

UNIVERSITY OF CALGARY

A Conditional Model of Knowledge Attributions

by

Julianne Chung

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES  
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE  
DEGREE OF MASTER OF ARTS.

DEPARTMENT OF PHILOSOPHY

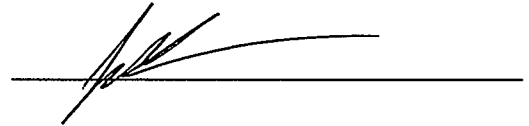
CALGARY, ALBERTA

AUGUST, 2009

© Julianne Chung 2009

UNIVERSITY OF CALGARY  
FACULTY OF GRADUATE STUDIES

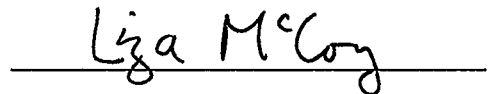
The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "A Conditional Model of Knowledge Attributions" submitted by Julianne Chung in partial fulfilment of the requirements for the degree of Master of Arts.



Supervisor, Jeremy Fantl, Department of Philosophy



Dennis McKerlie, Department of Philosophy



Liza McCoy, Department of Sociology

August 13<sup>th</sup>, 2009

Date

ABSTRACT: *A Conditional Model of Knowledge Attributions*

(Julianne Chung)

My primary purpose in this thesis is to argue that, even if skepticism is true, we can nonetheless be warranted in attributing knowledge. In chapter one, I survey two kinds of skeptical argument, both of which have been influential in recent years: one from the supposed infallibility of knowledge, the other from what are known as “lottery problems”. In chapter two, I examine and evaluate three recently proposed theories of knowledge attributions that are motivated, at least in part, by the arguments surveyed in the first chapter: contextualism, contrastivism, and an account sketched by Gilbert Harman and Brett Sherman. Finally, in chapter three, I advance a conditional model of knowledge attributions that I think manages to satisfy several important desiderata for a plausible theory of knowledge attributions, and argue that this conditional model of knowledge attributions is preferable to the three accounts discussed in chapter two.

# A Conditional Model of Knowledge Attributions

Approval Page.....	ii
Abstract.....	iii
Table of Contents.....	iv
Introduction.....	1
Chapter One: Skepticism, Infallibilism, and Lottery Problems	
I. Skeptical Arguments from the Infallibility of Knowledge.....	3
i) Infallibilism and Fallibilism.....	3
ii) David Lewis.....	4
iii) Peter Unger.....	8
II. Skeptical Arguments from Lottery Problems.....	21
i) Lottery Problems.....	21
ii) Closure Principles.....	25
III. Five Desiderata for a Theory of Knowledge Attributions.....	26
Chapter Two: Recent Theories of Knowledge Attributions	
I. Contextualist Theories of Knowledge Attributions.....	27
i) A Brief Exposition.....	27
ii) Potential Problems.....	34
II. Contrastivist Theories of Knowledge Attributions.....	37
i) A Brief Exposition.....	37
ii) Potential Problems.....	41
III. Harman and Sherman's Account.....	45
iii) A Brief Exposition.....	45
iv) Potential Problems.....	47
IV. Two Additional Desiderata for a Theory of Knowledge Attributions.....	48
V. The Seven Desiderata and These Three Accounts.....	49
Chapter Three: A Conditional Model of Knowledge Attributions.....	51
I. Initial Exposition.....	52
II. Restricting Assumptions.....	56
i) Epistemic Priority.....	57
ii) Reasonability.....	59
iii) Raised Standards.....	61
a) Stakes.....	62
b) Salience.....	62
III. Why Think That We Can Imply So Much?.....	64
IV. What Differentiates This Account?.....	71
V. Why Endorse This Account?.....	71
Conclusion.....	77
Bibliography.....	78

## Introduction

Skepticism in epistemology (henceforth simply, “skepticism”) can be defined as the view that no, or at most very few, propositions are known. Skepticism remains a highly influential philosophical position that is perhaps as irksome as it is compelling; for, though it has endured, it is not easy to fully endorse. After all, we quite often attribute knowledge to ourselves and to others. We like to think that we know rather mundane things such as the names of our family members and friends, where we live, and what we’ll be doing over the course of the next day. We also like to think that we know much that is quite a bit more interesting and contentious.<sup>1</sup> Further, we frequently speak of things such as “scientific knowledge,” “common knowledge,” “encyclopedic knowledge” and “bodies of knowledge”. Knowledge, it seems, is everywhere, and because of this, it might seem absurd to even entertain a view on which we know nothing, or next to nothing. Additionally, arguments in favor of skepticism tend to imply not only that very little is known, but also that very little *can be* known. Thus, skepticism might seem potentially troubling in addition to merely seeming absurd; for if the prospects for acquiring knowledge are that dim, just what is the point of engaging in inquiry? Don’t we all inquire precisely to acquire knowledge? If we cannot know much, then we are doomed to live in a state of ignorance, and this seems like a decidedly negative thing indeed!

My primary purpose in this paper is to argue that, even if skepticism is true, we can nonetheless be warranted in attributing knowledge. In chapter one, I survey two

---

<sup>1</sup> David Lewis puts the point similarly in, “Elusive Knowledge,” *Australasian Journal of Philosophy* 74 (1996), 549.

kinds of skeptical argument, both of which have been influential in recent years: one from the supposed infallibility of knowledge, the other from what are known as “lottery problems”. In chapter two, I examine and evaluate three recently proposed theories of knowledge attributions that are motivated, at least in part, by the arguments surveyed in the first chapter: contextualism, contrastivism, and an account sketched by Gilbert Harman and Brett Sherman. Finally, in chapter three, I advance a conditional model of knowledge attributions that I think manages to satisfy several important desiderata for a plausible theory of knowledge attributions, and argue that this conditional model of knowledge attributions is preferable to the three accounts discussed in chapter two.

## Chapter 1: Skepticism, Infallibilism, and Lottery Problems

In this chapter, I survey two kinds of skeptical argument, both of which have garnered considerable attention in recent years. The first kind of skeptical argument stems from the supposed infallibility of knowledge. The second kind of skeptical argument stems from what have come to be known as “lottery problems”. By examining and evaluating these two kinds of skeptical argument, I hope to illuminate five desiderata that any plausible theory of knowledge attributions should satisfy.

### I) Skeptical Arguments from the Infallibility of Knowledge

#### (i) *Infallibilism and Fallibilism*

Roughly stated, infallibilism (with respect to knowledge) is the view that knowledge entails the impossibility of error; that is, on infallibilist theories of knowledge, it is not possible for an individual to know a proposition *p* if it is possible that he or she is wrong about *p*. Fallibilism is the negation of this claim; thus, according to fallibilist theories, it is possible for an agent to know a proposition *p* even if it is possible that he or she is wrong about *p*.

It is not too difficult to come up with linguistic evidence that seems to support infallibilism. Consider the following propositions:

- 1) “I know that Dostoyevsky wrote *Crime and Punishment*, but it is possible that I am wrong.”
- 2) “He knows that Washington, D.C., is the capital of the United States of America, but he could be wrong.”

3) “She knows that World War II ended in 1945, but she cannot be completely sure about that”.

Because sentences such as these sound inconsistent, they suggest that it is not possible for an agent S to know a proposition p fallibly. It seems right to say that, if it is possible that an individual is wrong about a proposition p (that is, if her evidence for or method of determining that p does not entail p), then that individual does not know that p. This is the core infallibilist intuition.

The problem with this, however, is that it is relatively easy to imagine ways in which it is possible that one is wrong about most, if not all, propositions that one purports to know. Let us suppose that I claim to know that Dostoyevsky wrote *Crime and Punishment*. There are many ways in which I could be wrong. For instance, perhaps memory might not be serving me correctly in this instance; it is possible that, though I think that I know that Dostoyevsky wrote *Crime and Punishment*, I am actually confusing him with some other author. Perhaps, the source from which I gleaned that information is a purveyor of erroneous information. Perhaps Dostoyevsky stole the manuscript. Or, perhaps *Crime and Punishment* does not even exist; rather, it is just a figment of my imagination!

Skeptical arguments along these lines have been advanced by a number of philosophers; in recent years, however, the most prominent proponents of such arguments have been David Lewis and Peter Unger.

(ii) *David Lewis*

In “Elusive Knowledge,” David Lewis presents the following argument:



(P1) It seems as if knowledge is by definition infallible.

(P2) If knowledge is infallible, then for any proposition *p* that an agent *S* takes him or herself to know, *S* must have eliminated all possibilities in which not-*p*.

(P3) This is not the case (indeed, it is perhaps not even possible/practical) for most (if not all) propositions *p* that are taken to be known.

These three premises entail the conclusion that (C) we know nothing, or at any rate, next to nothing.<sup>2</sup>

According to Lewis, speaking of fallible knowledge (that is, of knowledge despite un-eliminated possibilities) just sounds inconsistent. After all, if one claims that an individual *S* knows a proposition *P*, and one grants that *S* cannot eliminate a certain possibility in which not-*P*, it seems as if one has granted that *S* does not know *P* after all.<sup>3</sup>

The trouble here, of course, is that it seems as if there is always (or at any rate, almost always) some un-eliminated possibility of error, even with respect to mundane knowledge attributions. As Lewis notes, these un-eliminated possibilities of error may often sound far-fetched, but they are possibilities nonetheless. And, as such, they might be thought to undermine even our most basic knowledge attributions. As Lewis writes,

Let your paranoid fantasies rip – CIA plots, hallucinogens in the tap water, conspiracies to deceive, old Nick himself – and soon you find that uneliminated

---

<sup>2</sup> Lewis, D. "Elusive Knowledge," *Australasian Journal of Philosophy* 74:4 (1996), 549

<sup>3</sup> *Ibid.*, 549. Taking the proposition, "S knows that Dostoyevsky wrote *Crime and Punishment*," as an example, according to Lewis, if one grants that *S* cannot eliminate the possibility (say) that someone else wrote the manuscript (and Dostoyevsky simply stole it), then it seems as if one has granted that *S* does not know that Dostoyevsky wrote *Crime and Punishment*.

possibilities of error are everywhere. Those possibilities of error are far-fetched, of course, but possibilities all the same.<sup>4</sup>

So, says Lewis, we think that we both know a lot and that knowledge must be infallible. But, we either have fallible knowledge or at best, next to none, and at worst, none at all. As he puts it, we “are caught between the rock of fallibilism and the whirlpool of skepticism. Both are mad!”<sup>5</sup>

Lewis points out that fallibilism has been taken to the less objectionable view, because it seems to many that fallibilism simply requires a less frequent correction of what we typically say.<sup>6</sup> So, when forced to choose between skepticism and fallibilism, most choose fallibilism. Lewis urges contented fallibilists, however, to try to hear the following phrase with “fresh ears”: “He knows, yet he has not eliminated all possibilities of error”. As Lewis puts it, even if one has “numbed one’s ears”, this overt, explicit fallibilism still sounds wrong, as it sounds contradictory.<sup>7</sup>

Lewis’s argument here captures the core infallibilist intuition because, if one can eliminate all possibilities in which *p* does not obtain, then it is not possible that one is wrong about *p*. Lewis does not say much more in defense of infallibilism; rather, he seems to take it to be intuitive enough so as not to require further argument. It is worth noting, however, that though Lewis endorses infallibilism, he ultimately rejects skepticism (in fact, he thinks that it is a “Moorean fact” that a lot is known, which is to say that he is more sure that a lot is known than he is of any premises of an argument

---

<sup>4</sup> Ibid., 549

<sup>5</sup> Ibid., 550

<sup>6</sup> Because if infallibilism is true, in most cases in which the verb “knows” is applied, it is applied incorrectly.

<sup>7</sup> Ibid., 550

to the contrary).<sup>8</sup> That said, his argument in favor of infallibilism could still be construed as providing the framework for an argument in favor of skepticism. After all, one might argue that it is not possible for agents to rule out every possibility that not-*p* for many if not all of the propositions *p* that we take ourselves to know. One might maintain that certain possibilities cannot be ruled out even in principle; for example, one perhaps cannot, even in principle, rule out the possibility that one is a disembodied brain-in-a-vat who is hooked up to equipment programmed to provide it with visual, auditory, tactile, and other sensory inputs such that one thinks that one is experiencing a whole variety of things, but actually is not. If one cannot rule out this possibility, it is plausible that one cannot know anything about the world; after all, one cannot rule out the possibility that everything that they perceive to be happening is illusory. Other possibilities that also arguably cannot be ruled out in principle include the possibility that one is being deceived by an evil demon and the possibility that one has spent his entire life in a lifelong dream-state in which nothing is real.<sup>9</sup> Thus, a fairly widespread skepticism arguably follows if knowledge is infallible in Lewis's sense. Though Lewis does attempt to address these sorts of concerns by positing a contextualist theory of knowledge attributions (I discuss contextualist theories in the next chapter), insofar as one finds such a solution wanting, one might find herself inclined towards skepticism if one finds Lewis's rather intuitive argument in favor of infallibilism compelling.

---

<sup>8</sup> Ibid., 549

<sup>9</sup> These examples are first due to Rene Descartes. See meditation one in Descartes, R. *Meditations on First Philosophy* (Cambridge: 1996)

(iii) *Peter Unger*

Peter Unger, one of the most influential skeptics of the 20<sup>th</sup> century, presents a similar, but far more elaborate argument for the infallibility of knowledge in “An Argument for Skepticism.” Unlike Lewis, Unger claims that, because knowledge is infallible, no one ever knows *anything at all* to be the case. That is, in his view, no one knows anything about, say, the past or the present, about external objects or his or her own experiences, about complicated contingencies, or even about the simplest mathematical truths.<sup>10</sup>

Unger’s argument has two premises. The first is the proposition:

(P1) If someone knows something to be so, then it is all right for the person to be absolutely certain that it is so.

