-pack, but I wouldn't put myself in danger of financial ruin. On the other hand, if I lose out on \$100,000, it would result in great financial difficulty. Even though gaining \$100,000 would surely be a great boon, the benefits that it would produce wouldn't come near to offsetting the problems incurred by the loss.

So, when assessing the ratio between the outcomes, it clearly isn't just about the monetary value because a multitude of other factors will come into play.<sup>98</sup> For instance, I might take the risk of losing \$100,000 if gaining it would have some important implications. Perhaps, if I get that \$100,000, then I would be able to save the lives of 10 people who would die otherwise. This change might alter whether the action is practically rational, but even if it does, *it would clearly alter the ratio as well*. Also, we might discuss the practicality of a billionaire risking \$100,000 in contrast to myself. With a billionaire, the practicality quickly changes since she is, relative to our financial standing, risking far less than I am. Further examples verify the general point that external factors are relevant to the ratio between

---

<sup>98</sup> Kahneman and Tversky make a similar point: "...the derived value(utility) function of an individual does not always reflect "pure" attitudes to money, since it could be affected by additional consequences associated with specific amounts."(1979, 279)

payoff and penalty.<sup>99</sup> So Fantl and McGrath are not helped by looking at cases in which the monetary ratio stays the same, but the amounts increase.

They might also argue that cases like 10K and Nickel aren't relevant to the truth of KJ – the KJ principle is entirely compatible with the idea that our practical intuitions track the payoff/penalty ratio. So, it's possible that the 'Worth It' thesis is true, but the subject also loses knowledge in high stakes/bad ratio cases. I grant that these cases don't provide a straightforward rejection of KJ. Rather, the criticism comes in terms of the better explanation that Worth It *and* intellectualism provide over practicalism. In that regard, we can say a bit more and show how intellectualism should be favored over KJ.

#### ***§4.9 KJ and Conversational Oddities***

First, if JB-Know and 10K are offered one after the other, we get some odd results by interpreting the outcome according to KJ. If KJ is true, then Jack doesn't know in JB-Know because he isn't justified to answer and so his claim to know would be false. However, once the payoff changes in 10K he is justified to answer which implies that he does know that *p*. Why would that be? The only change between the two cases is the increased payoff. So, it is reasonable to conclude that this should be our explanation if KJ is true – Jack acquires knowledge because the payoff increased. But clearly this is an odd sort of explanation. People don't acquire knowledge because it would turn out better for them if they are successful

---

<sup>99</sup> To make a fair assessment we would need to know, among other things, the subject's socioeconomic status, what they stand to gain with the money and what they would lose if they had to endure the penalty.

in acting on that knowledge. On the other hand, if we understand these cases according to the intellectualist view, we get no such oddities. Jack knows that  $p$  in both cases and the only difference is that it isn't worth it for him to act in the first instance.

In a similar way, an equally odd explanation occurs with the two cases involved in Nickel. At first the subject knows when the ratios are even, but once the payoff is decreased and it is no longer practical to act on  $p$ , the subject loses knowledge according to KJ. Again, no other factor changes except the payoff implying that one can lose knowledge simply because the payoff of a successful act *decreases*. Again, this is an odd result. People don't lose knowledge because the outcome of successfully acting on that knowledge becomes less lucrative. As before, intellectualism provides a clean and intuitive explanation—Jack knows in both, but he is only practically justified to act in the case in which it is worth it for him to do so.

Following the first point, if KJ is true, it would permit explaining the acquisition and loss of knowledge in counterintuitive ways. In correspondence with Fantl and McGrath, Mikkel Gerken points out that KJ permits citing practical reasons to defend knowledge attributions and denials and this leads to unordinary claims.<sup>100</sup> For instance, someone could say they don't know *because* the stakes are too high which seems counterintuitive.

---

<sup>100</sup> Gerken (2012)

A similar point applies to cases like 10K and Nickel. In the latter case, the payoff increased which resulted in making the action rational. If KJ is true, then Jack doesn't know  $p$  when the payoff is low, but he does know  $p$  when the payoff is increased. This implies that he acquired knowledge as a result of an increase to the payoff. So, it would be appropriate to cite such a reason if asked to explain how the subject acquired knowledge between JB-Know and 10K. But clearly it would be odd for someone to say "Well I didn't know  $p$  because the payoff was so low, but now the payoff increased. So, I do know that  $p$ ."

The same issue arises when looking at the Nickel cases. If KJ is true, and the person loses knowledge when the payoff decreases, then it would be appropriate for someone to explain the loss of knowledge as due to a decrease in the payoff. But again, it would be odd for someone to say "I did know  $p$ , but the payoff decreased. I guess now I don't know that  $p$ ." Though similar to the issue in the above paragraph, it presents a further problem for KJ because Fantl and McGrath take that principle to be in line with our standard knowledge-attributing/denying practices. So, the fact that KJ permits odd explanations of knowledge acquisition and loss creates a significant problem for their view.

