# THE IMPORTANCE OF KNOWLEDGE PER SE

by

JEFFREY GLICK

A Dissertation Submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Professor Peter Klein

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

October, 2007

ABSTRACT OF THE DISSERTATION

The Importance of Knowledge Per Se

By JEFFREY GLICK

Dissertation Director:

Peter Klein

A traditional line of inquiry in epistemology tried to analyze the concept of knowledge into its constituent components. In virtue of understanding these alleged more basic concepts, such as truth, justification, and belief, it was hoped that a complete and informative theory of knowledge would emerge. According to the revolutionary approach advocated here, one which originates in Timothy Williamson's Knowledge and Its Limits, better success can be achieved by reversing this conceptual analysis structure by taking knowledge as the fundamental explanatory tool in epistemological theorizing. I defend the view that this knowledge-theoretic approach exceeds the explanatory value of its conceptual analysis competitor in the sense that the best explanations of epistemologically significant phenomena are appropriately expressed in terms of knowledge per se.

# DEDICATION

This dissertation represents the combined diachronic efforts of scores of people. Although I wrote the words, the training, encouragement, and support I received from my family, friends, and colleagues was necessary to complete the task. I owe a debt of gratitude to so many people that it is difficult to tally them all up. So let me begin by offering a general thank you to all of you who have stood by me during this ordeal. I am truly lucky.

To begin I want to acknowledge the love and forward-looking generosity of my great-grandmother Nora Gould. She and my great-grandfather Dan Gould were instrumental in facilitating my transition to and persistence in graduate school. My Mother, Andrea Clark and my Father, Stephen Glick, each deserve special recognition for the love, encouragement, and, whether they realize it or not, practical philosophical insight which they gave to me freely. It is my joy to have in my life such a wonderful step-mother Didi and step-father David who have been every bit my family for years. My sister Jennifer has consistently been a vocal advocate of my continuing education and an excellent source of positive reinforcement.

My Grandma Elaine, the warmest person I have ever known, will never let me, or any of our family, forget that she gave me my first bath. I am certain that the event had a significant impact on my life, and without it I should surely have failed. To my Grandpa Jerry I owe more than I can compress into so meager a dedication. He is the Jewish Mahayana Bodhisattva. There is hardly an aspect of my life he has not profoundly

influenced in a positive direction. My Grandparents are special people. I love them and thank them.

Philosophical debate is not restricted to the classroom and conference. I enjoyed and profited from the many heated discussions involving my Aunts Debbie, Sandi, and Pam. These are smart, assertive women who contributed to an environment where questions and arguments were encouraged. That kind of context allowed me to express my philosophical leanings long before I began graduate school, and the effect resonates onward.

California State University at Northridge, where I did my undergraduate work, equipped me with an outstanding set of philosophical tools for graduate school. I most humbly and sincerely thank Takashi Yagisawa, Ron McIntyre, Jim Kellenberger, Bonnie Paller, Cindy Stern, Frank McGuinness, Jacob Hale, Jeff Sicha, Charles Crittenden, and Donald Salter. My friends Steve Kaufmann and Robert Boyle, each a competent philosopher, were wonderful conversational companions. Jim Tomberlin, who passed away while I was finishing my graduate coursework, provided an ideal role model for the kind of serious academic work I am now privileged to do. I still enjoy teaching logic out of the book he used.

Rutgers University granted me exactly the kind of education I needed from a set of professors with unmatched skill. I thank Howard McGary, Stephen Stich, Barry Loewer, Ted Sider, Dean Zimmerman, Alvin Goldman, Tim Maudlin, Larry Temkin, Ernie

Lepore, Ruth Chang, Frankie Egan, Jeff McMahon, Jerry Fodor, and Frank Arntzenius. The graduate student population at Rutgers is humbling to say the least. I thank my brilliant and helpful friends for their encouragement and support: Angel Pinillos, Brendan Neufeld, Kate Devitt, Iris Oved, Justin Sharber, Adam Swenson, Adam Sennet, Dennis Whitcomb, Sam Cumming, Carl Feierabend, Dan Kelly, Alex Jackson, Jacob Ross, and Jason Turner. The amazing administrative staff at Rutgers deserve acknowledgement. Thank you Susan Viola, Pauline Mitchell, and Mercedes Diaz in the philosophy department, and Alex Bachmann and Barbara Sirman in the graduate school. I also want to single out David Benfield from Montclair State to thank for several productive exchanges. It's a lot of names, and the compression does a disservice to the magnitude of each of their contributions.

I owe special thanks to Bernard Molyneux and Doug Parvin. These are two of the smartest people I know. Living with them shaped my mind in more ways than I can number. Most of my philosophical conversations in graduate school were with them and I am sure that was to my credit. Their input was and continues to be more valuable than gold.

The head of my dissertation committee, Peter Klein, has been inspirational. He has always pushed me to excel and reach a higher standard. I endeavor to do justice to the exemplary job he has done in guiding me towards my goal as a professional philosopher. I extend my heartfelt gratitude to the other members of my committee, Ernie Sosa, John

Hawthorne, and Igal Kvart for conversation, education, and criticism which I desperately needed and from which I was able to craft a dissertation worth reading.

Of all the people who have helped me achieve this goal, of all the people who struggled with me, laughed with me, and encouraged me towards success, it is my most loving wife Kristy who deserves the highest appreciation. It would be all too easy to leave the thread of the dedication and expound my love for this exceptional person. But more germane to our present purposes, I want to impress upon any reader the immense magnitude of the contribution she has made to my life. She listened to my ramblings and complaints, she allowed me to test her intuitions mercilessly, and she gave me substantive philosophical advice during the writing process. She is a perfect companion and I love, respect, and thank her.

TABLE OF CONTENTS

# Chapter One

## The Project and Its Motivations

The principal aim of what follows is to defend the explanatory importance of the concept of knowledge. Much of human activity and circumstance can best be understood by theories that quantify over knowledge. Attempts to provide comparably good explanations of such activity that do not rely on the concept of knowledge are arguably unsuccessful. Some of the following chapters provide accounts of such explanations and demonstrate the failures of their non-knowledge-based competitors. The view that knowledge is important to philosophy has a venerable lineage, perhaps beginning with Plato's Meno. But more recently, Timothy Williamson's Knowledge and Its Limits provides a systematic epistemological theory that takes the primacy of knowledge as its fundamental motivation. This work is revolutionary in the sense that it opposes the vast majority of the epistemological theorizing over the last century. Some of the following chapters provide analyses of the particulars of his theory. Since the theoretical background motivating certain moves is not always easy to characterize *in situ*, the sensible place to begin is with an overview of some of Williamson's epistemological commitments and their arguments. Once they are established, it will be easier to understand how the other chapters develop the broader project of defending the theoretical importance of knowledge.

### Conceptual Analysis

One of the most significant aspects of Williamson's approach to epistemology is his rejection of the project of analyzing the concept of knowledge into constituent concepts. He thus distances himself from the vast majority of 20th century epistemology in favor of treating knowledge as a primitive concept. Since the concept of knowledge is both primitive and important, we should expect it to figure in the explanation and elucidation of other related epistemic concepts. The more traditional position is that epistemologists should be looking for necessary and sufficient conditions for knowledge in an attempt to decompose that concept into its constituent concepts. A trivial example of this is to analyze the concept *brown cow* into the constituents *brown* and *cow*. A slightly more sophisticated example is the analysis of the concept *bachelor* into the cluster of concepts *unmarried*, *adult*, and *male.*[1] But Williamson inverts this explanatory structure. Instead of gaining insight into knowledge by breaking it down into better understood parts, we can take knowledge as a primitive concept and try to understand these putative constituent concepts in terms of this primitive. This procedure can also be extended to gain insight into concepts which were never presumed to be part of an analysis of knowledge. So for example, we should expect knowledge to play a significant role in the best explanation of epistemic concepts like evidence and introspection. Williamson does argue along these lines and indeed displays the utility of the concept of knowledge by showing how it figures in a correct understanding of less obviously epistemic phenomena such as the giving of testimony. Yet despite the reversal of a traditional epistemological

[1] It is interesting to note that even this simple decomposition is not entirely uncontroversial. Is a newborn baby boy a bachelor? Is the pope a bachelor? The availability of plausible counterexamples to many proposed non-circular analyses is an indicator that such a project is problematic. For additional arguments against the view that most concepts can be analyzed into their conceptual constituents, in the form of definitions, see Jerry Fodor's Concepts; Where Cognitive Science Went Wrong, Oxford University Press, 1998.

taxonomy, this new epistemological theory does not and need not reject all of the insights which epistemologists have brought to the fore in their diligent, albeit somewhat misguided, investigations. Other epistemologists were not simply wasting their time. Rather, what is required is a reorientation of their claims. In several cases, suggestions for analytic decompositions of knowledge may instead be regarded, inside Williamson's epistemological program, as revelatory of how knowledge works, its limits, and its scope, all with the caveat that these are non-decompositional in nature. A key question which arises at this point is this: What is the status of such claims?

Williamson is somewhat coy on this issue. He does speak to it briefly though:

> "The working hypothesis should be that the concept *knows* cannot be analyzed into more basic concepts. But to say that is not to say that no reflective understanding of it is possible." [2]

Apparently what epistemological investigation is now concerned with is gaining "reflective understanding" of the concept of knowledge. Such an understanding may even result in our awareness that knowledge has certain necessary entailments. But what is to be avoided is the urge to regard such an understanding as progress toward an analysis of knowledge into other, less complex concepts.[3]

One way to understand the value of a reflective understanding of knowledge is by recognizing that any knowledge claim carries with it certain metaphysical entailments. If

---

[2] Williamson, Timothy. Knowledge and Its Limits. Oxford University Press, 2000, pg 33. (Hereafter KAIL)

[3] The traditional project of uncovering a complete, informative, and non-circular analysis of knowledge has many supporters. But there is no general consensus on what the right analysis is. The literature is replete with attempts to provide such analyses along with the consequent counterexamples and reformulations. For a recapitulation of this dialectic, Williamson recommends Robert Shope's The Analysis of Knowledge: A Decade of Research, Princeton University Press, 1983. The present theory takes the failure of this decompositional program as its starting point.

S knows that P, then, among other things, P is true, S believes that P, and S has some justification for P. But this may seem unduly weak. One may take a view of knowledge whereby one does not have the concept of knowledge unless one also recognizes the truth of such entailments. Such a criterion would be too strict, since certainly children know various things and even take themselves to know various things and yet may not have a good grasp of what justification is or whether they have it in particular cases. Nevertheless, if one is not completely convinced of Quine's well-known thesis, one could regard statements like "If S knows that P, then S has justification for P" as analytic, that is, true in virtue of the meaning of the component terms. Taking such a view does not commit one to the further view that the concept of knowledge is analyzable or decomposable into component concepts. Consider, for example, the statement "If X is maroon, then X is red." Such a statement is analytically true, but one need not have a concept of red in order to have a concept of maroon. There is no requirement on concept possession that one recognizes every general category under which the specific concept falls. Nevertheless, recognizing that the statement is an analytic truth is a step further than recognizing that there is an entailment relation between being maroon and being red.

Part of Williamson's project is to defeat the hypothesis that the concept of knowledge has a correct analysis. The attack is negative insofar as it examines the strongest arguments available in favor of such a view and finds them wanting. Consider for example, the following inference:

(E): If every correct application of a concept C is also a correct and non-trivial application of a concept C*, then C* is an analytic component of C.

(The phrase 'non-trivial' is ineliminable in (E). Without it, any concept which applied to everything would be an analytic component of every concept, an unhappy result.) Inference (E) appears to be at work in traditional epistemology. It is, and has been, widely held that knowledge implies, among other things, belief and truth. In fact, the entailments between knowledge and belief as well as knowledge and truth are useful as tools to test whether a given mental state qualifies as a knowledge state. Thus we are strongly inclined to deny that an agent S knows some proposition P if P is false, and we are likewise strongly inclined to deny that an agent S knows that P if S does not believe that P.[4] These two implications have been cited as evidence for the view that knowledge is a composite concept combining internal, psychological elements (belief) as well as external, semantic elements (truth). But there is reason to question the underlying general pattern of inference embodied in (E) and exemplified in each of these inferences.

Consider the following triad:[5]

1. X is an equilateral triangle.
2. X is a triangle whose sides are indiscriminable in length to the human eye.
3. X is a triangle.

[4] There have been some attempts to block this entailment. Consider the following claim: "Susan did not believe that Karen ate the last cookie, she knew it." Now one could treat this claim as elliptical for "Susan did not *just* believe that Karen ate the last cookie, she knew it.", though that sort of revisionist move typically induces controversy. This tangent will not be addressed further here.

[5] Drawn from KAIL, pg 28

While it is true that 1 implies 2 and 2 implies 3, it is nonetheless improper to suggest that the indiscriminability of the sides of an equilateral triangle is somehow an analytic component of the concept *equilateral triangle*. Equilateral triangles could exist without human eyes, therefore the concept *equilateral triangle* cannot be decomposable into any cluster of concepts containing, as an ineliminable component, the concept *human eye*. Since the inference from 1 to 2 holds and since such an inference does not require that the concept *indiscriminable in length to the human eye* is an analytic component of the concept *equilateral triangle*, it follows that (E) is false.

The above triad can also teach us another lesson. Both the properties *equilateral triangle* and *triangle* are geometric properties. They each can be specified solely within a theory of geometry. But the property *triangle whose sides are indiscriminable in length to the human eye* is a non-geometric property. It cannot be specified solely within a theory of geometry. Therefore, the following inference does not hold either:

(E*): For properties C, C*, and C** where each is non-equivalent to any other, if every occurrence of C is an occurrence of C*, and every occurrence of C* is an occurrence of C**, then if C and C** are both properties of type T, then C* is also a property of type T.

Williamson puts it this way, "Something sandwiched between two geometrical properties need not itself be a geometrical property."[6]

Knowledge as a Mental State

[6] KAIL, pg 28

Anyone interested in preserving the spirit of the historical investigation of knowledge should find Williamson's views appealing in at least the following respect. If traditional epistemology is valuable, it is only because knowledge itself is valuable. A good way to demonstrate the value of knowledge is to reveal what roles it plays in our network of explanatory mechanisms. For Williamson, epistemology is integral to the philosophy of mind. One misses an important aspect of our mental lives if one ignores epistemological concerns. Knowledge and Its Limits provides a defense of the view that knowledge is a mental state in the same sense that believing that the sum of two and two is four is a mental state or that feeling pain is a mental state. In its strongest form, the view is that knowledge is the most general factive mental state, the factive mental state that one is in if one is in any other factive mental state. The effect is that knowing that p is important if any of a host of other attitudes are, including seeing that p, remembering that p, regretting that p, and proving that p. A correct and complete understanding of factive mental states will have to advert to knowledge. (E*) can be used to justify the claim that knowledge is not a mental state. The argument proceeds as follows.

> If S knows that P, then S believes truly that P. If S believes truly that P, then S believes that P. Suppose for reductio that knowing is a mental state. Believing is uncontroversially a mental state. If (E*) holds, then believing truly is a mental state. Since believing truly is clearly not a mental state, it follows that knowledge is not a mental state.

What the triad shows is that (E*) cannot be correct. Therefore the proposed reductio does not go through and the question of whether knowledge is a mental state remains open.

Another source of resistance to the claim that knowing is a mental state is the view that one's own mental states must be open to first-person access in a way that knowledge states clearly are not. Consider the state of being in pain. If one is in pain, so the view goes, one is in a position to know that one is in pain. A great deal is built into the restriction "being in a position to know" over and above "know". For example, the person in question must possess the right concepts, must not be under the influence of some mind-altering pharmaceutical, and must be considering the question of whether she is in pain. Such first-person knowledge of one's own mental states is, on this view, a requirement of any genuine mental state. The state one is in is not mental unless one is in a position to know that one is in it. A host of states whose status as mental is uncontroversial seem to meet this requirement. One is normally in a position to know whether one believes that P for any P. One is normally in a position to know whether one desires that P for any P. Phenomenological states such as being in pain or being itchy also seem to fit this pattern. Conversely, whether one does not believe that P, does not desire that P, is not in pain, or is not itchy also seem to be open to a special kind of infallible first-person access.

Even if belief, desire, regret and the like are mental states and are open to the desired kind of infallible first person access, knowledge looks to be on less sure footing. One is clearly not always in a position to know whether or not one knows that p, for all p. One

way to argue for this conclusion is to note the aforementioned concern that one may not have the mental capacity to grasp claims of iterated knowledge. But aside from that line of thought, there is the more straightforward conclusion one can draw from the Gettier cases. Nigel is given strong evidence which he trusts and which comes from a reliable source that Kristy is in the United States. Nigel is a competent logician and deduces from his justified belief that either Kristy is in the United States or she is in Holland. As luck would have it, Kristy has recently and unbeknownst to Nigel's source taken a trip to Holland and is presently in that country. Nigel is correct about his deduced belief. It is true that either Kristy is in the United States or she is in Holland. What's more, Nigel is justified in believing that either Kristy is in the United States or she is in Holland. However, there is a strong intuitive appeal to the claim that Nigel does not know that either Kristy is in the United States or she is in Holland. His lack of knowledge is not open to first-person access. Nigel does not know that he does not know that either Kristy is in the United States or she is in Holland. If knowledge needs to be perfectly open to first-person access in order to qualify as a mental state, then knowledge is not a mental state.

But such perfect access is not required of mental states. A weaker condition which captures only the positive case is more palatable. Call any state which is open to the relevant kind of first-person access 'luminous'. A state is luminous if and only if whenever one is in that state, then one knows that one is in that state (without requiring the converse). While it is easy to see that one may not always know one is not in S when one is not in S, Williamson's claim is that only trivial mental states are luminous. All the

interesting mental states, belief, desire, feelings of pain or itchiness, alertness, drunkenness and the like, are all non-luminous. If non-trivial, uncontroversially mental states are non-luminous, then there is no argument to be made against the claim that knowledge is a mental state which relies on the premise that all mental states are luminous. Chapter three discusses the anti-luminosity argument in detail. The gist of that argument is that non-trivial mental states can all be put on a continuum of evaluation. At each end of the continuum, instances of the mental state under consideration are transparent; one can know one is in them when one is in them and one can know one is not in them when one is not in them. But for any non-trivial mental state, there exist borderline cases which manifest a necessary ignorance for creatures of limited discriminatory powers such as ourselves. If no non-trivial, uncontroversially mental state is luminous, then it is no argument against knowledge being a mental state that knowledge is not luminous.

Externalism[7]

Another reason that one might reject Williamson's account of knowledge as a mental state is that one has been convinced of an internalist picture of the mind. According to the internalist, there are metaphysically discrete states whose conjunction entails a knowledge state. One of these states is an external, environmental condition. The other is an internal, mental condition. For any proposition p about the external world and some agent S, if S knows that P, then P is true given that knowledge is factive. The truth component, claims the internalist, is an external, non-mental part of the knowledge state.

[7] Much of the argumentation here is drawn from KAIL, chapter three 'Primeness'.

If S knows that P, then S believes that P. The belief component is an internal, mental part of the knowledge state. When one is in the correct belief state, and the environment corresponds to the belief state, and the belief state was entered into in a correct, that is, justified manner, then (other things being equal) S knows that P. The internalist finds such a taxonomy inviting since it keeps all mental causation spatially close to the agent. Beliefs are in the head. Beliefs are mental. Beliefs are the right kinds of entities to bear the responsibility of unmediated mental causation. Even if one does not accept this last claim, one could still appreciate that belief states appear to be *better* suited to the task of unmediated mental causation than knowledge states which require that the external world be a certain way in addition to one's internal states. Causation is local; no action at a distance. Since the truth component of knowledge routinely involves objects, events, or states which are distant from the brain, if knowledge does in fact break down the way the internalist says it does, then knowledge is not as good a candidate for mental causation as belief is. Furthermore, if knowledge is specifiable as a conjunction of internal and external states, then knowledge appears to be causally parasitic on the implied belief state. All of the causal-explanatory work which might be done by treating knowledge as a mental state can be done at least as well or better by deploying the separate aspects of a (clearly mental) belief state on the one hand and an environmental state on the other.

We need a way to taxonomize states which are essentially combinations of purely internal states and purely external states. Believing truly that Julius Caesar was emperor of Rome is an example of such a state. It involves the purely internal state of believing that Julius Caesar was emperor of Rome and the obtaining of a certain purely external

state consisting of Julius Caesar's having been emperor of Rome. The dividing line between internal and external is usually drawn near the skull. Call states which exhibit the conjunction of one internal and one external state "composite". Let a "prime" state be one which is not composite.

If a state S is really composite, then it should be possible to construct the following group. Let a and b be cases in which S obtains. Let c be a case which is internally like a and externally like b. If S is composite, then S obtains in c as well. Believing truly is a typical composite state. Put to work on the previous example, Anne and Bill each believe truly that Julius Caesar was emperor of Rome. Consider a person Charlie who is internally like Anne and externally like Bill. Charlie thus believes that Julius Caesar is emperor of Rome (Anne's internal state), and he also is situated in a world where Julius Caesar was emperor of Rome (Bill's external state). In such a case, Charlie believes truly that Julius Caesar was emperor of Rome, just as the theory predicts.

Compare the case of believing truly with the case of knowing. If internalism about the mental is correct, then knowing that p is metaphysically reducible to a conjunction of an internal, mental state and an external, non-mental state. But knowledge is not reducible in that way; knowledge is prime. To see this, we can establish a triad of cases, a, b, and c, in which a and b manifest knowledge while case c does not. If case c is constructed in the right way, the result will establish that knowledge is not composite.

Here is an example from Williamson.[8] Let case a be a case in which Larry knows via testimony that p. In case a, Smith tells Larry that p, Smith is objectively a trustworthy person, and Larry trusts Smith and thereby comes to believe that p. Also in case a, Jones, an objectively untrustworthy person, tells Larry that p. But Larry doesn't trust Jones in a. Let case b be a case in which Larry knows via testimony that p. In case b, Jones tells Larry that p, Jones is objectively trustworthy, and Larry trusts Jones and thereby comes to believe that p. Also in case b, Smith, an objectively untrustworthy person, tells Larry that p. But Larry doesn't trust Smith in b. Finally, let case c be the conjunction of the internal state of Larry in case a and the external state of Larry in case b. So in case c, Larry trusts Smith, Smith tells Larry that p, but Smith is objectively untrustworthy. Larry comes to believe that p in c because c is internally like a, but Larry does not know that p in c because his informant is untrustworthy. Knowledge is not simply a conjunction of a type of internal state with a type of external state.

By crafting triads in the appropriate way, Williamson deploys this argumentative strategy to establish the primeness of states like seeing water, hearing noises, and having singular beliefs. If these states are mental, than so is knowing. If knowing is non-mental, then so are these other states. By tethering together a wide array of states so that they hang or fall as one, Williamson forces the internalist to either concede that knowledge is mental or that a substantial portion of our ordinary talk about the mental is misguided.

The picture under attack here, if a metaphor will do, is one in which people are infants immersed in water. They are surrounded by the external world, but are separate

[8] KAIL, pg. 72

from it, cut off in a fundamental way. By contrast, the view of the mental that Williamson endorses allows for vigorous interactions between the agent and the world in virtue of the irreducible interconnectedness of subject and her environment. It is beyond contention that for each case of knowing that p, there is some external state and some internal state that the knower is in. But that is to miss the explanatory significance of knowledge. Explanatory tools need to have a certain amount of generality. They need to be able to explain slightly different cases by appealing to common properties that unify them. That one's brain is in a certain state while the environment is in a certain state when one knows that p is not a particularly useful fact, unless there is some reason to think that these states are sensitive to one another. Intuitively, one of the features missing from the internalist conception of knowledge is the causation connection from world to belief. The ways in which the external world causes various beliefs, especially by stimulation of the sensory organs, is important for explaining human activity. Successful action is best explained by models which incorporate a deep and persistent connection between the subject and the world she inhabits.[9] Since those connections are severed in the internalist model, that is a reason to reject internalism about the mental.

[9] See chapter five for further discussion of this issue.

# Chapter Two

## Evidence

The Position:

Williamson's view of the nature of evidence is a departure from traditional accounts of evidence. His theory is simple to state; it is that one's evidence is identical with what one knows—E=K. At several points throughout Knowledge and Its Limits, Williamson argues in favor of reversing the standard accounts of knowledge which decompose that concept into its conceptual constituents: justification, truth, belief, and whatever else one regards as necessary for knowledge. Instead, he suggests that better progress can be made by understanding epistemic notions, such as justification, essentially in terms of knowledge. The view that one's evidence is identical with what one knows is more natural given that context. There is a short argument for this claim, and it is important to understand what that argument is and how Williamson defends its premises in order to address some of the criticisms which have been leveled against it. Following that, one can examine how a view that identifies one's evidence with one's knowledge comports with some contemporary accounts of adequate justification, such as foundationalism and reliabilism.

The argument in favor of E=K is clearly valid:

1. All evidence is propositional.
2. All propositional evidence is knowledge.
3. All knowledge is evidence.

Therefore, All and only knowledge is evidence.[1]

All Evidence is Propositional

Each of the premises is suspect on different grounds. In defense of premise 1, Williamson argues that a propositional account of evidence captures all the functional aspects of evidence that we pretheoretically desire. Those functional aspects include, but are not limited to, inferring to the best explanation, ruling out possible hypotheses, and updating our subjective probabilities.

It is important to pause and take note of the overarching themes of a knowledge-first view put into application here. This strategy of identifying crucial functions of a concept as a vehicle for gaining deeper insight into it is systematic in Knowledge and Its Limits. In articulating the view that knowledge is the norm of assertion, the argument derives its potency from ordinary observations about the use of knowledge-theoretic explanations in our judgments about discourse. Mary said more than she had the authority to say because she did not know her statement to be true. Chapter five argues that analogous reasoning shows that knowledge-theoretic explanations are ineliminable in certain moral cases. Here we are examining the role that evidence plays (or should play) in our epistemic practices and arguing, with Williamson, that knowledge is best suited to play that role.

[1] Knowledge and Its Limits Oxford University Press, 2000. (Hereafter, "KAIL") pg. 193

Each successful determination of the correct explanation of these phenomena lifts the veil, to some degree, and permits us to more fully understand what knowledge is. This by itself is inadequate to undermine the traditional project of decomposing the concept of knowledge into its putatively better understood constituents, but it should at least motivate one to examine the prospects of approaching the fundamental question in epistemology from a new perspective.

It would be unreasonable to exclude propositions from the extension of the concept of evidence. Those propositions must, of course, have a certain degree of epistemic plausibility. That the Earth is round is a proposition which could have usefully served as evidence to those cave-dwelling would-be astronomers who, understandably though incorrectly, regarded the Earth as flat. Since these distant ancestors of ours had no grounds for believing that the Earth is round, that fact prohibited this proposition from functioning as evidence for them. The failure was not the result of the propositional nature of the possible evidence in question, but instead was a result of the unavailability of justification for the proposition. Had these astronomers utilized better models and better instruments, they might have been able to hold that the Earth is round and utilize that proposition in their explanations of the celestial motions. What they did have available as evidence included the propositions 1) that the Earth appears from our perspective to be flat, 2) that the Sun appears to move around the Earth, and 3) that the Earth appears to be stationary. Indeed, when heliocentric models of the Solar system were proposed, it was this sort of evidence which needed to be explained in light of its apparent, but not actual, incompatibility with the new theories.

Nonetheless, there are surely other forms of evidence which do not, on the face of it, have a propositional component. Sensory experiences and memories of sensory experiences often contain so much information that any attempt to capture their totality in a single proposition would surely fail. Sometimes we have no words to express why it is that person A looks like person B. Yet the sensory information, whether occurrent or recalled, appears to confer some degree of epistemic merit on a wide array of propositions, often with the result that we know the proposition in question. That person A's face matches the face of person B in one's memory is usually sufficient for one to know that A is B. A substantial argument needs to be deployed to show that such an ordinary conception of evidence ought to be abandoned.