The second is the proposition:

(P2) It is never all right for anyone to be absolutely certain that anything is so.<sup>11</sup>

These two, taken together, entail the conclusion (C) that no one ever knows that anything is so.<sup>12</sup>

Unger’s Argument for P1:

Unger claims that to deny P1 is to, “do violence to the meaning of the word know,” and offers a variety of linguistic evidence in favor of it. As Unger notes, we

---

<sup>10</sup> Unger, P. “An Argument for Skepticism,” in *Arguing About Knowledge* ed. Duncan Pritchard and Ram Neta (Routledge: 2009), 467

<sup>11</sup> It is worth noting that Lewis’s brand of infallibilism can be construed as differing a bit from Unger’s. Unger characterizes knowledge in terms of it being “all right to have an attitude of certainty,” whereas Lewis characterizes knowledge in terms of “eliminating all possibilities of error”. They are not equivalent, however; after all, Unger think that one must at least have an attitude of certainty towards a proposition p if one is to know p; however, Lewis has no such requirement. Conceivably, one could eliminate all possibilities of error without having a corresponding psychological attitude of certainty.

<sup>12</sup> *Ibid.*, 467

often think that someone is certain of a proposition but shouldn't be. That is, we often think that, though a certain individual may *feel* certain of a proposition, he really ought not have this attitude of certainty. Thus, we might ask such an individual a question like, "How can you be certain of that?". As Unger sees it, in asking this question, we imply that it might not be all right for the individual to feel certain, and imply further that this is because he might not really know that proposition. If the individual could demonstrate that he does indeed know it, it would be natural to withdraw the question. Why is this? Unger suggests that we are able to imply so much because we all accept the idea that if one knows a proposition, then it is all right for one to feel certain of it, but if one doesn't then it isn't. This suggests that there is some analytic connection between knowing and being all right to feel certain.<sup>13</sup>

According to Unger, the idea that knowing entails its being all right to be certain is further suggested by the fact that knowing entails, at least, that one indeed *is* certain; in other words, one must at least have an *attitude of certainty* towards any proposition that one purports to know. This is made clear by the inconsistency of propositions like, "He really knew that it was raining, but he wasn't absolutely certain that it was".<sup>14</sup> Unger also contends that the entailment from knowing to being absolutely certain can be seen once one realizes the equivalence between an agent's knowing a proposition and his knowing it with absolute certainty. Though there may be cases in which we would more naturally react by claiming that the individual in question simply knew a proposition rather than knowing it for certain (such as a case

---

<sup>13</sup> Ibid., 467

<sup>14</sup> Ibid., 467

in which a man, while looking for his cuff-links, unerringly went to the spot where they were located, though he had doubts as to whether or not they were indeed located there), Unger claims that our readiness to attribute knowledge in such cases merely shows that we use the term “know” loosely in certain cases, and suggests that we are simply far more strict in our use of the term “know” when the term “certain” enters into the picture. Unger thinks this a more plausible hypothesis than the idea that there is no equivalence between knowing and knowing with absolute certainty, for propositions such as “He knew it, but he didn’t know it for certain,” and “He really knew it, but he didn’t know it with absolute certainty” are inconsistent.<sup>15</sup>

#### Unger’s Argument for P2:

So, why is it never all right for anyone to be certain that anything is so?

According to Unger, this is due to the fact that an agent’s being absolutely certain of a proposition involves that agent in having a severely negative attitude with respect to the matter of whether that proposition is so. This attitude is that no new information, evidence, or experience that an agent might ever have will be seriously considered by the agent to be at all relevant to any possible change in his or her thinking in the matter. To put it another way, Unger thinks that if one is absolutely certain of a given proposition, then one must have a corresponding attitude that any new information, evidence, or experience that would count against that proposition should be ignored.<sup>16</sup>

---

<sup>15</sup> Ibid., 468

<sup>16</sup> Ibid., 469

To show that this is so, Unger begins by offering some linguistic evidence. He asks the reader to consider a case from the work of Norman Malcolm, a case in which an agent knows the following proposition: "There is an ink-bottle before me". According to the preceding discussion of P1, says Unger, we have seen that knowledge of a proposition plausibly entails, at the least, that one has an attitude of certainty towards that proposition. But, according to Unger, one has an attitude of certainty towards the proposition that "There is an ink-bottle before me," if and only if one has the following attitude: "insofar as I care about being right about whether an ink- bottle is or was before me, no matter how things may seem to appear, I will not count as contrary evidence any possible sequence of events, however extraordinary". So, even if several people entered the room and declared that there was no ink- bottle there, a photograph of the desk showed no ink-bottle present, or, upon reaching out to grab it, one's hand instead passed through the ink-bottle, an agent should not count any of these experiences as evidence against the proposition that "There is an ink-bottle before me," if he indeed really is certain of that fact.<sup>17</sup>

At this juncture, it is important to note that Unger stresses that a proposition like, "I should regard nothing as evidence that there is no ink-bottle before me right now" must be regarded as an expression of a man's current attitude, and not as a prediction of what he will do under certain future circumstances. Thus one can allow as consistent a sentence like, "He is absolutely certain that there are automobiles, but he may change his mind should certain evidence come up". This is because, even if his present attitude is that he will not, things may not happen in accordance with that

---

<sup>17</sup> Ibid., 469

attitude; for instance, things might happen that cause him to become uncertain, or his attitude might simply evaporate. However, an inconsistent sentence will arise once it is ensured that the severely negative clause is embedded so that it is clear that the man's current attitude is at issue. A sentence like, "He is absolutely certain that there are automobiles, but his attitude is that he may change his mind should certain evidence come up," is always inconsistent.<sup>18</sup>

Now, Unger does not think that this linguistic evidence suffices to show that being certain, or the attitude in knowing, demands the severe attitude posited above. Rather, he thinks that what is needed is to fit this severe attitude into a more general account of things. Thus, he invokes his account of absolute terms. According to Unger, absolute adjectives like "flat," "useless," and "certain" purport to denote a limiting state or situation to which things may approximate more or less closely. Thus, with respect to such adjectives, modifiers such as "absolutely," "completely," and "perfectly" are redundant apart from being points of emphasis. Unger admits that various locutions that contain absolute terms such as "certain" will appear to indicate that certainty admits of degrees; however, he maintains that such locutions will always admit of a paraphrase where this appearance is dispelled in favor of a more explicit reference to an absolute limit. Thus, "That's pretty certain" can be better paraphrased as, "That's pretty close to being absolutely certain"; and "He is more certain of this than of that," can be better paraphrased as "He is absolutely certain of this but not of that or else he is closer to being absolutely certain of this than of that".

---

<sup>18</sup> Ibid., 471



Similar paraphrases are possible for any proposition containing an absolute adjective.<sup>19</sup>

Additionally, Unger thinks that absolute adjectives typically have contrasting terms that are relative adjectives: for instance, “certain” has “confident” and “doubtful”; “flat” has “bumpy” and “curved”; and “useless” has “useful” and “serviceable”. Unger points out that, unlike absolute terms, relative terms do in fact admit of matters of degree. Thus, the sentence “He is pretty confident,” cannot be paraphrased as “He is pretty close to being absolutely confident”, and the sentence “That is very useful,” cannot be paraphrased as, “That is very close to being absolutely useful”. Because such terms really do admit of matters of degree, it does not follow from the fact that a person is, say, confident, that no one else could ever be more confident. However, this is the case when it comes to any absolute term.<sup>20</sup>

It is also worth noting that, on Unger’s account of absolute terms, it is a necessary condition for the correct application of an absolute adjective that certain things denoted by relative adjectives are entirely absent. With respect to the adjective “certain,” if an individual is certain of something, then there must be a complete absence of doubt as to whether or not that thing obtains.<sup>21</sup>

Now, if one acknowledges this last claim, says Unger, something further should be recognized: any openness on the part of an individual to consider new experience or information as seriously relevant to the truth or falsity of a claim must also be absent if one is to be certain of that claim. To put it more formally, if S feels certain

---

<sup>19</sup> Ibid., 471

<sup>20</sup> Ibid., 472

<sup>21</sup> Ibid., 472

that p, then it follows that S is not at all open to considering any new experience or information as relevant to his thinking in the matter of whether p.<sup>22</sup>

So, given all this, why is it never all right to have this “severely negative” attitude? Just what is wrong with not allowing any new evidence to have any effect at all on one’s thinking on certain matters?

According to Unger, it is never all right to have such an attitude because having it is always dogmatic. This is because, for any proposition that one purports to know, there is always some experience or sequence of experiences that one could have that would, or at any rate, should cause one to begin to doubt, or become uncertain of, that proposition. With respect to the case of the ink-bottle mentioned above, there are a number of experiences that would cast some doubt on the proposition that, “There is an ink-bottle before me,” in addition to those already mentioned. For instance, one can imagine being approached, say, by government officials who seem to demonstrate that the object on one’s desk is a container of a material to poison the water supply, which somehow found its way out of government hands and into one’s home. This container is disguised to look like an ink-bottle, but it is seen to have many small structural features essential to such a container of poison but which no ink-bottles have. One might think, then, that this object is not really an ink-bottle, even if it were to contain ink; rather, it is something else. At any rate, an attitude that one would disregard such evidence as being relevant to the truth or falsity of the proposition

---

<sup>22</sup> Ibid., 472

that, "There is an ink-bottle before me," is dogmatic, as such an experience should not simply be disregarded should it occur.<sup>23</sup>

Unger notes that, at this point, one might wish to retreat, claiming that one can at least be absolutely certain that "There is something that looks like an ink-bottle before me," or perhaps something else along those lines. So, though the above experience might count as evidence against the proposition that, "There is an ink-bottle before me," it does not count against a number of other propositions that one might know instead.

Unger holds, however, that for any proposition that one might represent oneself as knowing, one can always imagine some experience that would cast doubt on it. Even in what he terms, the "hardest cases" (such as cases about one's direct experiences, mathematical or logically necessary truths, or about whether or not the world or oneself has existed for more than five minutes), one can always imagine some experience that could (and indeed, should) cause one to doubt propositions that really do seem (absolutely) certain at first glance. For instance, with respect to the matter of what the sum of 1 and 1 is, one can imagine being confronted by a group of mathematicians who insist that 1 and 1 is 3, rather than 2. With respect to necessary truths, one can imagine hearing a voice that explained that people oftentimes become confused by the meanings of certain terms. And so on. Though the situations are far-fetched, they are possible, and should they obtain, says Unger, we should find our confidence about certain propositions diminished.<sup>24</sup>

---

<sup>23</sup> Ibid., 474

<sup>24</sup> Ibid., 477

Unger admits that one might object to him at this point by claiming that, if any such far-fetched scenario occurred, a more natural response would be to think that one had gone mad, or at least had his capacity for judgment impaired in some way. As he notes, however, this attitude of prospective self-defeat is quite compatible with lessening one's confidence. Thus, even if one did think herself quite mad as a result of experiencing such a far-fetched scenario, as Unger sees it, one should still find herself in a position where she would become uncertain nonetheless. And, because knowledge entails its being all right to be certain, then one can never know anything; as it is clearly never all right to be certain.<sup>25</sup>

So, what are we to make of Unger's arguments? After all, if one accepts them, then we have no knowledge at all: an unacceptable conclusion for most. Thus it is hardly surprising that Unger's brand of skepticism has not exactly caught on. It simply seems to require too dramatic a revisal of people's linguistic behavior, for, if Unger is correct, people utter false statements all the time when they attribute knowledge. However, one might think that, in the absence of good reasons not to, people ought only say what is true. Unger, however, provides no such reasons here; thus his account might be taken as requiring us to change the way we talk and think. Unger also seems to claim that even thinking that one knows a certain proposition involves one in having a dogmatic attitude that is always unjustifiable, which might, at first pass, sound a bit strong.

That said, I think that there is much right with Unger's claims. And, though mounting a complete evaluation of his arguments is beyond the scope of this paper, I

---

<sup>25</sup> Ibid., 475

would like to briefly comment on what I take to be their strengths and weaknesses here.

First, the linguistic evidence that Unger brings up appears to strongly support the claim that knowledge requires at least that one *feels* certain. Any sentence of the form, “S knows that p, but S is not certain,” is bound to sound inconsistent. The linguistic evidence also seems to support the claim that knowledge requires that it is all right to be certain; after all, it is certainly not sufficient for knowledge that one merely feels certain. Rather, it seems as if one must be justified in feeling certain, so to speak, perhaps in virtue of the evidence one has (among other things).

Unger is also right, I think, in saying that we often employ the term “know” rather loosely. We often claim to have known, say, where certain things were located, what team would win the game, or what the outcome of a story would be, even though, upon reflection, we can often be persuaded (and I would venture, usually very easily) to admit that we did not really know these things at all.

I am also sympathetic to Unger’s account of absolute terms (though I am not sure that the term “flat” belongs in this category). In addition to the terms that Unger mentions (“certain”, “flat”, and “useless”), a number of other terms come to mind that could possibly be classified as absolute. For instance, any term that ends with the suffix, “less,” seems to qualify as absolute: for instance, consider “flawless,” “boundless,” and “flavorless”. “Perfect” might also be an absolute term (though it is perhaps synonymous with “flawless”). So might be the term, “dead”. And, I venture (along with Unger) that “certain” and “knows” are absolute as well: they plausibly do not admit of matters of degree. Though it seems as if knowledge might admit of

degrees in expressions like, "S knows better than L that p," following Unger, I think such statements are better paraphrased as, "S is closer to knowing that p than L is".<sup>26</sup>

There is also, I think, a compelling argument to be found in Unger's defense of P2 above (the claim that it is never all right for anyone to be absolutely certain that anything is so). Though one could simply deny that knowing that p involves one in having the attitude that Unger claims one must have, I think it entirely plausible that knowing a proposition p would indeed involve one in having such an attitude, especially if infallibilism is true. For, if one infallibly knows a proposition p, then it must be impossible that one is wrong about p; thus, there would be no reason to entertain evidence against p.