Interpreting these cases according to Worth It and intellectualism results in some common and plausible responses. For example, the subject might say, "I know the answer, but I'm not going to risk some mishap or mistake when I'd only get a jelly bean. It's not even worth trying." This sort of response is natural and intuitive. Given the disparity between the odd-sounding responses implied by KJ

and the natural-sounding responses of intellectualism, it is clear we should opt for the latter.

So, it looks like WIT provides the better interpretation of JB-Know. Most importantly, if WIT is correct, then there is no threat to intellectualism because we also have a plausible way to explain why Jack's response seems natural that avoids concluding his claim to know is false.

#### ***§4.10 Conclusion***

I have discussed Baron Reed's counterexample to KJ and argued that it is a more effective attack on KJ than Fantl and McGrath have taken it to be. Not only is it a counterexample to KJ, it can also be modified to attack Safe Reasons which is the second premise of their principled argument. Since KJ is not supported by any principled argument, the debate between the practicalist and intellectualist comes down to the explanatory plausibility of the views.

Together with WIT, intellectualism provides a cleaner explanation for why certain actions with higher risks are not practically rational. As I've argued, it is not about a lack of justification, but the simple fact that the payoff is not good enough in relation to the penalty. For an act to be practically rational, the payoff must be good enough. The examples discussed, 10K and Nickel, confirm this idea and, moreover, they lead to some odd problems for Fantl and McGrath's view; namely, if KJ is true, then one can gain or lose knowledge merely from an increase/decrease in the payoff alone. But this is counterintuitive and easily

avoidable if we hold to intellectualism. In sum, there doesn't seem to be much reason to adopt Fantl and McGrath's view over intellectualism.

## Bibliography

- Bach, Kent. "The Emperor's New 'Knows'." In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, by Gerhard Preyer, & Georg Peter, 51-89. Oxford University Press, 2005.
- Brown, Jessica. "Comparing Contextualism and Invariantism on the Correctness of Contextualist Intuitions." *Grazer Philosophische Studien*, 2005: 71-100.
- Buckwalter, Wesley. "Knowledge Isn't Closed on Saturday: A Study in Ordinary Language." *Review of Philosophy and Psychology*, 2010: 395-406.
- Buckwater, Wesley, and Jonathan Schaffer. "Knowledge, Stakes, and Mistakes." *Nous*, 2015: 201-234.
- Casullo, Albert. *A Priori Justification*. New York: Oxford University Press, 2003.
- Cohen, Stewart. "Contextualism, Skepticism, and the Structure of Reasons." *Philosophical Perspectives*, 1999: 57-89.
- . "Contextualism and Skepticism." *Philosophical Issues*, 2000: 94-107.
- DeRose, Keith. "Contextualism: An Explanation and Defense." In *The Blackwell Guide to Epistemology*, by John Greco, & Ernest Sosa, 187-205. Oxford: Blackwell, 1999.
- . "How Can We Know that We're Not Brains in Vats?" *The Southern Journal of Philosophy*, 2000: 121-148.
- . "The ordinary language basis for contextualism, and the new invariantism." *Philosophical Quarterly*, 2005: 172-198.
- . *The Case for Contextualism*. Oxford University Press, 2009.
- Dougherty, Trent. "Book Review: Knowledge in an Uncertain World." *Mind* 122, 2013: 1078-1085.
- Fantl, Jeremy. *Re: The reason-giving argument against pragmatic encroachment*. March 22, 2012. <http://certaindoubts.com/the-reason-giving-argument-against-pragmatic-encroachment/#comment-26990> (accessed January 13, 2015).
- Fantl, Jeremy, and Matthew McGrath. "Evidence, Pragmatics, and Justification." *The Philosophical Review*, 2002: 67-94.
- . *Knowledge in an Uncertain World*. Oxford University Press, 2009.
- . "Precis of Knowledge in an Uncertain World." *Philosophy and Phenomenological Research*, 2012a: 441-446.