Resistance to the first premise, that all evidence is propositional, is most perspicuously generated by those who wish to maintain a certain brand of foundationalism. Foundationalists hold that the nature of justification is such that every justified proposition has an evidential history which neither can extend infinitely nor can circle back so that, through a series of inferences, one is left with some proposition lending evidential support to itself by being in its own evidential ancestry. The only remaining option is that the chain must terminate in some basic, foundational beliefs. These foundational beliefs must be either justified or unjustified. If they are unjustified, then it is puzzling how they could lend epistemic support to propositions which they are meant to justify. How then is a foundational belief justified? Two answers suggest themselves. Either the proposition is justified in virtue of an inference from some other

proposition, or the proposition is justified but not by virtue of any inference. Since self-justifying propositions are suspect for precisely the same reasons that circular justification is rejected, the third possible answer (that the foundational beliefs are literally self-justifying) is a non-starter. And since all justified propositions need to be justified by something, there must be something which justifies the foundational beliefs without needing to be justified itself. That thing cannot be another proposition lest the same questions emerge.[2] This line of thought led some foundationalists to the conclusion that sensory experiences, phenomenal experiences, were the ultimate source of justification for all other empirical propositions.[3] Phenomenal experiences do not need to be justified, chiefly because they are not candidates to be justified. The question "What reason do you have for seeing that?", interpreted as a question about the epistemic status of the seeing and not as a question about the authority of the agent in question to engage in the seeing act, is nonsensical. Sensory episodes don't get justified by reasons the way beliefs do. But phenomenal experiences are not propositions. If all evidence is propositional, then phenomenal experiences cannot be evidence and this foundationalist model is undermined. The natural response from this sort of foundationalist is to reject Williamson's first premise.

In order to counter this line of thought, Williamson proposes to explain how it might appear that phenomenal experience, understood non-propositionally, acts as a justifier

---

[2] This is one formulation of the so-called Regress Argument. See Audi, Robert. 'Contemporary Foundationalism' and Bonjour, Laurence. 'A Critique of Foundationalism' in The Theory of Knowledge ed. Matthias Steup, Wadsworth Publishing Company 1999 for further discussion.

[3] Non- empirical propositions, such as those one often encounters in mathematics, logic, or semantics, are often left out of this discussion. A correct explanation of justification, evidence, knowledge, and other related phenomena may be quite different in the realm of the a priori than it is in the realm of the a posteriori. Here we shall confine ourselves to the standard practice of considering only empirical propositions.

when in fact it is propositions either about the phenomenal experience or propositions which are the content of the phenomenal experience[4] which are solely doing the justificatory work. This explanation revolves around the three important roles which evidence plays, namely legitimizing one's updating of subjective probabilities, the ruling out of competing hypotheses (by being incompatible with them), and suggesting ways in which one can infer to the best explanations. Non-propositional sensory experiences are ill-suited to these tasks. But before examining Williamson's argument for premise one, let us consider deploying an alternative strategy.

Williamson is intent on ruling out any non-propositional elements from the set of genuine evidence. But ordinary uses of the word 'evidence' tend to include in their extension both propositional and non-propositional members. Leaving aside the lofty examples upon which philosophers tend to draw, one need only look at the legal uses of the word 'evidence' to see that this is true. Particular items are routinely entered "into evidence". There is an evidence room in all police stations where items like guns, biological samples, and the contents of wallets are kept for use during trials. Attorneys can argue that evidence has been tampered with. All of these uses of the word 'evidence' fly in the face of premise one. Even if there are roles for evidence to play which cannot be played by anything other than propositions, there are different roles for evidence which cannot be satisfied by propositions. To force a change in the use of the word in order to satisfy premise one would do excessive violence to our semantic intuitions. And

[4] In 'Veridical Hallucination and Prosthetic Vision' Australian Journal of Philosophy, September 1980, David Lewis puts the point thusly: "The content of the experience is, roughly, the content of the belief it tends to produce." Since belief is an attitude toward a proposition if any attitude can be, the content of an experience is propositional.

any theory so revisionist in spirit had better have formidable reasons in its favor. Williamson himself may regard the justifications of premise one provided in Knowledge and Its Limits as sturdy enough to justify this linguistic alteration, but others may reasonably remain unconvinced of the need to forfeit the word to Williamson's interpretation. Certainly there is a propositional sense of the word 'evidence', but this is not the only sense of the word, nor the only one important to epistemology.

Williamson is aware of this difficulty. He says, "Why should all evidence be propositional? It would not be on a broad interpretation of 'evidence'. In the courts, a bloodied knife is evidence."[5] One might remark that the "broad interpretation" of 'evidence' is only so-called because Williamson's interpretation of evidence is conspicuously narrow. But Williamson need not try to convince us that *the* correct view of evidence is the one which satisfies the three aforementioned roles. To insist on this line of thought is tantamount to denying that the judicial roles for evidence are important and central ones to our understanding of that concept. The quarrel is merely verbal and there is no need to be proprietary about the word 'evidence'. Such a dispute takes us away from what is interesting about Williamson's theory. A slight modification will help. Let there be a distinction between propositional evidence (PE) and physical evidence (PhE). We may then replace Williamson's central equation 'E=K' with 'PE=K' and leave an analysis of PhE aside. Certainly there is an intimate connection between PhE and PE. For instance one cannot have PE without some PhE[6], at least when PhE is

[5] KAIL, pg. 194
[6] And perhaps, further, the converse holds.

construed broadly enough to encompass things like brain states[7]. The propositional content of PE is necessarily tied to the existence of particular PhE. For example, the evidence a jury member has for her belief that Larry is the murderer is a function of, among other things, the existence of a bloody knife, his fingerprints on the knife, a video recording of Larry committing the crime, as well as the various utterances made by experts in court. This distinction between propositional and physical evidence requires a modification of premise one, a modification which makes that premise trivially true. Instead of beginning with the claim that all evidence is propositional, we instead begin with the claim that all PE is propositional. But the triviality of the first premise does not go very far in diminishing the work which needs to be done to substantiate the view that PE=K. What the distinction does is to prevent one from being distracted by issues which are not actually relevant to Williamson's thesis.

What *is* relevant to Williamson's thesis is that certain central roles for evidence can only be fulfilled by equating evidence with knowledge. This suggests a second alternative way to view Williamson's original argument. Instead of equating evidence in general with knowledge, we instead equate evidence* with knowledge, where evidence* is that evidence which is used (or could be used) to update subjective probabilities, rule out competing hypotheses, and infer to the best explanation. Since physical evidence alone is, arguably, poorly suited to these roles, only propositions will qualify as evidence*. This second alternative also has the virtue of better preserving the original formulation of the argument, however it is still worthwhile to mark the distinction

[7] This condition is required in order to accommodate the idea that mathematical knowledge, among other things, is often justified for one when one has a proof in mind. Similarly, the typical justification of one's knowledge that one is in pain may be seen as a function of what state one's brain is in.

between propositional evidence and physical evidence. Hereafter we shall adopt the first alternative, though the remaining discussion could be modified to accommodate the second alternative.

Evidence in Action

While the truth of the modified premise one, that all propositional evidence is propositional, cannot usefully be disputed, it is still worthwhile to investigate the ways in which specifically propositional evidence does serve as evidence. As already mentioned, the three chief roles which propositional evidence plays are to update one's subjective probabilities, to rule out competing hypotheses, and to infer to the best explanation.

With respect to the first role, the standard Bayesian formula for updating one's subjective probabilities is to conditionalize on the available evidence. Since the only candidates for having probabilities are propositions, it follows that the standard Bayesian picture already requires that propositional evidence be a genuine and important form of evidence. Consider the following situation: one is drawing playing cards one at a time from a standard fifty-two card deck. In the first ten draws, one does not draw any aces. Since there are four aces in the deck, the probability that the next card drawn will be an ace is four in forty-two or just over nine and one-half percent. With each successive draw that does not result reveal an ace, the probability of pulling an ace on the next draw increases. But the collection of drawn cards themselves, the physical evidence, is insufficient to carry the probabilistic weight which allows us to conditionalize on

evidence. It is *that* each successive card drawn is not an ace that renders it more likely that an ace will be drawn next. In fact, the hypothesis itself is stated in explicitly propositional terms. *That* an ace will be drawn next is what is rendered more probable by the evidence.

This last point, that hypotheses are propositional, hardly bears mentioning given the factivity of knowledge. If hypotheses are the things which scientists and others are attempting to know (or know the negation of), then the factivity of knowledge demands that hypotheses be propositional. Only propositions are truth-evaluable. Both what confers evidential weight and what receives evidential weight are propositional.[8]

Similar remarks apply to the ruling out of hypotheses. Consider the hypothesis that there are no people in the White House at time t. We investigate this hypothesis by aiming a telescope at a White House window. We observe a person inside the building at t. That we have observed a person inside the White House at t is inconsistent with the hypothesis that there are no people inside the White House at t. To avoid misunderstanding the view, it is crucial in this context to distinguish between "seeing" and "seeing that".[9] To see a person in the White House is to have a visual experience which is appropriately causally dependent on a person being in the White House, but is such that it does not require one to have a concept of either *people* or *the White House*. To see that there is a person in the White House requires both seeing a person in the

[8] A very similar argument occurs in section 4, *A Coherentist Critique of Foundationalism* in Ernest Sosa's "The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge", Midwest Studies in Philosophy, Vol. 5: Studies in Epistemology, Minneapolis: University of Minnesota Press, 1980.
[9] KAIL chapter one, especially pg. 38

White House and grasping the proposition that there is a person in the White House. Additionally, Williamson argues that the experience must be such that one comes to know that there is a person in the White House on the basis of said experience, but the knowledge component is not necessary at this point in the discussion. For the purposes of this argument, let it be the case that if one sees that there is a person in the White House, then one must grasp[10] the proposition that there is a person in the White House and endorse[11] such a proposition on the basis of the experience. If one merely sees a person in the White House at t without seeing that there is a person in the White House at t, then one cannot legitimately rule out the hypothesis that there are no people in the White House at t. In such a circumstance, while one would have the right physical evidence to rule out the hypothesis, one would not have the right propositional evidence to rule it out. Physical evidence alone is insufficient to rule out hypotheses; propositional evidence derived from the physical evidence must be deployed.

This characterization of using evidence to rule out hypotheses may suggest an implausible picture. Scientists certainly do not throw out whole theories on the basis of a single piece of evidence, no matter how glaring the inconsistency between that evidence and an accepted and well-confirmed theory. If an otherwise successful, simple, coherent, and explanatory theory is marred only by one observational anomaly, then that calls for further investigation and experimentation, not the rejection of the established view. But

[10] To grasp a proposition is to 1) possess the concepts which the proposition comprises, 2) possess the computational capacity to organize those concepts in the required manner, and 3) to deploy one's computational capacity by bringing to mind the proposition in question. Perhaps this formulation is theoretically inelegant given that condition 3 entails conditions 1 and 2. Nevertheless, redundancy can be a vehicle for clarity.

[11] 'endorse' is used here as a neutral term which either comprises or is necessary for belief, strong belief, and knowledge as well as other suitable affirmative attitudes which one could hold with respect to a proposition.

this is not a fair parsing of the original claim. The position here is simply that *if* one is to rule out a hypothesis, or indeed an established theory, *then* one should do it on the basis of what one knows. The surprising conclusion, if there is one, is that sometimes we are mistaken about what our evidence includes. The reason that scientists are unwilling to give up a powerful theory in the face of limited evidence is because they have concerns about the extent of that evidence. If theory T entails that event E will not occur and event E is genuinely observed[12] to occur, then T must be strictly false and the observer is in a position to dispense with it. But if T is sufficiently well-established, that will undermine the epistemic status of any putative observation that E occurred. There will be doubts as to whether E was in fact observed. In Williamson's terminology, the scientist would have seen E occur but would not have seen *that* E occurred. Only the latter entails that the scientist knows that E occurred and therefore that E is legitimately part of her evidence. But much of what we would intuitively count as evidence remains evidence in this problematic scenario. The scientist has as part of her evidence that she seemed to observe that E occurred, that it looked as though E occurred, etc. Given a preponderance of this kind of limited evidence, the original theory could be overturned precisely because the observation that E occurred would have risen to the status of evidence and would therefore properly rule out any hypothesis which entailed its negation.

Finally, only propositional evidence can allow one to infer to the best explanation of a class of phenomena. Inferring to the best explanation is a process in which one endorses a hypothesis because it does a better job of explaining the evidence than any of its available competitors. But one cannot explain merely physical evidence. Only

[12] In the sense of "seeing that" rather than plain "seeing".

propositional evidence is an adequate candidate for explanation. Linguistic evidence bears this out.[13] Explanations are what follow 'because' in sentences of the form "X because Y." For example, World War I began because Archduke Ferdinand was assassinated. Non-propositional entities do not make for grammatically correct substitutions for either X or Y in that formula. "World War I because Archduke Ferdinand's traumatized corpse" is not proper English. Although his corpse is good physical evidence, it is only the propositional evidence derived from it which serves the explanatory role. There is, however, the construction "X because of Y" which does allow for non-propositional entities to serve as substitution instances of Y, as in, "Bob did not go to the party because of Nancy", but the phrase "because of" is semantically fused. In any event, Nancy is not an explanation. That Nancy was present at the party and that Bob does not like Nancy is an explanation, but such constructions are clearly propositional.

The foundationalist who regards sense-data or sensory experiences as the final grounding of all correct justification for a posteriori reasoning may find this emphasis on propositional evidence disconcerting. In their view, experience plays an ineliminable role in the justification of beliefs, a role which is not adequately served by propositions. However, even they must admit that for the purposes of updating subjective probabilities, ruling out hypotheses, and inferring to the best explanation, it is propositional evidence which carries the workload. The defender of the centrality of propositional evidence need not be dismissive of the foundationalist's concerns. In fact, it is open to a supporter

[13] See KAIL pg. 195. Williamson draws our attention to concurring support from Peter Unger in Ignorance a Case for Scepticism, 1975 Clarendon Press, pg. 204-206.

of Williamson's theory to argue that the physical evidence supplied by sensory experiences, that is, the raw experience or the qualia, are ineliminable sources of the only adequate propositional evidence to serve as the foundation for all further justification. Nevertheless, there may be cause for concern.

Ineffable Experience

One problem is that there appear to be sensory experiences which cannot be encapsulated by any description. Many of our experiences are too multivariate, too amorphous, and too eclectic to be captured in their entirety by even the most sophisticated literatus or the most gifted poet. One's particular experience of the Grand Canyon, in all its visual, audible, and even olfactory majesty, outruns one's vocabulary. Since, the objector will claim, there is no proposition which can serve as the justifier for one's belief that, say, the Grand Canyon is lovely, the only remaining candidate to serve the evidential role is the experience of the Grand Canyon itself. While this is physical evidence for the belief that the Grand Canyon is lovely, it is not propositional evidence that the Grand Canyon is lovely. Since, the objection continues, it is physical evidence which is playing the justificatory role in this case and not any propositional evidence, it follows that any account of evidence which rules out non-propositional elements from serving the role of updating one's subjective probabilities is incorrect. Since Williamson provides such an account, his view is therefore mistaken.

There is an adequate response to this concern.[14] The objection presupposes that the required propositional evidence must contain explicit descriptive components which when combined constitute a complete semantic "picture" of what is experienced. The idea is that the only viable propositional evidence for the belief that the Grand Canyon is lovely would have to take the form of a complete description of every bit of minutiae that the experience comprised. Since there are serious concerns regarding the limitations of the descriptive abilities of human beings, our conceptual capacities with respect to grasping large propositions, and simply our inattentiveness to every relevant detail, it cannot reasonably be supposed that we do, in fact, even entertain such complicated propositions let alone rely on them for the justification of our beliefs. What the objection ignores is the feasibility of using demonstratives. While there are many propositions which can be articulated with varying degrees of difficulty, there are others which can only be grasped by having a particular experience and then referring to it by way of demonstration. For example, one can believe that one is having a visual experience of a hunter green circle, two feet in diameter, on a white background. The content of that belief would have among its components the concepts of *hunter green* and *circle*. But there are those who do not have the concept *hunter green*. They have the same visual capacities as those who do have that concept, but they cannot grasp propositions of the aforementioned variety precisely because they do not have the requisite concepts. Nevertheless, there is a way that such a conceptually impoverished individual can express her belief. She can believe that there is a circle of *that* color.

[14] KAIL, pg. 197-198

In the case of the belief that the Grand Canyon is lovely, the proposition that is doing the evidential work is the proposition that the Grand Canyon looked, sounded, and smelled like *that*. It is true that without actually experiencing the particular complex experience that the agent underwent, others cannot grasp the proposition which serves as the evidence for her belief that the Grand Canyon is lovely. But such a proposition does exist, even if it cannot be articulated without using the demonstrative 'that', and even if it cannot be grasped without undergoing a particular complex experience. Since the objection rested on the unavailability of any such proposition, the objection is refuted.

Propositional Evidence is Knowledge

The second premise of Williamson's argument, that all propositional evidence is knowledge, also finds plausibility by satisfying the three aforementioned central roles for knowledge. With respect to updating one's subjective probabilities, it would be illicit to do so unless the evidence was known. But might not one object that it is perfectly reasonable to update one's subjective probabilities in an hypothesis h provided that one has a justified belief that p, where the probability of h conditional on p is greater than the probability of h alone? Yes, one rationally ought to update one's subjective probability that h is true if one is justified in believing that p, where p renders it more likely that h is true than the initial probability for h. But in that case, the proposition that is one's evidence is that one is justified in believing that p is true, not that p is true. The degree to which p is justified in one's justified belief that p will and should factor into any

alteration of one's subjective belief in h. But p itself is only part of one's evidence if it is known.

Explosive Evidence

Williamson offers a further argument against the view that merely justified true beliefs[15] can serve as evidence[16]. If merely justified true beliefs were to count as evidence, then the propositions which they justified would themselves count as evidence. This larger set of evidence could then justify still more propositions which then count as evidence and so on. A cascade effect would grant the status of evidence to far more propositions than we in fact want to countenance as genuine evidence. In Williamson's words, "The result would be very different from our present conception of evidence" [17]. While the propositions which a justified true belief justifies are thereby, of course, justified, it should not follow that the propositions so justified are themselves evidence. What should count as evidence is that the propositions so justified are justified propositions. But there is a difference between the proposition p and the proposition that p is justified. This difference is blurred if any justified belief can count as evidence.

An example will help illustrate the point. Williamson does not describe such an example in Knowledge and Its Limits, though he has provided one in conversation[18]. Suppose that one knows that Jones, a man in his early fifties, is alive today and in fine

[15] That is, justified true belief that is not knowledge.
[16] KAIL, pg. 201
[17] KAIL, pg. 201
[18] Email correspondence, March 17th 2006

health. Uncontroversially, knowing p is sufficient for having p as a component of one's evidence. Such evidence justifies the belief that Jones will be alive tomorrow. Crucially it does not render the belief known. Instead, if one draws the conclusion that Jones will be alive tomorrow, then, given the previously mentioned evidential basis of such an inference, one will have a justified belief that Jones will be alive tomorrow. Let us also grant that the inference is correct; Jones will be alive tomorrow. One therefore has a justified true belief that Jones will be alive tomorrow. If that is sufficient for the belief to count as evidence, then in addition to one's evidence including the proposition that Jones is alive today, one will also possess as part one's evidence the proposition that Jones will be alive tomorrow. But this latter proposition will serve as evidence for the conclusion that Jones will be alive the day after tomorrow. It will do so with the same justificatory weight as the original belief lent to the first inference. "Obviously a sorites threatens", says Williamson. At some point, one will reach a day far into the future where one could not reasonably be justified in believing that Jones will survive one more day. But if any justified true belief can count as evidence, then one should be justified in believing that Jones will last that additional day. Since one is not, evidence cannot consist merely in one's true justified beliefs.

In contrast, the knowledge account of evidence has no such unwelcome consequences. Suppose one draws the conclusion that Jones will be alive tomorrow. One's current knowledge that Jones is alive and in good health makes it likely that Jones will be alive tomorrow. Absent any concerns to the contrary, one has a justified belief that Jones will be alive tomorrow. Grant that Jones, in fact, will be alive tomorrow.

One's justified true belief that Jones will be alive tomorrow does not, on the knowledge account, increase one's evidence set to include the proposition that Jones will be alive tomorrow. Perhaps if one considers one's inferences, one can come to the conclusion that one has a justified belief that Jones will be alive tomorrow. If one's methods are good, one can know that one has a justified belief that Jones will be alive tomorrow. But such knowledge does not threaten a sorites as it does in the previous view. One's knowledge that Jones is alive today and one's justified belief that Jones will be alive tomorrow (even allowing that it is a true justified belief) make it likely that Jones will be alive the day after tomorrow. But even if one is reflective about one's inferences and recognizes their status as either known or merely justified, one's evidence is far more limited and qualified than in the previously mentioned view.

One issue to consider here is how independent this argument is for the position that only knowledge can serve as evidence. Even if the example works as stated, it is not sufficiently powerful on its own to validate Williamson's contentious E=K thesis. One obstacle is that it does not rule out views in which some condition stronger than being known is placed on propositions in one's evidence set. Any view of that sort would also survive the explosive evidence objection. Instead, we only have an argument that weaker conditions, like mere justification coupled with truth[19], will not serve as adequate theories of the nature of evidence.

Relevant and Irrelevant Evidence

[19] The 'mere' has wide scope over the 'coupled with'. It rules out cases in which one's justified beliefs consist in one's knowledge.

The third premise in Williamson's argument, that all knowledge is evidence, is perhaps the least contentious. At a minimum, any knowledge which is directly relevant to a given hypothesis should form part of the evidence in favor of or opposed to it. Either way, any knowledge relevant to a given hypothesis is evidence. But Williamson's third premise is more general. The claim is that even knowledge which is apparently irrelevant to a hypothesis should count as evidence. Since it is not obviously evidence either for or against the given hypothesis, the case needs to be made that it is, nevertheless, evidence.

The case in favor of the view that even knowledge irrelevant to a given proposition nevertheless counts as evidence can follow at least two avenues. The first is that not all evidence must be evidence either for or against a given proposition; such a demand is a false dilemma. It is permissible to allow for a neutral category. Evidence which is neutral with respect to a proposition appropriately neither increases nor decreases one's subjective credence in that proposition. It neither confirms nor contradicts a given hypothesis. It neither promotes nor detracts from a proposed explanation. Yet it does satisfy the rest of the requirements for evidence. It is propositional; it is grasped; and it is endorsed. Of course it must also be true.

The second avenue is to note that the justificatory relevance of a piece of knowledge to a given proposition is a contextually dependent matter. In particular, whether one proposition is relevant to another can change over time. Note, for example, Peter Klein's comment in discussing Dretske's Zebra Case:

> "Would it be relevant evidence *for* the claim that they are zebras that they are not cleverly disguised aliens from a recently discovered planet outside our solar system? Or that they are not newly invented super-robots? Or that they are not members of the lost tribe of Israel who have been hiding out disguised as zebras? To repeat, that they are not painted mules would become relevant evidence only if there were some reason to suspect that the animals are painted mules."[20]

The distinction to which Klein is drawing attention is between relevant and irrelevant evidence, crucially not between evidence and non-evidence. Since known propositions have all the other elements required to be legitimate evidence, to exclude them as evidence on the grounds that currently they lack relevance to the proposition at hand is to be naively shortsighted. It is often the case that we cannot tell what will be relevant in the future. Better to hold all our currently or seemingly irrelevant knowledge in reserve than to cast it off as non-evidence. Returning to the legal examples, it would be criminally negligent to discard any bit of bagged physical evidence before a case is resolved. Even then, physical evidence is typically kept for a time in the event that new information is uncovered which would make that physical evidence, and the consequent propositional evidence derived from it, relevant to the case. As a purely methodological point, it is better to allow any knowledge to count as evidence regardless of its relevance to the question at hand than to force it to switch from evidence to non-evidence because one wants to be proprietary about the word 'evidence'. Letting all knowledge count as evidence has no downside whereas denying that all evidence counts as knowledge has several.

## Reliabilism

[20] "Contextualism and Academic Skepticism", in Skepticism edited by Ernest Sosa and Enrique Villanueva, Blackwell Publishers 2000.

Reliabilism has an illustrious history with many prominent defenders. Without getting into the specifics of a particular version of reliabilism, there are some features common to this view of justification. A belief is reliably justified only if the process or the method from which the belief is acquired reliably yields truths with respect to the content of the belief so acquired. This condition is circular insofar as it deploys the explanandum *reliably* in the explanans. But it can be expanded to overcome this difficulty. A process or method reliably produces truths only if it tends to produce them with few exceptions. The exceptions themselves should be explainable if the process is to be reliable. Consider the belief that there is a red ball. The process which produced the belief utilized a visual system comprising such components as one's eye, one's optical nerve, and various portions of one's brain. When presented with an unobstructed view of a red ball in sufficient light, a normal visual system will tend to produce, indeed, will nearly unerringly produce the belief that there is a red ball, provided that the viewer has sufficient conceptual resources to grasp such a proposition. Of course, such a system can fail to produce such a belief or produce it in inappropriate circumstances. Sometimes one will mistake a red ball for an orange ball. But such errors or omissions should be explainable in terms of things like a defect in the optical nerve, the interference produced by drugs in one's system, or the level of fatigue of the viewer. Nevertheless, in the main, beliefs about the presence and color of medium-sized objects which are causally dependant on the visual system of an ordinary human being are true. Reliabilists regard such justified true beliefs as constituting knowledge.

One standard difficulty for reliabilism is the generality problem.[21] The issue is that any reliabilist account of an adequately justified belief must specify the process which is causally responsible for the belief in question. After all, it is the reliability of a given belief-forming process which is central for the reliabilist account. But for any belief, the causal process from which it originates is debatable. Since there are a variety of candidates for *the* causal process, and since not all the candidates are in fact reliable processes, there appears to be no non-arbitrary way of specifying the relevant causal process which is in play when a belief is said to be reliably produced, i.e. a good candidate for knowledge. Consider the earlier example of coming to believe that there is a red ball. Only a vague sketch of the process was there articulated. Notice how many different processes could reasonably be considered. 1) Utilizing one's visual apparatus to observe a red ball. 2) Utilizing one's visual apparatus to observe a red ball while paying close attention. 3) Utilizing one's visual apparatus to observe a red ball on Tuesdays. 4) Utilizing one's visual apparatus to observe a red ball in New York. The list is ever so extendable. To cite just one potential complication, when the level of specificity of the process reaches a point where the observational event is characterized uniquely (that is, in such a way that no other observational process is type identical with it), then how is the reliability of such a process determined? It cannot be that such a process is reliable if and only if it always produces true beliefs (though this may seem the strongest form of reliability), for then reliability comes too cheaply. *Any* true belief has some process which is uniquely specifiable and such that it always produces true beliefs. Taking a modal perspective will not remedy the problem. Processes can be specified

[21] Conee, E. and Feldman, R. 1998. 'The Generality Problem for Reliabilism' *Philosophical Studies*, 89 1-29

with reference to the world in which they are deployed. At a minimum, any actual process can be specified with reference to the actual world. Once that element is added, there is no point in wondering how the process would behave in merely possible, relevantly possible, or counterfactual circumstances, for such cases do not arise. The generality problem may well be insoluble. If it is, then reliabilism suffers a mortal blow.

Even if reliabilism turns out to be defunct as a theory about the elusive fourth condition for an analysis of knowledge (or as a theory of the nature of justification necessary for knowledge in a standard justified, true belief account), there is much to be said in favor of its inclusion in our explanations of why particular justified true beliefs fail to qualify as knowledge. Williamson holds this view:

> "Let us concede for the sake of argument that the generality problem is indeed insoluble. It does not follow that appeals to reliability in epistemology should be abandoned. For the insolubility of the generality problem means that the concept of reliability cannot be defined in independent terms; it does not mean that the concept is incoherent. Most words express indefinable concepts; 'reliable' is not special in that respect. Irrespective of any relation to the concept *knows*, we clearly do have a workable concept *is reliable*…The concept is certainly vague, but most words express vague concepts; 'reliable' is not special in that respect either. The concept *is reliable* need not be precise to be related to the concept *knows*; it need only be vague in ways that correspond to the vagueness in *knows*." [22]

The examples where deploying concerns about reliability assists in the explanation of the degree of epistemic success are numerous and easy to construct. Consider a poker player deliberating over whether to call a particular bet. Many novice players trust in what may loosely be termed a "gut feeling" in order to determine whether the play is prudent. Such feelings are notoriously unreliable and the success of knowledgeable players at the expense of such amateurs is evidence for it. Veteran players utilize a range of

[22] KAIL, pg. 100

information including but not limited to their knowledge of the relevant probabilities of card distributions, their previous experience with the particular opponent, and subtle bodily cues which help reveal the attitude of the opposing player.

Rarely if ever does either player actually know what cards the other is holding. But the beliefs of the experienced player are better justified (and more likely to be true as a result) than those of the novice. A direct and appropriate explanation of such epistemic success can be given in terms of the reliability of the relative processes from which each player's belief is derived. To omit any mention of reliability on the grounds that the generality problem, or some other technical problem, complicates a complete theory of reliability is to miss out on available and useful data. Explanations which advert to reliability *are* useful. Even if reliabilism per se is problematic, the concept of reliability should not be abandoned.