However, I have some concerns about the claim that it is always wrong or dogmatic to ignore evidence against certain propositions even given infallibilism. With respect to necessary truths, for instance, it might be all right to not seriously consider evidence against them, insofar as one considers it at all. For instance, even if thousands of people were to show up at my office door tomorrow, each waiting for their turn to inform me that there are some married bachelors out there in the world, I cannot imagine allowing this to count against my knowing that "All bachelors are unmarried". I can imagine thinking that perhaps everyone else had started using the term "bachelor" in a different way, or that the whole world had gone crazy, as in some science-fiction film. But I cannot imagine letting this count against my knowledge of that proposition on my current understanding of the terms involved. I am not so sure

---

<sup>26</sup> Even if there are better or worse instance of knowledge, this does not show that knowledge necessarily admits of degrees. In an example courtesy of Jeremy Fantl, home runs do not admit of matters of degree, but there can be better or worse ones.

that my confidence would or even should be shaken any. Unger thinks that any potential scenario in which one could imagine one's confidence being shaken about a proposition *p* should suffice to show that one does not know *p*; however, I think that this is far from clear.

It is worth pointing out that other philosophers have argued that knowing a proposition *p* might involve one in having a wrongly dogmatic attitude. For instance, a problem due to Saul Kripke turns upon the closure of knowledge under implication. The problem is as follows: *p* implies that any evidence against *p* is misleading. So, whenever you know that *p*, you know that any evidence against *p* is misleading. It follows that, if you know *p*, you should ignore any evidence against *p*. The problem, of course, is that this appears to be plainly dogmatic: you should never just ignore evidence against *p*.<sup>27</sup> Thus the problem is often referred to as, "the dogmatism paradox".

Philosophers have responded to the dogmatism paradox in other ways. In *Thought*, Harman writes,

The argument for paradox overlooks the way actually having evidence can make a difference. Since I now know that Tom stole the book, I now know that any evidence that appears to indicate something else is misleading. That does not warrant me in simply disregarding any further evidence, since getting that further evidence can change what I know. In particular, after I get such further evidence I may no longer know that it is misleading. For having the new evidence can make it true that I no longer know that Tom stole the book; if I no longer know that, I no longer know that new evidence is misleading.<sup>28</sup> (1973, 149)

---

<sup>27</sup> Lewis, D. Elusive Knowledge," *Australasian Journal of Philosophy* 74:4 (1996), 564

<sup>28</sup> Harman, G. *Thought* (Princeton: 1973), 149

In effect, Harman denies what Roy Sorenson calls “the hardness principle” with respect to knowledge. According to this principle, one knows a proposition *p* only if there is no evidence such that if one knew about the evidence one would not be justified in believing *p*.<sup>29</sup> This has as a consequence that knowledge cannot be undermined by new evidence. Harman, however, denies this, as do other philosophers including Sorenson.<sup>30</sup>

However, even if non-skeptics such as Harman have at least plausible responses to the ‘dogmatism paradox,’<sup>31</sup> I think that it is important to note that, as far as the skeptic is concerned, no potentially paradoxical situation arises in the first place. If one knows that *p*, one can ignore evidence against *p*, sure; but one (at least almost) never knows that *p*, so one is (at least almost) never justified in simply ignoring evidence against *p*. It is worth emphasizing as well that my goal in this chapter is not to show that skepticism must be true, but to discuss arguments in favor of the view that it is at least plausible. Because skepticism can be seen as potentially safeguarding us against the threat of dogmatism, this suggests a way of viewing skepticism that is less perhaps less negative than has been often thought. Instead of seeing the skeptic’s arguments as “knowledge destroying” (which they surely are – but this isn’t a bad thing, necessarily), one might be able to instead regard them as “mind-opening”. That is, skepticism might be construed as protecting us from the

---

<sup>29</sup> Sorenson, R. “Dogmatism, Junk Knowledge, and Conditionals,” *The Philosophical Quarterly* 38 (1988), 437

<sup>30</sup> *Ibid.*, 433-454

<sup>31</sup> It is important to note, however, that a solution like this likely only works if one endorses fallibilism. Insofar as one finds infallibilism compelling, such a route is arguably not an option. For again, if infallibilism is true, then one can only know a proposition *p* if one’s evidence or method of determining that *p* entails *p*. And, if this is the case, no new evidence should make a difference.



dangers of dogmatism rather than simply plunging us into the darkness of ignorance.

David Hume appears to be sympathetic to this sort of skepticism. In, “An Enquiry Concerning Human Understanding,” Hume writes,

The greater part of mankind are naturally apt to be affirmative and dogmatical in their opinions; and while they see objects only on one side, and have no idea of any counterpoising argument, they throw themselves precipitately into the principles, to which they are inclined; nor have they any indulgence for those who entertain opposite sentiments. [...] But could such dogmatical reasoners become sensible of the strange infirmities of human understanding, even in its most perfect state, and when most accurate and cautious in its determinations; such a reflection would naturally inspire them with more modesty and reserve, and diminish their fond opinion of themselves, and their prejudice against antagonists. [...] In general, there is a degree of doubt, and caution, and modesty, which, in all kinds of scrutiny and decision, ought for-ever to accompany a just reasoner.<sup>32</sup>

So, though mounting a complete defense of infallibilism (and the skepticism that arguably results from endorsing it) is beyond the purview of this work, because the intuitions that underlie infallibilism are so strong, I take it to be a significant point in a theory of knowledge attributions’ favor if it is consistent with it. Thus I think that it is entirely reasonable to hold that any theory that could both explain why we attribute knowledge so often and remain consistent with infallibilism would have the advantage over its fallibilist counterparts.

## II. Skeptical Arguments from Lottery Problems

### *(i) Lottery Problems*

John Hawthorne opens his book, *Knowledge and Lotteries*, with a brief statement of the “lottery problem”:

---

<sup>32</sup> Hume, D. *Enquiries Concerning Human Understanding and Concerning the Principles of Morals* ed. L.A. Shelby-Brigge and P.H. Nidditch (Oxford:1975), 161

Suppose someone of modest means announces that he knows he will now have enough money to go on an African safari this year. We are inclined to treat such a judgment as true, notwithstanding various far-fetched possibilities in which that person suddenly acquires a great deal of money. We are at some level aware that people of modest means buy lottery tickets from time to time, and very occasionally win. And we are aware that there have been occasions when a person of modest means occasionally suddenly inherits a great deal of money from a relative from whom he had no reason to expect a large inheritance. But, despite all this, many normal people of modest means will be willing, under normal circumstances, to judge that they know that they will not have enough money to go on an African safari in the near future. And under normal circumstances, their conversational partners will be willing to accept that judgment as correct.

However, were that person to announce that he knew that he would not win a major prize in a lottery this year, we would be far less inclined to accept his judgment as true. We do not suppose that people know in advance of a lottery drawing whether they will win or lose. But what is going on here? The proposition that the person will not have enough money to go on an African safari this year entails that he will not win a major prize in a lottery. If the person knows the former, then isn't he at least in a position to know the latter by performing a simple deduction?<sup>33</sup>

As pointed out by Jonathan Vogel, the problem generalizes. For instance, I am inclined to think that I will be spending some time in Canada next year. But, once the question arises, I am not inclined to think that I will not be among the unfortunate people to die suddenly (say, in an accident or of a heart attack) before then. However, if I know that I will be spending some time in Canada next year, I should be able to deduce that I know that I will not die suddenly before then in an accident or of a heart attack. This seems unacceptable: there does not seem to be any way that I could know that!<sup>34</sup>

---

<sup>33</sup> Hawthorne, J. *Knowledge and Lotteries* (Oxford: 2004), 1-2

<sup>34</sup> *Ibid.*, 3

Further, as Hawthorne notes, the problem also generalizes to cases that do not involve the future (and indeed, that do not simply involve lotteries in which the draw has already occurred but the results are unknown). For instance, I am inclined to think that I know where my car is parked right now. But, once the question arises, I am not inclined to think that I know whether or not I am one of the unlucky people whose car has been stolen during the last few hours. And, this will obviously generalize to all cases of non-observational beliefs: who the current president of the United States is (he may have died in the past five minutes), whether my refrigerator is running (there may have been an electrical outage in my neighborhood, or it may have broken), or whether my favorite sports team won last night (perhaps there was a misprint in the newspaper).<sup>35</sup>

And, indeed, the problem potentially extends beyond cases involving putative knowledge of unobserved objects or events to cases involving putative perceptual knowledge as well. Hawthorne presents an admittedly far-fetched problem designed to show this: a person takes a pill from a bucket of 10,000 pills. One pill out of the 10,000 induces a blue-green color reversal, which is to say that it will make blue things look green and green things look blue to anyone who ingests it. The remaining pills are inert. The pill-taker knows this. As it happens, he takes an inert pill, leaving his color perception mechanisms undisturbed. He looks at a blue patch and forms the belief that it is blue. We may be inclined to count him as knowing that he is seeing a

---

<sup>35</sup> Ibid., 4

blue patch. But, we will be far less open to the suggestion that he can deduce and come to know that the pill he took was inert.<sup>36</sup>

In each of the above cases, the structure of the problem is the same: there is an ordinary proposition (a proposition of the sort that people ordinarily take themselves to know) and a lottery proposition (a proposition of the sort that, while highly likely, is a proposition that people are disinclined to take themselves to know). But, in each case, the ordinary proposition entails the lottery proposition.<sup>37</sup>

Notice too that the problem does not hinge on whether or not knowledge is infallible. That is, the lottery problem is potentially just as big a problem for fallibilists as it is for infallibilists. For, even fallibilists may find themselves uncomfortable with claiming that they know that a certain lottery ticket will not win, that they know they will not die in a car accident or of a heart attack in the next year, that their car has not been stolen, that the pill that they took was inert, and so on (despite the overwhelming probability that these propositions are true).

As Hawthorne notes, these sorts of cases generate pressure to accept a fairly extreme form of skepticism for several reasons. First, there is a strong intuitive pull to keep one's judgment about the lottery proposition and reject one's judgment about the ordinary proposition. And, even if one does not feel this intuitive pull, and thinks that one can know the relevant lottery proposition, a problem arises as a result of conjunction introduction. After all, if I can know of myself, say, that since I know I will not be able to afford a safari in the near future, I know I will not win the lottery, I can

---

<sup>36</sup> Ibid., 4

<sup>37</sup> Ibid., 5

presumably know such things about various acquaintances as well. Consider 1,000 such acquaintances: if I know of each of them that they will not be able to afford a safari, I can know of each of them that they will not win a lottery. But, assuming that I can extend my knowledge by conjunction introduction, I can now know of ALL of them that they will lose. Soon enough, I will take myself to know of a large chunk of lottery-ticket holders that they will all lose! And, if there are 1,001 ticket holders, and I can know that 1,000 will not be able to afford a safari in the near future (and thus not the lottery either), I can even know the winner by deduction! So, in an effort to keep us from knowing too little, an individual who tries to maintain that we can know the relevant lottery propositions risks allowing that we know far too much.<sup>38</sup>

## *(ii) Closure Principles*

Lottery problems rely on the extremely plausible (though not wholly uncontroversial) idea that knowledge transfers over entailments; that is, that some sort of “closure principle” holds for knowledge. Simply stated, the basic idea is this:

Basic Closure Principle: If S knows that p and S knows that if p, then q, S is in a position to deduce and thereby come to know that q.<sup>39</sup>

---

<sup>38</sup> Ibid., 6-7

<sup>39</sup> My formulation of closure principles is similar to Timothy Williamson's. On page 32 of *Knowledge and Lotteries*, Hawthorne puts Williamson's suggestion as follows: “knowing  $p_1 \dots p_n$ , competently deducing q, and therefore coming to believe q is a general way of coming to know q.” Jonathan Kvanvig too thinks that such an approach is promising: see “Contextualism, Contrastivism, Relevant Alternatives, and Closure,” *Philosophical Studies* 134 (2007), 132. Though an exact formulation of a basic closure principle that avoids counterexample has been a bit difficult to formulate, as Kvanvig puts it, “Nonetheless, it is still plausible to think that some closure principle is true. The approach I find most promising here is to strengthen the antecedent in the way suggested by Williamson and Hawthorne to

According to Timothy Williamson, closure principles articulate what is an extremely intuitive idea; namely, that deduction is a way of extending one's knowledge.<sup>40</sup> Closure principles have also been taken to be a "Moorean fact".

### III. Five Desiderata for a Theory of Knowledge Attributions

Due to the force of the skeptical arguments presented above, any theory of knowledge attributions must engage with them to some extent. Given this, I think that five desiderata for an adequate theory of knowledge attributions can be advanced. A theory of knowledge attributions should, if possible:

- 1) Either be consistent with skepticism or be able to adequately explain our skeptical intuitions
- 2) Either be consistent with the claim that much is known or be able to adequately explain our tendency to attribute knowledge fairly often
- 3) Either be consistent with infallibilism or be able to adequately explain infallibilist intuitions (though, as noted, any theory that is consistent with infallibilism is arguably preferable).
- 4) Solve the lottery problem
- 5) Be consistent with some form of closure principle

---

include the idea that q comes to be believed on the basis of a competent deduction from p. Still more qualification will be needed - e.g., it will be important that additional defeating information regarding p and regarding q is not acquired in the process of deduction - but the prospects are good for a successful principle along these lines."

<sup>40</sup> Ibid., 33

## **Chapter Two: Recent Theories of Knowledge Attributions**

In this chapter, I look at three recently proposed accounts that are motivated, at least in part, by the issues discussed in chapter one: contextualism, contrastivism, and an account advanced by Gilbert Harman and Brett Sherman. Though there have been many theories of knowledge attributions offered in response to skeptical arguments such as those presented in the previous chapter, I have chosen to discuss these in particular because they are similar to, but importantly different from, the conditional model that I sketch in the next chapter.

By examining and evaluating these accounts, I hope to illuminate two additional desiderata that any plausible theory of knowledge attributions should satisfy. I then note that, though each account is compelling in certain ways, they all fail to satisfy at least one of the seven desiderata.