- . "Replies to Cohen, Neta, and Reed." *Philosophy and Phenomenological Research*, 2012b: 473-490.
- Feldman, Richard. "Having Evidence." In *Philosophical Analysis*, by D.F. Austin (ed.), 83-104. Kluwer Academic Publishers, 1988.
- Feldman, Richard, and Earl Conee. *Evidentialism*. Oxford University Press, 2004.
- Feltz, Adam, and Chris Zarpentine. "Do You Know More When it Matters Less?" *Philosophical Psychology*, 2010: 683-706.
- Gerken, Mikkel. *Re: The reason-giving argument against pragmatic encroachment*. March 22, 2012. <http://certaindoubts.com/the-reason-giving-argument-against-pragmatic-encroachment/#comment-26990> (accessed January 13, 2015).
- Hawthorne, John. *Knowledge and Lotteries*. Oxford University Press, 2004.
- Jr., Mylan Engel. "Is Epistemic Luck Compatible with Knowledge?" *Southern Journal of Philosophy*, 1992: 59-75.
- Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, 1979: 263-291.
- Lewis, David. "Elusive knowledge." *Australasian Journal of Philosophy*, 1996: 549-567.
- Nagel, Jennifer. "Knowledge ascriptions and the psychological consequences of changing stakes." *Australian Journal of Philosophy*, 2008: 179-294.
- . "Epistemic Anxiety and Adaptive Invariantism." *Philosophical Perspectives*, 2010: 407-435.
- Neta, Ram. "Contextualism and the Problem of the External World." *International Phenomenological Society*, 2003: 1-31.
- Pinillos, N Angel. "Knowledge, Experiments, and Practical Interests." In *Knowledge Ascriptions*, edited by Jessica Brown, & Mikel Gerken, 192-221. Oxford University Press, 2012.
- Pinillos, N Angel. "Some Recent Work in Experimental Epistemology." *Philosophy Compass*, 2011: 675-688.
- Reed, Baron. "A Defense of Stable Invariantism." *Nous*, 2010: 224-244.
- . "Resisting Encroachment." *Philosophy and Phenomenological Research*, (2012): 465-472.



- . "Practical Matters Do Not Affect Whether You Know." In *Contemporary Debates in Epistemology*, edited by Matthias Steup, John Turri, & Ernest Sosa, 95-106. Wiley-Blackwell, 2013.
- Russell, Bruce. "How to Be an Anti-Skeptic and a Noncontextualist." *Erkenntnis*, 2004: 245-255.
- . "Contextualism on a Pragmatic, Not a Skeptical, Footing." *Acta Analytica*, 2005: 26-37.
- Rysiew, Patrick. "The Context-Sensitivity of Knowledge Attributions." *Nous* 35, no. 2001: 477-514.
- . "Speaking of Knowing." *Nous*, 2007: 627-662.
- Schaffer, Jonathan. "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism." *Philosophical Studies*, 2006: 87-107.
- Sripada, Chandra, and Jason Stanley. "Empirical Tests of Interest-Relative Invariantism." *Episteme*, 2012: 3-26.
- Stanley, Jason. "On the Linguistic Basis for Contextualism." *Philosophical Studies*, 2004: 119-146.
- . *Knowledge and Practical Interests*. Oxford University Press, 2005.
- Stine, Gail. "Skepticism, Relevant Alternatives, and Deductive Closure." *Philosophical Studies*, 1976: 249-261.
- Turri, John. "Epistemic Contextualism: An Idle Hypothesis." *Australasian Journal of Philosophy*, 2016: 141-156.
- Williamson, Timothy. *Knowledge and Its Limits*. Oxford University Press, 2000.

**ABSTRACT****A DEFENSE OF EPISTEMIC INVARIANTISM AND INTELLECTUALISM**

by

**LUCAS RYAN LOCKARD****August 2018****Advisor:** Dr. Bruce Russell**Major:** Philosophy**Degree:** Doctor of Philosophy

The traditional view of knowledge includes two plausible claims. The first claim is that the truth-conditions required to know some proposition  $p$  do not vary. So, for any two subjects who have the same epistemic position regarding  $p$ , either both know that  $p$  or both do not know that  $p$ . An implication of this invariantist position is that any difference in knowledge will be explained by a difference in their epistemic position. The second claim is that only epistemic or 'truth-conducive' reasons are relevant to considerations of knowledge. This intellectualist position has recently come under attack from those who argue that whether someone knows that  $p$  is partly restricted by their practical situation.

The first task of this dissertation is to defend invariantism from the contextualist view that the truth-conditions required to know that  $p$  are context-dependent. Contextualists argue that their view best explains our ordinary and intuitively correct knowledge-attributing practices. In response, I argue that they do not adequately address other related claims that are equally ordinary and

intuitive. The debate largely hinges on identifying a mistake in the relevant cases and to that end I argue a closer analysis of contextualist cases reveals just such a mistake. In short, the ‘knowledge-denial’ rests on a false belief.

The second task is to defend intellectualism from the practicalist view of Jeremy Fantl and Matthew McGrath. Initially, I argue that the argument from their central KJ principle is susceptible to a counter-example first used by Baron Reed to attack the principle itself. This counter-example is much stronger than first realized as it also applies to the SafeReasons principle used as a premise in the argument for KJ. After showing the argument to be unsound, I discuss further examples which reveal a second condition on rational action. This condition suggests that knowing  $p$  is not sufficient to justify acting on  $p$ —as is implied by KJ—the act must also be worth it.

## AUTOBIOGRAPHICAL STATEMENT

Lucas began his academic career at Grand Rapids Community College as a classical guitarist and then moved on to Grand Valley State University where he received a BA in Music. After his time at GVSU, he headed off to Calvin Theological Seminary where he received an MA of Divinity. Teaching was, and still is, his main interest. To that end, he began his pursuit of a doctorate at Western Michigan University where he received an MA in Philosophy. Upon graduating from WMU, he was accepted at Wayne State where he completed his PhD in 2018.