On a more exegetical note, Williamson himself is committed to the view that reliability is necessary for knowledge. The earlier quotation occurs during his discussion of luminosity and his argument against it.[23] The anti-luminosity argument itself is crucial for several other positions Williamson adopts in other places, so if it relies on a tenable notion of reliability, then so too do those other positions.

Of course, to say that reliability is necessary for knowledge, that an unreliable belief-forming process compromises the formed beliefs' status as knowledge, is not to say that reliability is a component in a non-circular analysis of knowledge. Such a conclusion is

[23] See chapter three for further discussion of the anti-luminosity argument.

an *infamnia*. Many concepts are such that the existence of their extensions entails the existence of members of other concepts. But that entailment relation is not equivalent to nor as strong as a constitution relation. For Williamson, knowledge has no correct, non-circular analysis in the sense of 'analysis' in which the concept analyzed is decomposed into constitutive members. *A fortiori*, knowledge is not correctly and non-circularly analyzable into some complex containing reliability as a constituent. To rehearse just one counterexample[24], equilateral triangles are always such as to have sides indiscernible in length to the human eye, but being indiscernible in length to the human eye is not a constitutive member of a correct analysis of being an equilateral triangle. At a minimum, there is a clear difference between this relationship and the one that holds between *bachelor* and *male*. Even though the second of each pair of concepts is entailed by the first, only in the latter case is the connection a constitutive one.

Nevertheless, Williamson needs knowledge to entail reliability, at least in the sense required for the anti-luminosity argument. Such a reliance on reliability reveals an overarching virtue of the Williamsonian approach. Despite rejecting so much of what traditional epistemology considered invaluable, like the justified, true belief account of the analysis of knowledge, his theory is still in a position to utilize much of the philosophical work done by these earlier, and contemporary, philosophers.[25] Williamson can incorporate the work of reliabilists into his own theory by acknowledging that reliability is necessary for knowledge. The particular placement of a reliability condition inside a general epistemological program will still be a point of difference between

[24] KAIL, pg. 28
[25] See chapter four for further discussion of this point.

Williamson and the reliabilist. For the reliabilist, the reliability condition on knowledge will be a fourth component in an analysis of knowledge while for Williamson the reliability constraint is necessary for knowledge without being a conceptual component. But whenever a reliabilist is in a position to deny that a person S knows that p on the grounds that the process by which S acquired the belief that p is unreliable, Williamson will be in a position to agree. The upshot is that a reliabilist account of propositional evidence, where reliable belief-forming processes are necessary for the content of a belief to count as evidence, is compatible in most important respects with Williamson's view that one's evidence is identical with what one knows.

# Chapter Three

## Anti-Luminosity and Its Critics

Timothy Williamson argues at length in Knowledge and Its Limits[1] that knowing is a mental state. He takes pains to emphasize that knowing is a purely mental state in the sense that for every knowable proposition p, there is a mental state type that is both necessary and sufficient for knowing p. For Williamson, knowing that p is at least as mental as hoping that p, believing that p, feeling warm, or seeming to see a bicycle. His defense of the claim that knowing is a mental state primarily consists in defeating the chief objections directed against it, rather than providing positive support for the position (though that is also pursued). Chapter four of the book is dedicated to the objection that knowing cannot be a purely mental state on the grounds that purely mental states are such that they are open to a special kind of first-person access. This special kind of first-person access, which Williamson calls 'luminosity', consists in the satisfaction of the following schema by a condition C:

(L) For every case α, if in α C obtains, then in α one is in a position to know that C obtains.

[1] Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press. , hereafter KAIL.

where a condition either obtains or fails to obtain in a case, and a case depends on a subject, a time, and a possible world. Williamson maintains that no non-trivial conditions satisfy (L), and consequently there are no conditions which form a central core of purely mental states from which knowledge is allegedly to be excluded. His argument against (L) is largely based on considerations of reliability and empirical observations of the limited discriminatory capacities of ordinary human beings. While reliability is importantly related to knowledge, the particular spin on reliability which Williamson utilizes in the anti-luminosity argument is not the only permissible understanding of that property. That point, if correct, may seem to cordon off the effects of the anti-luminosity argument, limiting its scope to those conditions knowledge of which is contingent on the type of reliability specified in the anti-luminosity argument. Some authors have pursued that line, but their arguments are subject to strong criticism. The conclusion one should draw is that these opponents of Williamson's anti-luminosity thesis must embrace an unmotivated notion of reliability as it applies to the gradation structure in the original anti-luminosity argument. A different set of authors have advocated a modified version of the luminosity thesis that avoids Williamson's argument and that captures the intuitive appeal of the relationship between genuine mental states and first-person accessibility. Against these authors the response should be that the uses to which a luminosity thesis can be put are not ones to which the modified version of luminosity can be put.

The dialectic begins with Williamson's assertion that knowing is a purely mental state. Since one is in an unusually perspicuous epistemic position with respect to one's own mental states, for some mental states S, the fact that one is in S should be a sufficient

basis for one's knowledge that one is in S. Given a sufficient basis for knowing that p, an agent is "in a position to know that p" if whenever the agent considers whether p is the case, then the agent knows that p is the case. His opponent holds that some mental states are such that whenever one is in them, then one is at least in a position to know that one is in them. A few further background conditions are that she be free from the influence of psychotropic chemicals, fatigue, and other cognitively debilitating constraints, and that she have the conceptual resources to consider whether she is in the condition in question. Feeling cold, smelling sweetness, hearing a tone, and other conditions of that sort are thought to be open to that type of access in a way that "knowing that p" is not. This difference is thought to constitute a rejection of knowing as a central mental state and it is this conclusion that is the main target of the anti-luminosity argument.

It is an easy mistake to infer from the anti-luminosity argument that Williamson thinks that we have rather limited access to whether we know p, for any p. In fact, Williamson endorses the general sentiment that mental states are open to a special kind of first-person epistemic access, and he does not recommend any exceptions for knowledge. Two examples:

> "Our extensive but not unlimited ability to know without further observation whether we know something is what enables us to use knowledge as evidence. It constitutes an extensive but not unlimited ability to know without further acquisition of evidence whether something is part of our present evidence."[2]

> "Perhaps failures of transparency could not be the normal case, although that claim would require extensive argument. A more plausible claim is that we have some non-observational knowledge of our own mental states and not of the mental states of others. But then the same may be said of knowing:

[2] KAIL, pg. 15

> we have some non-observational knowledge of our own knowledge and ignorance and not of the knowledge and ignorance of others."[3]

So while Williamson is arguing that no non-trivial mental state is luminous, he is not arguing that non-trivial mental states are not especially epistemically accessible by the subject.

Mental states that satisfy (L) are philosophically interesting. They permit an avenue by which one can be saved from complete skepticism. That is not the only way that one might oppose the skeptic, but it is a notable one. In virtue of our direct contact with certain mental states, we can be assured that at least some facts are knowable because some facts are known. Consider Smith who is looking at the Mona Lisa. Smith may wonder whether the object she is looking at is in fact the Mona Lisa or merely a mock version that is sometimes presented as the real artwork (for reasons of security). Smith is not always in a position to know whether what she sees is in fact the Mona Lisa. Contrast that with Smith's seeming to see the Mona Lisa. Provided that Smith has an adequate understanding of the salient properties of the Mona Lisa, it appears that she is in a position to know that she seems to see the Mona Lisa merely by paying careful attention to her internal states. "Seeming to see" is open to first-person access in a way that plain "seeing" is not. While it is possible for one to fail to attend to the question of whether one seems to see the Mona Lisa and thereby fail to know that one seems to see the Mona Lisa, the claim is simply that if the agent does attend to the question of whether

[3] KAIL, pg. 25

a luminous condition obtains, then if that condition does obtain, then the agent knows that the condition obtains. This type of position opposes Williamson's argument.

Williamson begins his critique of (L) by choosing an arbitrary mental state that appears initially to have as good a chance at being luminous as any mental state. He selects the condition of feeling cold since if there is a class of mental states that is open to first-person access of the required kind, feeling cold is prima facie likely to qualify. Next he constructs a case in which an individual, Smith, initially clearly feels cold, then warms up slowly over a long period of time until she finally clearly does not feel cold. Suppose that at t0 it is zero degrees Celsius. Over the course of the day the temperature gradually rises until at tN the temperature is 35 degrees Celsius. During this time, Smith is continuously attending to the question of whether she feels cold.

Smith is a normal human being with ordinary discriminatory capacities with respect to her phenomenal states. At t0 Smith clearly feels cold. At tN Smith clearly does not feel cold. As time progresses beginning from t0, Smith is increasingly less confident in her judgments of whether she feels cold. It is important for Williamson's example that the amount of time which passes between any successive t(x) and t(x+1) be quite small. In particular, since the normal discriminatory capacities of human beings with respect to their phenomenal states do not permit Smith to notice the gradual qualitative changes between any two arbitrarily close time periods t(x) and t(x+1), there must be a finite unit of time into which the period t0 through tN can be divided such that Smith cannot discern whether the qualitative state experienced at t(x) is the same or different from the

qualitative state experienced at t(x+1). Williamson settles on increments of one millisecond for each t(x), but one could just as easily run the argument with smaller units of time.

The next step in the argument is to defend the following thesis (the "feeling cold thesis"):

(FCT): For all times t(x), if at t(x) Smith knows that she feels cold, then at t(x+1) Smith feels cold.

Is (FCT) uncompelling in this context?

Before exploring Williamson's reasons for endorsing (FCT), it is important to note that there is a way of understanding (FCT) such that its invocation in an argument against (L) is infelicitous. Provided that the condition of feeling cold satisfies (L), the moment that Smith changes from feeling cold to not feeling cold is the very same moment that Smith changes from being in a position to know that she feels cold to not being in a position to know that she feels cold. If we are also to regard the condition of feeling cold as either obtaining or failing to obtain at every time t(x), then there is a particular t(i) such that at t(i) the condition that Smith feels cold obtains while at t(i+1) the condition that Smith feels cold does not obtain. If feeling cold satisfies (L) (i.e. if feeling cold is a luminous condition), then at t(i) Smith knows that she feels cold. But by stipulation t(i+1) is such that the condition of feeling cold does not obtain then. So there is a time

t(x) such that at t(x) Smith knows that she feels cold but at t(x+1) Smith does not feel cold which is in conflict with (FCT). Therefore no consistent defender of (L) would deign to grant (FCT), especially not in any acceptable argument against (L).

In the strongest language, perhaps too strong, the criticism here is that Williamson is begging the question against defenders of the luminosity of feeling cold if he invokes (FCT) in his critique. But no such commitment to a charge of begging the question is required. It is enough for the defender of (L) to point out that the account of reliability embodied in (FCT) deploys an uncompelling version of that epistemic virtue, one which is either insufficiently substantiated or inappropriate in the dialectic given the theoretical presuppositions of the conversational participants. Resisting the urge to accuse Williamson of begging the question frees the defender of (L) from the burdensome if not impossible task of clarifying what exactly the fallacy of begging the question amounts to. By doing so, the focus of the discussion is properly located in the substantiation of the reliability principle (FCT), rather than the satisfaction of one or another accounts of a fallacy which itself has a substantial literature behind it.[4]

If Williamson were only appealing to (FCT) as an ad hoc device to undermine (L) as it applies to feeling cold, then the charge that the invocation of (FCT) in this context is in some way illicit might have some bite. As it happens, Williamson provides independent motivation for (FCT). If such grounds are genuinely independent, then that should exonerate Williamson from the present worry. But a strong, independent defense of

[4] See, for example, the discussions and bibliographies in Sorensen, R. ‘‘P, therefore P’ without Circularity’ *Journal of Philosophy* 88/5 May 1991 and ‘Unbeggable Questions’ *Analysis* 56/1 January 1996.

(FCT) is required, especially since it is a relatively easy series of logical steps from (FCT) and permissible background assumptions to the conclusion that luminosity fails for the condition of feeling cold. And since feeling cold is paradigmatic of putatively epistemically open mental states, the failure of this condition to be luminous augurs poorly for any defense of the view that (L) is satisfied by some non-trivial mental conditions.

As Smith experiences changes with respect to her phenomenal states from t0 to tN, she begins, at t0, with maximal confidence that she does in fact feel cold. With the passage of time, she gradually loses confidence in this judgment. At some point during the day, her confidence in the judgment that she feels cold will be entirely gone and in its place will be an increasing amount of confidence in the judgment that she is not feeling cold until finally she has maximal confidence in the judgment that she does not feel cold. If one knows that p, then one must be at least reasonably confident that p. Lacking confidence in one's beliefs is a prima facie impediment to knowledge. But merely having reasonable confidence that some condition holds is by itself insufficient to know that the condition holds; one's confidence must be reliably based if it is to be appropriate for knowledge. Unreliably based confidence, such as the sort derived from wishful thinking, is inadequate for knowledge. Let us grant that the condition of feeling cold satisfies (L). If the condition of feeling cold satisfies (L), then there will be some point during the day which is the last moment during which Smith feels cold and is simultaneously the last moment during which she knows that she feels cold. Consider the critical times t(i) and t(i+1) such that at t(i) Smith (allegedly) knows that she feels cold

(because she does feel cold) while at t(i+1) she does not know that she feels cold (because she does not feel cold). Since the changes in Smith's degree of confidence with respect to the judgment that she feels cold differ only slightly from one millisecond to the next (since it is based on her limited discriminatory abilities), if at t(i) Smith knows that she feels cold while at t(i+1) she does not feel cold, then she will have a significant degree of confidence in the judgment that she does feel cold at t(i+1) despite the fact that she does not feel cold at t(i+1).

More importantly, the degree of confidence which Smith possesses at t(i) that she feels cold is nearly, though not perfectly, identical to the degree of confidence she possesses at t(i+1) that she feels cold. She must have some sensible measure of confidence in either case in order to make the judgment at all. The trouble is that processes which give an agent reasonable degrees of confidence in a judgment that a condition obtains during circumstances in which that condition obtains, but which also give nearly as much confidence in the same judgment during circumstances in which the condition does not obtain are too unreliable to justify (in the sense required for knowledge) the confidence one has in the judgment during the times that the condition obtains. But since that is exactly what is going on in Williamson's example, Smith's confidence in her judgment at t(i) that she feels cold is based on a process which is unreliable and therefore does not allow the judgment to count as knowledge.

This understanding of what reliability amounts to constitutes the independent motivation for (FCT). In order for Smith to know that she feels cold, it must be the case

that she at least feels cold when she has similar levels of confidence in the judgment that she feels cold, at least when the basis for the confidence is virtually the same in each case. Thus if Smith knows at t(x) that she is cold, she must in fact feel cold whenever she has nearly as much confidence that she does feel cold, which in this case occurs at t(x+1).

Taken together, (FCT), (L), and the fact that at t0 Smith feels cold entail that at tN Smith feels cold. To see this, consider an instance of (L):

(L*) If at t(x) Smith feels cold, then at t(x+1) Smith knows that she feels cold.

By stipulation, at t0 Smith feels cold. By (L*) if at t0 Smith feels cold, then at t1 Smith knows that she feels cold. So by Modus Ponens at t1 Smith knows that she feels cold. By (FCT) if Smith knows at t1 that she feels cold, then at t2 Smith feels cold. But then by (L*) Smith also knows that she feels cold at t3 and the reasoning proceeds in a similar fashion all the way to tN. So, at tN Smith knows that she feels cold. But ex hypothesi Smith does not feel cold at tN, rather she is highly confident in her knowledge at tN that she does not feel cold.

Since the example clearly requires that Smith does not feel cold at tN while the combination of (FCT), (L*), and the fact that at t0 Smith feels cold clearly requires that Smith does feel cold at tN, some part of the argument must be faulty. Since each inference utilizes only valid rules (universal instantiation and Modus Ponens), at least one of the premises must be false. It is hopeless to argue that Smith does not feel cold at t0.

Such a move would still leave one open to a similar argument that instead begins with Smith not feeling cold and then gradually decreases the temperature until she clearly does feel cold. So either (FCT) or (L*) must be false. Since Williamson has apparently provided independent justification for (FCT), it follows that (L*) must be false. This completes the main line of Williamson's argument against the luminosity of non-trivial properties. All that remains is to make the generalization from the fact that the selection of "feeling cold" as the condition to be tested was entirely arbitrary, which is obvious; there are no non-trivial conditions that satisfy (L).

Since there is some cause for concern that by invoking (FCT) Williamson is invoking an overly controversial premise against defenders of (L), it is prudent to more carefully examine his arguments in favor of (FCT). As mentioned before, (FCT) is supposed to be an application of a reliability constraint. Invocations of reliability as a constraint on knowledge are in themselves harmless, since a defender of (L) has no special reason to deny the intuitive claim that one's belief-forming processes should be reliable if the beliefs so-formed are to constitute knowledge. The real question is how reliable that process needs to be and in what ways the reliability of a process can be evaluated. If the motivations for the particular reliability constraint that Williamson has in mind are independent of an attempt to defeat (L), then that will undermine the criticism that (FCT) is intrinsically unacceptable to a defender of (L).

Williamson says the following in connection with justifying his reliability constraint: "The basis on which one judges that one feels cold need not change suddenly as one

gradually becomes colder."[5] Since it is more germane to the example Williamson gives, we can interpret this claim as applying equally well to cases where one gradually becomes less cold. The basis in question here is the sequence of Smith's phenomenal states associated with active introspection during a period in which Smith begins by feeling cold and then gradually becomes less cold. Provided that Smith is, to the best of her abilities, actively and continuously considering whether she feels cold, there is a phenomenal state p(x) associated with every t(x) such that the basis for Smith's belief (or lack of belief) that she feels cold as well as the basis for Smith's confidence in the belief (or lack of belief) that she feels cold at t(x) is p(x). Each p(x) is only slightly different from p(x+1). Consider the two times t(y) and t(y+1) such that at t(y) Smith feels cold while at t(y+1) Smith does not feel cold. If there were a sharp contrast in the phenomenal states associated with t(y) and t(y+1), then that would be sufficient to explain the reliability of Smith's beliefs at t(y) even though the same judgment at t(y+1) would be unreliable. The two judgments would be based on sufficiently distinct phenomenal states and therefore the unreliability of one of them would not affect the reliability of the other. However, since the underlying basis for Smith's judgments is a sequence of only slightly different phenomenal states, ex hypothesi indiscriminable to Smith, there is no sharp difference between the two bases which a defender of (L) might exploit in opposition to Williamson's criticisms.

One can derive (FCT) from Williamson's notion of safety. A belief is safe, in his sense, only if one could not easily have made a mistaken judgment on the grounds of a very similar basis. To know that p, one's belief must be safe in Williamson's sense. For

[5] KAIL. pg. 99

a pair of indiscriminable (but non-identical) states, like the temporally close pairs in the earlier example, if one believes that one is in the phenomenal state p(x) occurring at t(x), then one could easily have believed that one was in the phenomenal state p(x+1) occurring at t(x+1). Consider two adjacent times on the cutoff of feeling cold, t(y) (where one does feel cold) and t(y+1) (where one does not feel cold). Suppose for reductio that S knows at t(y) that S feels cold. Since the experiences at the two times are indiscriminable, S could easily have believed that S felt cold at t(y+1). But such a belief at t(y+1) would be mistaken, so S's belief at t(y) isn't safe. The argument relies on the empirical observation that human beings cannot distinguish between phenomenal states which are sufficiently similar.

If this is an accurate picture of what Williamson has in mind when he said "The basis on which one judges that one feels cold need not change suddenly as one gradually becomes colder", then it appears that a defender of (L) (in the form of (L*)) has a retort. In the example involving Smith, it was stipulated that the divisions of time would be finite, but arbitrarily small. It is important for Williamson that the divisions be small enough to lend support to the claim that between the two moments[6] which flank the boundary of feeling cold t(y) and not feeling cold t(y+1), the basis on which Smith judges that she feels cold when she feels cold (at t(y)) is only slightly different from the basis

[6] This way of characterizing the temporal aspect of the situation suggests an implausibly discrete model of time. If the change in temperature is continuous, then even within very small time intervals the phenomenal "state" of the subject will be in flux. Trying to specify a last or first moment of time in any finite interval in order to compare the states of those moments is not trivial. So it is a great convenience to speak as though there was but one steady phenomenal state for each t(x). Since the time intervals are so small, it is harmless for present purposes to speak in this way. The reason is that the argument could be recast in terms of the pairwise comparison of any pair of moments (a,b) such that a is a member of (tx) and b is a member of (tx+1), where those time intervals are sufficiently small. For such pairs (a,b), the phenomenal state in each will be subjectively indiscriminable from the other.

she has access to when she does not feel cold (at t(y+1)). The intuitive appeal of Williamson's objection rests on the idea that there is a gradual change in temperature. Suppose that Smith were allowed to experience whatever phenomenal state is associated with t(x) for a length of time sufficient for her to become familiar with the sensation. If she were then permitted to experience the phenomenal state associated with t(x+1) and subsequently become familiar with it, she would still not be in a position to tell whether the state at t(x) was qualitatively identical to the one at t(x+1). Recall that increments of one millisecond were merely a convenience; Williamson could have run the argument using increments of one picosecond or smaller. A defender of (L) (in the form of (L*)) might take the following position. Yes, the changes in phenomenal states are gradual. Between any two successive phenomenal states there is little difference. Perhaps so little difference that Smith is not in a position to know that she is in one rather than the other. Nevertheless, the states themselves are different. When Smith feels cold, that condition is *luminous*, it "shines brightly enough" that careful consideration by Smith as to whether she feels cold will always permit her to tell that she is cold when she is cold. Certainly there are some tricky cases. The two times t(y) and t(y+1) which flank the border between feeling cold and not feeling cold appear problematic. But despite the minute differences between the phenomenal states p(y) and p(y+1), Smith's belief-forming process is in fact reliable. It is reliable precisely because the condition of feeling cold is luminous. When Smith feels cold, she is an excellent gauge of that fact. Her judgments on the matter are entirely reliable owing to the special kind of first-person access that she has to her own mental states.

One could, in this way, maintain the view that Williamson has the wrong notion of reliability in play here. Perhaps (FCT) does not capture what is important about reliability as it relates to knowledge of feeling cold under the circumstances so described. But Williamson has provided an argument in favor of that account of reliability, so if one wishes to reject it, then one should either provide grounds for thinking that the argument is unsound, or offer a competing account of the requisite type of reliability that has comparable merits in areas like explanatory value, intuitive appeal, and theoretical cohesion. Without such details, the rejection of (FCT) looks unmotivated. Merely noticing the logical possibility of a competing account of reliability is not enough to seriously threaten a standing account that displays a number of desirable theoretical virtues.

With respect to the issue of whether invoking (FCT) is illegitimate in an argument against a defender of (L), we appear to have reached an impasse. If Williamson's characterization of reliability in terms of proportions of confidence in one's judgments is correct, then (L) cannot be right. If (L) is correct, then Williamson's characterization of reliability cannot be right. Further frustration can be encountered when one realizes that Williamson is not necessarily resting his discussion of reliability on any general principle.

> "The use of the concept *is reliable* here is a way of drawing attention to an aspect of the case relevant to the application of the concept *knows*, just as one might use the concept *is reliable* in arguing that a machine ill serves its purpose. The aim is not to establish a universal generalization but to construct a counterexample to one, the luminosity principle (L). As with counterexamples to proposed analyses of concepts, we are not required to derive our judgment as to whether the concept applies in a particular case from general principles"[7]

[7] KAIL, pg. 101

He is more inclined to evaluate particular cases using pre-theoretic intuitions.[8] Since there is no general principle of reliability to be evaluated here, the discussion degrades into intuition mongering. If one thinks that some conditions satisfy (L), then one will not be inclined to accept Williamson's evaluations of what is and is not reliable. If one accepts Williamson's evaluations of reliability, then (L) looks unappealing. However, the fact that one does not accept Williamson's notion of what reliability is does not entail that the notion of reliability has no role to play in our understanding of knowledge. There are other notions of reliability which a defender of (L) could exploit in order to accommodate the plausible intuition that knowledge should be the result of, among other things, a reliable belief-forming process. These alternatives would have to be spelled out in detail and judged on their explanatory merits. Of course the work has yet to be done for this opponent of Williamson.

Brueckner and Fiocco

Anthony Brueckner and M. Oreste Fiocco, in their 2002 paper "Williamson's Anti-Luminosity Argument"[9], also take issue with the characterization of reliability which

[8] This strategy is not uncommon, especially in epistemology. Brian Weatherson, in his "What Good are Counterexamples" *Philosophical Studies* 2003, concurs. In any case, issues of whether it is preferable to rely more on strong intuitions about particular cases than on theoretically virtuous theories (one's with a great deal of coherence, explanatory scope, simplicity, etc.) or vice versa is independent of Williamson's discussion as results there will affect a wide variety of otherwise unrelated theories. For more on this, see Roderick Chisholm's "The Problem of the Criterion" from The Foundations of Knowing, University of Minneapolis Press, Minneapolis, 1982, reprinted in The Theory of Knowledge 2nd edition, edited by Louis Pojman, Wadsworth Publishing Company, 1999.

[9] Brueckner, A and Fiocco, M. 'Williamson's Anti-Luminosity Argument' *Philosophical Studies* 110, 2002

Williamson offers. In particular, they focus on a generalization of (FCT) which they label (R):

(R) For all $\alpha i$ such that $0 \leq i \leq n$, if KC in $\alpha i$, then C in $\alpha i+1$.

where '$\alpha i$' stands for the case at time i, 'C' stands for the subject of the case feeling cold, and 'KC' stands for the subject of the case knowing that she feels cold.[10] They find this reliability principle wanting since, according to them, it delivers the wrong verdict about whether one knows in a case which Williamson himself has described.

The case they have in mind involves the death of Abraham Lincoln. Consider the pair of times $\alpha i$ and $\alpha i+1$ which flank the moment of death of the former president. At $\alpha i+1$ Lincoln has already died. Since he is dead, he can no longer be president, and consequently no one can know at $\alpha i+1$ that Lincoln is president. But at $\alpha i$, while Lincoln is still alive, a common person's belief (at that time) that Lincoln is president is based on ordinary evidence like witnessing Lincoln's inauguration, or reading about it in a newspaper, or hearing from a reliable source that he was elected president. These sources of evidence would have resulted in the belief that Lincoln is president long before he entered Ford's Theatre. And in the absence of any evidence to the contrary, the belief that Lincoln is president (while he was president) is well-justified for the ordinary person. Since the evidence in support of the belief that Lincoln is president is the same at $\alpha i$ as it was for a substantial amount of time before $\alpha i$, Brueckner and Fiocco conclude that if

[10] ibid. pg. 285

someone knew that Lincoln was president on the basis of the aforementioned evidence, say, a week before the assassination, then that person also knew that Lincoln was president at αi. But it is precisely this fact, that at αi someone knows that Lincoln is president and at αi+1 Lincoln is not president, which serves as a counterexample to a substitution instance of (R), where the condition of feeling cold is replaced by the condition that Lincoln is president. If reliability demands that things which are known must be true in the next millisecond, then the Lincoln case should not be possible; yet it is.

Before addressing this argument directly, it is useful to examine the textual interpretation of Williamson which is operating here. Brueckner and Fiocco argue that Williamson has committed himself to the claim that at αi a common person can know that Lincoln is president on the basis of the evidence she earlier accepted. But this is more than what Williamson explicitly states. Williamson utilizes the Lincoln example, which Brueckner and Fiocco quote in their paper, to illustrate the failure of a strengthened KK thesis, that whenever one knows that p, one knows that one knows that p, and whenever one does not know that p, one knows that one does not know that p. Such failures of perfect first-person epistemic access to one's knowledge is one type of avenue by which Williamson's opponents might reject the claim that knowing is a mental state. For completeness, the text reads as follows:

> "One is surely not always in a position to know whether one knows p (for almost any proposition p), however alert and conceptually sophisticated one is. The point is most vivid when the subject believes p falsely. Consider, for example, the situation of a generally well-informed citizen N.N. who has not

> yet heard the news from the theatre where Lincoln has just been assassinated.[11] Since Lincoln is dead, he is no longer president, so N.N. no longer knows that Lincoln is president (knowing is factive). However, N.N. is in no position to know that anything is amiss. He continues reasonably to believe that Lincoln is president; moreover, this seems to him to be just another item of general knowledge. N.N. continues reasonably to believe that he knows that Lincoln is president. Although N.N. does not know that Lincoln is president, he is in no position to know that he does not know that Lincoln is president." [12]

As one can see, Williamson is not in that passage discussing issues of reliability. Furthermore, there is no commitment here to any position on whether N.N. knows that Lincoln is president a moment before Lincoln dies.