### **I. Contextualist Accounts of Knowledge Attributions**

#### *(i) A Brief Exposition*

Prominent proponents of contextualist theories include David Lewis, Keith DeRose, and Stewart Cohen. Though particular brands of it are nuanced, contextualism, broadly construed, is the view that the proposition expressed by a given knowledge attribution depends at least in part on the context of the agent attributing knowledge. Contextualist theories thus all share in common the view that the truth-conditions of knowledge attributions are at least in part determined by the attributor's context. It is the "attributor" part that makes the view novel and interesting (not to mention contentious): again, "context" here is to be understood as

referring to particular features of the knowledge *attributor's* psychology or conversational/practical situation (for this reason, some choose to refer to contextualist views as being “attributor-contextualist” for clarity). This has as a consequence that the truth-values of the propositions expressed by “S knows that p,” can vary in potentially surprising ways (as a result of this variability in truth-conditions resulting from changes in the attributor’s psychology or conversational/practical interests).

It is worth emphasizing at this point that, on contextualist theories of knowledge attributions, it is only relative to contextually determined standards that knowledge sentences can express complete propositions at all. According to these theories, knowledge attributions are taken to resemble utterances involving uncontroversially context-sensitive terms, including indexicals such as “he,” “she,” “I,” “that,” and “here”. This means that utterances of a sentence of the form, “S knows that p,” can express interestingly different propositions<sup>41</sup> depending on the context in which they are uttered as a result of contextual parameters associated with the verb “know” itself.<sup>42</sup>

An example will help to illustrate how this is supposed to work. In his seminal paper, “Contextualism and Knowledge Attributions,” Keith DeRose presents two cases, which he calls “Bank Case A,” and “Bank Case B”:

---

<sup>41</sup> Notice that because knowledge attributions often feature indexicals, a given knowledge attribution can express ‘uninterestingly’ different propositions depending on the context in which it is uttered. For instance, the sentence “I know that I’ll be there tomorrow,” expresses different propositions depending on who utters it, when, and about what location. This is not the phenomenon that contextualists are interested in; rather, they are interested in variability in truth conditions as a result of the use of the verb “know” and not some obvious indexical contained within the knowledge attribution. Jason Stanley also makes this point in *Knowledge and Practical Interests* (Oxford: 2005).

<sup>42</sup> Hawthorne, J. *Knowledge and Lotteries* (Oxford: 2004), 53.



Bank Case A. My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." I reply, "No, I know it'll be open. I was just there two weeks ago on Saturday. It's open until noon."

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, "Banks do change their hours. Do you know the bank will be open tomorrow?" Remaining as confident as I was before that the bank will be open then, still, I reply, "Well, no. I'd better go in and make sure."<sup>43</sup>

De Rose asks the reader to assume i) in both cases, the bank will indeed be open on Saturday, and ii) there is nothing odd going on that is not mentioned in the above description of the cases. According to De Rose, in Case A, when he attributes knowledge of the bank's being open tomorrow to himself, he says something true. That is, in Case A, the proposition expressed by (1), "I (Keith De Rose) know that the bank will be open tomorrow," is true. However, De Rose claims that, in Case B, when he *denies* knowledge of the bank's being open tomorrow to himself, he says something true as well! That is, in Case B, the proposition expressed by (2), "I (Keith De Rose) do not know that the bank will be open tomorrow," is true. But, he has no more or better evidence in Case A than he has in Case B. And, he admits, it seems quite natural to say

---

<sup>43</sup> DeRose, K. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (1992), 913

that, as a result, (3) if he knows that the bank will be open in Case A, then he knows that the bank will be open in Case B. De Rose claims that he can indeed consistently assert (3) without having to say something in addition that he thinks would be obviously inconsistent with what he says in (1) and (2); namely, that (4) if the proposition expressed by “I (Keith De Rose) know that the bank will be open on Saturday,” is true in Case A, then the proposition expressed by “I (Keith De Rose) do not know that the bank will be open on Saturday” is false in Case B.<sup>44</sup>

How is this possible? According to De Rose (and, indeed, other defenders of contextualist theories), it is possible because the stakes and practical interests of the agent attributing knowledge (as well as what is salient to the attributor) are such that the standards for knowledge are lower in Case A than they are in Case B. As such, the proposition expressed by the sentence, “I know it’ll be open,” has different truth-conditions in Case A than it does in Case B because of contextual parameters associated with the verb “knows”. In effect, the word “know” means something different in each of the two cases. And, as a result, the proposition expressed by the sentence, “I know it’ll be open” can be true in Case A and not in Case B, without pain of inconsistency.

Similar ideas can be found in the work of others who defend contextualist theories. For instance, in “Elusive Knowledge,” David Lewis states that, “I say S knows that P iff P holds in every possibility left uneliminated by S’s evidence – Psst! – except for those possibilities that *we* are properly ignoring. ‘We’ means: the speaker and hearers of a given context; that is, those of us who are discussing S’s knowledge

---

<sup>44</sup> Ibid., 913-914

together".<sup>45</sup> Stewart Cohen also shares the view that the knowledge attributor's psychology or conversational/practical situation can affect whether or not one speaks truly when one attributes (or denies) knowledge, and, like De Rose, presents his reader with a case to consider:

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes I know it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really know that the plane will stop in Chicago. They decide to check with the airline agent.<sup>46</sup>

According to Cohen, Smith says something true when he attributes knowledge of the plane's stopping in Chicago to himself. However, Mary and John also say something true when they deny knowledge of the plane's stopping in Chicago to Smith. Cohen justifies his claims in much the same way as De Rose does: by claiming that there is no one standard for the application of the term "knowledge" that holds across all situations. Rather, there are many different standards for the application of the term, "knowledge", and it is the attributor's context that determines which standards are appropriate in a given case.<sup>47</sup>

In addition to providing linguistic evidence (such as the cases presented above), proponents of contextualist theories of knowledge attributions also argue for their position by explaining how contextualism can combat skeptical arguments

---

<sup>45</sup> Lewis, D. "Elusive Knowledge," *Australasian Journal of Philosophy* 74:4 (1996), 554

<sup>46</sup> Cohen, S. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13, 58

<sup>47</sup> *Ibid.*, 59

(including though not limited to arguments such as those presented in the previous chapter) and solve skeptical puzzles such as the lottery problems. Though contextualists of different persuasions offer slightly different solutions to these puzzles, their responses tend to take a general form.

In response to skeptical arguments from the infallibility of knowledge, contextualists generally deny that the truth-conditions for knowledge attributions are always as strict as skeptics like Unger propose. Rather, they hold that, though in some contexts, the truth conditions for knowledge attributions are that strict, in other contexts they are not. According to contextualists, there is something about skeptical arguments, or the possibilities that they involve, which effect or reflect a shift in context, such that there is a corresponding shift in the relevant standards.<sup>48</sup> To put it another way, contextualists generally argue that the skeptic somehow manages to change the context from low-standards to high-standards in the course of her argument. As Keith DeRose says,

The fact that the skeptic can thus install very high standards which we don't live up to has no tendency to show that we don't satisfy the more relaxed standards that are in place in ordinary conversations. Thus, it is hoped, our ordinary claims to know will be safeguarded from the apparently powerful attacks of the skeptic, while, at the same time, the persuasiveness of the skeptical arguments is explained.<sup>49</sup>

Thus, contextualists maintain that the skeptic is being equivocal in her arguments. As Jonathan Schaffer puts it,

---

<sup>48</sup> Rysiew, P. "Epistemic Contextualism," in the *Stanford Encyclopedia of Philosophy* (2007). <http://plato.stanford.edu/entries/contextualism-epistemology/>. Accessed October 24, 2008.

<sup>49</sup> De Rose, K. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (1992), 917

The idea here is [that it is right] in ordinary context C1 to say (i) 'Moore knows that he has hands' while the skeptic who brings up extraordinary possibilities such as being a brain-in-a-vat is right in such an extraordinary context C2 to conclude (ii) 'Moore does not know that he has hands'. The skeptic is just wrong to think that these conflict. There is no more conflict between (i) and (ii) then there would be between Ann's saying (iii) 'I am at the Haymarket café' and Ben's saying (iv) 'I am not at the Haymarket café'. Any appearance of conflict is due to an equivocation on the reference of a term [that behaves like or is] an indexical (be it 'knows' or 'I').<sup>50</sup>

Thus, the contextualist can solve the lottery problem as follows: Though there is a sense in which the statement "S knows that he will not be able to afford a safari next year," expresses a truth, there is another sense in which that statement does not express a truth. If the individual attributing knowledge to S is not seriously considering the possibility that S will, say, win the lottery, then the proposition expressed by "S knows that he will not be able to afford a safari next year," is true. However, once the attributor seriously considers that possibility, we have a context shift in which the verb "know" expresses a new relation such that the proposition expressed by "S knows that he will not be able to afford a safari next year," is false.<sup>51</sup>

Contextualists think that their view has another significant advantage: in addition to solving the skeptical puzzles, they say, contextualism squares well with our everyday knowledge-attributing practices. For, it does seem as if contextual factors are relevant to whether or not we are willing to attribute knowledge in ordinary cases. They cite examples such as the Bank and Airport Cases presented above as evidence for this, noting that contextualism has the resources to explain our intuitions about such cases, whereas many other theories of knowledge attributions

---

<sup>50</sup> Schaffer, J. "From Contextualism to Contrastivism," *Philosophical Studies* 119 (2004), 75

<sup>51</sup> Hawthorne, J. *Knowledge and Lotteries* (Oxford: 2004), 80-81

cannot adequately account for them. That is, contextualism can explain the at times shifty acceptability of knowledge attributions.

## *ii) Potential Problems*

Though a thorough exposition of the various potential problems with contextualist accounts of knowledge attributions (and the responses to them) is beyond the scope of this paper, it is worth highlighting several of what have been taken to be the most serious problems with contextualism (if only to motivate the account of knowledge attributions that I offer in the next chapter).

Among the most serious objections to contextualism is that “knows” simply does not seem behave like indexicals or other context-sensitive terms. This is evidenced by the fact that we do not have much trouble using and understanding propositions containing other indexicals or less controversially context-dependant expressions. Consider the following exchange: I am walking down the street with a friend, and we pass a gentleman of considerable height. After passing him, I mention to my friend that the man “sure was tall.” I venture that not a single friend of mine (or indeed, of anyone else’s) would ever respond to this by saying something like, “That man was not tall at all. After all, his height was only a tiny fraction of the height of the Sears Tower”. Such confusions do not seem to occur often when it comes to indexicals or less-controversially context-sensitive terms like “tall”, “short”, “flat”, “he”, “she”, “I”, and the like. However, if contextualist theories are correct, then this is the sort of mistake that we’re making concerning the term “knows” all the time. So, if knowledge is a context-sensitive term in the way that contextualists specify, why do we have such

trouble recognizing it as such, when we do not have the same sort of trouble with less-controversially context sensitive terms that to which it is purported to be similar?

Philosophers such as Jason Stanley and John Hawthorne have pressed similar lines of argument.<sup>52</sup> As Stanley notes in, *Knowledge and Practical Interests*, in the following exchange, every statement expresses a truth:

A: (Looking at a zebra in a normal zoo). I know that is a zebra.

B: But can you rule out its being a cleverly painted mule?

A: I guess I can't rule that out.

B: So, you admit that you don't know that's a zebra, and so you were wrong earlier?

A: I didn't say I did. I wasn't considering the possibility that it could be a cleverly painted mule.

As Stanley points out,

A's final utterance, according to the contextualist's semantics [expresses a proposition that] is perfectly true. But this seems a very strange result. It is extremely difficult to make sense of A's denial except as a lie. But instead the contextualist predicts that it [expresses something] true.<sup>53</sup>

Contextualists generally respond to this sort of objection by positing an error theory of some kind, claiming that we fail to fully appreciate the contextualist semantics or faithfully track shifts in context. In the words of Keith DeRose, we are "bamboozled by our own words".<sup>54</sup> According to De Rose, "

---

<sup>52</sup> See Hawthorne, J. *Knowledge and Lotteries* (Oxford: 2004), 51-112 and Stanley, J. *Knowledge and Practical Interests* (Oxford: 2005), 16-84 for their discussions of contextualism and its difficulties.

<sup>53</sup> Stanley, J. *Knowledge and Practical Interests* (Oxford: 2005), 52

<sup>54</sup> DeRose, K. "Bamboozled by Our Own Words," *Semantic Blindness and Some Arguments Against Contextualism*, *Philosophy and Phenomenological Research*, 72: 316

[...] When it looks as if the contextualist has to say something strongly counter-intuitive, what he must say turns out to be, on the contrary, something fairly theoretical concerning the truth-conditions of certain sentences. Do we really have strong intuitions about such things?<sup>55</sup>

Many philosophers, however, find this approach highly implausible.<sup>56</sup> Though perhaps these concerns do not comprise knock-down objections to contextualism, the fact that “knows” just does not seem to be context-sensitive in the way that contextualists claim is a major point against it: A theory that could solve the same problems and account for the same intuitions without having to resort to claiming that “knows” is context-sensitive in this way would be preferable.<sup>57</sup>

Contextualist theories of knowledge attributions have a number of other counterintuitive consequences. As its proponents often admit, if anything, we seem to be intuitive invariantists. As Patrick Rysiew puts it, “it is very difficult to for many non-contextualists to accept that there could ever be a case in which an agent S has the same psychology and evidence as an agent R with respect to a proposition p, and the sentence “S knows that p” expresses a truth, whereas “R knows that p” does not, simply because S and R are being evaluated from within different contexts that have to do with the conversational/practical interests and psychology of the individual *attributing* knowledge.”<sup>58</sup>

---

<sup>55</sup> DeRose, K. “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52 (1992), 927

<sup>56</sup> Including but not limited to, Jason Stanley, Jonathan Schaffer, and John Hawthorne.

<sup>57</sup> Hawthorne, too, seems to share this sentiment; see *Knowledge and Lotteries*, 80.