There are reasons to think that N.N. does not know a moment before Lincoln dies that Lincoln is president. For one thing, it is true that some moments before that there was a man pointing a loaded gun at Lincoln. That fact itself seems to defeat the justification N.N. has for his belief that Lincoln is president, at least in some relevant sense of 'defeat'.[13] Since knowledge is a binary state, one either has it or fails to have it with respect to a given proposition, there must be some pair of successive milliseconds t and t* such that at t N.N. knows that Lincoln is dead and at t* N.N. does not know that Lincoln is dead. But t and t* need not coincide with $\alpha i$ and $\alpha i+1$, the times flanking Lincoln's death. It is true that across the time period from $\alpha i$ to $\alpha i+1$ there is a change in the situation which is relevant to whether anyone can know that Lincoln is alive. One can't know what isn't so. But whether at $\alpha i$ N.N. is justified, in the sense relevant and necessary for knowledge, in believing that Lincoln is alive has yet to be established. Just

---

[11] In truth, Williamson misrepresents the facts here. If assassination requires death, then Lincoln was only mortally wounded at Ford's Theatre. He died nine hours after the shot inside the Peterson House across the street. And yet, "Lincoln was assassinated inside the Peterson House" also seems false.

[12] KAIL, pg. 23

[13] See, for example, Lehrer, K. and Paxson, T. 'Knowledge: Undefeated Justified True Belief' *The Journal of Philosophy* 66/8 1969. For confirmation that considerations of defeasibility are useful inside a Williamsonian framework, see chapter four.

because one's evidence was good enough for one to know that p at some time t does not imply that the same set of evidence is good enough for one to know that p at some later time t*. This is especially apparent inside a theory that embraces externalistic virtues. Since, for Williamson, one's evidence is identical with what one knows, and since what one knows is responsive to factors independent of one's beliefs, the justificatory status of one's beliefs can change when one's environment changes.

But there are more simple-minded observations to notice which help undermine the relationship between justification and knowledge that Brueckner and Fiocco seem to rely upon. Consider Kristy and her cat Nigel. At t, Bernard knows by a combination of sight, testimony, and inference that Kristy owns Nigel. But shortly after t, Bernard moves across the country. The fact that he is no longer in proximity to Kristy and Nigel does not immediately undermine his justification, and thereby his knowledge, that Kristy owns Nigel. Without being given evidence to the contrary or the emergence of some pernicious defeater, there is no reason to suppose that Bernard immediately loses his knowledge. His memory and inductive evidence about the persistence of ownership relations between pet-owners and pets sustains his knowledge even at a distance. But the effect is only temporary. As time passes and his evidence is not updated, those facts begin to erode his earlier justification for the belief that Kristy owns Nigel. After all, Kristy might unexpectedly decide to give Nigel away as a gift. Nigel might die. Kristy might die. The more time that passes without confirmation of his earlier beliefs, the more subjectively likely it is that the belief is false. At some point, the accumulated corrosion of his justification will obviate his knowledge that Kristy owns Nigel. This is similar to

the situation described by Brueckner and Fiocco. A generally well-informed citizen who has not yet been informed of Lincoln's assassination (shortly after it occurred) is not necessarily in a position to know that Lincoln is president moments before he is assassinated. The testimonial and observational justification which the citizen earlier possessed regarding the current president are not necessarily enough to sustain his knowledge that Lincoln is president at all times between when the justification was originally acquired and when Lincoln dies.

Even if Williamson were committed to the position that at αi N.N. does know that Lincoln is dead while at αi+1 N.N. does not know that Lincoln is dead, that alone would not be enough to show that the reliability principle (R) is false. Williamson himself may be responsible for the confusion here, for he is not as clear as he could be in the original text about the scope of principle (R).[14] He is using it in his anti-luminosity argument in order to capture something epistemically relevant about the feeling cold case *as described*. It is a fact about the process of warming up from dawn until noon that, when one is measuring in temporal increments as small as a millisecond, one's discriminatory capacities with respect to one's own feelings are not perfectly sensitive to the minute differences between any two adjacent points. Any confidence which one has in one's judgments about whether one feels cold at a given time is only minutely different from the amount of confidence that one has a moment later, too small a difference to allow for knowledge in the one case and false belief in the other. The issue is intimately connected to our limited discriminatory capacities, we are not perfect observers even of ourselves. For our beliefs to be reliably based *in the case as stated* it must be the case that whenever

[14] KAIL, pg. 97. There the principle is called '(Ii)'.

one knows that one is cold, in the next millisecond one is cold. The fact that in some other cases it is possible for one to know that p at t and then have it be false that p a millisecond later at t* is not relevant. Brueckner and Fiocco are not attacking the right principle.

If they really wish to undermine (R) with the Lincoln example, then they must construct a case which matches the feeling cold case in terms of its reliance on our limited discriminatory capacities. Consider Dr. Joseph Barnes, one of the doctors present at Lincoln's death. Suppose Dr. Barnes paid close attention to the heart rate and breathing of his dying patient continuously until a few minutes after Lincoln was dead. This situation is now in line with the feeling cold case. Presumably Lincoln's condition deteriorated, his heart rate fluctuated, his breathing grew more labored, and then all was still and quiet. But Dr. Barnes goes on listening nonetheless, searching for some sign of life from his beloved president. Only when he is sure that Lincoln has passed does he declare him deceased. There was a point $\alpha i$ such that at $\alpha i$ Lincoln was alive and a millisecond later, at $\alpha i+1$, Lincoln was dead. But did Dr. Barnes know at $\alpha i$ that Lincoln was president? Allowing that after 1861 Lincoln is president if and only if he is still alive, it is doubtful that Dr. Barnes had such knowledge. Dr. Barnes doesn't know that Lincoln is dead until he has witnessed the effects of death for a period of time longer than a millisecond. But as Lincoln's death-like symptoms become more and more obvious, each cycle between heartbeats represents an area of ignorance. Dr. Barnes is not a perfect doctor; still less is Lincoln's death luminous for him, no matter how attentive the

good doctor is. The Lincoln case does not pose the problems for Williamson that Brueckner and Fiocco suggest.

But there are further proposed counterexamples from Brueckner and Fiocco to overcome. Here, in response to a desire to see their counterexample more closely parallel Williamson's own condition of feeling cold, they offer the dead parrot example:

> "Suppose that S is staring at a dead parrot for five hours and correctly believes that he sees a dead parrot throughout this interval. At the time t at which the interval ends, S sees the dead parrot and then blinks. One millisecond later, at t+1, S opens his eyes and sees a dead-parrot-hologram. At t+1, S mistakenly believes P (= S sees a dead parrot).
>
> Let us assume that prior to t+1, the Deception Squad had been completely unable to produce any holograms. The hologram-producer finally goes briefly on-line at t+1, and it is linked to a hologram-placer that randomly places holograms. It just so happens that the Squad's first and only success is a dead-parrot-hologram that winds up being placed before S at t+1, unbeknownst to the Squad (who were trying to produce a live-chihuahua-hologram). So prior to t+1, S was not at any time hanging fire on, or in any sense in imminent danger of, being deceived by means of holograms."[15]

This proposed counterexample offers a different reason against (R) than the Lincoln example did. In the dead parrot case, it seems that at t there is no question that S knows that S sees a dead parrot, whereas Dr. Barnes was not in a comparable position in the Lincoln case. While staring at the parrot for five hours before t, S knows by sight that there is a dead parrot before him; and by introspection and inference S knows that S sees a dead parrot. Since at t+1 S does not see that there is a parrot before him, S does not know that there is a parrot before him, and therefore S does not know that S sees a dead parrot. But if at t+1 S does not know that S sees a dead parrot, then, by an application of (R), S does not know at t that S sees a dead parrot. But there is nothing to distinguish what is taking place at t from what took place during the hours just prior to t, and during

[15] Brueckner and Fiocco 2002, pg. 289

those hours S clearly knew that S saw a dead parrot. The strategy here is similar to the Lincoln case. Brueckner and Fiocco begin by illustrating a case in which a subject switches from knowing to not knowing in the space of a millisecond, in such a way as to violate condition (R).

But this second proposed counterexample to (R) succumbs to essentially the same criticism as the Lincoln case. Brueckner and Fiocco have picked the wrong target. The reliability principle embodied in (R) is not a general reliability condition but only a characterization of what reliability amounts to when considering conditions which are under observation by beings of limited discriminatory abilities, and then only when the underlying basis for one's confidence in the judgment changes at most slightly from one interval to the next. That may strike one as an overly restrictive thesis. But it is possible to subject any non-trivial condition to the same kind of anti-luminosity argument used for the condition of feeling cold. If all non-trivial conditions can be substituted into the schema of the original anti-luminosity argument, then there is no significant restriction on the scope of the anti-luminosity thesis.

Consider an aspirant condition which may not at first appear to be a good candidate for the anti-luminosity argument's structure. Scientists are routinely in a position where they must rely on the readings of their instruments. Suppose that there is a computer monitor that gives a digital read out of the temperature of water in a given container. The instrument is only fine-grained enough to detect changes of whole degrees. Even the condition "the screen reads 20 degrees" is subject to the anti-luminosity argument. What

is necessary is to create a scenario in which the screen changes from 19 to 20 degrees over a long period of time. The interval can be broken down into small units. If there are only minute changes in the way the screen looks, say because only one pixel at a time is changed or the pixels themselves dim in and out, then whenever S knows that the screen reads 20 degrees, it must in fact read 20 degree in the next moment or else the reliability of S's confidence in her judgment that the screen reads 20 degrees is not reliably based. That conditions are often, even routinely such that whenever one is in them one can know one is in them is not good enough for the condition to be luminous. A higher degree of transparency is required and that is what Williamson rejects. Other conditions are not likely to fare better than this one as candidate exceptions from the scope of the anti-luminosity argument.

Hawthorne and DeRose

John Hawthorne and Keith DeRose have each offered a similar criticism of Williamson's anti-luminosity argument.[16] Their criticism is not that the anti-luminosity argument is unsuccessful, but that even if it is successful the desired conclusion does not follow. While the anti-luminosity argument is put to use for other purposes, one of its stated goals is to oppose the view that reasons for denying that knowledge is a purely mental state include the fact that genuine mental states are open to a special kind of first-person epistemic access. Williamson characterizes this level of access as perfect luminosity; for any genuine mental state, whenever one is in it, one is in a position to

[16] Hawthorne, J. "Knowledge and Evidence" Philosophy and Phenomenological Research, March 2005 and DeRose, K. "Review of Knowledge and Its Limits" British Journal of Philosophy of Science 53, 2002.

know that one is in it. But this is not the only way that one may quantify the appropriate degree of epistemic access required by genuine mental states. Hawthorne and DeRose offer similar accounts of such a competing quantification.

The motivation for an alternative account of the degree to which one must have access to a condition for such a condition to count as genuinely mental stems from a desire to preserve the "Cartesian Theatre" model of our mental lives. There was no need for the anti-luminosity argument to demonstrate that knowledge was not perfectly open to first-person access; ordinary facts, like the earlier quoted passage about N.N., are enough to show that it is false that for any p, whenever we know that p, we know that we know that p. But there is still hope that a narrow set of mental conditions will serve as the genuinely mental central core of mental states whose distinguishing feature will be that the members of the set would all be such that a subject enjoyed special epistemic access to them, at least in the sense that whenever one was in one of them, one would be in a position to know that one was in one of them. An alternative account of the level of required accessibility for admittance into this set is useful if it can accept that no non-trivial condition is luminous, and yet explain why people thought that there was a central core of luminous mental conditions by showing that there is another account of the degree of accessibility required of a genuinely mental condition. In that way, the main conclusion of Williamson's anti-luminosity argument is preserved, but there is theoretical room to accommodate the intuitions of those who see knowledge as importantly different from more pure mental states. Perfect luminosity would have been nice, but if knowledge differs significantly enough from a natural class of uncontroversially mental

conditions with respect to how accessible such conditions are to the subject, then Williamson will be in a poor position to defend his view that knowledge is a genuinely mental state.

DeRose adopts this strategy first by granting the conclusion of Williamson's anti-luminosity argument. He then explains that the cases which are counterexamples to luminosity all occur at the border area between the regions where the condition clearly applies and the region where the condition clearly does not apply. The closer one is to the border, the more unreliable are one's judgments of whether the condition obtains. Ignorance is ever-present, but only at the edges.

Williamson has defended his position on the necessity of reliability for knowledge by adverting to safety.[17] Roughly, for Williamson, and others[18], one knows that p only if one's belief that p is safe. An exact characterization of what 'safe' means in this context is somewhat elusive. But Williamson does have several important connections to draw between the concepts of safety, reliability, robustness, and other related notions. A belief is safe if it is not in danger of being false. In counterfactual cases which are close enough to the actual world, one's belief is not false.[19] DeRose draws on this reliance when he offers a principle of weak luminosity (WL) which will achieve the second purpose of securing the intuitions of those who initially opposed the anti-luminosity argument.

---

[17] KAIL, pg. 123-127
[18] Sosa, E. 'How to Defeat Opposition to Moore' *Philosophical Perspectives* 13 Epistemology, 1999
[19] See, for example, Luper-Foy, S. *The Possibility of Knowledge: Nozick and his Critics* Totowa, N.J.: Rowman and Littlefield 1987

(WL) For every case α, if in α C safely obtains, then in α one is in a position to know that C obtains.[20]

Given that a central core of mental states are weakly luminous in the sense of (WL), Williamson's anti-luminosity argument can no longer be used to object to the view that knowledge is not a mental state because knowledge is not as accessible as a central core of other, more obviously mental states. Simple phenomenal conditions, like feeling cold, can satisfy (WL) in a way that knowing that p cannot.

Hawthorne suggest an argumentative strategy that proceeds along similar lines. He begins by observing that one motivation for rejecting luminosity is that a phenomenal conception of evidence, which Williamson rejects and some skeptics have defended, is more plausible if the phenomenal is a realm utterly open to epistemic access. If no class of conditions are luminous, then there is no special reason to hold phenomenal evidence in higher regard than other sorts evidence. Those who wish to hold on to a solely phenomenal conception of evidence are thus well-advised to look for fault in Williamson's anti-luminosity thesis. Failing that, however, there are still avenues of retreat. A defender of the purely phenomenal conception of evidence could grant that Williamson's anti-luminosity argument works, but nevertheless there remains a principled and important epistemic asymmetry between a central core of purely mental

[20] DeRose, 2002.

states and the hybrid state of knowledge.[21] The central core is *cozy* while knowledge is not. Hawthorne defines coziness as follows:

"A condition C is cozy iff in every case α in which, determinately, C obtains, one is in a position to know in α that C obtains."[22]

The idea is that any ignorance in the anti-luminosity argument is essentially due to vagueness. It is important to note that a condition can fail to obtain determinately and yet obtain nonetheless. If that vagueness is removed by only considering that region which is not near the border area between the condition obtaining and not obtaining, then there are several phenomenal conditions which appear to satisfy coziness. Since knowledge does not satisfy coziness, we here have an epistemic property which can serve as a dividing line between genuinely mental states and imposter mental states.

There is a complication here that Hawthorne is eager to avoid. One might think that coziness is a somewhat trivial property which is easily had by a variety of conditions not all of which are good candidates for membership in a central core of genuinely mental states.[23] If one is allowed to pick out conditions by way of a demonstrative, then one will be in a position to know that one is in a particular phenomenal condition simply by thinking of it as *this* condition. "I feel thus" will pick out whatever phenomenal condition one is in and there is no worry about ignorance here. However one feels, one

[21] Hawthorne himself does not endorse this view, but instead points out that Williamson has yet to overcome it.
[22] "Knowledge and Evidence" Hawthorne 2005 pg. 453
[23] ibid. pg. 453, the examples are from there.

feels whatever one feels. But non-mental conditions satisfy this as well. Whatever the temperature is, one is nearly always in a position to know that the temperature is *this*, so the condition that the temperature is *this* will be cozy. In fact, this point goes beyond coziness. Admitting conditions specified by demonstratives will permit many conditions so described to not only be cozy but luminous. Neither Hawthorne nor Williamson desires that result.

To combat this difficulty, Hawthorne adjusts his account of coziness so that it applies to concepts of conditions rather than the conditions themselves. The revised account of coziness is this:

"A concept C, denoting condition c, is cozy iff in any case α in which C determinately holds, one is in a position to know in α that c obtains by an exercise of C."[24]

Given this new definition, there is no longer a worry about the triviality of coziness. Concepts of extra-mental conditions, like the temperature being 72 degrees Fahrenheit, are not cozy. Even when and where the concept "temperature of 72 degrees Fahrenheit" applies determinately, one need not even be in a position to know that the temperature is 72 degrees by an application of that concept. The new definition of coziness also renders plausible the original idea that there are a central core of concepts of phenomenal conditions which are open to a special kind of epistemic access. Namely, that when the concept determinately applies, one is in a position to know that the condition denoted by the concept obtains. Since the concept of knowledge does not satisfy coziness, we have

[24] ibid. pg. 454

here a criterion by which one can separate the concepts of genuinely mental conditions from the concepts of not genuinely mental conditions.

Hawthorne is conscious of the fact that not all phenomenal conditions are cozy. For example, "highly complex phenomenal conditions (such as how many phenomenal lines there are before me) are not, intuitively cozy."[25] He urges[26] that the fact that there are some phenomenal conditions which do not satisfy coziness is no impediment to the existence of a central core of mental states which do satisfy coziness. Perhaps that is the case.

There are a number of ways one might attempt to respond to the Hawthorne/DeRose position. If there are genuinely mental conditions which do not satisfy coziness, or if there are cozy conditions which are not actually mental, then the value of the Hawthorne/DeRose position is injured. One of Williamson's stated goals is to block attempts to demarcate a class of *genuinely* mental states from states which are less than fully mental if such a demarcation puts knowledge on the non-genuinely mental side. So one issue to consider is whether the DeRose/Hawthorne position achieves something which Williamson's theory is in a position to deny. Hawthorne has already conceded that there are phenomenal conditions which are not cozy. And the conditions in question, being phenomenal, are uncontroversially mental. So the best that the position can achieve is to establish a class of cozy mental conditions which form a basic core, but which exclude knowledge from the core. But in what sense are these conditions basic?

---

[25] ibid pg. 453

[26] Again, Hawthorne does not overtly endorse this line of thought, but sees it as an unresolved problem for Williamson's view.

One way they could be basic is by being constituents of the more complex, non-cozy phenomenal conditions. The directness of our epistemic access to complex phenomenal conditions would then be a function of our access to the basic conditions coupled with restrictions based on our finite cognitive abilities.

If this is the sense of "basic" in which a class of cozy phenomenal conditions is basic, then it involves more complex theoretical commitments than one might expect. It requires an account of why the basic core is important, and the explanation of its importance cannot simply be that knowledge ought to be counted as non-mental in some sense. It also requires a principled way to delineate between basic and non-basic concepts. The resolution of these theoretical concerns should precede the use of their consequences in an argument against Williamson. It is neither clear that such commitments can be legitimized, nor that either Hawthorne or DeRose would endorse such a position.

But even if we grant that there is a class of basic phenomenal conditions, that is not enough to disregard the importance of the anti-luminosity thesis. Williamson uses the anti-luminosity thesis in a number of places for a variety of purposes. The Hawthorne/DeRose position is designed to complicate only one of the uses to which it is put. Here is an unresolved worry for anyone who is convinced of their thesis. Williamson has a provocative argument against skepticism which begins by constructing an argument on the skeptic's behalf and then shows that parallel reasoning leads one to a

clearly false conclusion. But that argument for skepticism begins with a premise which commits the skeptic to the denial of the anti-luminosity thesis. Namely:

(1) For any appropriate property π, in any case in which one's evidence has π, one knows that one's evidence has π.[27]

The argument for the skeptic will not be valid if (1) is instead replaced with the following:

(1*) For any appropriate property π, in any case in which one's evidence determinately has π, one knows that one's evidence has π.

To see this, we can examine each step in the original argument and add references to coziness where appropriate.[28] To begin, the first premise is (1*) itself. Next, premises (2) and (3) can be used without modification, since their justifications are independent of whether the appropriate properties of one's evidence are cozy.

(2) For any appropriate property π[29], if in the good case one's evidence lacks π, then in the bad case one knows that in the good case one's evidence lacks π.

(3) It is consistent with what one knows in the bad case that one is in the good case.

---

[27] KAIL, pg. 171
[28] The numbering and the unmodified premises themselves follow Williamson, KAIL pg. 171-172.
[29] In order to capture as wide a set of properties as possible, Williamson holds that appropriateness of a property π is closed under complementation. That is, if π is an appropriate property, then so is not-π.

Since the goal of this argument is to show that one's evidence is the same in both the good and bad cases, we need a given fact about one's evidence in one case and a reductio assumption concerning the other case that is in conflict with that given fact.

(4) In the bad case one's evidence has π.

And assume for reductio that

(5) In the good case one's evidence lacks π.

Since (2) and (5) are identical to the original premises, they continue to entail:

(6) In the bad case one knows that in the good case one's evidence lacks π.

At this point the argument hits a snag. In the original argument, premises (1) and (4) were used to entail that in the bad case one knows that one's evidence has π. But since the new argument only has access to (1*), the best that it can manage when combined with (4) is:

(7*) If one's evidence in the bad case determinately has π, then in the bad case one knows that one's evidence has π.

From (5) and what one knows in the bad case as a consequence of (7*), one can deduce that one is not in the good case, but such a deduction can be accomplished only if the π in question applies determinately. Therefore, from within the bad case,

(8*) If π applies determinately, then it is inconsistent with what one knows in the bad case that one is in the good case.

When the π in question does in fact apply determinately, then the consequent of (8*) contradicts (3). In such a case, we can discharge the reductio assumption (5) and deduce its complement:

(9) In the good case one's evidence has π.

What has been demonstrated is that one's evidence in both the good and bad cases is the same, but that demonstration is possible only if the evidence in question has its relevant properties determinately. If the evidence in question does not have the relevant properties determinately, then (8*) remains inert and the conclusion that one's evidence is the same in both the good and bad cases is blocked. It would be folly for the skeptic to rest her position on such a thin argument, for if all it takes for one's evidence to fail to be the same in both the good and bad cases is that the relevant properties of that evidence do not determinately apply, then there is no real threat on this front from skepticism at all.[30]

[30] Williamson says of the original argument for the sameness of evidence, "If something like this argument is not the reason for which skeptics and others think that one has the same evidence in the two cases, it is not at all clear what is." (KAIL, pg. 173)

Plenty of appropriate properties of one's evidence do not determinately apply. To restrict what counts as an appropriate property of one's evidence to just those properties which only apply if they determinately apply is to invoke an ad hoc and unduly strong condition on one's evidence. Without the luminosity premise (1), Williamson's work against the skeptic is already finished.

# Chapter Four

## Testimony and Knowledge

Much of what we know is known by way of testimony[1]. Whether in the form of verbal communication, the written word, or more subtle cues, the giving and receiving of testimony is inextricably connected to our social network of inter-doxastic support. So much is granted by Jennifer Lackey in "Testimonial Knowledge and Transmission"[2]. What she argues in that paper is that it is possible for knowledge to be acquired solely on the basis of testimony (and a few uncontroversial background beliefs), even if the individual giving the testimony does not herself know her statements to be true. Lackey's cases in favor of this view are focused around two central properties: the failure of defeaters to transfer to a testimonial recipient even when those defeaters prevent the testifier from knowing, and a possible epistemic asymmetry between a testifier and a hearer with respect to the influence of unresolved yet pernicious skeptical worries. The cases themselves appear plausible. Unfortunately, Lackey's desired conclusion does not follow. This paper will argue that Lackey's proposed counterexamples to the general principle (2) are spurious:

[1] The word 'testimony' is used by Lackey, and many other epistemologists, somewhat differently than it is used in common parlance. Typically, to give testimony that p requires that there be an appropriate authority present. In its present, philosophical usage, the term is meant to capture the same general activity as telling and asserting. The testifier of p is at least representing that one could come to know that p on the basis of her testimony.

[2] Lackey, J. 'Testimonial Knowledge and Transmission', *The Philosophical Quarterly*, Vol. 49 (Oct. 1999)

"(2): For every speaker S and hearer H, if H comes to know that p via S's testifying that p, then S must know that p."[3]

For while Lackey has shown that not all defeaters transfer when testimony occurs, nevertheless, whenever a speaker does testify to something which she does not know, there will always be some defeater for the listener.[4]

The principle embodied in (2) has intuitive merit. Testimony is open to problems that do not encumber other modes for knowledge acquisition, such as perception. In a case of seeing that p, for example, one is faced with standard skeptical concerns about malevolent demons and mad scientists, as well as more nuanced Cartesian worries over the reliability of our epistemic faculties given that they have, occasionally, misled us in the past. Sight is not a perfect guide to what is true. But whatever difficulties might beset one's attempts to know that p via sight, those problems are often compounded in the case of testimony, for testimony itself must eventually be grounded in some other epistemic mode; a chain of reasons cannot invoke testimony at every stage. And with each stage in a string of testimony, the connection with the original non-testimonial information becomes ever more distant and tenuous. On this point, Hume says:

> "Our evidence, then, for, the truth of the Christian religion is less than the evidence for the truth of our senses; because, even in the first authors of our religion, it was no greater; and it is evident it must diminish in passing from them to their disciples; nor can any one rest such confidence in their testimony, as in the immediate object of his senses."[5]

---

[3] ibid pg. 473

[4] More rigorously, whenever a speaker S testifies that p while she does not know that p, there will always be a defeater d for any hearer H with respect to H acquiring knowledge that p solely on the basis of S's testimony that p and uncontroversial background assumptions.

[5] Hume, David Enquiry Concerning Human Understanding, "Of Miracles"

The danger of error demands a high degree of epistemic integrity in order for the testimony to be an avenue to knowledge for the listener. The principle in (2) captures that intuition by requiring that testimony be known by the speaker before it can be a source from which others can know. If testimony must be known whenever it serves as the basis for another's knowledge, then there is no intrinsic drop in the epistemic status of the information conferred. In that sense, a solid chain of testimony, made solid in virtue of being known at each stage, links one's present knowledge to some earlier non-testimonial knowledge. If knowing that p via testimony required anything less, then the epistemic status of the putative knowledge that p would be impugned.

Lackey is unconvinced. According to her, there are cases which clearly show that a person B can acquire knowledge that p via testimony that p from a person A even if A does not know that p. She blames adherence to (2) on the view that defeaters are transmitted via testimony. Once we divest ourselves of the opinion that defeaters are always transmitted in this way, we are free to acknowledge that a hearer may have better epistemic contact with the content of the testimony than the speaker does. For instance, a hearer may be safe from defeaters which attack the speaker. Since, in her view, a hearer can come to know that p via testimony even when the first testifier in the chain of testimony lacks knowledge, testimony can be a way of *generating* knowledge instead of merely *transmitting* it.

The plausibility of Lackey's position is strengthened by the independent, prima facie clarity of our judgments regarding the proposed counterexamples to (2). Thus the case is here presented in full:

> "Suppose that a Catholic elementary school requires that all teachers include sections on evolutionary theory in their science classes and that the teachers conceal their own personal beliefs regarding this subject matter. Mrs. Smith, a teacher at the school in question, goes to the library, researches this literature from reliable sources, and on this basis develops a set of reliable lecture notes from which she will teach the material to her students. Despite this, however, Mrs. Smith is herself a devout creationist and hence does not believe that evolutionary theory is true, but she none the less follows the requirement to teach the theory to her students. Now assuming that evolutionary theory is true, in this case it seems reasonable to assume that Mrs. Smith's students can come to have knowledge via her testimony, despite the fact that she fails [to believe what she says] and hence does not have the knowledge in question herself. That is, it seems that she can give to her students what she does not have herself. For in spite of Mrs. Smith's failure to believe and therewith to know the propositions she is reporting to her students about evolution, she is a reliable testifier for this information, and on the basis of her testimony it seems that the students in question can come to have knowledge of evolutionary theory. I take it that similar considerations apply in cases where a Kantian teaches utilitarianism, a dualist teaches physicalism, an atheist teaches Christianity, and so on. If the theory in question is true and the hearer comes to believe it by means of the teacher's testimony, then, I would say, the hearer can acquire knowledge on this basis despite the failure [of the speaker to believe (and thereby know) what was said]." [6]

The bracketed portion might be seen to contain an ambiguity. Unless otherwise specified, let "know what was said" describe the situation in which the content of the assertion is also a member of the knowledge set. The issue is not about linguistic comprehension.

In this first proposed counterexample to (2), Lackey articulates a case in which (allegedly) Mrs. Smith does not know that evolutionary theory is true, Mrs. Smith gives testimony to the students to the effect that evolutionary theory is true, and the students come to know (or at least are in a position to know) that evolutionary theory is true on the basis of Mrs. Smith's testimony. By hypothesis, evolutionary theory is true. The

[6]Lackey 1999, pg. 477

students have every reason to trust Mrs. Smith, and the information which Mrs. Smith is conveying was in fact derived from reliable sources. It appears that the belief that evolutionary theory is true is justified for the students. The students routinely acquire beliefs on the basis of Mrs. Smith's testimony, and this case is no exception. Thus the students have a justified, true belief that evolutionary theory is true.

Before going further, it is useful to distinguish two possible interpretations of the Mrs. Smith case.[7] Under the first interpretation, when Mrs. Smith presents the material on evolution, she does so by proffering evidence in its favor which she gleaned from reliable scientific sources. Mrs. Smith herself does not believe that evolutionary theory is true, but she does believe that the individual pieces of evidence which she presents to her students, such as the fossil records, the analyses of geological strata, and the carbon dating results, are all correct. Her resistance to the conclusion that evolutionary theory is true is a function of her dogmatic commitment to anti-evolutionary canon. In this situation, if the students are reflective, they can come to know that evolutionary theory is true on the basis of some of Mrs. Smith's testimony because they can deduce it from the evidence she provided in her testimony.

Under the second interpretation, Mrs. Smith does not provide the students with sufficient evidence of the sort discussed in the first interpretation from which they would be in a position to deduce that evolutionary theory is true. Instead, she provides a detailed characterization of what evolutionary theory is and then asserts that it is true. Alternatively, she does provide the students with the sort of evidence discussed in the

[7] Igal Kvart pointed out this ambiguity in conversation.

first interpretation, but the students are unreflective about the logical or abductive relationship between the evidence and the conclusion that evolutionary theory is true. They take her word for it at every step.

It is the second interpretation of the Mrs. Smith example that is relevant to Lackey's discussion. What is at issue is whether a hearer can acquire knowledge that p solely by way of testimony that p from a speaker who does not know that p. Of course it is impossible to eliminate all instances of inference and background knowledge from cases of testimony. The students must utilize other epistemic resources to secure the connection, for example, between the sounds that Mrs. Smith makes when she speaks and the propositions that she wishes to convey. They must infer that when Mrs. Smith asserts something in class that they are expected to take it seriously. These are just a few of the many features of testimonial activity which are supposed to be bracketed off for the present debate. The single most important feature of the Mrs. Smith example is that the students' primary epistemic contact with the content of Mrs. Smith's assertions is a function of their accepting what Mrs. Smith said because she said it.

Similar considerations apply to interpreting the word 'via' in (2). What is important about the principle is the connection between the testifier and hearer as mediated by the expression, interpretation, and acceptance of a given utterance on the basis of the word of the speaker. What is not supposed to be included are unusual causal connections that may satisfy the "via" clause but fall outside the philosophically relevant context. For example, a testifier might assert that p with such a loud voice that it causes an avalanche.

If the avalanche hits a hearer just right, the hearer could come to believe that p as a causal consequence of the testifier asserting that p. But that would not be an instance of the communicative activity that is presently under scrutiny.[8]

Even though the students unquestionably have a justified true belief, Gettier's famous 1963 paper[9] demonstrated that justified, true belief is alone insufficient for knowledge. Defeasibility theory is one prominent view designed to account for the remaining elements in the correct analysis of knowledge. Lackey characterizes two defeasibility principles. Each share a common feature which is best illustrated by the weaker of the two:

"(JP) If S reports that p to H and H has no defeaters for S's report that p, then H is justified in accepting that p on the basis of S's testimony."[10]

Lackey apparently endorses some such principle, at least to the extent that no case of knowledge acquisition by way of testimony can succeed if the hearer has any (undefeated) defeaters for p. The exact meaning of "has no defeaters" turns out to be an important issue for Lackey's thesis. It is also the locus of the argument against Lackey's position.

[8] Specifying the individuation conditions for cases of pure testimonial knowledge is itself a daunting enterprise and too expansive to handle in this context. It is hoped that legitimate philosophical headway can be achieved here without providing a perfectly robust and specific account of that social practice.
[9] Edmund Gettier, 'Is Justified True Belief Knowledge?', *Analysis*, Vol. 23 (1963)
[10] Lackey 1999, pg. 474

Lackey offers three views on how defeaters function. She divides them into doxastic defeaters, normative defeaters, and factual defeaters. A doxastic defeater functions "by virtue of being *believed*, regardless of [its] truth-value".[11] Absent being believed, a proposition cannot defeat the justification for a belief even if it is a true proposition. But a belief can defeat the justification for another belief even if the first belief is in fact false. Doxastic defeat has a remarkably internalistic character. A normative defeater functions "by virtue of being [a proposition] that S *should believe* given the evidence which is available to S."[12] Normative defeaters need not be believed in order to defeat the justification for a particular belief. But a normative defeater must not be totally disconnected from the subject's doxastic system either. While not as internalistic as doxastic defeat, normative defeat is still subject to concerns about what is epistemically available to a person given her evidence. Finally, a factual defeater functions "by virtue of being *true*".[13] This is an externalistic conception of defeat. A belief B is defeated by a factual defeater D only if D satisfies the following: a) D is true, and b) if D were added to the set of one's beliefs, then the resulting set would no longer justify B. While factual defeaters must be true, they need not be believed nor need they be within the easy epistemic grasp of a subject given the subject's evidence.

While Lackey does have things to say about both normative and factual defeaters, doxastic defeaters and their failure to always transfer via testimony are at the core of her objections to (2). But it is precisely this narrow focus on doxastic defeaters which perpetuates the mistake in Lackey's argument. For although in Lackey's example there

[11] ibid, pg. 474
[12] ibid, pg 475
[13] ibid, pg 475

is, arguably, no transfer of the doxastic defeater from Mrs. Smith to her students, there is a factual defeater which complicates the acquisition of knowledge and in fact makes it impossible for the students to know that evolutionary theory is true if their sole grounds for that belief is Mrs. Smith's testimony. Thankfully, Lackey does some of the work laying the foundations for this line of objection to her view, but she fails to carry the reasoning all the way through. That problem will be corrected below.

If one is to deny that Mrs. Smith's testimony caused the children to know that evolutionary theory is true (via that testimony), then, as Lackey points out[14], one should object to some or all of the claims that (A) Mrs. Smith's testimony is the source of the children's knowledge, (B) the children come to have knowledge rather than merely justified true beliefs, and (C) the way the children come to know in this case is via testimony. Lackey provides reasons to support all three claims. Unfortunately, her defense of claim (B) is inadequate.

In defending (B) Lackey actually provides the grounds for rejecting it. She initiates the discussion by remarking that to deny that the children come to *know* in the Mrs. Smith example, rather than merely having a justified true belief, to her "seems quite implausible"[15]. But, to her credit, she provides various reasons why one might take that line despite its (apparently) initially unpalatable veneer. Regarding rejections of B she says[16]:

[14] ibid, pg. 478
[15] ibid. pg. 480
[16] ibid. pg. 480

> "…[P]erhaps one way to support it is to argue that testimony cannot be a source of knowledge if the speaker in question is lying or falsely testifying. For if one reports that P but does not believe that P, then, it might be argued, it is merely an accident if a hearer comes to believe truly that P on this basis. And since it is widely assumed that there must be some non-accidental connection between a subject's belief that P and the fact that P, in order to rule out Gettier-type cases, one cannot know that P via a speaker's false testimony. Given this, one might argue that, broadly speaking, Mrs. Smith is lying to her students, since she is reporting what she herself does not believe, and hence even if her students acquire true beliefs via her testimony, this is merely an accident."[17]

This is, roughly, the right objection to Lackey's proposed counterexample to (2). But the issue is unnecessarily obscured by bringing to bear the concept of accidentality. Non-accidental belief is surely a desideratum of an adequate account of testimonial knowledge, but there are other virtues which a belief must have in order to be properly categorized as knowledge. In particular, according to Lackey herself, there must not be any undefeated defeaters for the hearer for the belief so formed. By fixating on the reliability of Mrs. Smith's statements as an answer to any concerns about accidentality, Lackey misses the more important question of whether the children have any defeaters *that result from the testimony being given*. More on this in a moment.

Lackey defends B by arguing that in the example Mrs. Smith used reliable processes to formulate her articulations in class. That is, she retrieved the information from books which are considered authoritative by the relevant scientific community. In this way, she secures the non-accidental connection between Mrs. Smith's statements about evolution and the facts about evolution. Since the reliability of Mrs. Smith as an information conferring agent is unthreatened, Lackey concludes that there is no intrinsic accidentality in this case. If the lingering concerns regarding (A) and (C) are discharged, then there remain no interesting grounds for resisting Lackey's central point, viz. it is possible for a

[17] In passing we may note that "one cannot know that P via a speaker's false testimony" has a reading under which it is trivially true. To interpret Lackey charitably, it is best to treat "false testimony" as elliptical for "testimony which the testifier does not herself believe."

hearer to know that p via testimony from a speaker even if the speaker does not know that p.

The trouble with that line of thought is that it places a great deal of theoretical weight on the notion of reliability. That in itself is not an unacceptable move. Reliabilism has a distinguished history in the epistemological literature. But Lackey has already located the problem with (2) in the possible failure of defeaters to transfer from speaker to hearer. She is fence-straddling a pair of theoretical backgrounds, reliabilism and defeasibility theory, which have traditionally been at odds with one another over the correct analysis of knowledge. Again, that in itself is not an unacceptable move. Acquiring a true belief in a non-accidental way is a plausible condition for having knowledge. But if non-accidentality were the only condition necessary for turning true belief into knowledge, then there would be no need to discuss the failure of defeaters to transfer since whether a particular belief is defeated by a doxastic, normative, or factual defeater would be entirely irrelevant to whether a hearer can know that p on the basis of a speaker's testimony that p.

For the reliabilist, what is importantly accidental is the link between the truth of the proposition and the belief about it. If the relationship between the truth of p and the belief that p is perniciously accidental, then knowledge that p cannot result. The reliabilist thus demands that such connections be secured by reliable processes linking the belief to the truth. With certain modifications, this is the sense of accidentality that Lackey is exploiting. She argues that the relationship between the statements proffered

by Mrs. Smith and their truth is non-accidental. The non-accidentality in question is clearly a matter of what process was used to transmit the claims about evolution. Lackey puts it thusly, "What is crucial for hearers in order to acquire knowledge from speakers is that the proffered *statements* be reliable, that is, that they be truth-conducive in some way." [18]

Contrast that perspective with a defeasibility theorist who identifies a crucial element for knowledge in the relationship between justification and truth. Consider a set E of all relevant evidence both for and against a proposition p. Of the members of that set, a certain portion will also belong to the belief set B of a given individual S. If an individual S has enough of the right members in B, S may qualify as knowing that p on the basis of that evidence. If there is some member d of E such that d conjoined with the members of B would not justify p, then, in a certain sense, it is accidental that S has a justified true belief that p. What is important is that S could have had a different portion of E in B and then S would not have been justified in believing that p. The existence of such a d, a defeater, is a prima facie impediment to knowledge. But the sense of 'defeat' used here is not the doxastic defeat which Lackey emphasizes, but the broader notion of factual defeat. Lackey has not excluded factual defeat from the discussion and nothing in her paper reveals any inclination to do so. While Lackey has shown that Mrs. Smith's statements are non-accidental along one axis of evaluation, she has yet to show that Mrs. Smith's statements are non-accidental *simpliciter*.

[18] Ibid. pg. 481, Lackey's emphasis

What, then, is this factual defeater that overrides the students' attempts to learn about evolution from Mrs. Smith? A number of candidates suggest themselves depending on how one fleshes out the Mrs. Smith example. Suppose that Mrs. Smith has a positive belief that she knows that evolution is a false theory. Suppose that she believes it is false on the basis of sources she takes to be authoritative, but which in fact fall short of that mark. What will Mrs. Smith's assertions in class feel like to Mrs. Smith when she says them? If Mrs. Smith has a modicum of moral sensibility, they should feel like lies. Given her creationist loyalties, Mrs. Smith might in fact feel some genuine guilt over her decision to put school policy ahead of her obligation not to bear false witness. That Mrs. Smith believes that she lied is a factual defeater for the students' beliefs that evolutionary theory is true. If one conjoins "Mrs. Smith believes that she was lying when she taught evolution" with what the students knew at the end of the session on evolution, then the resulting conjunction does not justify the students in believing that evolution is true. Mrs. Smith's testimonial act has given rise to a factual defeater for her audience.

In response to this objection one might reply on Lackey's behalf that the cases can be refined in such a way as to block to proposed factual defeater. Perhaps Mrs. Smith has no strong convictions on this matter. Perhaps, despite being a creationist, she is open-minded enough to accept that opinions which oppose the scientific majority opinion are likely to be problematic. In such a case, the proposed defeater, that Mrs. Smith takes herself to be lying, disappears because Mrs. Smith does not really believe she knows whether creationism is true. She definitely doesn't believe that it is not true. But that's not enough to make her statements count as lies.

The retreat will only take one so far. For even if there is no defeater having to do with lies, there will still always be some additional defeater which arises in concert with cases of testimony in which the testifier doesn't know. The counterexamples that Lackey proposes against (2) all rely on the speaker not knowing the content of her assertions when she makes them. Yet it is that very fact, that the speaker does not know that p when she testified that p, which serves as the factual defeater in all cases.[19] It is a defeater which arises as a consequence of the testimonial act. To put the point in Lackey's terms, defeaters may be acquired by way of testimony not merely through transmission from the speaker, but by being generated by the testimonial act.

Suppose that Mrs. Smith's students were to learn, immediately after class, that Mrs. Smith didn't know that evolutionary theory is true. The rational thing to do in that case is for them to withhold their assent from what Mrs. Smith said. Other things being equal, knowing that a testifier does not know what she testified to undermines the justification one has for the claims. It is because of this fact that the proposed defeater has its defeating effect. The account, it should be emphasized, is not counterfactual in spirit. One does not consider the subjunctive conditional "If the students had believed that Mrs. Smith didn't know what she said, then they would not have known that evolutionary theory is correct." Instead, the view takes seriously the notion of propositional justification, a relation which holds between sets of propositions irrespective of whether they are believed. It then takes account of different sets of propositions and evaluates

[19] On page 165 of his book Certainty: A Refutation of Skepticism, University of Minnesota Press 1981, Peter Klein has a paragraph characterizing an argument similar to this.

whether the sets justify some particular belief. So in the present example, the defeater works because the set of propositions which is the students' evidence at the end of the class justifies their belief that evolutionary theory is true, while the same set with the addition of the proposed defeater fails to justify the belief that evolutionary theory is true.

Of course, even defeaters are defeasible. There may be a defeater of the proposed defeater, a "neutralizer", "restorer", or "defeater eater" depending on one's favored terminology, which can undermine the defeating effect of the initial defeater that Mrs. Smith does not know what she said. The students might learn that the reason that Mrs. Smith doesn't know that evolutionary theory is true is that she doesn't believe that evolutionary theory is true, but, nevertheless, Mrs. Smith is an accurate reporter of the information from scientists who do know that evolutionary theory is true. Given this new information, the students can have the knowledge that evolutionary theory is true. But this is now no longer a case of simple testimony, and indeed this new information is itself sufficient for the students to know that evolutionary theory is true.

If Mrs. Smith had initially reported to the students that although she herself does not believe in evolution, respected authorities do, then the students can learn that evolutionary theory is true because now their beliefs are based on a belief in Mrs. Smith's ability to accurately convey information from respected authorities. But this is not the sort of case that Lackey wants. What she needs are examples in which the students come to learn that p solely by testimony from someone who does not know that p.

The focus of the discussion now shifts to questions about whether a proposed defeater is genuine or misleading.[20] Genuine defeaters are not defeated by other defeaters, while misleading defeaters are defeated by other defeaters. If a defeater is misleading, then several conditions apply to it. They include conditions which apply to defeaters generally, such as being true and being such that their conjunction with the original justification fails to justify the proposition in question. Additionally, misleading defeaters have their defeating effect solely in virtue of rendering plausible[21] some false proposition. That is, had the defeater not rendered plausible the false proposition, then it would not defeat the original justification at all. Only genuine defeaters are problematic for one's justification. The fact that Mrs. Smith does not know that evolutionary theory is true is not a misleading defeater. There is no proposition f which it justifies that is both false and such that the defeating effect of this defeater is parasitic on rendering f plausible.

One might suggest, in opposition to this last claim, that there is in fact a false proposition which is rendered plausible by the proposed defeater that Mrs. Smith doesn't know that evolutionary theory is true. It is f*:

f*: Evolutionary theory is false.

[20] See Peter Klein, 'Misleading Evidence and the Restoration of Justification', *Philosophical Studies*, 37 (1980), pp. 81-89 and "Misleading 'Misleading Defeaters'" *Journal of Philosophy*, 75 (1979), pp. 382-386, both of which Lackey cites (pg. 475) as good examples of a defeasibility theory in action.
[21] Roughly, for a proposition to be rendered plausible is for its truth to be more likely than not.

If the proposed defeater did render f* plausible and was such that it defeated the justification the students had for their beliefs that evolutionary theory is true solely in virtue of the fact that it rendered f* plausible, then indeed the proposed defeater would be misleading and Lackey's counterexamples to (2) would be safe. But neither conjunct holds. Since the latter conjunct will not hold if the former does not, it is sufficient to demonstrate that the proposed defeater does not render f* plausible. To see that it does not, one must recognize that Mrs. Smith's failure to know that evolutionary theory is true could be explained, and entailed, by any of the following: a) It is false that evolutionary theory is true, b) Mrs. Smith does not believe that evolutionary theory is true, c) Mrs. Smith lacks adequate justification for the belief that evolutionary theory is true, or d) Mrs. Smith is subject to some version of a Gettier problem. Of these, f* is equivalent to option a. But evidence which is sufficient to render a disjunction plausible need not render each disjunct plausible. As a limiting case of this, consider a set of evidence in favor of some proposition p. That evidence will be equally compelling in favor of the disjunction (p v (a & ~a)). But that surely does not imply that the evidence in question renders the contradictory right disjunct plausible. Since there is no special reason to think that Mrs. Smith's failure to know is a consequence of f* rather than anything else, there are no grounds for thinking that the proposed defeater (that Mrs. Smith does not know what she said) renders f* plausible.

It is further of no use to volunteer (R) as a potential restorer:

(R): The only reason that Mrs. Smith does not know that evolutionary theory is true is that she does not believe that evolutionary theory is true.

For (R) does not neutralize the defeating effect of the original defeater. Instead, it provides an entirely new justification for the students' belief that evolutionary theory is true, since (R) entails that evolutionary theory is true. The words 'the only' are doing a great deal of work. By invoking the definite description, (R) implies that there is no other component of the knowledge relation which is missing other than the belief. But an uncontroversial element of knowledge that p is that p is true. When the issue is whether one's justification for a belief that p is defeated, it is illicit to invoke either p or anything that entails p as a defeater defeater. The point cannot be overemphasized. If such propositions were allowed to serve as neutralizers for defeaters, it would undermine all of defeasibility theory. All justified true beliefs would turn out to be undefeated, which is an unacceptable consequence for anyone who thinks that defeasibility theory is even remotely on the right track. To see this, consider a set of evidence that E that justifies ones' belief that p. Now consider a defeater d which, when coupled with E, does not justify one in believing that p. If we are now allowed to undercut the defeating effect of d by adding the true proposition p (or anything that entails p) then the resulting set does justify one's belief that p, since the probability of p conditional on p is one. Since Lackey has already endorsed the view that the presence of defeaters can undermine the justification one has for one's beliefs, she is similarly committed to rejecting (R) as a genuine restorer.

Notice how this systematic strategy of testimonial defeat applies naturally to Lackey's other main counterexample to (2). Again, to preserve the integrity and initial plausibility of the case, it is here presented in full:

> "Jane is *currently* in the grips of skeptical worries which are so strong that she can barely be said to know anything at all. (I here emphasize 'currently' to capture the idea in contextualist views of knowledge that skeptical doubts may undermine knowledge while those doubts are being entertained, even if they need not undermine knowledge in ordinary or everyday contexts.) That is, her belief that she could now be the victim of an evil demon is strong enough to defeat the justification she has for many of her ordinary beliefs and, moreover, it is currently an undefeated defeater. Jim, a passerby, approaches her, asks her where the café is, and she reports that it is around the corner, but does not report her skeptical worries to Jim. Now Jim has never considered any skeptical possibilities at all, and hence he does not have any doxastic defeaters for his ordinary beliefs. Furthermore, he does have positive reasons for accepting Jane's report, e.g., he has perceived a general conformity between facts and the reports of many speakers in these types of contexts, and he has inductively inferred that speakers are generally reliable when they are giving directions, and Jane does not exhibit any behavior which indicates a lack of sincerity or competence with respect to her report. So Jim forms the true belief that there is a café around the corner on the basis of Jane's testimony…Given that Jane has an undefeated defeater which Jim does not have, he has knowledge which she lacks. Yet at the same time it seems possible for Jim to come to know that the café is around the corner via Jane's testimony even though her skeptical doubts currently undermine her knowing this."[22]

Lackey's technique is by now clear. Identify the defeater which is blocking the speaker from knowing, and then show how that defeater fails to transfer in a case of testimony. Since it is beyond serious doubt that a hearer can acquire a true belief on the basis of testimony, even if the speaker does not know whether the testimony is true, Jim is in a reasonably strong epistemic position with respect to Jane's claims. What remains is to demonstrate a lack of accidentality and then Jim can be said to know. But the method which Lackey uses to adjudicate cases of accidentality is by way of reliability. However, there is an unresolved issue of accidentality revolving around the testimonial defeater (that Jane doesn't know where the café is). So even though Jane's doxastic defeater (her skeptical worry) doesn't transfer to Jim when she tells Jim where the café is, and even

[22] Lackey, 1999: pp. 484

though Jim has a true belief as a consequence of Jane's testimony, Jim's newly formed belief is still susceptible to epistemic undermining by the factual defeater that Jane doesn't know where the café is.

Let us compare the case to a slightly different one. Consider the situation in which Jane knows she is in a particular kind of skeptical world. Suppose, somehow, that she finds out that since thirty minutes ago there really are malevolent demons in the area who delight in tricking people into believing that cafés are in one place when they are actually in another, though they do not always exert their influence. Her epistemic position in this case is similar to her epistemic position in the original case. In both she does not know where the café is and she has a belief to that effect. In both she believes that the café is around the corner, but she does not have adequate justification for her belief. Jim asks Jane where the café is and she reports that it is around the corner without reporting her knowledge of the demons' activities. Jim has all the same reasons for accepting Jane's statement, "e.g., he has perceived a general conformity between facts and the reports of many speakers in these types of contexts, and he has inductively inferred that speakers are generally reliable when they are giving directions, and Jane does not exhibit any behavior which indicates a lack of sincerity or competence with respect to her report." If the demons happen not to be engaging in any trickery with respect to the location of the café which Jane has in mind, then Jim will form a true belief about the location of the café. Since people in Jim's experience are generally trustworthy with respect to these matters, he has a justified true belief about the location of the café. But his belief can

hardly be called knowledge. This is just a version of the standard phony barn cases[23] which are generally uncontroversially regarded as undermining knowledge.[24] Since Jane is in the grip of skeptical worries in the original case, the basis for her assertions should seem to her similar to the present case. From Jim's perspective, nothing is different in the two cases. By analogy, Jim does not know where the café is in the original case.

Information Transfer Response, an Alternative to Lackey:

Intuitions regarding what is required for knowledge to result from testimony can vary widely. One might grant that the critique of Lackey proposed here is a good one. Given that Lackey embraces a defeasibility condition on knowledge, and given that the operative defeater (that the speaker does not know what she said) does indeed undermine the justification that the hearer has for the belief so acquired, it follows that Lackey's criticism of (2) is nullified. But so what? With some modification, the spirit of Lackey's position can be recast in a model that avoids the complicating defeater and reinforces the intuition that (2) is false.

According to this new view, we should understand testimony in terms of information transferal.[25] Information is not a set of beliefs since the content of beliefs must be grasped by the believer whereas the information conveyed by a speaker need not be

[23] Alvin Goldman, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* 73 (1976)
[24] For an interesting discussion of the instability of intuitions regarding phony barn cases, see Tamar Szabo Gendler and John Hawthorne, 'The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies' *Philosophical Studies*, 124 (June 2005).
[25] For articulate opposition to this view, see Lehrer, K. 'A Critique of Externalism' from *Theory of Knowledge* 1990 reprinted in *The Theory of Knowledge*, ed. Louis Pojman, Wadworth Publishing Company 1999.

grasped by the speaker. Information has propositional content since it can be either correct or incorrect information based on whether the contained propositions are true or false. It can be carried by a speaker either consciously, as when a speaker either believes or knows the information spoken to be true, or it can be carried without conscious involvement, as when a person verbalizes a message the contents of which are beyond her lexical understanding. But information is not like a justified belief in that it can disappear in the face of defeaters. Information is impervious to defeaters; it is simply either correct or incorrect.

If we reframe the Mrs. Smith case in light of these considerations, we reach the following explanation. Various scientists spent the better part of their lives devoted to understanding the biological developmental history of the varieties of life on this planet. They codified their research into various texts from which others might come to learn about evolutionary theory. These texts, *ex hypothesi*, contain the correct information about evolutionary history. This information has an impeccable etiology that ensures no epistemic worries. If one reads these books and articles, one will be in a position to know that evolutionary theory is true. Mrs. Smith is a diligent and competent teacher. She knows how to investigate the subjects she is required to teach. When she reads the reports of the scientists, despite her failure to believe the contents, she is absorbing the information contained therein. When she speaks to her students about evolutionary theory, she rehearses the scientific information in her own words, but with an conscious intention to be true to the authors' purpose. When the students walk away from class with a firm belief that evolutionary theory is true, that belief is based on a reliable source

conveying correct information garnered at each stage in the most epistemically robust ways. To withhold the label "knowledge" from what the students have is to be blinded by fancy theoretical motivations. The students have a justified true belief achieved by the transfer of correct information through a reliable process. When they tell their friends about what they learned in class, their friends can come to know that evolutionary theory is true too. Correct information is doing all the interesting work here; knowledge is just correct information acquired in the right ways.

That picture of testimony stands opposed to the one endorsed here. On the present view, knowledge and information come apart. Although one cannot transfer knowledge by testimony without also transferring correct information, it is possible, and perhaps all too common, to transfer correct information without transferring knowledge. This is a trivial point though, since idle and arbitrary gossip manifests this feature. But if we add that the correct information must be derived in a reliable way, then the present claim is still that, in testimony, properly derived correct information can come apart from knowledge. An analogy will help display the reasons why.

Jerry Fodor once commented that God can't make a dime. His point was not about God's omnipotence nor about God's fiscal prospects. Dimes are artifacts which are partly individuated in terms of their causal history. The United States Mint is the only body authorized to stamp dimes. The authorization they carry is legal, but it carries a deeper metaphysical import. If an object was not created by the U.S. Mint, then that object is not a dime. Fake dimes are not dimes. The fact that fake dimes look like real

dimes is part of what justifies the budget for Secret Service investigations of counterfeiting rings. While dimes may not in fact be worth counterfeiting, someone might take it upon herself to refine the skill of duplicating the perceptible qualities of a dime. But no matter how good these creations matched a dime in color, shape, texture, and other important features, the fact that they were manufactured outside the U.S. Mint means that they are not dimes and the duplicator has committed a crime. Presumably God could make an object which even the experts at the U.S. Mint would be unable to differentiate from one of their own. This object could enter circulation, be passed from person to person to satisfy various transactions, and finally wend its way into someone's coin collection, all the while being treated like a real dime and functioning in every relevant respect like a real dime. But it would still be counterfeit. The numismatic, metallurgical, and artistic talents of God notwithstanding, any object not created by the U.S. Mint is not a dime.