<sup>58</sup> Rysiew, P. “Epistemic Contextualism,” in the *Stanford Encyclopedia of Philosophy* (2007). <http://plato.stanford.edu/entries/contextualism-epistemology/>. Accessed October 24, 2008.



Because of considerations such as those presented above, I venture that a major point against contextualism is that it does not seem to get the phenomenology of knowledge attributions right. What we want a theory of knowledge attributions to do, ideally, is accurately capture what seems to be going on when we attribute knowledge. Indeed, this seems to be what we want out of any philosophical theory that is about something we “do”. Though this perhaps isn’t a necessary constraint on a satisfactory account of knowledge attributions, I think it reasonable to hold that any theory that accurately captured the phenomenology of knowledge attributions while satisfying other desiderata for a theory of knowledge attributions would be preferable to one that did not.

The last line of objection that I will mention here concerns the possibility that contextualists fail to respond to the skeptic effectively. This is because contextualist theories are silent as to whether or not we *really ever know* anything; for, the theory only addresses the question of whether or not we ever speak truly when we utter sentences of the form “S knows that p”.

## **II. Contrastivist Accounts of Knowledge Attributions**

### *(i) A Brief Exposition*

Jonathan Schaffer, the main proponent of contrastivism, offers both a negative and positive characterization of the view. On the negative characterization, contrastivism is understood as denying the fairly widespread assumption that “knows” denotes a binary relation of the form  $K_{sp}$  (S knows that p). On traditional views, it is generally supposed that there is one  $K_{sp}$  relation; on contextualist views, it

is held that there are many. On contrastivist views, however, no such relation exists; rather, it is held that the assumption that knowledge is a binary relation is an error due to the seductive pull of the surface grammar of utterances involving the verb “knows”. On the positive characterization, contrastivism is the view that knowledge is a ternary relation of the form  $Kspq$  (S knows that p rather than q), where q is a contrast proposition.<sup>59</sup>

As Schaffer notes, this characterization raises several clarificatory questions, including the following: how does q factor into knowledge attributions? Also, why believe that q actually exists?<sup>60</sup>

In response to the first question, Schaffer holds that the relation between “knows” and “ $Kspq$ ” is that “knows” always denotes a three-place relation of the form “ $Kspq$ ” even when “knows” appears to denote only a two-place relation of the form  $Ksp$ . Schaffer acknowledges that some may wish to dismiss the hypothesis of a hidden contrast variable out of hand. He, however, argues that q is syntactically real by offering several diagnostics for the existence of covert variables, and showing how each suggests that “knows” is indeed ternary. The strongest arguments in favor of this view are what Schaffer calls the argument from contrastive ascriptions and the argument from focus.

With respect to the argument from contrastive ascriptions, according to Schaffer,

One diagnostic for covert variables is the existence of overt counterparts that articulate the same lexical items plus an extra argument place. For instance,

---

<sup>59</sup> Schaffer, J. “From Contextualism to Contrastivism,” *Philosophical Studies* 119 (2004), 77

<sup>60</sup> *Ibid.*, 77

one reason for thinking that “prefers” denotes a ternary relation despite the existence of binary preference attributions such as “Ann prefers chocolate” is the existence of overtly contrastive preference attributions such as “Ann prefers chocolate rather than vanilla”.<sup>61</sup>

Knowledge attributions, he points out, have overtly contrastive forms too.

Though there are binary ascriptions such as “Holmes knows that Mary stole the bicycle,” there are contrastive ascriptions such as “Holmes knows that Mary stole the bicycle rather than the wagon”. In his view, to consider only the binary ascriptions is to commit a sampling error.<sup>62</sup>

With respect to the argument from focus, Schaffer writes,

Yet another diagnostic for [the presence of] covert variables is the existence of semantic sensitivity to focal differences. For instance, ‘Ed prefers *drinking* tea’ and ‘Ed prefers drinking *tea*’ differ in truth conditions: if Ed’s overall preference ranking is drinking coffee, then drinking tea, then bathing in tea, then the former is true but the latter false.<sup>63</sup>

An example due to Fred Dretske reveals that knowledge ascriptions can also generate semantical focus sensitivity. Consider the knowledge attribution, “S knows that Clyde sold his typewriter to Alex”. If the emphasis is on sold, then S knows that Clyde sold his typewriter to Alex rather than gave it or lent it to him, for instance. But, if the emphasis is on Alex, all S must know is that Clyde sold his typewriter to Alex rather than to someone else.<sup>64</sup>

In addition to arguing that the best linguistic treatment of knowledge attributions may require that *q* is syntactically real, Schaffer argues for the

---

<sup>61</sup> Ibid., 78

<sup>62</sup> Ibid., 78

<sup>63</sup> Ibid., 79

<sup>64</sup> Ibid., 79

contrastivist position by claiming that the presence of  $q$  allows for a solution to skeptical problems. For instance, when faced with the question, “Does Moore know that he has hands,” the skeptic will answer, “no,” whereas others might answer, “yes”. If  $q$  is recognized, these answers are potentially consistent as a result of differences in the  $q$ -values expressed: the skeptic is correct in saying “no” because Moore does not know that he has hands rather than vat-images of hands, and others are correct in saying “yes” because Moore does know that he has hands rather than stumps. Notice too that the lottery problem can be solved in much the same way. Though  $S$  might know that he will not be able to afford a safari rather than having his paychecks cover the costs of the trip in addition to his bills and other expenses,  $S$  might not know that he will not be able to afford a safari rather than win the lottery.

Schaffer also points out that the existence of a covert variable  $q$  could explain acceptability differences among knowledge ascriptions, including the at times shifty acceptability of binary knowledge ascriptions. Consider, for instance, the statement ‘I know that it is 3:15’. As Schaffer sees it, such a statement is ordinarily acceptable on the basis of a glance at a clock, but not when the accuracy of the clock is in question; the explanation for this, on a contrastivist view, is that in the ordinary case, the contrast does not include the possibility that the clock is not keeping the correct time, whereas the introduction of the question of accuracy shifts the implicit contrast to include this possibility.<sup>65</sup>

It should be pointed out that contrastivism is quite similar to, but importantly different from, contextualism. Contrastivism is similar to contextualism in that, on

---

<sup>65</sup> Ibid., 81

both views, knowledge attributions have context-dependent truth conditions that are factored into the truth-conditions, and many of our ordinary knowledge attributions are shielded from the threat of skeptical doubt. However, there are several key differences, the most important of which has to do with how the truth-conditions of knowledge attributions can vary as a result of contextual parameters associated with the verb “know”. On contextualist theories, the truth-conditions of knowledge attributions can vary as a result of the occurrence of “knows” because “knows” is treated as an indexical; on contrastivist theories, however, truth-conditions are context- dependant because “knows” denotes a ternary relation of the form,  $K_{spq}$  in which the contrast variable,  $q$ , is determined contextually. And, as explained above, contextualism and contrastivism solve skeptical puzzles differently.<sup>66</sup>

### *(ii) Potential Problems*

Among the most serious objections to contrastivism is that “knows” simply does not seem to denote a ternary relation. Consider the following exchange:

S: “Ann prefers chocolate ice cream.”

J: “To what?”

S: “Oh, to vanilla ice cream.”

I venture that the reason it seems natural for J to ask S what Ann prefers chocolate ice cream to is because the statement that S uttered just does not seem complete without more information because there is an unidentified covert variable in the mix. Compare this to:

---

<sup>66</sup> Ibid., 82

S: "Holmes knows that Mary stole the bicycle"

J: "Rather than what?"

S: "Uh..."

S's statement in the second exchange does not seem to stand in need of further clarification; rather, it seems natural to respond to J's question with bewilderment. Though some of Schaffer's examples of knowledge attributions seem to fit a ternary model for "knows", for many knowledge attributions, it arguably would take quite a lot of thinking to come up with a relevant contrast proposition, if indeed it is possible to come up with one at all. For instance, consider a case of knowledge pertaining to a mathematical truth. If I purport to know that  $5 + 5 = 10$ , do I really purport to know this in contrast to something else? What would the contrast expression be here? Am I claiming to know that  $5 + 5 = 10$  rather than  $5 + 5 = n$  for any  $n$  not identical to 10? This seems strange. Or, take the sentence that, "I know my name is Julianne." Do I really purport to know that "I know my name is Julianne rather than (any number of names)?" I think not. For many knowledge attributions, there will arguably not be a plausible contrast expression in play.

If one is sympathetic to this sort of objection, one might wonder what we ought to make of Schaffer's arguments in favor of treating "knows" as syntactically ternary. I suggest that the linguistic evidence that he provides can be arguably accounted for on a less controversial binary model. Thus, "Holmes knows that Mary stole the bicycle rather than the wagon," is perhaps better translated as, "Holmes knows that Mary stole the bicycle and that she did not steal the wagon", and "S knows that Clyde sold

his typewriter to Alex,” (with the emphasis on sold) becomes, “S knows that Clyde sold his typewriter to Alex and that he did not give or loan it to him”.

Another potential objection is that contrastivists are arguably committed to an error theory similar to the one advanced by contextualists. After all, we arguably do not have a lot of trouble using and understanding other ternary expressions, so why all the confusion surrounding, “knows”?

As a result of these sorts of concerns, I venture that, like contextualism, contrastivism just does not seem to get the phenomenology of knowledge attributions right. That is, it just does not seem to accurately capture what it is that we are doing when we attribute knowledge. Though it must be emphasized that this is not a knockdown objection to the view, I take it to be a serious drawback.

Another potential problem with contrastivist accounts of knowledge attributions is that they may have to deny closure principles or risk their account not being able to effectively respond to the skeptic. Jonathan Kvanvig argues for this in, “Contextualism, Contrastivism, Relevant Alternatives, and Closure”:

Suppose it is known that Jack is a brother rather than a sister of Jill. Being a sister entails not being an apparition produced by an evil genius, but contrastivism will not be able to mimic the anti-skeptical stance of contextualism if this entailment is allowed to play a role in some closure principle. First, one cannot supplant the contrast claim involving sisterhood with lack of apparitionhood, because being a brother does not contrast with lack of apparitionhood. If supplanting is to occur, it will have to be with the negation of the entailment, so that knowing that Jack is a brother rather than a sister implies knowing that Jack is a brother rather than an apparition. Schaffer explicitly rejects expanding or replacing the contrast claim, and the reason is clear: one can only mirror the virtue of contextualism with respect to skepticism by keeping the contrast claim the same or strengthening it. So, knowing that Jack is a brother rather than a sister of Jill can imply knowing that Jack is a brother rather than an older sister of Jill, but it cannot imply knowing

something about Jack in contrast either to experiencing some apparition or not experiencing such an apparition.<sup>67</sup>

Kvanvig then argues that, because of this, the contrastivist cannot both respond to the skeptic and save closure. His argument depends on three assumptions: first, a contrast proposition is a logical contrary, second, no closure principle can allow a weakening of the contrast claim, and third, any closure principle will allow weakening of the knowledge claim.<sup>68</sup>

According to Kvanvig, these assumptions yield the following simple contrastive closure principle:

Contrastivist Closure Principle: If S knows p rather than q, p entails r and t entails q, then S knows r rather than t.<sup>69</sup>

But, says Kvanvig, this is incompatible with contrastivism's being an anti-skeptical view. This is because not all contraries to a claim are contraries to claims implied by the original claim. Kvanvig uses the following example to show this: "Jack knows that his pet is a dog rather than a cat." As Kvanvig notes, being a dog implies being an animal. But, the contrastive knowledge attribution can't be "Jack knows that his pet is an animal rather than a cat," because being an animal does not contrast with being a cat. Further, since every concept stronger than being a cat (like being a Persian cat) will also entail being an animal, the contrast claim has to be either weaker or logically unrelated. However, as noted above, such a maneuver would result in

---

<sup>67</sup> Kvanvig, J. "Contextualism, Contrastivism, Relevant Alternatives, and Closure," *Philosophical Studies* 134 (2007), 133

<sup>68</sup> One might also worry that contrastivist accounts fail to respond to the skeptic effectively for other reasons. For, one can imagine the skeptic resisting even the claim that, say, "Moore knows that he has hands rather than stumps".

<sup>69</sup> Ibid., 134



contrastivism's being unable to counter the skeptic's arguments.<sup>70</sup>

### III. Harman and Sherman's Account

#### *(i) A Brief Exposition*

In, "Knowledge, Assumptions, Lotteries," Harman and Sherman note that they endorse two theses about knowledge, one of which is the following thesis, which they term, "T1":

T1: What one knows can and usually does rest on assumptions one justifiably takes for granted without knowing them to be true.<sup>71</sup>

Harman and Sherman note that T1 seems highly intuitive to them, and provide an example similar to the examples outlined in the first part of this paper. They ask the reader to consider the case of an ordinary individual that they will call Sam, an individual who knows that he owns a car that is presently parked outside in front of his house. As Harman and Sherman put it, Sam's knowledge rests on various assumptions that he does not know but justifiably takes for granted: that there is an external world including cars and houses, that he is not a brain in a vat who simply imagines that he has a car and a house, and that no one has taken his car away since he parked it in front of his house an hour ago.<sup>72</sup>

As Harman and Sherman note, though part of the explanation of Sam's knowing that his car is presently parked outside in front of his house is that he justifiably (and truly) takes it for granted that the car hasn't been stolen, the truth value of the

---

<sup>70</sup> Ibid., 135

<sup>71</sup> Harman, G. and Sherman, B. "Knowledge, Assumptions, Lotteries" *Philosophical Issues* 14 (2004), 492

<sup>72</sup> Ibid., 493

knowledge attribution is not relative to the things that Sam “takes for granted”. That is, the claim that, “Sam knows his car is parked outside,” is not elliptical for something more complex.

Further, Harman and Sherman note that, on their view, knowledge that *p* which rests on justifiably taking it for granted that *q* is not just knowledge that, if *q* then *p*. This is presumably because they think that this would make knowledge too easy. They write, “Sam knows that if he is not a brain in a vat, he is not a brain in a vat. But he is not justified in coming to believe that he is not a brain in a vat and he cannot come to know that he is not a brain in a vat just because he is taking that for granted”<sup>73</sup>.