Knowledge is like information currency. One of the purposes of testimony is to cause others to believe truths. True beliefs help people navigate the world. Without true beliefs, it would be quite difficult to satisfy most desires. Our species is and has been in the business of documenting and verbalizing true beliefs as a way to endure and improve over time. Sometimes what we thought was a true belief is learned not to be true, and then we do our best to update in the light of the discovery. Knowledge entails true belief. But beyond that, it entails at least two further elements. First, a connection between the belief and the truth which is not accidental (as the reliabilists remind us). Second, the justifying reasons for the belief must not be overridden by complicating facts relevant to

the content of the belief (as the defeasibility theorists remind us). These two extra elements constrain the permissible testimonial interactions. Now there may be further elements which we would want any proper testimony to have. The simplest, most convenient, and most explanatorily satisfactory way to capture all of the desiderata of ideal testimony is that the content of the testimony should be what the testifier knows. Testifying to what one does not know is epistemically criminal. It is like passing along a counterfeit note to an unsuspecting cashier. If the note is of high enough quality, no one will be the wiser. The note will go on circulating, just as a real note would, all while satisfying the right monetary functions. Similarly, one can report true information, and that information, if properly structured, will perform all the same functions as knowledge.[26] Yet we value the strong tie between belief and truth that knowledge provides. When one asserts what one does not know, something has gone wrong with the testimonial process.

Two Further Rebuttal Attempts

1) A defender of Lackey's position might respond that the proposed factual defeater, that the testifier does not know what she said, is illegitimate because it deploys an epistemic term. But this is too strong a restriction for defeaters. Suppose that Jones tells Smith that all of Smith's beliefs are false and suppose further that Smith has good general

[26] This view has ancient origins, "…true opinion is as good a guide as knowledge for the purpose of acting rightly. …[T]rue opinions are a fine thing and do all sorts of good so long as they stay in their place, but they will not stay long. They run away from a man's mind; so they are not worth much until you tether them by working out the reason. … Once they are tied down, they become knowledge, and are stable. That is why knowledge is something more valuable than right opinion. What distinguishes one from the other is the tether." Plato. "Meno." *The Collected Dialogues of Plato*. Hamilton, Edith and Huntington Cairns, ed. Princeton U: Princeton. 1961. 381-2

reasons to trust Jones and as a consequence unreflectively accepts what Jones said. If Smith believes that water freezes at zero degrees centigrade at sea-level, and then Smith conducts several controlled experiments which all point to the conclusion that water does freeze at zero degrees centigrade at sea-level, then Smith's belief should serve as a defeater for the belief she initially acquired from Jones. To block this on the grounds that epistemic concepts cannot be deployed in defeaters is unmotivated.

Similar considerations apply to restricting defeaters by not allowing the concept of knowledge to play any role in them. The idea is that while epistemic concepts in general can be parts of defeaters, knowledge per se cannot. This new proposed restriction has the virtue of avoiding the unintuitive consequences discussed in the previous paragraph while still blocking the principal objection to Lackey's putative counterexamples to (2). But even if the restriction only applied to cases of testimony it would still not be a good one since the existence of Gettier cases clearly impedes the transfer of knowledge from speaker to hearer. Lackey is in no position to dispute this. If S tells H that p, but S has a justified true belief that p which is not knowledge due to a Gettier case, then that fact is a factual defeater that undercuts H's justification for p. It is likely non-trivial to provide precise necessary and sufficient conditions for being a Gettier case, but a sensible starting point is that a case C is a Gettier case only if in C one has a justified true belief that p without knowing that p. Since the presence of Gettier-type complications for a speaker undermine a hearer's ability to come to know that p on the basis of that speaker's testimony that p, and since any plausible analysis of what a Gettier case is will invoke the

concept of knowledge, it follows that knowledge cannot be cordoned off from the set of legitimate and permissible constituents of defeaters.

2) Another way a defender of Lackey's position might respond is to argue that the proposed factual defeater, that the testifier does not know what she said, simply begs the question against Lackey's position. It would be dialectically inappropriate[27] to suggest that the lack of knowledge on the part of the speaker can serve as a defeater in this context because that is tantamount to the negation of Lackey's view. It might be true, but it is not a good argument against Lackey's position.

This is a delicate issue since it does seem that testimonial defeat entails that (2) is true and the conclusion of Lackey's argument is that (2) is false. But the fact that a criticism C entails the negation of the conclusion of an argument A does not imply that in arguing for C one *eo ipso* begs the question against a proponent of A. Usually what is required in a case of begging the question is that in arguing for C one uses as a premise something equivalent to the negation of the conclusion of A. But that is not taking place in the present criticism. Notice that (2) makes no mention of testimonial defeat. Additionally, we have seen that there are reasons to accept the existence of testimonial defeat which are independent of concerns about whether the general principle (2) is correct. More speculatively, we can observe that not all cases of begging the question are equally pernicious, so even if the charge holds, the effect may not be devastating to the criticism.

[27] For a discussion of begging the question and why it cannot be a formal matter, see Roy Sorensen 'Unbeggable Questions', *Analysis*, (Jan. 1996) and ''P, therefore P' without Circularity', *Journal of Philosophy*, (May 1991).

Consider someone who takes the view that there are no good reasons for anything.[28] Any attempt to provide good reasons to the contrary will inevitably succumb to a charge of begging the question. But the situation does not appear desperate for the defender of good reasons. For a defender of Lackey's view to deploy this objection to the present criticism requires strong evidence that in fact the criticism does beg the question and it does so to its detriment. It is not obvious that both tasks can be accomplished.

The most straightforward way to beg the question against Lackey is to use (2) itself as a premise in an argument against Lackey's proposed counterexamples. Since that is not taking place, we must look elsewhere for the premise that supposedly begs the question against Lackey's view. The natural place to look is to the following premise (D):

(D): If H hears from S that p while S does not know that p, and H comes to believe that p on the basis of S's testimony that p, then the justification that H has for p is defeated by the fact that S did not know that p when S said that p.

It has already been established that Lackey cannot reject (D) due to the presence of epistemic terms. More specifically, she cannot reject (D) due to the presence of the word 'know'. She cannot reject (D) on the grounds that it employs the notion of defeat, since her objection to (2) is predicated on the failure of defeaters to necessarily transfer from speakers to hearers. She further cannot object to (D) on the grounds that it employs the wrong notion of defeat. In her discussion of the distinction between doxastic, normative,

[28] This case is presented by Timothy Williamson, *Knowledge and Its Limits*, (Oxford University Press, 2000), pp. 188. Hereafter, KAIL.

and factual defeaters, she refers, approvingly, to a body of literature in which the notion of factual defeat is cashed out in terms of the effect of adding a proposed defeater to a set of evidence for p and observing whether the newly formed set in fact justifies p.[29] In a broad sense, the criticism of Lackey's proposed counterexamples to (2) is that it is a consequence of an inconsistent view. The explanation of why it seems that (D) begs the question against Lackey, if one is needed, is that (D) is the most important premise in an argument demonstrating the inconsistency of her view. Of course she may still wish to resist (D), but the inclination to resist a devastating premise is not evidence that the premise begs the question.

If successful, this strategy for handling the Lackey counterexamples to (2) has the virtue of systematically undermining them. All that Lackey really needs for her attack on (2) is for there to be one case which violates it. A refutation of such numerous counterexamples, especially considering the diversity in what features are exploited to deny knowledge to the speaker, is unlikely to succeed by engaging in a "one-off" strategy of handling counterexamples one at a time in different ways depending on the nature of the case. A better approach, like the one utilized here, is to draw on the background conditions common to all the proposed counterexamples, and use those to undermine whatever initial plausibility the counterexamples may have had. The fact that in each testimonial act in Lackey's proposed counterexamples there is a factual defeater based on the failure of the speaker to know suggests that (2) is a principle with quite a bit to recommend it. Minimally, Lackey's criticism of the requirement that a testifier must

[29] For example, Keith Lehrer and Thomas Paxson 'Knowledge: Undefeated Justified True Belief', *Journal of Philosophy*, 66 (1969) and Peter Klein 'A Proposed Definition of Propositional Knowledge', *Journal of Philosophy*, 68 (1971).

know that p in order for her audience to come to know that p via her testimony is abrogated.

Williamsonian Connections

Lackey's suggestion that an epistemically virtuous chain of testimony can exist despite the fact that various testifiers do not know what they assert has ramifications for Timothy Williamson's general epistemological view. Williamson argues for a reversal in the standard explanatory arrangements. Instead of trying to analyze the concept of knowledge in order to ferret out the various conceptual components (justification, belief, etc.), Williamson instead argues that these notions are better understood by treating knowledge as a primitive by which other epistemic notions are explained. The more problems and phenomena that succumb to a knowledge-based explanation, the better Williamson's theory fares.

One area where knowledge-based theorizing has a high level of explanatory value is the general practice of assertion. Assertion is for Williamson what testimony is for Lackey. If there is a genuine distinction to be made here, it appears that neither is making it. Williamson's view is that the practice of assertion has but one constitutive rule[30]:

(The Knowledge Rule) One must: assert p only if one knows p.

[30] KAIL, pg. 243

The rule is individuating for assertion, no socio-linguistic practice is more closely associated with knowing than assertion is. The rule also has the virtue of explaining many of the more familiar Gricean maxims, such as the maxim to make one's contribution to the conversation true.[31]

In order to capture the normative element inherent in (The Knowledge Rule), hereafter '(TKR)', Williamson limns a view in which evidential authority to assert that p is transferred via assertion from speakers to hearers. Asserting that p when one lacks the evidential authority to assert that p leaves one open to legitimate criticism. It is a kind of cheating which is prohibited by (TKR). One need not obey the rule in all cases of assertion, but tacit responsiveness to the normative force of the rule is necessary to engage in the practice of assertion. In an idealized case, an initial person comes to know that p some way other than by testimony, say, by perception. Then that person tells a second individual that p and on the basis of the first's testimony the second comes to know that p. The testimonial chain can proceed in this way indefinitely and, if things go right, at every step the hearer is in a position to come to know that p on the basis of the testimony of the speaker. The evidential authority conferred across this chain of testifiers is secure because the knowledge rule ensures that in cases of correct assertion, assertions which are in compliance with (TKR), the testifier knows what she is saying.

So at bottom Williamson's theory is in a position to endorse (2). In fact, it is a natural fit. If knowledge is to be the fundamental conceptual tool for epistemological theorizing, then it had better appear prominently in the rules for practices which are sensitive to

[31] ibid. pg 242

epistemological factors. For these reasons, it is inappropriate and even a little disappointing that Williamson twice explicitly concedes Lackey's position:

> "Someone who does not know p lacks the authority to assert p, no matter how plausibly he gives me the impression that he has done so. Although there are special cases in which someone comes to know p by hearing someone who does not know p assert p (Lackey 1999), the normal procedure by which the hearer comes to know p requires the speaker to know p too."[32]

> "In normal circumstances, when the hearer knows that the speaker asserted p, the speaker has no reputation for unreliability, and so on, a speaker who asserts p thereby puts a hearer in a position to know p if (and only if) the speaker knows p (see Lackey 1999 for some qualifications)."[33]

By conceding Lackey's general thesis, that knowledge can be acquired by way of testimony without the testifier knowing what she said, Williamson weakens the connection between knowing and asserting and thereby opens his theory to counterexamples on those grounds. The less explanatory generality knowledge-based theories have for actions and mental states, the less plausible is Williamson's inversion of the traditional epistemological approach to the explication of knowledge in terms of other concepts.

Viewing the matter with a charitable eye, one might account for Williamson's failure to counter Lackey's position by suggesting that perhaps the problem is not as invidious as it might seem. For even if there are special circumstances which falsify (2), there is still a wide enough range of cases which do fit the traditional, perhaps idealized account of testimonial chains. These chains need explaining and the knowledge-based account of assertion is in a position to explain them. If some cases of assertion happen to fall

[32] Ibid. pg. 257
[33] Ibid. pg. 267

outside the scope of the explanatory model, then that is a limitation, but not a damning one.

While that is a polite response on Williamson's behalf, it is not theoretically satisfactory. Williamson's arguments in his chapter on assertion often revolve around the failure of suggested alternative rules for assertion to capture enough cases to satisfy the background demand of maximal explanatory generality. For example, in comparing the following two rules:[34]

(The Truth Rule) One must: assert p only if p is true.
(The Warrant Rule) One must: assert p only if one has warrant to assert p.

Williamson notes that the latter rule is superior to the former insofar as the latter prohibits assertions based on lucky guesses while the former does not. The fact that (The Warrant Rule) covers more cases than (The Truth Rule) is reason to favor (The Warrant Rule). For comparable reasons, Williamson should be wary of allowing any exceptions to his final proposal for the constitutive rule of assertion (TKR). A theory with complete generality over its data set and offering no chance for exceptions in the future is more methodologically appealing than one which admits of exceptions. Just as Occam's Razor recommends parsimony, so too is full generality a methodological desideratum for any theory. Williamson should not concede the issue to Lackey.

[34] Ibid. pg. 242

We have seen that Williamson has good reason to reject Lackey's proposed counterexamples to (2), but we need to establish that his theory is strong enough to legitimize such a rejection. First, note that the discussion of testimonial defeat made no appeal to a Williamsonian theory. The criticism offered earlier against Lackey was based on features internal to defeasibility theory, a theory that Lackey, at least in some sense, endorses. Depending on how one characterizes the defeasibility theory in question, Williamson is actually in a position to endorse many of the arguments which deny knowledge to an individual on the basis of the existence of undefeated defeaters. Second, once the type of defeasibility itself is spelled out, there are ways of viewing defeasibility from inside a purely Williamsonian theory which support the criticism of Lackey's position that was offered above.

Since the testimonial defeat criticism against Lackey's view is not Williamsonian in origin, it is useful to see how it might be annexed into his broader theory. One way to see whether the criticism is compatible with Williamson's view is to examine the individual theoretical commitments inherent in the criticism and check whether they are commitments which are either shared by or consistent with Williamson's explicit position. The bulk of the work done by testimonial defeat relies on the falsity of knowledge ascriptions in cases where there is an undefeated defeater. Williamson's epistemology is certainly amenable to a principled distinction between cases of knowledge and ignorance, and nothing in Knowledge and Its Limits suggests that the presence or absence of defeaters is irrelevant to the truth of knowledge ascriptions.

Of course there are deep differences between Williamson and endorsers of defeasibility theory, especially with respect to what position any defeasibility thesis holds in a broader theoretical framework. For a typical defeasibility theorist, the defeasibility condition is a fourth condition of a justified, true belief analysis of knowledge. Its presence is demanded by the Gettier cases and the issues that that literature raised. Williamson reverses this structure and lets knowledge be a primitive, unanalyzable concept which serves as the primary implement of epistemological theorizing. But these major structural differences need not interfere with a common argumentative strategy. Each theory can make use of the existence of defeaters in order to substantiate the theory's judgments of knowledge or ignorance. For the traditional defeasibility theorist, a defeater complicates putative cases of knowledge by being inconsistent with the definition of knowledge. For a Williamsonian, the presence of a defeater complicates putative cases of knowledge because, like unreliability, the existence of an undefeated defeater is incompatible with knowledge (although this is not a definitional or analytic matter). There is no need in principle for Williamson to reject the conclusions that defeasibility theory reaches with respect to whether a given case is a genuine case of knowledge. Especially when there are grounds to deny that a given individual knows that p, when something necessary for knowledge is missing, whenever a defeasibility theorist would say that the individual does not know that p, a Williamsonian is in a position to agree.

The real elegance of this structure is apparent once it is coupled with the fact that Williamson is in a comparable relationship with reliabilism. When it comes to denying

that an individual knows that p, whenever a reliabilist would so deny, Williamson is also in a position to deny. Reliabilism and defeasibility theory are each manifestations of different, but equally relevant, intuitions about knowledge. The confusion that they share, according to Williamson, is their treatment of the conceptually analytic relationship between the concepts of reliability and defeasibility respectively and the concept of knowledge. Once we divest ourselves of the old justified, true belief model of knowledge, we are free to utilize the more general, less theory-laden intuitions and subsequent efforts of those who endeavored to find a fourth condition in an analysis of knowledge. Studies of epistemic defeat and reliability help us better understand knowledge, not by progressively defining it, but by revealing the reliance of our extensive epistemic vocabulary on our familiarity with knowledge.

This structure may not be apparent at first glance. Williamson conspicuously does not address defeasibility theory per se in Knowledge and Its Limits. He instead makes some general remarks about the efforts of post-Gettier epistemologists:

> "Other analyses are circular rather than false; if someone insists that knowledge *is* justified true belief on an understanding of 'justified' strong enough to exclude Gettier cases but weak enough to include everyday empirical knowledge, the problem is likely to be that no standard of justification is supplied independent of knowledge itself. This book makes no attempt to survey even the most salient analyses of knowledge proposed in recent decades and the counterexamples to which they succumb…"[35]

Given this, one might conclude that Williamson has a counterexample against defeasibility theory in mind, but notice that there is another route permitted by the quoted passage. If it can be demonstrated that defeasibility theory deploys the concept of

[35] Ibid, pg. 4

knowledge in some ineliminable way in its account of defeasibility, then Williamson's complaint will be that the theory is circular. Circularity need not be a vice in all cases. Dictionaries are notoriously circular. What is important is that the judgments, especially the negative judgments, of defeasibility theory and reliabilism are trustworthy because they each focus on an important aspect of knowledge. Circularity need not be a problem which affects such judgments.

Here is one place where it appears that Williamson can criticize defeasibility theory for being circular. However the defeasibility condition on knowledge is spelled out, there will always be room for restorers. A restorer is a defeater that defeats another defeater. A justified belief for which there is a defeater can remain justified if there is a restorer to override the original defeater. Of course the process can repeat; there can be defeaters for restorers and then restorers for those as well. But any adequate account of restorers will need to restrict them in various ways. For example, suppose that p is true, S believes that p, and S has good reason to believe that p. Suppose also that there is a defeater d which defeats S's justification for p. Since p is true, it looks like p itself is a perfect candidate restorer to overcome whatever the defeater is. For however pernicious the defeater might be, if one were to add the fact that p to one's evidence set, even given the defeater, one would be justified in believing that p. This is clearly an unwelcome result since if the fact itself is allowed to serve as a restorer for a justified but provisionally defeated belief, then there will never be any true, justified, and defeated beliefs. More complex restorers which entail the fact in question are similarly suspect as admissible. Since there is no general consensus about what the limitations are for restorers, there is the potential for an

irresolvable interplay between various candidate defeaters, restorers, defeaters for those restorers, and so on. Judging whether on balance the defeaters outweigh the restorers can test the limits of one's intuitions.[36] But if Williamson is right, then this difficulty is best explained as a struggle of intuitions over whether the subject knows. If the subject knows, then the restorers outweigh the defeaters. If the subject doesn't know, then the defeaters outweigh the restorers. The charge is that the defeasibility theorist cannot escape from utilizing the concept of knowledge in the explanation of knowledge. While this is not enough to establish that defeasibility is circular, it does validate the legitimacy of thinking about defeaters from within a Williamsonian framework. The circularity itself is not worrisome since it occurs at a deeper level in the defeasibility theory than the particular judgments that theory gives as to whether one knows that p in a particular case. For those latter judgments, a defender of Williamson can happily agree, especially when those judgments are negative.

There is still further textual evidence to support the conclusion that Williamson intends defeasibility theory and reliabilism to be subsumed by his epistemology:

> "The present account of knowing makes no use of such concepts as *justified*, *caused*, and *reliable*. Yet knowing seems to be highly sensitive to such factors over wide ranges of cases. Any adequate account of knowing should enable one to understand these connections. This challenge is not limited to the present account: standard accounts of knowing in terms of justification must enable one to understand its sensitivity to causal factors, and standard accounts of knowing in terms of causal factors must enable one to understand its sensitivity to justification…"[37]

[36] In 'Misleading Evidence and the Restoration of Justification' *Philosophical Studies* 1980 pg 88, Peter Klein says, "…conflicting intuitions concerning whether a defeater is genuine or misleading can account for the conflicting intuitions concerning the status of S'[s] knowledge."
[37] KAIL, pg. 41

Since an integration of sensitivity to reliability and defeasibility is required of any adequate account of knowledge, Williamson's motivation for bringing the best parts of reliabilism and defeasibility theory together is independent of his broader epistemological picture. Given that Williamson has explicitly treated the notion of reliability as necessary for knowledge, while at the same time denying that reliability is a conjunct in an analysis of knowledge, it is natural to extend the same treatment to defeasibility. Once this is accomplished, the Williamsonian is in a position to run the same criticism of Lackey's proposed counterexamples to (2) as given earlier.

# Chapter Five

## Knowledge, Legality, and Morality

Epistemology is a peculiar area of philosophy in the sense that its value as a pursuit is often challenged. Certain ways of framing epistemological questions, such as "what is the nature of justification?" or "what is the nature of evidence?", are less susceptible to this concern. But when an epistemologist is engaged in the historically important quest for an understanding of knowledge per se, she is more likely to meet resistance. According to these naysayers, the theory of knowledge appears to have less bearing on contemporary concerns than the ancillary studies of justification and evidence. If the theory of knowledge is to remain a proper object of study for epistemologists, then a defense against these criticisms is compulsory. It is therefore useful to investigate the ways in which the study of knowledge over and above the study of these subsidiary concepts should be of interest. This chapter will explore the relationship between the epistemic and moral components of certain morally significant events as well as some legally significant events. The suggestion will be that Timothy Williamson's view that knowledge is the norm of assertion can be adapted to capture the importance of knowledge in the evaluation of certain morally significant events. If the connection is plausible, then this constitutes an extension to the overall project of regarding knowledge as a central and ineliminable element of a wide range of philosophical theorizing.[1]

[1] The salient proponent of which is Timothy Williamson in his Knowledge and Its Limits Oxford University Press, 2000. (Hereafter, "KAIL") Williamson's own work on this front includes his views that knowledge is a central and ineliminable component of our best theories of assertion, human behavior, and evidence.

The envisioned opposition to this view regard knowledge as, at best, a useful shorthand for the more complicated cluster of properties that are, at bottom, responsible for any theoretical work that knowledge performs. These include philosophers who are convinced that some sort of probabilistic account of justified belief will yield a satisfying explanatory apparatus for intentional action. Quite a bit of explanatory power can be achieved by a theory which talks just about beliefs, desires, and the way the world is. Still more can be achieved if the belief is justified. The challenge is to demonstrate that knowledge is worth studying given how much explanatory work can be done by talking about truth, belief, and justification separately.

There are two paths one can take in responding to this challenge. The two paths are not mutually exclusive. The first is to regard knowledge as intrinsically valuable. According to this response, the study of knowledge is on a par with the study of the good in meta-ethics or the study of the beautiful in aesthetics. "Everyone by nature desires to know" says Aristotle. This suggests that as human beings our cognitive goal is to understand our world. Mere true belief is insufficient for this task; we might accidentally have the right beliefs and in such a circumstance we could easily have been mistaken. Real understanding of our world and our place in it requires a stronger kind of fit between our cognitive states and the world. When one knows, one has the appropriate kind of fit.

Some may regard any attempt to capture the intrinsic value of knowledge as hopelessly question-begging. The problem is not unique to epistemology; meta-ethics shares this difficulty. When Kant or Mill attempt to persuade us to adopt their favored

criterion for intrinsic moral value, they do so by pumping our intuitions by suggesting cases where their moral view seems to give the intuitively correct answer. But such cases are rarely conclusive, especially since there appears to be any number of counterexamples to the pure Kantian deontological theories and to a pure utilitarian theory. A defense of anything as the ultimate bearer of intrinsic value, whether it be in an epistemic, moral, or aesthetic context, is best accomplished by displaying, in a systematic way, how all or most of our ordinary intuitions about the subject are captured by the theory which holds that thus-and-so has intrinsic value, all else having derivative value insofar as it promotes the thus-and-so.

Whether a case for the intrinsic value of knowledge can succeed, there is a second way to justify the study of knowledge. According to this alternative route, a correct understanding of knowledge will assist us in theory construction in areas other than those traditionally and most intimately associated with epistemology. Williamson endorses this second line of thought. He points out that we need the concept of knowledge in our explanations of human action. While a certain degree of success can be achieved by deploying the putatively less contentious notions of justification and evidence, better explanations of human behavior, ones which cover a broader range of cases with the highest predictive success, will be characterized in terms of knowledge.

It is worth noting that there is a third, hybrid position to take on the question of whether knowledge is intrinsically or derivatively valuable. In pursuing the line that knowledge has derivative value, one looks for ways in which theories that invoke

knowledge best capture the phenomena to be explained. We want good explanations where we can have them, they are intrinsically valuable. Knowledge is derivatively valuable insofar as it helps us achieve such explanations. The greater the explanatory power, the wider the range of applicability, and the more cohesive the knowledge-theoretic theories are, the greater the derivative value of knowledge. But if the research project which Williamson began is continuously successful on this front, then that suggests that there is a deeper reason for its success. As more and more phenomena are subsumed under knowledge-theoretic accounts, a threshold is reached where it is more theoretically satisfying to hold that knowledge itself is the important concept, the one with intrinsic value. In articulating the ways in which theories benefit from couching their explanations in terms of what is known (or should be known), the intrinsic value of knowledge is displayed. It is precisely because it is intrinsically valuable that it plays the roles it does in such a diverse array of areas.

One of the uses of the concept of knowledge is to characterize probabilistic differences between the likelihood of success under differing epistemic conditions. Kate wants a beer. In scenario type one she knows that there is beer in her fridge. In scenario type two she has a true belief that there is beer in her refrigerator, though she does not know that there is beer in her refrigerator. Kate is more likely to be drinking a beer in the near future conditional on her being in the first type of scenario than she is conditional on her being in the second scenario. The reason is that knowledge provides one with a robust connection to the facts that cannot be mimicked by weaker epistemic relations. This connection is what improves the chances for success of an action conditional on

knowing the relevant facts. Mere true belief that is not knowledge is susceptible to epistemic undermining by future evidence in a way that knowledge is not. If on the way to the refrigerator Kate meets Morgan who claims (falsely) that there was no beer in the refrigerator when he last looked, Kate is more likely to give up her belief that there is beer in the refrigerator if she did not before know that there was beer in the refrigerator than if she did know that there was beer in the refrigerator. Belief that there is beer in the refrigerator will suffice in the short term. Explanations of immediate behavior which are characterized without reference to what is known but only to what is believed fare as well as (though no better than) the knowledge-theoretic competitor. But when we leave the synchronic case and examine the diachronic case it is clear that having knowledge of the relevant facts in many, many ordinary situations raises the likelihood of success of a given action in that context over and above what true belief can render.

One might accept the previous line of thought and yet conclude that knowledge is not necessary if we can provide a supplement to true belief. If we allow that the propositions are justified as well, then it seems that one acquires the right kind of connection with the environment without having to advert to knowledge and without any intrinsic loss in the likelihood of success of a given action. But this modification will not succeed. One can be justified in believing that p, where p is true, even if one does not know that p. If one's justification is not sufficient to know that p (given the truth and belief components), then there will be a weakness in the justification which future counterevidence to p can exploit. Call weaknesses of that sort 'J-imperfections'. If one encounters counterevidence to p that reveals a J-imperfection, then one is liable to give up on a belief

that is essential to maximizing one's success in a given activity. Similar counterevidence will not have such a destructive effect on genuine knowledge, since in that case there is no J-imperfection to exploit.

One might be uncomfortable with the way the explanations have been given because they appear to require that propositions can be true at one time and then false at another. This conflicts with certain views in the philosophy of language that require propositions to hold their truth values eternally. Each proposition is presumably indexed to a time for maximal specificity. For expository reasons, the explanations given here will flout that convention. However, the problem is not pernicious. A translation mechanism exists in the following form. Let p* be a proposition that is known at time t1. The general reliability of physical law allows us to reliably form beliefs about the present based on earlier true beliefs about the past.[2] Propositions like "Bob's keys are in his desk drawer at 6/26/2007 10:42:58 AM", if known to be true for certain kinds of reasons, will make it entirely reasonable to believe a sequence of corresponding propositions which differ only in their time index.[3] Let p1, p2, p3,…pn be the sequence of propositions beginning with p* that differ from p* only in their time index; the higher the number, the later the time. Different attitudes toward p* can then be evaluated with respect to how far into the sequence one will believe truly.

Knowledge, then, is well-correlated with successful action especially in the diachronic case. But if sustaining a true belief over time is what makes knowledge

[2] This assumes, not uncontroversially, that there is no serious problem with conventional induction.
[3] As a limiting case, consider that the keys cannot move faster than the speed of light, so they could not leave the desk in the next picosecond.

superior to mere justified true belief, then there are other states which appear to do that job at least as well as knowledge does. Consider the state of dogmatically believing truly that p. If one dogmatically believes truly that p, then 1) p is true, 2) one believes that p, and 3) one's belief is not susceptible to undermining by future evidence. In a contest with knowledge, dogmatic true belief can be just as conducive to success. In fact, a dogmatically held belief can, in principle, be even more robust to counterevidence than an item of knowledge. Since success is at least as highly correlated with dogmatically held true belief as it is with knowledge, why take knowledge to hold a privileged status?[4]

One answer is that knowledge has a kind of generality of application that is missing in dogmatically held true belief. The goal here is to develop a theory which accounts for the wide range of human behavior in a way that unifies it. That requires an explanation which can cover a large range of cases in the same way. Knowledge is better suited to that task than dogmatically held true belief. Not all cases of knowledge that p involve a dogmatically held true belief that p.[5] The cases where successful action is brought about by a dogmatic belief in some true proposition p are relatively few in number. Most of what people believe is more pliant than that. Certainly explanations of successful action which rely on the dogmatic belief will explain their cases well, but those explanations will not cover nearly as much territory as explanations which rely on what the agent knows.

---

[4] Williamson's discussion of this point is unusually terse. He says, "Of course, profoundly dogmatic beliefs which are impervious to future evidence and do not constitute knowledge may be even more likely to persist than beliefs that are rationally sensitive to future evidence and do constitute knowledge, but then the subject's cognitive faculties are not in good working order. Since the difference between your present knowledge and your present true belief matters for predicting your future beliefs, it matters for predicting your future actions, because they will depend on your future beliefs." KAIL pg, 79

[5] In fact, the combination may not be possible. Dogmatism is not usually a good epistemic reason to hold a belief.

We can extend Williamson's argument for the utility of knowledge in the explanation of human behavior to cases in ethics. In certain ethically significant situations, optimal explanations of the moral culpability or moral praiseworthiness of an agent are characterized in terms of knowledge. An example will illustrate this point. Bob is walking down the street. He sees his friend Larry lying on the ground face down in a puddle. No one else is in sight. Instead of helping Larry, Bob passes by en route to his favorite antique store. Larry subsequently drowns.

In this case, we regard Bob as morally culpable. His actions are incompatible with correct moral behavior. Why so? A number of possible explanations suggest themselves.

One explanation could be that Bob's actions caused Larry pain, or that Larry's ability to experience life to the fullest was halted because of Bob's negligence. Such an account is consequentialist in spirit since it does not advert to any intention that Bob had (which would suggest a more deontological approach) nor does it mention Bob's lack of strong moral character (which would suggest a virtue ethics approach).

There are major if not conclusive objections to consequentialism. One such objection is that consequentialism gives the wrong verdict in a range of morally significant cases. Consider the so-called "lollipops for lives" case. There is a machine that produces lollipops which are delicious and are so well made that they may be licked indefinitely

without diminishing their size or quality. But in order to produce the confections, the machine requires certain raw materials. A necessary ingredient is that an innocent, fully conscious, adult human being must be liquefied in an exceedingly painful way. Failure to produce high levels of pain or using an individual who is not innocent results in defective lollipops. But once a lollipop is made, it can be licked over and over by an indefinite number of people. The innocent person suffers for the creation of the lollipop. The people who taste the lollipop gain a finite, but measurable amount of pleasure. If enough people lick the lollipop, then eventually there will be more pleasure generated by the licks than there is suffering generated by liquefying the innocent person. Since the net result is a gain in pleasure, consequentialists who regard maximization of pleasure as the defining characteristic of correct moral behavior will be forced to say that using the machine is a morally correct action. Provisions will have to be made which restrict the effects of the machine in order to secure this result. In particular it may be stipulated that no one will know where the lollipops come from and that the innocent person will not be missed. The upshot of this objection is that purely consequentialist explanations of moral culpability are less than transparently plausible and may in fact be entirely wrong-headed.

This counterexample to pure consequentialism is just one of many which is opposed to any account that takes morality to consist in the maximization of a given property whether it be physical pleasure, happiness, beauty, fairness, or what have you. It works less well or perhaps not at all against those consequentialist theories which regard certain consequences as categorically prohibited and it also works less well or perhaps not at all against those consequentialist theories which are in some ways equivalent to an intention-

based theory in which the maximization of actions with favorable intentions is central. But it is the former type of view which is here being rejected. The reason that the pure consequentialist theory needs to be rejected is that it suggests that the mental state of an individual during a morally significant action has no bearing on the moral status of the action in question. Such a view is at odds with perfectly ordinary intuitions about morality and threatens the knowledge-based account of morality which is developed below. While there are plausible counterexamples to a purely intension-based account of morality, that theory is not as menacing to a knowledge-based account of morality since that view at least permits the mental states of the agent a role in the explanation of moral actions.

Here is a second explanation of why Bob is acting in a morally incorrect way. Since Bob knew that Larry was face down in a puddle, and since Bob knew that Larry cannot survive if left face down in a puddle, and since Bob knows that no one else is likely to be able to help Larry, and since Bob knows that it is wrong to let his friend drown, Bob's failure to help Larry constitutes a moral crime. This explanation works equally well for the deontologist and the virtue ethicist. While we need not endorse either of those moral views in their entirety, there certainly is some intuitive force to this type of explanation. This is the type of moral explanation one learns at Mother's knee.

The point is that the explanation makes heavy use of the concept of knowledge to explain what is wrong with Bob's behavior. Conversely, denials of knowledge serve to abrogate Bob's moral culpability. If Bob did not know that Larry was drowning, then he

cannot be held responsible for his failure to prevent it. If Bob did not know (per impossible) that Larry could not breathe if face down in a puddle, then Bob cannot be held accountable for Larry's death. Further examples of ignorance have similar mitigating results.

Let us now consider five similar scenarios which will illustrate the moral difference made by varying degrees of epistemic success.[6] The examples form a hierarchy. At the bottom level is false unjustified belief and at the top is knowledge. In the middle are various combinations of justification and truth. The format of the examples is as follows.

Henry is a troubled individual. Despite the fact that he has been treated with adequate care and concern by his father, he nevertheless harbors a distinct hatred for his parent. Henry returns home one night and seems to see through a window that his father's den is illuminated. (Whether he actually sees that the den is illuminated is a variable in the following cases.) For whatever reason, he decides that he must act on his hatred. He enters the home, retrieves a loaded and fully functional .357 magnum from his room, and heads toward the den. The den is small, containing little more than a desk and chair. The chair is normally in front of the desk and when it is in that position the chair has its back to the door. All of this is common to each of the examples.

Case One: Henry believes that his father is in the den. In fact, his father is not in the den. Henry does not have any particular reason for maintaining his belief. He knows that often enough his father will leave the light on in the den even when he is not present. But

[6] I am grateful to Dennis Whitcomb for an early discussion of these cases.

Henry is an unusual fellow and on the basis of little more than an inadequately justified belief he believes that his father is the den. Moreover, Henry knows that he doesn't have a good reason for believing his father to be in the den but stubbornly maintains his belief anyway. Henry approaches the door to the den and without opening it he fires several shots through the door. Since no one is inside, no one is injured.

Case Two: Henry believes his father is in the den. In fact, his father is not in the den. However, Henry does have a good reason for believing that his father is in the den. The reason is that in this case, unlike in Case One, his father has never in the past left the light on in his den when he was not inside. This constitutes good evidence that Henry's father is now inside the den. Moreover, Henry is aware that he has good evidence that his father is in the den and from Henry's perspective it feels for all the world like he knows that his father is in the den. Henry approaches the door to the den and without opening it he fires several shots through the door. Since no one is inside, no one is injured.

Case Three: Henry believes his father is in the den. And indeed, his father is in the den sitting in the chair at the desk. Henry does not have any particular reason for maintaining his belief. He knows that often enough his father will leave the light on in the den even when he is not present. But Henry is an unusual fellow and on the basis of little more than an inadequately justified belief he believes that his father is the den. Moreover, Henry knows that he doesn't have a good reason for believing his father to be in the den but stubbornly maintains his belief anyway. Henry approaches the door to the

den and without opening it he fires several shots through the door. Henry's father is hit three times in the back and subsequently dies from the injuries.

Case Four: Henry believes his father is in the den. And indeed, his father is in the den sitting in the chair at the desk. Henry has plenty of reason to believe that his father is in the den since at no time in the past has his father left the light on in the den without actually being inside. Henry knows that it is at least possible that his father is not inside, i.e. he knows that there is no law of nature that requires his father's presence in the den whenever the light is on. Still, his experience has been that when the light is on, his father is inside. As it happens, an elaborate ruse is being carried out. Local hooligans have replaced the den window with a window-façade that emits light which looks for all the world as though the den light is on. Henry is, of course, unaware of this fact. He therefore has a true and justified belief that his father is in the den. Henry approaches the door to the den and without opening it he fires several shots through the door. Henry's father is hit three times in the back and subsequently dies from the injuries.

Case Five: Henry believes his father is in the den. And indeed, his father is in the den sitting in the chair at the desk. Henry has plenty of reason to believe that his father is in the den since at no time in the past has his father left the light on in the den without actually being inside. Henry knows that it is at least possible that his father is not inside, i.e. he knows that there is no law of nature that requires his father's presence in the den whenever the light is on. Still, his experience has been that when the light is on, his father is inside. Henry approaches the door to the den. Suddenly he vaguely recollects a

philosophy course he attended where they talked about Gettier cases involving hooligans with light-emitting window-façades and decides that he had better be sure his father is actually present. Henry slowly opens the door and looks inside the den. He sees his father busily typing at the desk. Henry quickly fires several shots and hits his father three times in the back. His father subsequently dies from the injuries.

Each of these cases involves a morally relevant situation with a differing epistemic component. In Case One, Henry has an unjustified false belief. In Case Two, Henry has a justified false belief. In Case Three, Henry has an unjustified true belief. In Case Four, Henry has a true justified belief that for Gettier-type reasons fails to constitute knowledge. In Case Five, Henry has knowledge. We need not endorse any particular moral view in order to examine the moral differences in the cases. It is enough to draw on the most popular views to reach some plausible conclusions.

As a departure point, we can agree that in Case Five Henry is more morally culpable than he is in Case One. After all, in Case Five there is an actual murder. In Case Five, there is an awareness on the part of the murderer that a killing is about to take place, one for which he will be causally responsible. Only on the strictest intention-based accounts of morality does Henry's actions in Case One constitute at least as great a moral crime as his actions in Case Five, and then so only with debatable plausibility. For even if Henry has the killing intention in Case One, he lacks the requisite information to connect his killing intentions with an actual killing. True, had his father been in the den in Case One, then Henry would have killed him. But Henry did not have a good reason to believe that

his father was in the den. His actions are not even partially caused by his confidence in their likely success in killing his father. But in Case Five things are different. In Case Five we have a cold-blooded murder. The assailant confirmed the presence of his victim and assured his death in a careful, methodical way. In Case Five, Henry would not have fired if he had not seen his father, a fact which is not implied by Case One.

Once we accept that Henry is less morally culpable in Case One than he is in Case Five, it is a short step to recognize that no case is such that in it Henry has greater moral culpability than he has in Case Five. Any element which might add to Henry's moral culpability is present in Case Five if it is present in any of the other cases. The real issue now is to ascertain the remaining hierarchical relationships.

Between Case One and Case Two a reasonable position would be that in Case Two Henry is at least as morally culpable as he is in Case One. One does not mitigate moral responsibility by acting on reasonable evidence in such a way that if the action is successful, negative consequences will result. In fact, quite the opposite holds true. Because in Case Two Henry has information which enables him to understand the likely consequences of his actions, and since further he understands that the consequences would be morally wrong and takes the action anyway, Henry must be at least as guilty on an intention-based moral account in Case Two as he is in Case One. We can go further and conclude that there are good grounds for saying that Henry is *more* morally culpable in Case Two than he is in Case One, but that further conclusion is not required here. From a consequentialist standpoint, only the rule-driven varieties will regard Henry's

actions as more morally wrong in Case Two than they are in Case One. A strict consequentialist would presumably see no significant moral difference between Case One and Case Two, but that is just to concede the original point that in Case Two Henry is at least as morally culpable as he is in Case One.

The relevant difference between Case One and Case Three is that in Case Three Henry's father is in fact in the den. From a strictly consequentialist standpoint, the issue is already decided. Other things being equal, it is wrong to actually commit murder. Since no murder takes place in Case One, a consequentialist has no recourse to claiming that Henry's action is Case Three is significantly worse than it is in Case One. But if one is a strict intention-based moral theorist, the case is not so clear. In one sense, it is merely an accident that Henry's father happens to be in the den at the time of the shooting. Henry's intention to successfully kill his father did not actually play a role in his decision to fire the gun. After all, Henry knows that often enough his father is actually not in the den when the light is on. If we were to ask Henry in Case Three whether he knew that his father was in the den, he might very well say "no". Without a doubt it is probably less than morally righteous to fire a gun into a room when one has no good reason for thinking that no one is in there. But from a strictly intention-based account of morality the situations are symmetrical. The same intentions are operating as the cause of the action in each respective case. In Case One, no one gets killed. In Case Three, someone gets killed. The difference in consequences legitimizes the claim that in Case Three Henry is at least as morally culpable as he is in Case One.

Between Case Two and Case Three the issue is whether it is worse to have an unjustified true belief which results in a killing or a false justified belief which does not result in a killing. If the killing is all that matters for moral evaluation, then Case Three is worse than Case Two. But there is some cause for hesitation here. While in both cases Henry can be described as having a killing intention, of the two it is only in Case Two where that intention is connected with other beliefs in such a way as to improve the chances of success of the action. But in Case Three Henry does not believe that his reason for believing his father is in the den is a good one, so Henry is not maximizing his chances for success. In Case Two, Henry has evidence which he has good reason to believe will improve the chances of a successful killing. Acting on such evidence, other things being equal, counts as a moral wrong. It is debatable just how wrong such an action is. The point is that between Case Two and Case Three it is not altogether clear which is worse. Both are worse than Case One, and both are better than Case Five, but beyond that the situation is less than clear.

Case Four constitutes a bona fide Gettier case. Henry has a true belief which is grounded in good evidence. Perhaps it is the deviant causal chain from the evidence to the truth of what is believed which prevents this from being a case of knowledge. Perhaps it is the fact that there is a defeater, namely that there is a light-emitting window façade in place, to Henry's evidence which prevents this from being a case of knowledge. We can assess Henry's moral culpability in this case along the two axes of consequence and intention. Since there is an actual killing in Case Four, the pure consequentialist will regard Henry as morally blameworthy. Further, Henry has the murderous intentions

characteristic of a contemptible figure. True, the evidence on which his action is based is unreliable in the sense that Henry could easily have been wrong about the likelihood of his father's presence and therefore the likelihood of successfully killing his father by shooting through the door of the den. But the unreliability of Henry's evidence does not appear to militate against his blameworthiness in Case Four. On both a consequentialist and an intention-based account of morality, Henry has done something wrong. Moreover, there does not appear to be any difference between the magnitude of the moral crimes in Case Four and Case Five. Knowledge, then, appears to be no better than justified true belief when it comes to the assignment of blame in these morally relevant situations. But it is certainly no worse.

In light of these cases, we have revealed that one's epistemic status does matter for the evaluation of one's moral status in certain morally relevant situations. What remains to be seen is whether there are moral circumstances where even having justified true belief is not equivalent to having knowledge. A case is needed where knowledge of some facts would add something to the moral situation which would not be present even if mere justified true belief of the same facts was present. If such cases exist, then while justified true belief may account for the particular degree of moral blameworthiness or praiseworthiness in many situations, a full theory of moral culpability, one which can cover every case, would be incomplete without adverting to knowledge.

As a first step towards such cases, one can examine the United States' code of laws. US code section 152 is just one of many sections which has a knowledge component written into it explicitly. To wit[7]:

**Section 152. Concealment of assets; false oaths and claims; bribery**

> A person who -
>  (1) knowingly and fraudulently conceals from a custodian, trustee, marshal, or other officer of the court charged with the control or custody of property, or, in connection with a case under title 11, from creditors or the United States Trustee, any property belonging to the estate of a debtor;
>  (2) knowingly and fraudulently makes a false oath or account in or in relation to any case under title 11;
>  (3) knowingly and fraudulently makes a false declaration, certificate, verification, or statement under penalty of perjury as permitted under section 1746 of title 28, in or in relation to any case under title 11;

The section continues on in similar fashion, routinely making use of the phrase "knowingly and fraudulently". The language is not particular to bankruptcy law. Even the penalties sections of the legal code often require that the illegal activity took place while the perpetrator had knowledge of her actions and their effects. Here is an example from the penalties section of an agriculture law, Title 7, Chapter 35, subchapter 2, part D:

**Section 1379i. Penalties**

> (a) Forfeitures; amount; civil action
>  Any person who knowingly violates or attempts to violate or who knowingly participates or aids in the violation of any of the provisions of subsection (b) of section 1379d of this title shall forfeit to the United States a sum equal to two times the face value of the marketing certificates involved in such violation. Such forfeiture shall be recoverable in a civil action brought in the name of the United States.
> (b) Misdemeanors; punishment
>  Any person, except a producer in his capacity as a producer, who knowingly violates or attempts to violate or who knowingly

[7] http://caselaw.lp.findlaw.com/casecode/uscodes/18/parts/i/chapters/9/sections/section%5F152.html

> participates or aids in the violation of any provision of this part, or of any regulation, governing the acquisition, disposition, or handling of marketing certificates or who knowingly fails to make any report or keep any record as required by section 1379h of this title shall be deemed guilty of a misdemeanor and upon conviction thereof shall be subject to a fine of not more than $5,000 for each violation.

What these examples show is that the legal system, with whatever wisdom and prudence it has, regards knowledge as a relevant detail in determining whether a crime has been committed and whether a punishment should be levied. Without getting bogged down in the particulars of legal minutiae, in certain cases of securities fraud such as "insider trading" there are at least two components. The first component is that the individual traded securities in a certain way, and the second component is that the individual had knowledge of a certain sort. Any competent investor is well-advised to investigate the activities and policies of a company she is considering investing in. But if the investor has information which is unavailable to the general public, i.e. if the investor knows that company A is going to merge with company B, it is her knowledge of such facts which undermines the legality of her otherwise legal action. In short, anyone can buy and sell shares of any publicly traded company as he or she sees fit, except those who are in a special position with respect to internal information about the companies. That Jane, a private investor, bought 10,000 shares of company B prior to a merger of company B with the more powerful company A is not illegal. That she did so as a result of a protracted study of the likelihood of companies to merge with one another and the subsequent results is shrewd, but not illegal. She may have learned as a result of her studies that it was very likely that company A would merge with company B and that such a merger would benefit the shareholders of company B. But if her actions are a

result of viewing internal documents of company A which were not publicly available, whereby she came to know that company A would merge with company B, or if she was told by an executive of either company that the merger was about to take place, then there is an issue of legality. Her having a justified true belief (of the sort specified) that company A would merge with company B is insufficient to warrant an investigation into her actions. But if she has genuine knowledge of the future actions of the companies, then she is in a position to be indicted.

Here is one additional example from a fourth circuit appellate court decision[8]. The original case involved the defendant, Carr, who was found guilty of arson (in an apartment building) and causing the death of an occupant. The basis for Carr's appeal was twofold. One cause was that the prosecution neglected to indict him using the exact wording of the arson law. The appellate court upheld that portion of the district court's decision. The second cause regarded the sentence Carr received. The appellate court remanded the case to the district court for resentencing on the grounds that "we are not sure that the district court properly distinguished between recklessness and knowledge". There are federal sentencing guidelines which regulate how sentences are to be apportioned by the court. These guidelines draw a distinction between deaths caused by those who genuinely know the consequences of their actions and those who are merely reckless.

The appellate decision goes into some detail on this point. Specifically, the appellate court found that the district court either (1) equated reckless indifference with knowledge

[8] United States of America v. Adam Nicklous Carr, decided September 11th 2002

or (2) regarded the evidence that Carr acted recklessly as sufficient evidence to find that he also acted knowingly. If situation (1) obtained, then the appellate court concluded that the district court made a legal error. If situation (2) obtained, then the appellate court concluded that "the [district] court did not make a finding that Carr actually *knew* that people were inside the building at the time of the fire" (the emphasis on 'knew' is in the original opinion). Interestingly, what the district court did conclude was that "the only reasonable belief" was that there were people in the building at the time. The appellate court thus held that there is a distinction between having knowledge and having a belief which is the only reasonable one to hold.

It is worth reading the exact language the appellate court used:

> Application note 1 to § 2A1.1, in discussing departure, focuses on the defendant's state of mind: "If the defendant did not cause the death intentionally or knowingly," but caused it through "recklessness or negligence," "a downward departure may be warranted." *U.S.S.G. § 2A1.1*, cmt. n.1. The term "knowingly" thus indicates a more culpable state of mind than does the [**18] term "recklessness."

One issue to be resolved is how the terms 'recklessness' and 'knowingly' are defined for purposes of the sentencing guidelines. It appeals to the difference in the state of mind of the individual at the time of the commission of the crime. A definition of 'reckless' occurs in an application note[9] to the sentencing guidelines:

> "Reckless" refers to a situation in which the defendant was aware of the risk created by his conduct and the risk was of such a nature and degree that to disregard that risk constituted a gross deviation from the standard of care that a reasonable person would exercise in such a situation.

[9] § 2A1.4 (Involuntary Manslaughter)

As one can see, whether one acted recklessly depends on the degree of risk one is aware of at the time of the crime. The fourth circuit court drew the definition of 'knowingly' from a Supreme Court decision[10]. In that case, the Supreme Court ruled that an individual knowingly commits a crime "when he knows that the result is practically certain to follow from his conduct". This definition is unabashedly circular, but the Model Penal Code clarifies the issue[11]. An individual acts knowingly when he "is aware that it is practically certain that his conduct will cause such a result". Since a homicide committed with knowledge carries a stronger sentence than a homicide committed only with recklessness, courts are legally obligated to find for one or the other type. These definitions and clarifications constituted sufficient grounds in the eyes of the fourth circuit court to remand the case to the district court for resentencing.

The upshot of this last case is that the legal system distinguishes between degrees of awareness on the part of the perpetrator. Having a high degree of awareness, a well-justified and true belief if you will, of the likely consequences of one's actions does not in itself equate to having knowledge of the likely consequences of one's actions. If this legal distinction is a manifestation of our collective moral attitudes, then this constitutes just the sort of case we were interested in locating at the end of the discussion of Henry and his father.

The fact that the word 'knowledge' and its cognates occurs in the language of our legal system is a reason unto itself for studying knowledge. One cannot simply abandon

[10] *United States v. U.S. Gypsum Co., 438 U.S. 422, 445, 57 L. Ed. 2d 854, 98 S. Ct. 2864 (1978)*
[11] Model Penal Code § 2.02(2)(b)(ii)

the pursuit of an adequate theory of knowledge in favor of merely studying justification on pain of ignoring this salient aspect of our social lives.

But there is an additional conclusion one can draw from this that mimics certain comments that Williamson has made about knowledge and its relationship to assertion. According to Williamson, knowledge is the norm of assertion. He views the activity of making assertions as governed by rules that have normative, though not moral, force. To make an assertion that one does not know is to violate a standard of the social practice. Of course people will routinely violate this rule owing to the fact that we cannot always know that we know the things we assert. Nevertheless, the rule is in place and it acts as a norm against which we can judge and regulate our conversant activities. We can extend this rule to cover moral cases as well. Knowledge, for certain moral actions, is the norm of behavior. This means that there are governing principles for some activities that assign blame or praise in accordance with whether one knows various things. Recall the example of Bob viewing his friend Larry who is face-down in a puddle. The explanation which seems to accord most naturally with our intuitions is characterizable most broadly in terms of what Bob knows. It is his knowledge that Larry is likely to drown that morally obligates Bob to take action. The rule that Bob is obliged to follow is a norm of moral action. He may, on occasion violate this norm because, for example, he may not know that he knows that allowing his friend to drown is morally wrong. Nevertheless, the norm is present. Since he has the required knowledge about the likely consequences of neglecting his friend, he is thereby obligated to act to prevent harm.

Other moral cases serve to illustrate this point further. Consider this principle regarding how one may conduct a just war. A preemptive strike against one's enemy is warranted provided that one knows, among other things, that one is about to be attacked. Failure to know that one is about to be attacked potentially robs the preemptive attacker of her moral high-ground. There are, of course, moral rules which need not advert to knowledge in their characterizations. Perhaps one only needs a reasonably well-justified belief about the circumstances of one's neighbors in order to acquire an obligation to lend aid to those in need. But some moral rules, particularly where the potential consequences are substantial, demand a level of epistemic authority which can only be met by knowledge. In some circumstances one will falsely believe that one knows that a preemptive strike is about to take place but will believe so for good reasons. In such cases, we may regard the preemptive strike as justifiable and even morally permissible. But the norm will still be that one ought to have knowledge before one undertakes a preemptive strike. One's actions ought to be guided by these knowledge requirements.

Insofar as the norms for some moral actions mimic the structure of Williamson's norm of assertion, we have additional evidence in favor of Williamson's view about the essential importance of knowledge. For a perfect correlation to obtain between the moral case and the assertion case, it would be necessary to show that the relevant moral cases have the same structure with respect to the role of norms as the assertion case has with respect to the role of norms. If there is no relevant similarity in structure, then either there is not enough similarity between moral cases and assertion cases to unify the explanation of the importance of knowledge in both, or there is some other way of

understanding the role of knowledge in these moral cases which, while different from the assertion case, is similar enough to allow that Williamson's theory is extendable to the moral case if it fits the assertion case.

Williamson's theory of the role of knowledge in instances of assertion is easy to state. It is simply that knowledge is the norm of assertion. When one asserts that P, one is representing oneself as knowing that P. If S asserts that P without knowing that P, then that ought to prompt others to either correct S's assertion so that it is qualified by, for example, the prefix "I believe that", or at least it should demand further explanation. Claims of the form "P, but I don't know that P" are in a sense deviant. Of course there is no claim here to grammatical impropriety or outright contradiction. Instead, Williamson envisions assertion as an essentially social activity which is governed by various rules. Many are happy to agree that these rules include grammatical rules which should not be violated on pain of loss of comprehensibility. But beyond these grammatical rules are more pragmatic rules. Again, many are happy to concur that such pragmatic rules exist. The standard Gricean conversational maxims are widely regarded as responsible for the generation of additional levels of meaning beyond the mere semantic content of the given sentence. One way to understand Williamson's idea is that if assertion is to serve as a legitimate basis for acquiring knowledge, then the content of the assertions of a speaker ought to be known by that speaker.

The status of the rule that knowledge is the norm of assertion is that it is constitutive of the practice of assertion. One is properly engaging in the social practice of

conversation only if one knows the content of what one says. Many practices have constitutive rules. When one is playing chess, it is a constitutive rule that one may only move a bishop along the diagonal. Failure to adhere to this rule should enjoin others to correct one's behavior. Nevertheless, there is no claim that the violation of this rule implies that one is no longer playing chess. If that were so, then there would be no sense in correcting the individual. The norms being what they are, if one asserts that P without knowing that P, then one is subject to criticism on those grounds. One ought to say only what one knows. Since there are myriad ways to weaken one's claims with qualifications like "I believe that", "It is likely that", or "I heard from a reliable source that", it is incumbent on the speaker to do so if that is the limit of what the speaker actually knows to be the case.

By the luminosity argument[12], one is not always in a position to know that one knows the content of one's assertions. Therefore one is not always in a position to know whether one is obeying the norm of only asserting what one knows. This may be cause for pause. How can one be obligated to obey a norm which one is not always in a position to know whether one is violating it? One might think that such a condition on the correct practice of assertion would place undue obligations on conversational participants. But the fact remains that only very trivial conditions, such as those which are impossible, are such that if they obtain one is always in a position to know that they obtain. For all other conditions, there will be always be some element of risk, some degree to which the conditions' obtaining are not perfectly certain. Such is the predicament of agents with imperfect epistemic faculties. Nonetheless, standards for

[12] See chapter three for a discussion of luminosity.

correct action, whether in assertion, moral action, or simply obeying the rules of a game, are a routine part of our cognitive and social lives. To insist on perfect epistemic accessibility is, despite any initial opposition to the contrary, an implausible standard.