It is worth noting that Harman and Sherman do not seem overly concerned about skeptical arguments from the infallibility of knowledge; rather, their focus here is on the lottery problem. However, it is not hard to see how they might respond to such arguments: they might say that one is simply assuming that certain far-fetched scenarios do not obtain. Thus, perhaps one can “infallibly” know a proposition *p* in the sense that (for example) his or her evidence (or the method that she uses to conclude that *p*) plus her assumptions might make it impossible that he or she could be wrong.

Harman and Sherman solve the lottery problem as follows: *S* knows that he will not be able to go on safari next year because *S* is assuming that he does not win the lottery (among other things). But, he does not know that he will not win the lottery because there is nothing licit that he could assume in this case.

---

<sup>73</sup> Ibid., 493

*(ii) Potential Problems with Harman and Sherman's Account*

The principal difficulty with Harman and Sherman's account is that, in order for it to be tenable, they must deny closure principles of the sort first mentioned in chapter one. Recall the basic closure principle (BCP): if S knows that p, and S knows that p entails q, then S is in a position to deduce that q and thereby come to know that q.

Harman and Sherman need to deny closure principles for several reasons. First, one would be able to derive knowledge of the propositions that he or she is only entitled to assume if strong closure principles were true. For instance, Sam knows that his car is parked outside (assuming that it hasn't been stolen). However, Sam should also know that if his car is parked outside, then it has not been stolen. Thus, if closure principles are true, Sam also must be able to know that his car has not been stolen. However, it does not seem right that one should be able to come to know a proposition merely by assuming it. Second, unless Harman and Sherman denied closure, their account would be unable to provide a solution to the lottery problem: after all, if S really knows that he will not be able to afford a safari (which on Harman and Sherman's account, he does), if a closure principle of the sort just discussed is true, S should also be able to deduce that he will not win the lottery and thus come to know that as well (which on Harman and Sherman's account, he cannot; further, that he cannot know this is one of the intuitions at the heart of the problem).

It is worth noting that Harman and Sherman argue against closure principles, and do not take it to be a bad thing that their view is inconsistent with them. However,

given the intuitive plausibility of such principles, any theory that could accommodate them would arguably be preferable than one that could not, all things considered.

One might also worry that Harman and Sherman need to tell a more complete story about the nature of these assumptions. They do note that one must justifiably and truly take the propositions that she is assuming for granted, but they do not delve much more into what exactly this might involve. They also do not think that these assumptions are beliefs, though they must be justified. Thus, these assumptions can't just be the kind of assumptions that one might make when doing proofs in math or logic, or certain problems in physics, nor can they be beliefs about what one thinks obtains or will obtain (without knowing). In other words, for Harman and Sherman, the kinds of assumptions on which our knowledge rests aren't just propositions that are merely stipulated or taken for granted without preponderance of the facts; for they must be justified. But, they aren't beliefs either. More needs to be said about just what they are, how they work, and how we work with them.

#### **IV. Two Additional Desiderata for a Theory of Knowledge Attributions**

Given some of the issues addressed in this chapter, at this point two more desiderata for a theory of knowledge attributions can be advanced. In addition to the desiderata presented in the previous chapter, a theory of knowledge attributions should, if possible:

- 6) Accurately capture the phenomenology of knowledge attributions
- 7) Explain the at times shifty acceptability of knowledge attributions

## V: The Seven Desiderata and These Three Accounts

At this point, it might be useful to briefly make note of how each theory measures up with respect to the seven desiderata I have suggested.

Recall that a theory of knowledge attributions should, if possible:

- 1) Either be consistent with skepticism or be able to adequately explain skeptical intuitions
- 2) Either be consistent with the claim that much is known or be able to adequately explain our tendency to attribute knowledge fairly often
- 3) Either be consistent with infallibilism or be able to adequately explain infallibilist intuitions (though, as noted, any theory that is consistent with infallibilism is arguably preferable).
- 4) Solve the lottery problem
- 5) Be consistent with some form of closure principle
- 6) Accurately capture the phenomenology of knowledge attributions
- 7) Explain the at times shifty acceptability of knowledge attributions

As argued above, contextualism fails to satisfy 6, and may also not adequately satisfy 1. Dedicated infallibilists such as Unger may also feel that it does not adequately satisfy 3 either, as, on contextualist theories, there are indeed many circumstances in which propositions of the form “S knows that p” can express truths even if it is possible that S could be wrong about p. Contrastivism also arguably fails to satisfy 6, and may also have some difficulties with 1, 3, and 5. Finally, Harman and Sherman’s account does not satisfy 5, and is arguably too vague to decisively satisfy

many of the others, especially 3 and 7. Infallibilists in particular will likely be dissatisfied with the account, as it would only be, at best, “infallibilism on the cheap”; for, if consistent with infallibilism, on Harman and Sherman’s account one would have to be able to *infallibly know* propositions on the basis of assumptions that they are not even justified in fully believing!

### Chapter Three: A Conditional Model of Knowledge Attributions

In this chapter, I advance what I call a “conditional model” of knowledge attributions. This account is consistent with it being the case that, whenever an agent claims knowledge for him or herself or attributes knowledge to another in the form, “I know that p,” or “She (or he) knows that p,” this person is usually, *simpliciter*, stating something false. This is because I take the arguments surveyed in favor of a fairly widespread skepticism in the first chapter to be persuasive enough so as to warrant this maneuver. Thus, this theory of knowledge attributions does not aim to contest skeptical arguments; rather, it aims to explain when we are warranted in asserting sentences containing the verb “knows” even though we might be persistently using the term in error. For, as discussed in chapter one, one of the main problems with endorsing skepticism is that we quite often attribute knowledge to agents. So, if skepticism is indeed true, we say false things constantly, and this in turn seems to demand too drastic a revision of what we currently do and say. However, I think that one can acknowledge the truth of skepticism (if it’s true) without demanding that we radically alter our current linguistic practices and knowledge attributing behavior.

It is worth emphasizing that this account is intended to be consistent with *both* skeptical and non-skeptical theories of knowledge. For the remainder of the chapter, however, I will proceed as if skepticism is true to show that the account I propose is indeed consistent with the view.

## (I) Initial Exposition

Philosophers have commonly distinguished between what is said, strictly speaking, by a sentence, and what is implied. This account makes use of this distinction. On this account it is held that, when an individual utters a sentence of the form, “S knows that p,” even if the individual is usually, *simpliciter*, stating something false, he or she can successfully imply a proposition of the form, “S knows that if  $q_1 \dots q_n$ , then p” that may be true. This is because we often assume a number of propositions as conditions for the propositions that we purport to know. And, a more complex knowledge attribution that contains these assumptions can be implied by our simpler knowledge attributions. Further, a more complex knowledge attribution that contains these assumptions can be true. So, when we attribute knowledge to ourselves or to others, even if we, *simpliciter*, speak falsely, it is nonetheless possible for us to imply truths (or at least, what we might take to be truths).

Let us revisit an example from chapter one. If I say that, “I know that Dostoyevsky wrote *Crime and Punishment*,” it is plausible that, *simpliciter*, I utter a falsehood. As discussed in chapter one, there are many ways in which I could be wrong about that. Perhaps I am mistaking Dostoyevsky for another author because my memory is failing me in this case, perhaps the sources that I gleaned that information from are in error, perhaps Dostoyevsky stole the manuscript, or perhaps *Crime and Punishment* does not even exist (and so on).

However, when I claim that, “I know that Dostoyevsky wrote *Crime and Punishment*,” I do not claim to know this proposition *simpliciter*. Rather, I claim to know it *given a number of assumptions*. So, when I claim that, “I know that



Dostoyevsky wrote *Crime and Punishment*,” what I am plausibly implying is the claim that I know that Dostoyevsky wrote *Crime and Punishment*, *assuming* that memory serves me correctly, the sources that I gleaned the information from are not in error, Dostoyevsky did not steal the manuscript, that I’m not just imagining that the book exists, (and so on). This means that I am plausibly implying the claim that I know that, if memory serves me correctly, the sources that I gleaned the information from are not in error, Dostoyevsky did not steal the manuscript, that I’m not just imagining that the book exists, (and so on), then Dostoyevsky wrote *Crime and Punishment*.

At his point, a few clarifactory remarks are in order. First, a person attributing knowledge does not have to be conscious of a given assumption at the time of the attribution for the assumption to count as been having made or a knowledge attribution containing the assumption to be implied. This is because assumptions need only be held as *dispositional beliefs*. So, an agent can be counted as assuming a proposition *p* at a time *t* if the agent *dispositionally believes* *p* at *t*. This means that an agent can be plausibly credited with assuming many propositions at once, even if it is implausible that the agent consciously considered each or even any of these propositions. An example may help to illustrate why this is so. When I claim that I know that I have two hands, I am plausibly assuming a number of propositions, including but by no means limited to the following: I am not being deceived by a demon, I am not dreaming, and I am not hallucinating. However, it is implausible that, at least in most circumstances, I consciously consider these propositions as assumptions. Rather, they are the sort of thing that, if asked, I could sincerely say that I believed; that is, I *dispositionally* believe them. Thus it seems fair to credit me with

having assumed them, and with implying them in my knowledge attribution; for without them, my attribution is false (and I either believe that or can likely be persuaded), but I think my attribution is true.

Second, one might notice that the examples that have been provided so far have been from a first person perspective. What happens when we attribute knowledge from a second or third person one? Whose assumptions matter; the attributor's, or the subject's?

Because this account is intended to specify when we are warranted in asserting propositions containing the verb, "knows," it is the attributor's assumptions that matter. However, when an agent S attributes knowledge to another individual R, S is and in fact must be assuming things about what that R is assuming in order for his or her knowledge attribution to be warrantably assertible. For instance, if I say of an agent R, "R knows that Dostoyevsky wrote *Crime and Punishment*," what I imply is a statement of the form, "R knows that Dostoyevsky wrote *Crime and Punishment*, assuming that memory serves him correctly, that the sources that he gleaned the information from are not in error, Dostoyevsky did not steal the manuscript, that he's not just imagining that the book exists, (and so on), and assuming that R is assuming all that as well". After all, if one did not think that R is assuming that, say, the manuscript wasn't merely stolen by Dostoyevsky, one would be reluctant to attribute knowledge that Dostoyevsky wrote the book to R. Consider the following exchange between S and J about R:

S: "R knows that Dostoyevsky wrote *Crime and Punishment*."

J: "Can R rule out the possibility that he stole the manuscript from someone else?"

S: "I'm assuming that R is assuming that Dostoyevsky did no such thing."

Compare the following:

S: "R knows that Dostoyevsky wrote *Crime and Punishment*."

J: "Can R rule out the possibility that he stole the manuscript from someone else?"

S: "I have no idea."

I suggest that, in the first exchange, S was potentially warranted in asserting that R knows that Dostoyevsky wrote *Crime and Punishment*, but in the second exchange, he was not. If an agent S is to assert of another agent R that "R knows that p," S must be assuming that R is indeed assuming a sufficient number of propositions q as conditions for p in addition to assuming those propositions himself.

Third, one might wonder where one's evidence comes into the picture. On this account, evidence comes in as follows:

An agent S is warranted in asserting that an agent R knows a proposition p iff S can successfully imply a statement of the form, "R's evidence plus assumptions entail p." (In first person cases, S=R).

It should be stressed at this juncture that this conditional model of knowledge attributions is primarily an error theory if combined with skepticism. For, on this account, oftentimes, whenever an individual utters a phrase of the form, "I know that p," this individual is potentially saying something that is false, *simpliciter*. The reason

that this individual believes that he or she is saying something true is due to the fact that the individual is assuming a number of propositions in addition, which, if stated, might allow one to utter a true knowledge attribution. However, it is granted that it often will likely not be possible to credit agents with assuming a sufficient number of propositions. So, we are possibly still not going to be warranted in asserting many of our knowledge attributions, though we may think that we are.

## **(II) Restricting Assumptions**

At this point, it might be objected that the conditional model makes warrantably attributing knowledge too easy; for it seems that one can be virtually always warranted in attributing knowledge if certain assumptions are made.

It is perhaps worth emphasizing that, on this conditional model of knowledge attributions, it is possible that very few propositions are known *simpliciter*, though it is granted that many propositions are known *conditional on assumptions*. So, an agent can fail to be warranted in attributing knowledge just in case that agent's assumptions are insufficient to guarantee the truth of the implied attribution. This can happen in one of two ways. As noted above, the first way that this could happen is that the agent simply cannot be credited with having made the assumptions required to render his or her knowledge attribution true. For instance, if I were to declare that I know that I won't be able to afford a safari next year, and a friend responds with, "What if you win the lottery?", and I, instead of responding that I was assuming that I will not, exclaim, "Oh, I had not thought of that!", it seems that I plausibly should not have asserted that I knew I would not be able to afford the safari. For, not only did I not know that I

would not be able to afford the safari *simpliciter* in this case, it seems that I did not even dispositionally believe that I will not win the lottery. And the assumption that I will not win the lottery is plausibly required if I am to warrantably-assert that one knows that I will not be able to afford the safari.

The second way is that at least one of the requisite assumptions cannot be successfully implied. That is, some assumptions are impermissible because there are rules that we implicitly follow that govern the propositions that one can assume as conditions. So, though it may be the case that any knowledge claim can be made true if certain assumptions are made, it is not the case that one is allowed to just assume *any* proposition one wishes. In this section, I set out some of the ways in which one can go about attacking the permissibility of a given assumption. I describe four rules that govern whether or not a given assumption is allowed: the rule of epistemic priority, the rule of reasonability, the rule of raised standards, and the rule of truth. These rules are not simply *ad hoc*; rather, they capture certain intuitions regarding which propositions are impermissible to assume and why.

### *(i) Epistemic Priority*

The first restriction on assumptions is the rule of *epistemic priority*: in order for a proposition *q* to be a permissible condition for a proposition *p*, *q* must be epistemically prior to *p*. A proposition *q* is epistemically prior to a proposition *p* just in case, if one were to know that *p*, one could deduce that *q* and thereby get knowledge that *q*.

Consider the following statements: “I know that the Earth is spherical, assuming that I am not wrong,” and, “I know that the Earth is spherical, assuming that it is not the case that I do not know that fact”. It seems that I cannot assert that I know that the Earth is spherical given either of these assumptions, for the assumptions simply restate what I purport to know.

With this in mind, let us consider the statement, “I know that the Earth is spherical, assuming that the books that I have read which contain this assertion are not all purveyors of erroneous information about the Earth’s shape, the Earth has not recently morphed into some other geometrical form unbeknownst to me, memory serves me correctly, (and so on).” Unlike the statements above, I venture that this statement does not have a dissonant ring. Why is this? I suggest that it is at least in part because the assumptions in this case are not mere restatements of the knowledge attribution; rather, they seem to have other content; content that is *epistemically prior* to the proposition that one purports to know.<sup>74</sup>

So, the rule of epistemic priority will restrict all assumptions that simply assert the truth of or deny the falsity of knowledge claims in propositions like, “I know that p assuming that I am not wrong” and, “I know that p assuming that it is not the case that I do not know that p”; for in each of these claims, the assumption q is not epistemically prior to the proposition p that I purport to know. In these cases, the assumptions are not prior to the propositions conditional on them; rather, they are epistemically on

---

<sup>74</sup> It should perhaps be emphasized that, though adherence to the rule of epistemic priority is part of the reason that the preceding statement seems acceptable, it is, by no means, the entire story. For instance, the next rule, the rule of reasonable belief, plays a role in this as well.

par. As it turns out, however, other kinds of assumptions will also be restricted by this rule. To see this, let us examine another example.

Suppose that a friend and I are attending a luncheon for local singles. We spot a copy of the guest list and decide to take a look. The list is not organized in any way that demarcates the males and the females; it is simply an alphabetically-ordered list of names. My friend and I take a glance at the list, and she says, “I guess I know that Terry is a bachelor.” I ask her how she knows such a thing and she says, “Well, assuming that Terry is male, of course”. Such an exchange does not seem too odd. However, consider the following exchange: my friend takes a look at the list and says, “I guess I know that Terry is male.” I ask her how she knows that, and she says, “Well, assuming that Terry is a bachelor”. Now this exchange *does* seem odd. And it seems to me that the discrepancy is because of the rule of epistemic priority.

## *(ii) Reasonability*

The second restriction on assumptions is the rule of *reasonability*. This restriction is in place to disallow assumptions on the grounds that agent has more reason to believe that they are false than the agent has to believe that they are true.

An example should help to clarify why this rule is in place. Say that a very wealthy, miserly, and intensely frugal friend says to me, “I know that I will not be able to afford a new vehicle this year.” When I ask him how he knows that, he says, “Well, I am assuming that I will lose all my money”. It seems that a natural reaction would be that of incredulity. My friend should not say that he knows that he will be unable to afford the vehicle because it seems unacceptable for this person to assume such a

thing. So, in this case, the assumption is unacceptable because there are so many reasons to believe, in the case of this individual, that such an assumption is false, and very few, if any, to believe that they are true!

Conversely, if the same friend were to say, "I know that I will be able to afford a new vehicle this year," and, when I ask him how he knows this, he says, "Well, I am assuming that I do not lose all my money, the price of vehicles does not skyrocket, (and so forth)," this seems permissible. There is at least as much reason to believe that these assumptions are true as there is to believe that they are false. So it seems that my friend can warrantably assert that he will be able to afford the vehicle.

It may be interesting to note that the rule of reasonability permits at least some "neutral" assumptions. That is, the rule permits an agent to assume propositions that the agent has no reason to believe are false, but no reason to believe are true either. One might wonder at this point if these sorts of assumptions should be prohibited. After all, don't we usually require that an agent have at least some reason to believe that her assumptions are true?

The answer to this is that we sometimes do, and we sometimes do not. For instance, it seems intuitively correct that we can assume that we are not brains-in-vats, even though arguably, there is no evidence for or against this proposition.<sup>75</sup> However, in other cases, neutral assumptions plausibly are not permissible. For

---

<sup>75</sup> Likewise, it seems permissible to assume that you are not dreaming, and that there is no demon deceiving you. However, it does not seem permissible to assume that you are a BIV, that you are dreaming, or that there is a demon, which are arguably "neutral" in the sense discussed above as well. However, assuming those things would violate the rule of raised standards discussed in the next section. There are stakes concerns here, for assuming that you are a BIV (for instance) could have potentially ruinous consequences. Further, the possibility that you are not a BIV is arguably a constantly salient one.



instance, say that a friend, having recently become pregnant, tells me, “I know that I will be getting a lot of little girls’ clothes soon!” I ask her how she knows this, and she says, “Well, I am assuming that I have a girl!” It is much too early for anyone to know the gender of the fetus, and there is no other evidence to suggest that she will have a girl rather than a boy. Upon hearing this, I would be inclined to say to her, “Now, now, you don’t know that you’ll be getting a lot of little girls’ clothes soon. After all, you cannot just assume that you’ll have a girl, as you just as easily could have a boy instead.”<sup>76</sup>

Why the discrepancy here? I venture that it is at least in part because we cannot in principle get any evidence that would suggest anything about our BIV-status. However, we can in principle get evidence that will help us to determine the gender of my friend’s “future child”. Thus, it seems fair to say that neutral assumptions are permissible if we generally think that it is not possible in principle to get evidence for or against them, and not if we think that it is possible in principle to get evidence for or against them.

### *(iii) Raised Standards*

The third restriction on assumptions results from elevating the standards for knowledge. Standards plausibly determine which assumptions are permissible and which are not. It seems that, if the standards for knowledge are raised high enough, most or all but the strongest assumptions are restricted. Thus I call this restriction the rule of *raised standards*. This rule will have the consequence that, when standards are

---

<sup>76</sup> I am indebted to Jeremy Fantl for this example.

raised, it is almost impossible to assert that one knows certain propositions even given assumptions, because many assumptions will be rendered impermissible by this rule.<sup>77</sup>

Standards can be commonly raised in one of two ways. The first is by changing/increasing the stakes; the second, by increasing the salience of certain considerations.

a) Stakes

Keith DeRose's famous "bank cases"<sup>78</sup> presented in chapter two provides an excellent example of a scenario in which stakes seem to affect what can plausibly said be known.<sup>79</sup> Recall that, in Bank Case A, Keith and his wife are driving by the bank and spot a rather long line one Friday afternoon. It is not all that important that they deposit their cheques that day, so Keith suggests that they just deposit their checks Saturday morning; after all, he was in the bank just two weeks ago on Saturday, and it was open until noon. So, he knows that the bank is open tomorrow morning.

In this case, what Keith in this example is plausibly implying is that he knows that the bank is open tomorrow, assuming that memory serves him correctly, the bank has not changed its hours in the past week, the bank is not destroyed by tomorrow, the bank is not playing some cruel joke on its customers, the bank is not closed due to some accident or freak occurrence, and so on. Since these assumptions seem to be the kinds of things that are plausibly held as dispositional beliefs, are epistemically prior

---

<sup>77</sup> This rule bears similarity to David Lewis's "Rule of Attention" advanced in *Elusive Knowledge*.

<sup>78</sup> DeRose, Keith. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, Vol. 52, No. 4. (Dec., 1992), 913-914

<sup>79</sup> Though I agree with DeRose on this point, I obviously think stakes affect knowledge in a different way than DeRose does.

to the knowledge claim, and are reasonable, the assumptions seem permissible so far. Thus, Keith's claim that, "I know that the bank is open tomorrow," plausibly implies a true knowledge attribution.

In Case B, however, Keith and his wife *absolutely require* that their cheques be deposited no later than Saturday. Thus, the assumptions that Keith is entitled to make change. If enough depends on the cheque being deposited, one would want there to be as little chance of error as possible. Thus it seems plausible that, because of the elevation of potential costs, one is no longer allowed to assume that the bank has not changed its hours. Though there may not be more reason to believe that one's assumptions are false rather than true, there is no reason to believe that they are *sufficiently likely* to be true. In the "low-stakes" scenario it does not seem that there is reason to demand that assumptions be extremely likely to be true, or certainly true, in order to assume them, because nothing important really depends on their being true. However, in the "high-stakes" scenario there does seem to be such reason. High-Stakes cases, in general, seem to require a higher degree of confidence in one's assumptions. They seem to demand that you either have excellent reason to believe that your assumptions are true or that you are certain that your assumptions are true; to put it another way, you know your assumptions *simpliciter*.<sup>80</sup> But this is very unlikely to be the case for most assumptions. As a result, many assumptions will be restricted in high-stakes scenarios.

#### b) Salience

---

<sup>80</sup> Which is the case will likely depend on how high the stakes are.

Raising certain considerations to salience can also elevate standards. If a certain consideration becomes salient, it is no longer permissible to simply assume that it does or does not obtain. An excellent example of a case in which the salience of certain considerations is raised is an epistemology seminar on skepticism. Normally, when one attributes knowledge, one is (and is entitled to be) assuming that one is not a BIV. However, when this consideration is made salient in the context of the epistemology classroom, it seems that one is no longer entitled to simply assume that one is not a BIV. Thus it becomes very difficult to know anything even given assumptions in such a situation (since, of course, one can no longer assume the requisite propositions). Indeed, knowledge is difficult in any situation in which certain considerations are made salient and thus it no longer becomes permissible to simply assume them.

### **(III) Why Think That We Can Imply So Much?**

In chapter two of his seminal work, *Studies in the Way of Words*, H.P. Grice notes that in uttering certain sentences, we can imply much that is not, strictly speaking, part of the meaning of the sentences themselves in several ways. He calls this phenomena, “implicature”. Grice notes that there is a distinction between conventional and non-conventional implicature. A proposition *p* is conventionally implied by a sentence *t* if it is implied in virtue of the conventional meaning of *t*.<sup>81</sup> By contrast, if *p* is not implied in virtue of the conventional meaning of *t*, but by some other mechanism, it is non-conventionally implied. Grice is particularly interested in a

---

<sup>81</sup> Grice, H.P. *Studies in the Way of Words* (Harvard: 1989), 25-26

particular subclass of non-conventional implicatures that he calls conversational implicatures. On Grice's view, a proposition *p* is conversationally implied by a sentence *t* if it is implied in virtue of particular features of the particular conversational context/situation in which *t* occurs. Consider the following exchange:

S: "I am out of gasoline."

B: "There is a gas station around the corner."<sup>82</sup>

According to Grice, B's statement implies that the gas station is open and operational. However, it does not conventionally imply that the station is indeed open and operational; for it is not in virtue of the meaning of the sentence that such a thing could be implied. Rather, it is due to the conversational context/situation in which such a statement occurs.

According to Grice, conversational implicatures are essentially connected with certain general features of discourse; thus, his next goal is to try to say what these features are. Because our conversations are largely co-operative efforts in which each participant recognizes in them, to some extent, a common purpose or set of purposes or mutually accepted direction, Grice postulates "The Cooperative Principle" as a rough general principle that participants in conversations are expected to observe:

The Cooperative Principle (Co P): Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the exchange in which you are engaged.<sup>83</sup>

---

<sup>82</sup> Ibid., 32

<sup>83</sup> Ibid., 26

On the assumption that such a general principle is acceptable, Grice continues, four categories can be distinguished. More specific maxims fall under these categories, the following of which, Grice says, yields results in accordance with the Cooperative Principle.

According to Grice, the categories and maxims are as follows<sup>84</sup>:

The category of Quantity relates to the quantity of information to be provided in the course of conversation. These two maxims fall under it:

1. Make your contribution as informative as required for the current purposes of the exchange.
2. Do not make your contribution more informative than is required.<sup>85</sup>

The category of Quality has a super-maxim, "Try to make your contribution one that is true" and two more specific maxims:

1. Do not say what you believe to be false.
2. Do not say that for which you lack adequate evidence.

The category of Relation has only one maxim fall under it:

1. Be relevant.

The category of Manner differs from the first three categories in that it relates to how one should say things, rather than relating to what one should say. A supermaxim, "Be perspicuous," falls under this category, as do various maxims such as:

1. Avoid obscurity of expression.

---

<sup>84</sup> Ibid., 26-27

<sup>85</sup> Grice notes that this second maxim is debatable, but plausible nonetheless. Ibid., 26-27.

2. Avoid ambiguity.
3. Be brief.
4. Be orderly.

At this point, we can now look at Grice's characterization of the notion of conversational implicature: An individual who, by saying that p has implicated that q, may be said to have conversationally implicated that q provided that (1) this individual is observing the conversational maxims, or at least the Cooperative Principle, (2) this individual is aware that, or thinks that, q is required in order to make his saying that p consistent with this first presumption, and (3) this individual thinks that it is within the competence of the hearer to work out that q is required in the way mentioned in (2).<sup>86</sup>

With the above in mind, let us revisit an example of a sentence containing the verb, "knows" presented earlier: "S knows that Dostoyevsky wrote *Crime and Punishment*". On the conditional model of knowledge attributions that I am proposing, in uttering this sentence, S plausibly implies that, "I know that Dostoyevsky wrote *Crime and Punishment*, assuming that memory serves me correctly, the sources that I gleaned the information from are not in error, Dostoyevsky did not steal the manuscript, that I'm not just imagining that the book exists, (and so on)". I think that it is clear enough from the above discourse that this is a case of conversational rather than conventional implicature, as this proposition is implied in virtue of particular features of the particular conversational context/situation in which it occurs, rather than by the conventional meaning of the sentence, "I know that Dostoyevsky wrote

---

<sup>86</sup> Ibid., 31

*Crime and Punishment*". In fact, I suggest that this is the case for most, if not all, complex knowledge attributions that are implied by simpler ones; that is, they are conversationally rather than conventionally implied.

One might notice that, on the conditional model that I propose, we might be thought to consistently violate two of Grice's maxims when we attribute knowledge: the super-maxim, "try to make your contribution one that is true," and the maxim "do not say that which you believe to be false". After all, this account is intended to be consistent with skepticism. And, if skepticism is true, when we utter sentences of the form, "S knows that p," we, strictly speaking, utter falsehoods.

It is worth pointing out in response that on this account though we may, strictly speaking, utter falsehoods, we can imply truths. Moreover, we may say things that are false to avoid misleading others. For instance, compare:

"I know that Dostoyevsky wrote *Crime and Punishment*."

with

"I know that Dostoyevsky wrote *Crime and Punishment*, assuming that memory serves me correctly, the sources from which I gleaned that information are not in error, Dostoyevsky did not steal the manuscript, that I'm not just imagining that the book exists, (and so on)".

Though the latter proposition might be true and the former is false, actually uttering the latter might be quite misleading. Why? In explicitly uttering one's assumptions, one might imply that one is concerned that they obtain. It would perhaps be natural for one's conversational partner to inquire as to whether or not you had any compelling reason (that is, something much more than just the possibility



that it might not be) to think that memory might not be serving you correctly, that your sources were in error, that Dostoyevsky stole the manuscript, that the book might not really exist, and so forth. Thus, we may have the beginnings of another maxim: don't say things that you think will be misleading.<sup>87</sup>

Further, if one explicitly stated his or her assumptions every time he or she attributed knowledge, other maxims arguably would be violated. It is easy to see how the third maxim under the category of manner (be brief) would be violated if one were to start explicitly listing his or her assumptions in most situations. Perhaps most pressingly, both maxims under the category of "quantity" would plausibly be violated (make your contribution as informative as required and do not make your contribution more informative than is required). After all, there are few conversational situations in which one is going to be seriously interested in hearing about the far-fetched possibilities of error that one is assuming do not obtain when one attributes knowledge – for most of us arguably go about assuming (in the sense mentioned earlier in the chapter) that they do not obtain as well, and in the absence of good reasons to do so, we refrain from bringing them up. For this reason, the maxim, "Be relevant," would also plausibly be violated in many situations if one stated his or her assumptions explicitly.

Thus, I think that linguistic evidence gives at least initial plausibility to this sort of assessment. Consider another potential exchange between a doctor and a patient

---

<sup>87</sup> In fact, it might be the case that, sometimes, when one utters a truth, one might imply a falsehood that is so misleading that one should not have uttered the truth.

(let's call her Julianne) who has recently suffered a bit of a nasty spill and bumped her head:

Doctor: "Good afternoon. I'm Dr. Pierce. Do you know who you are? Do you know your name?"

Julianne: "Well, I don't really know."

Doctor: "No?"

Julianne: "Well, I don't know those things *simpliciter*, though perhaps I know them given certain assumptions. As far as I can recall, people have been calling me "Julianne" all my life, but there are lots of ways I could be wrong about that. After all, I could have been switched at birth, or I could be just imagining that my name is Julianne (and so on). So, yeah, I know my name: it's Julianne, assuming that those states of affairs do not obtain."

Doctor: "Right... maybe we should run some more tests."

I propose that any bewildered response on the part of the doctor in this exchange might be due to the fact that various conversational maxims are being flouted, and not due to the fact that anything false is being said. Thus, in this case, it is better to respond with the potentially false, "Sure; I know my name is Julianne," in order to avoid being confusing, misleading, and otherwise time-wasting. And, in most instances, I venture, it is generally better to stick with potentially strictly-speaking false utterances of "S knows that p," for similar reasons.

#### **(IV) What Differentiates This Account?**

As noted earlier, this account is importantly different from the accounts presented in the previous chapter. It primarily differs from contextualism (and contrastivism) in that it is not about the semantics of knowledge attributions, but rather, about the pragmatics of such attributions.<sup>88</sup> Additionally, it differs from contrastivism in that “knows” is not taken to denote a ternary relation. Finally, it differs from Harman and Sherman’s account primarily in that, on this conditional model, knowledge attributions are elliptical for more complex statements in the way that Harman and Sherman deny. Additionally, as a result of these core differences, the conditional model that I propose differs from these three accounts in the ways that it might be thought to satisfy the seven desiderata advanced in chapters one and two.

#### **(V) Why Endorse This Account?**

Recall that any plausible theory of knowledge attributions should satisfy the following desiderata:

A theory of knowledge attributions should, if possible:

- 1) Either be consistent with skepticism or be able to adequately explain our skeptical intuitions
- 2) Either be consistent with the claim that much is known or be able to adequately explain our tendency to attribute knowledge fairly often

---

<sup>88</sup> This account admittedly, however, might turn out to be equivalent to contextualism at the warranted-assertibility level.

- 3) Either be consistent with infallibilism or be able to adequately explain infallibilist intuitions (though, as noted, any theory that is consistent with infallibilism is arguably preferable).
- 4) Solve the lottery problem
- 5) Be consistent with some form of closure principle
- 6) Accurately capture the phenomenology of knowledge attributions
- 7) Explain the at times shifty acceptability of knowledge attributions

The conditional model just advanced, I think, can satisfy each of these desiderata. It should be easy enough to see how 1-3 are satisfied by this account: with respect to 1, I have tried to proceed with the exposition of this account in a manner that makes it clear how this account is consistent with skepticism's being true, if it is. And, since this account explains how and why we are warranted in attributing knowledge even if skepticism is true, 2 is satisfied as well. Further, since the account is neutral as to whether infallibilism or fallibilism is true, 3 is satisfied.

With respect to 4, the lottery problem can be solved on this account as follows: when an agent S attributes knowledge to an agent R of a proposition like "R knows that he or she will not be able to afford a safari next year," among the propositions being assumed as conditions is the proposition that S will not win the lottery. Thus, S is plausibly warranted in asserting that R will not be able to afford a safari next year. However, the reason that S is plausibly not warranted in asserting that R knows that

he will not win the lottery is that there is possibly nothing that can be permissibly assumed as a condition for that.<sup>89</sup>

With respect to 5, a proponent of this conditional model need not deny closure principles. Why not? Well, she can explain why one cannot assert knowledge of one's assumptions (simply by assuming them) another way: by referencing the rule of epistemic priority. Again, this rule states that in order for an agent to be allowed to assume a proposition *q* as a condition for a proposition *p*, *q* must be epistemically prior to *p*. A proposition *q* is epistemically prior to a proposition *p* just in case, if one were to know that *p*, one could deduce that *q* and thereby get knowledge that *q*. So, to use the "Sam and his car" example from Harman and Sherman, though Sam can potentially warrantably-assert that he knows his car is parked outside (on the condition that it has not been stolen), he cannot assert that his car has not been stolen simply by assuming that it has not, as these two propositions are clearly epistemically on par.

So, this account does not have to deny closure for those sorts of reasons, but does it in the end succeed in being consistent with it? I suggest that it does succeed in doing this; for the conditional model that I propose is consistent with its being the case that a proposition *p* is rarely known *simpliciter*. Thus, even if one knows that if *p* then *q* (*simpliciter*), since one will likely not know *p* *simpliciter*, one will not be in a position to derive knowledge of *q* *simpliciter*. Let us return to the example of Sam and

---

<sup>89</sup> Though "assuming that something really improbable does not happen" may be a candidate, the mention of the lottery arguably makes salient the possibility that something improbable could happen, and thus, *R* may not be able to assume that it does not in such cases. This also is arguably not epistemically prior to the proposition that one will not win the lottery.

his car. On this conditional model, Sam plausibly knows that “his car is parked outside, given that it has not been stolen (among other things)”. He also plausibly knows that if his car is parked outside, then it has not been stolen. However, he does not know that his car is parked outside *simpliciter*. He only knows that, “if his car has not been stolen (among other things), then it is parked outside”. And, it is not generally the case that  $((\text{if } q \text{ then } p) \text{ then } q)$ . That is, it does not follow from the fact that Sam knows that “if his car has not been stolen (among other things) then it is parked outside” that “his car has not been stolen”. Thus, he cannot make the inference from knowing that, “his car is parked outside given that his car has not been stolen” to knowing that, “his car has not been stolen”.

One might still worry that a proponent of this conditional model must deny some plausible closure principle. After all, on this account, Sam can assert that his car is parked outside (conditional on the assumption that it has not been stolen), and Sam arguably also can assert that if his car is parked outside, then it has not been stolen. So, if some form of closure principle holds for warranted assertibility, Sam should be able to warrantably-assert that he knows that his car has not been stolen as well.

The rule of raised standards, however, makes warrantably-asserting these two propositions (Sam knows his car is parked outside and If Sam’s car is parked outside, then it has not been stolen) at the same time impossible, however. For, if Sam asserts the conditional, “If my car is parked outside, then it has not been stolen,” he raises the possibility that it may have been stolen to salience with the result that it is no longer permissible for him just to assume it. And, since he cannot assume that his car has not been stolen, he cannot plainly assert that he knows his car is parked outside.

With respect to 6, I think that the conditional model that I advance here gets the phenomenology of knowledge attributions right, or at least, comes a lot closer to getting it right than do competing theories such as contextualism and contrastivism. That is, I think that it accurately describes what we are doing when we attribute knowledge. I take this to be the primary advantage of this account over contextualism and contrastivism. To see this, let us revisit DeRose's Bank Cases. What do contextualists have to say on the matter? They maintain that in Case A, the proposition expressed by "I know the bank is open," is true, as is the proposition expressed by "I know the bank is open," in Case B, because "know" means something different in each case as a result of changes in the attributor's context. How about contrastivists? They maintain that in Case A, the proposition expressed by, "DeRose knows that the bank is open tomorrow" is perhaps "DeRose knows that the bank is open tomorrow rather than that his memory is failing him," but in Case B, the contrast class is shifted such that "DeRose knows that the bank is open tomorrow" expresses the proposition, "DeRose knows the bank is open tomorrow rather than having changed their hours in the past two weeks such that it is no longer open on Saturdays." So, yes, DeRose does know that the bank is open tomorrow rather than that his memory is failing him, but he does not know that the bank is open rather than having changed its hours in the past week such that it is no longer open on Saturdays. And, a proponent of the conditional model just advanced? She claims that, in Case A, DeRose is warranted in asserting that the bank is open tomorrow because he can successfully imply that he is assuming that it hasn't changed its hours such that it is no longer open on Saturday. In Case B, because the consideration that the bank has

changed its hours in the past week is raised to salience as a result in the elevation of the stakes, DeRose is no longer entitled to simply assume it, and thus, he cannot warrantably-assert that he knows that it's open. I venture that this is a FAR more natural way of looking at things. Thus, the conditional model arguably more accurately captures what we are doing when we attribute knowledge than do the contextualist and contrastivist accounts.

Finally, with respect to 7, one can use cases like DeRose's Bank Cases and the case (presented in the previous chapter) involving the possibility that a zebra may be a painted mule to illustrate how a conditional model can explain the at times shifty acceptability of knowledge attributions.

Thus, I propose that a conditional model of knowledge attributions can better satisfy the seven desiderata advanced here than can the three accounts discussed in chapter two. Though there may be additional desiderata that a theory of knowledge attributions should satisfy, if possible, I think that the above should suffice to show that the conditional model that I advance is a plausible theory of knowledge attributions that deserves additional investigation.



## Conclusion

Since the conditional model that I sketch above is designed to explain when we are warranted in asserting knowledge attributions, it could be appended onto many theories, including those surveyed in chapter two. For, the conditional model says nothing about the conditions under which knowledge attributions are true or false. Thus, on its own, it is not a competitor to views such as contextualism, contrastivism, and Harman and Sherman's account. That said, the conditional model provides a means to argue against the views presented in chapter two when combined with certain other theories of knowledge. It is perhaps most interesting when combined with skepticism, for it can be used to explain why (and when) we are warranted in asserting knowledge attributions despite the fact that on the skeptic's view, those assertions are rarely, if ever, true. It also allows the skeptic to account for the core intuitions behind the views presented in chapter two.

## Bibliography

- Cohen, S. "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13, 57-89
- De Rose, K. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52 (1992), 913-929
- DeRose, K. "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments Against Contextualism," *Philosophy and Phenomenological Research* 72 (2006), 316-338
- Descartes, R. *Meditations on First Philosophy* (Cambridge: 1996)
- Grice, H.P. *Studies in the Way of Words* (Harvard: 1989)
- Harman, G. *Thought* (Princeton: 1973)
- Harman, G. and Sherman, B. "Knowledge, Assumptions, Lotteries" *Philosophical Issues* 14 (2004), 492-500
- Hawthorne, J. *Knowledge and Lotteries* (Oxford: 2004)
- Hume, D. *Enquiries Concerning Human Understanding and Concerning the Principles of Morals* ed. L.A. Shelby-Brigge and P.H. Nidditch (Oxford: 1975)
- Kvanvig, J. "Contextualism, Contrastivism, Relevant Alternatives, and Closure," *Philosophical Studies* 134 (2007), 131-140
- Lewis, D. "Elusive Knowledge," *Australasian Journal of Philosophy* 74 (1996), 549-567
- Rysiew, P. "Epistemic Contextualism," in the *Stanford Encyclopedia of Philosophy*, 2007 <http://plato.stanford.edu/entries/contextualism-epistemology>
- Sorenson, R. "Dogmatism, Junk Knowledge, and Conditionals," *The Philosophical Quarterly* 38 (1988), 433-454
- Schaffer, J. "From Contextualism to Contrastivism," *Philosophical Studies* 119 (2004), 73-103
- Stanley, J. *Knowledge and Practical Interests* (Oxford: 2005)
- <http://plato.stanford.edu/entries/knowledge-analysis/supplement.html>
- Unger, P. "An Argument for Skepticism," in *Arguing About Knowledge* ed. Duncan Pritchard and Ram Neta (Routledge: 2009), 466-478