If moral action had constitutive rules in the way that assertion does, then an easy parallel between the two could be accomplished. The view would be that the central constitutive rule of moral action, at least with respect to a significant portion of moral cases, is to act only if one knows the likely consequences of one's action. Consider the following case. An aging skyscraper is to be demolished. The dynamite has been placed and the immediate area around the building is cordoned off. The chief demolition engineer is ready to press the button which triggers the detonation. Just before the explosive event, a newspaper reporter asks the engineer whether there are people in the building. This is an important question. The potential loss of life represents a grave moral wrong and it is relatively easy to confirm that there are no people inside. Suppose that the engineer answers in the negative and then triggers the dynamite. As a matter of fact, the engineer is correct; there are no people in the building. The reporter is particularly inquisitive and asks the engineer—after the dust has settled—what grounds he had for saying that there were no people in the building. If the engineer responds that he was merely guessing, then the reporter will certainly write about that. In such a case the engineer acted negligently, even recklessly, and is properly subject to strong criticism. If the engineer responds that the building was unfit for shelter due to numerous leaks, the presence of asbestos, the pernicious rat and cockroach infestations, and that because of this it was highly unlikely that there were people in the building, then he is

equally subject to strong criticism. At best the evidence that the engineer had warranted the claim that it was highly unlikely that there were people in the building. But if instead the engineer responds that he and his staff toured the building just prior to the demolition, visually checking every floor and blocking off potential entrances as each floor was cleared, then this constitutes a reasonable and responsible pattern of action which satisfies the moral obligations that the engineer had. In this last case, the engineer knew that there were no people in the building. The epistemic standard for correct moral action in this case is that the engineer knows. Anything short of that would fail to adhere to the proposed constitutive rule of moral action.

Objection: Suppose that the engineer had acted in the proposed appropriate way but, despite his best efforts, there was a homeless person hiding in the ventilation system. The person dies in the demolition. Since knowledge is factive the engineer did not know that there were no people in the building. But the engineer did take appropriate action; he discharged his duties by engaging in a practice which, in the main, leads one to know that there are no people inside. Surely he is not to be blamed for the subsequent loss of life. He may agonize about it, he may even speculate about further measures he could have taken to ensure that no one was in the building, but only the most strict consequentialists would hold him morally accountable for the death of the homeless person (and then so only on the condition that the long-term consequences were in fact negative). More common moral intuitions hold that while the loss of life is unfortunate, the engineer is not to blame. But if the constitutive rule of this morally significant action is to trigger the

dynamite only if one knows that there are no people inside the building, then the constitutive rule is wrong.

Response: In the case of assertion, the constitutive rule is to assert that p only if one knows that p. But it often happens that someone will assert that p when she does not know that p, not through any fault of her own, but because she thought she had evidence which she in fact did not have. A woman reads a newspaper which contains an error. It says that the Yankees won the game last night when in fact they lost the game. She reports to her friend that the Yankees won the game. She has, in this case, asserted something which she did not know. She is therefore properly subject to correction, though she is not properly subject to criticism. Her intention was to assert what she knew, but, for reasons given by the luminosity argument, she is not always in a position to know what she knows or does not know. She has, in this case, an excuse. Moreover, it is a fairly good excuse since she attempted to act in accordance with the constitutive rule of assertion (though she would likely not put it quite that way). Similarly, in the case of the demolition engineer, he attempted to act in accordance with the constitutive rule of that morally significant action. He thought that he knew that there were no people in the building. The method that he used (viz. careful observation) to try to acquire that information does, in the main, lead to knowledge. He is therefore properly subject to correction, though he is not properly subject to criticism. Of course if the method he used more than infrequently led to such mishaps, then that would be cause for altering the procedures used in the future to secure the knowledge that there are no people in buildings which are about to be demolished. This is especially important considering the

gravity of the consequences which might occur. The upshot of this line of thought is that there are different appropriate reactions to violations of constitutive rules. The intent of the agent matters when determining what the appropriate response should be. Since the engineer made a good faith attempt to adhere to the rule (though he would likely not put it that way), the response should not be one of moral condemnation.

This analogy between assertion and certain morally significant situations only goes so far. Williamson's view of assertion as having constitutive rules was modeled on a relatively non-controversial view of the nature of games. Consider Dan Dennett's discussion in "Higher Order Truths about Chmess"[13]. This is a game which is like ordinary chess in every way with the exception that the king may move two squares rather than just one. The point is that by altering just one movement rule, the game deserved a new name. Chmess is not chess, although many of the strategies that work well in chess will likely work well in chmess. That the king may only move one square at a time is a rule of chess. The status of this rule fits Williamson's model as a constitutive rule, at least if Dennett's sentiments are to be trusted. If one alters the movement rules of chess, then one is no longer playing chess. Of course, one might be playing chess and either mistakenly or deliberately violate such rules. However such violations do not in themselves imply that one is no longer playing chess. If it did, then the normal response to such violations would seem odd. If I sit down to play chess with an opponent and she moves her king two spaces, I would and should remind her that a king may only move one space. Certainly there is a jocular tone I might adopt where I

[13] Available online: http://www.ephilosopher.com/modules.php?op=modload&name=Sections&file=index&req=viewarticle&artid=9

respond that I did not realize we were playing *chmess*, but that is a special case and may be put to one side here. It would further be inappropriate, even rude, to question my opponent as to what game we are in fact playing. Charity demands that we treat the move as an error. But in order for it to be an error, we must actually be playing chess.

Williamson holds that the constitutive rule of the conversational act of assertion is the following[14]:

(The knowledge rule) One must: assert p only if one knows that p.

The rest of that chapter is a defense of the knowledge rule. This rule cannot be altered in the way that a conversational convention may be altered without forsaking the act of assertion.[15] Of course people will routinely violate the knowledge rule whether by accident, malicious intent, or some other means. But an appropriate response to putative violations of the knowledge rule is to ask how one knows (the content of) what one just said. Even the question "How do you know?" presupposes that the asserter knows the content of her assertion.[16] If the asserter does not know (the content of) her assertion, then something is wrong with her having asserted it. She overstepped her evidential authority.

[14] KAIL, pg. 243
[15] Williamson says, "It is pointless to ask why the knowledge rule is the rule of assertion. It could not have been otherwise." KAIL pg. 267
[16] KAIL, pg. 252-253

Similar thinking applies in the case of the newspaper reporter asking the engineer how he knows that there are no people inside the building. This question presupposes that he does know that there are no people inside the building. After all, if he does not know that there are not people inside, then something is dreadfully wrong with his demolishing the building. For a demolition event to take place which is, ceteris paribus, free of problematic moral implications, it is necessary that the demolition engineer, whose job it is to ensure that there are no people in the building, knows that there are no people in the building. This suggests a simple constitutive rule:

(KRMD) One must: perform D only if one knows that not-p.

The knowledge rule for moral demolitions (KRMD) takes 'D' to stand for a given demolition and 'p' to stand for people being in the building to be demolished. If one is not (at least tacitly) sensitive to this rule, then one is not conforming to the moral demands of the demolition engineer occupation. Unlike the knowledge rule for assertion, the 'must' here does carry moral force. Such a rule is certainly sufficient to capture a central component of the duties of a demolition engineer. At least three objections can be leveled against it. First, it requires too much of engineers. Second, it has trouble accounting for Gettier-type cases. Third, its status as a constitutive rule is questionable.

The first objection suggests that there are lesser rules which might more adequately capture the moral obligation of the demolition engineer. One proposal could be the following:

(RBKRMD) One must: perform D only if one rationally believes that one knows that not-p.

The rationally believes oneself to know rule for moral demolitions (RBKRMD) also takes 'D' to stand for a given demolition and 'p' to stand for people being in the building to be demolished. The advantage of such a rule is that it seems to more adequately capture our moral intuitions in the case where there is a homeless person hiding in the ventilation system. Since the engineer did rationally believe that he knew that there were no people in the building, it follows that he discharged his relevant duty. (RBKRMD) explains that fact better, or at least in fewer steps, than (KRMD).

Perhaps this is an improvement on (KRMD), but if so it is not obvious. If the defender of (RBKRMD) favors it because of its first-person epistemic accessibility, then she will be disappointed. Human beings do enjoy some special first-person access to what they rationally believe themselves to know, but the access is not perfect. In particular, it has not been shown that such access is superior to that of (KRMD). Both knowing and rationally believing oneself to know are conditions which are susceptible to Williamson's anti-luminosity argument. For each of them it is not always the case that when one is in that condition one knows that one is in that condition. Therefore there will still be cases where one thinks oneself to rationally believe that one knows that there are no people in the building when in fact one is mistaken. Such cases will require

additional explanations about what constitutes an adequate excuse. Those explanations will mimic the explanations given for (KRMD).

(RBKRMD) is a parallel of an alternative rule of assertion that Williamson considers, namely RBK[17]:

(The RBK rule) One must: assert p only if one rationally believes that one knows p.

One of the criticisms which Williamson gives against RBK is that it too easily allows one to transfer the authority to assert p even when one knows that not-p. The idea is that an individual A who knows that p is false may create misleading appearances which convince another individual B that A in fact knows p. If A is trustworthy and the circumstances are right, then B will rationally believe that he knows that p. Williamson points out that B has not in fact acquired the authority to assert p, despite appearances. By analogy, one may object that (RBKRMD) too easily allows one to transfer moral authority to demolish a building even when one knows that there are people inside. For if A knows that there are people inside but creates misleading appearances which convince B that A in fact knows that there are not people inside, then B may come to rationally believe that he in fact knows that there are no people inside. But in this case, B does not, despite appearances, have the moral/evidential authority to demolish the building even though B does rationally believe that he knows that there are no people in the building. Of course, any blame should be placed on A's shoulders, but there is still something problematic about B's action.

[17] KAIL, pg. 261

But beyond this, (RBKRMD) still has need of the concept of knowledge. In that sense, even if it is a better rule than (KRMD) it still supports the general Williamsonian idea that knowledge is a central and ineliminable concept in the explanations of various phenomena. Substituting the concept of highly justified belief for knowledge simply will not do. Suppose there is a lottery in which every building in the area is assigned a number. Let us suppose that there are 10,000 such buildings. There is one number drawn and the corresponding building will have one person inside while all other buildings will be evacuated. Without knowing the results of the lottery, the demolition engineer has very good probabilistic evidence that the building he is about to demolish will be evacuated. Nevertheless, it is morally incumbent on him to find out what number was drawn.[18]

The second objection is that the present account cannot adequately handle Gettier-style variations. Suppose that the demolition engineer surveys each floor of the building carefully, but unbeknownst to him, there is a person hiding in a ventilation duct. Seconds before the demolition occurs, a passing alien spaceship transports the person from the ventilation duct onto the safety of the deck of the ship. The engineer thus has a justified true belief that there is no one in the building, but given the circumstances he does not know that no one is in the building. If (KRMD) is correct, then the engineer has done something wrong. But surely he has not done anything wrong. He engaged in a process

[18] Some may hold that it is further morally incumbent on him to examine the building for himself. After all, lives are at stake.

which tends to lead to knowledge, but because of highly unusual conditions he failed to know. No one was actually harmed, and the engineer seems to be blame-free.

One way to respond to this issue is to try to bring the Gettier-type cases into the theory by special accommodation. We could then propose a new rule (KRMD*) which reads:

(KRMD*): One must: perform D only if either one knows that not-p or one does not know that not-p but only for Gettier-type reasons.

Such a principle has all the feel of an ad hoc patch up. It looks like (KRMD) was faced with a serious counterexample and, in order to overcome it, made the problem go away without any theoretical motivation. Ad hoc appendices are methodological pariahs. But the situation is not nearly as grim as it may appear. If the modification can be substantiated by a previous theoretical commitment, and not just a quick response to a proposed counterexample, then it dodges the criticism that it is ad hoc.

Such a reconciliation is possible. The worrisome Gettier-type cases have certain commonalities. They are all cases in which the demolition engineer has a justified true belief that there are no people in the building. They are all cases in which the demolition engineer does not know that there are no people in the building. Since the failure of the engineer to know that there are no people in the building is not parasitic on the absence of belief, justification, or truth, we must look elsewhere for the complicating factor. Stating

what this missing factor is has been the holy grail of many epistemologists since 1963. But whatever the condition turns out to be, the description of the Gettier cases need not change. Further, Williamson has argued that any account of such a fourth condition will either fail by being either false or circular. Succinctly, Gettier cases are cases in which an individual has a justified true belief which is not knowledge. The structure of the Gettier cases is thus not characterizable without adverting to knowledge. This itself is a substantial concession to the general project of using knowledge as a primary explanatory tool. If knowledge is to be the primary implement of epistemological investigation, then it is no surprise that the characterization of other epistemic concepts should ineliminably advert to knowledge. Gettier cases are certainly special; they destroy knowledge in ways not obvious to overcome. But Gettier cases are closely related to ordinary cases of knowledge. In particular they each demand justification, truth, and belief. Given the facts about probabilistic evidence and how it will never be sufficient for knowledge[19], the justification in question in a Gettier case is, generally speaking, the right sort of justification for knowledge. Resistance to (KRMD) on the grounds that Gettier-type cases are counterexamples to it stems from the deep similarities between those cases and ordinary cases of knowledge.

The third objection to (KRMD) is more pressing. Morality, though it does have rules (at least in some sense of the word 'rule'), is hardly like a game. It is at best a strain on the analogy to suggest that morality has constitutive rules. What is more plausible is that morally significant actions have standards that one must strive to meet. In the case of demolishing a building, one may do so only if one knows that there are no people inside.

[19] This is due to the problem of lotteries.

When one instructs apprentice demolition engineers, one should teach them that they need to know that the building to be demolished is entirely vacated prior to detonating the explosives. For a demolition event to be morally unproblematic it must be the case that the engineer knows that there are no people in the building. Failure to know such information explains why the action is wrong. This line of thought paves the way for an additional rebuttal to the second objection. It would be unsatisfying if every proposed norm for moral action which was initially articulated in terms of what the agent knows had to be recast in terms of the disjunction of knowing or being in a Gettier case. That move is correct for (KRMD*), but it need not be deployed in every case. When developing guidelines to express the moral obligations of an individual with a particular social role, often the best expressions of such obligations are knowledge-theoretic. Demolition engineers, as the designated parties responsible for destroying buildings, must strive to know that there are no people inside the buildings they destroy. Trying to do anything less, such as trying to either know or be in a Gettier case, is to aim at too weak a goal given how high the stakes are. Both particular actions and dispositions to act can be evaluated along a moral dimension. The point here is that some of those evaluations are best explained by referring to what the agent knows, tried to know, or should have known.

The move from treating knowledge as a component of a constitutive rule for certain moral actions to treating it as the standard for certain moral actions preserves the points of the previous discussion. While it would perhaps have been more theoretically satisfying to take knowledge as a constitutive rule, the spirit of that account can be

sustained by treating knowledge as a necessary and ineliminable element in our best explanations of some types of moral action. Such a view constitutes an extension to the Williamsonian project of taking knowledge as a central theoretical concept.

# Chapter Six

## Contextualism

The revolutionary epistemological theory that Timothy Williamson advocates in Knowledge and Its Limits offers novel solutions to a variety of epistemological problems. The account addresses most of the major topics in epistemology over the last fifty years, topics which often trace their philosophical lineage to questions raised by the Ancients. But certain theoretical possibilities never arose for Plato, Aristotle, Sextus Empiricus, or even Descartes or Hume. In recent years, a group of epistemologists have advanced an account of the semantics of knowledge ascriptions according to which differences in context yield different relations for the semantic value of 'know' and its cognates. This contextualist program seeks to dissolve one of the most paradoxical epistemological arguments on offer. Williamson has argued, in a separate work[1], that contextualism is incorrect on internal grounds. But there is another series of criticisms of contextualism that are more directly consequences of the epistemological commitments of Williamson's book. These appear in "Knowledge, Context, and the Agent's Point of View"[2]. It is not a goal of this chapter to defend contextualism in general. Any argument here which appears to have that goal should be interpreted as providing support for the compatibility of contextualism with a Williamsonian account of the nature and extent of our knowledge and its important explanatory roles. This chapter will first characterize a version of a paradoxical argument which simultaneously brings out the uncomfortable commitments

[1] Williamson, Timothy 'Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge', *The Philosophical Quarterly* 55, 219 2005
[2] Williamson, Timothy 'Knowledge, Context, and the Agent's Point of View', in Contextualism in Philosophy eds. Preyer and Peter, Oxford University Press 2005

of both a highly skeptical position and an inflexibly dogmatic one. Next it will explain the contextualist resolution of the argument that avoids the unwanted consequences of either of those extreme viewpoints. Following that, it will provide an argument against contextualism that relies on premises whose primary motivations are, not always obviously, connected to issues raised in Knowledge and Its Limits. Once those objections are on the table, it will be argued that they are not decisive against contextualism. An important moral to draw from this is that the neutrality of the Williamsonian view with respect to contextualism is a methodological asset.

Victoria is standing at a bus stop at noon waiting for the number four bus. It is raining, so she is eager to board. However, the rain does not significantly impair her vision. The properly labeled number four bus slowly rolls to a stop in front of her. Victoria is an ordinary person with an ordinary visual and conceptual system, and she is attending to the question of whether the bus is before her. She is not in "phony bus country", nor are there malevolent demons or deranged, mind-tinkering scientists working mischief in the vicinity. It seems that in this situation the following is true:

(1) Victoria knows that the number four bus has arrived.

From any sentence S, an infinite number of sentences are validly derivable. In addition to the straightforward logical implications, there are an infinite number of metaphysical entailments that obtain. While Victoria does not know all the entailments of any claim, she does have enough conceptual sophistication to recognize some of the

more obvious facts which follow from what she knows. Specifically, the following seems plausible:

(2) If Victoria knows that the number four bus has arrived, then she knows that there is not an eccentric scientist causing her to falsely believe that the number four bus has arrived.

After all, if there were an eccentric scientist causing her to falsely believe that the number four bus had arrived, then it would be false that the number four bus had arrived. Since she knows that the number four bus has arrived, and since knowledge is factive[3], she can perform the deduction and conclude the consequent of (2) which is:

(3) Victoria knows that there is not an eccentric scientist causing her to falsely believe that the number four bus has arrived.

Unfortunately, this last claim is uncomfortably strong. It requires that a competing hypothesis be ruled out on the basis of evidence which does not seem robust enough to accomplish the task. If an eccentric scientist had caused Victoria to falsely believe that the number four bus had arrived, then she would be having the same kind of visual experience[4] that she in fact was having.[5] The two experiences are qualitatively

[3] For rare opposition to this view, see Hazlett, Alan 'The Myth of Factive Verbs' *Philosophy and Phenomenological Research* forthcoming.
[4] Characterizing the situation as one in which Victoria has the same *evidence* in the skeptical possibility as she has in the non-skeptical actuality may be a more natural characterization. However, that terminology prejudices the case against Williamson's view of evidence. See chapter two of this dissertation and KAIL chapters eight and nine for discussion of this view of evidence and its role in avoiding skepticism.

indiscriminable from her perspective. It is hubris for her, or anyone else, to hold that Victoria can correctly decide between these closely competing explanations of her experiential state, even when one is (arguably) more common and less complicated, and the other (arguably) less likely and deeply contrived. When the probability of a hypothesis is quite low on one's evidence, that alone is insufficient to know that the hypothesis is false. Suppose there is a lottery with a billion tickets and only one winner. If one has purchased a single ticket, the chance that one has the winning ticket is exceptionally low. But despite the low probability, one cannot know that one's ticket is a loser if one's evidence is solely probabilistic, even if one's ticket is in fact a loser. But, so the argument goes, Victoria has no non-probabilistic basis on which to distinguish between a veridical experience of seeing the number four bus arrive and an illusory experience of the same phenomenal type. She is, like all of us, an imperfect epistemic agent.

Given the truth of (1) and (2), it follows by Modus Ponens that (3) is true. Given that (3) is false and (2) is true, it follows by Modus Tollens that (1) is false. The intuitive allure of (1) and (2) coupled with the implausibility of (3) create a tense paradox that requires a philosophical resolution. Since any consistent position with respect to (1), (2), and (3) will require abandoning some pre-theoretic intuition, one neutral demand to place on any adequate resolution of the paradox is that the source of the competing intuition be explained.

---

[5] For a discussion of the virtues of a subjunctive conditional account of the implausibility of (3), see DeRose, Keith 'Solving the Skeptical Problem', *The Philosophical Review* 104 (1995).

There are two straightforward views that can resolve the paradox, albeit with less than fully satisfying accounts of the pull of the competing view. The first is to adopt an aggressive dogmatism. By accepting (1) as common sense and embracing the instance of closure captured in (2), the dogmatist bites the bullet and accepts (3) as an unexpected but nevertheless fully justified conclusion. Resistance to (3) is blamed on an overly "philosophical" mind that is inappropriately troubled by remote possibilities. The second is to adopt a pessimistic skepticism. By rejecting (3) as pure *chutzpah* and embracing the instance of closure captured in (2), the skeptic bites the bullet and rejects (1) as an unexpected[6] but nevertheless fully justified conclusion. The initial attractiveness of (1) is blamed on an insufficiently "philosophical" mind that is inappropriately ignoring undischarged and incompatible possibilities. In both cases, the opposing position is seen as resting on ludicrous foundations and the proper resolution of the dispute will involve the conversion of one camp to the point of view of the other.

Contextualism is a view that attempts to both resolve the paradox and provide an explanation of the initial intuitive appeal of each member of the paradoxical group. By doing so, contextualism legitimizes the claims of the apparently opposing parties by regarding each as true in its own proper sphere. The way it does so is by relativizing the knowledge ascriptions to a particular context. An occurrence of 'know' or 'knows' in one context need not express the same relation in a distinct context. Different contexts import different demands on what qualifies as knowledge. This sensitivity to context results in a different function from sentences to propositions for each differing context. While on the surface there may appear to be a genuine dispute between the epistemic

[6] Not unexpected for the skeptic herself perhaps, but certainly unexpected for most.

liberal and conservative, in point of fact they are speaking past each other. Since their sentences equivocate on the word 'knows' across contexts, there is no incompatibility between their orthographically similar utterances.

When one considers (1) by itself, the context is one in which the standards for knowledge are not particularly high. So (1) turns out to be true because one is not required, by the epistemic standards in that context, to rule out certain contraries. When one considers (3) by itself, the context is one in which the standards for knowledge are higher since, at a first approximation, certain relevant but undischarged contraries are salient. So (3) turns out to be false since Victoria's evidence is not strong enough to rule out the possibility described there. Once the possibility in (3) has been raised, the context has shifted so that ordinary knowledge ascriptions like (1) are rendered false. The skeptical position is explained because once the extreme possibility is part of the context, knowledge of mundane facts is harder to acquire. The dogmatist position is explained because if an extreme possibility has not been raised, there is no need for it to be part of the context and therefore the word 'knows' gains a wider correct application. Paying attention to what context is in question clears up certain epistemological confusions.

Claim (2) is structurally interesting. In that conditional, there are two instances of 'knows', each demanding a relativization to a context. When both the antecedent and the consequent are governed by the same context, (2) is plausible on the general methodological grounds that wherever possible one should preserve closure of

knowledge under known deduction.[7] With no difference in context between the two occurrences of 'knows', there is no obvious worry about (2). If one knows against low standards that p, and one knows against low standards that p entails q, then one is in a position to know[8] against low standards that q. This is correct even when q is a skeptical hypothesis like the one in (3). A typical skeptic is not concerned with whether we know against low standards that we are not brains in vats. Knowing against low standards is not really a philosophically interesting relation for a skeptic.[9] It is knowing against high standards that matters, and it is this relation that the skeptic thinks one cannot have to various skeptical possibilities and by consequence to mundane propositions about the external world. If one knows against high standards that p, and one knows against high standards that p entails q, then one is in a position to know that q. This is correct even when q is a skeptical hypothesis like the one in (3). One's knowledge against a high standard has sufficient evidential backing to discharge various skeptical contraries; that is what meeting a high standard for knowledge requires.

However, if there is a shift in the contexts of 'knows' from antecedent to consequent in (2), then the entailment is not so obvious. Certainly if the context for the antecedent is one where low standards for correct applications of 'knows' apply and the context for the consequent is one where high standards for correct applications of 'knows' apply, then

[7] Not all epistemologists endorse even a suitably mollified version of closure. For articulate discussion of the issues involved, see John Hawthorne's and Fred Dretske's contributions to *Contemporary Debates in Epistemology* ed. Matthias Steup and Ernest Sosa, Blackwell Publishing 2005.
[8] Certain other criteria are required here as well, largely to avoid technical objections. For instance, one must maintain knowledge that p after deducing q in the relevant type of closure.
[9] There are, of course, differences in the strengths of various skeptical positions. One might be a skeptic about the external world, of the past or future, of the a priori, or even of the legitimacy of deductive reasoning. Obviously this last type of skeptic falls outside the scope of the present discussion, since if all deductive reasoning is called into question, there can be no non-question begging way of arguing in favor of its correctness.

there is little if any reason to believe that (2) holds. But one might reason that if the context for the antecedent is one in which the standards for knowledge are high, then if the shift in context produces a lower standard for knowledge in the consequent, then the entailment does hold. The high standard of the first context provides enough epistemic support to warrant the low standard ascription in consequent. But this simple view ignores the fact that standards of evaluation can make demands on the correct application of 'knows' that function along multiple and independent dimensions. The cognitive accessibility of defeaters, the existence and proximity of misleading evidence, and the facts about various counterfactuals may all vary depending on the context. A high standard for knowledge may demand an array of epistemic merits for a correct application of 'knows' while a low standard may demand fewer but different merits for the correct application of 'knows'. Without a substantive account of the relationship between high and low standards, the mixed contexts case for the evaluation of (2) is complicated. Thankfully, the present discussion does not require any resolution of this problem, since the dogmatist, the skeptic, and the contextualist are all happy to countenance (2) true, and to regard the context as fixed over the entire conditional.

The context sensitivity of sentences is a perfectly general phenomenon of which epistemic contextualism, the context sensitivity of 'knows', 'justified', and the like, is a species.[10] The most familiar form of context sensitivity occurs in the use of indexical words like 'I', 'now', and 'here'. The semantic contribution each of these words makes when deployed in a sentence in a given situation is sensitive to the identity of the speaker,

[10] For discussion, see Conee, Earl 'Contextualism Contested' in *Contemporary Debates in Epistemology*, ed. Steup and Sosa, Blackwell Publishing 2005.

the time of the utterance, and the location of the utterance respectively. In an important sense, different speakers will mean different things by these words. Of course there is something that remains the same across contexts, it is the *character* of these words in David Kaplan's sense.[11] The character of the word 'I' is that it picks out the speaker. It does so in a directly referential way, unmediated by a descriptive element. But the *content* of the sentence which contains these words shifts with each use. For sentences like 'I am here now', an unmodified disquotational account of truth is unsatisfactory. For most sentences s, 's' is true if and only if s. But it is false that every occurrence of 'I am here now' is true if and only if I am here now. Because of contextual variation, the referents of 'I', 'here', and 'now' on the right of the biconditional need not match the referents of their counterparts in the named sentence on the left.

While some context sensitivity in our language is uncontroversially present, there are words which are less clearly context sensitive. Some have suggested that vague predicates are context sensitive and that recognizing their context sensitivity resolves certain uncomfortable philosophical dilemmas. Bob is six feet, one inch in height. Compared to his jockey friends, it seems correct to classify Bob as tall. Compared to his professional basketball playing friends, it seems correct to classify Bob as not tall. To simplify the case, assume that the sentences used to express these claims differ only with respect to the word 'not'. Yet here the resistance to a contextualist resolution is at least contentious.[12] It requires a strong philosophical argument to demonstrate that there is a single property of tallness, referred to in each of the competing claims, which either

[11] Kaplan, David 'Demonstratives' in *Themes from Kaplan* ed. Almog, et al. Oxford University Press 1989
[12] The most famous example of a philosopher who argues in favor of sharp boundaries for all vague terms is Timothy Williamson himself in *Vagueness* Routledge, 1994.

applies or does not apply to Bob depending on whether his height falls within the range of that property. On that view, one of the two claims must be false. A more natural resolution is that there is no real dispute going on here. The correct application of 'tall' is different in the differing contexts because the word expresses distinct properties depending on the comparison class. The superficial surface structure may suggest an inconsistency, but when the contextualist semantics are applied, the tension disappears. There is no need for one to give up on either claim because there is no dispute in the first place. One need only pay attention to the context in which the sentence is being evaluated in order to correctly interpret whether the sentence is true or false in that context. It is folly to look for a uniform cross-contextual interpretation of vague terms like 'tall', however theoretically satisfying it might be to find one.[13]

Words like 'knows', 'justified', 'warranted' and the like are vague. There are borderline cases for each where it is unclear whether a given case falls inside or outside the range of correct applications of the terms. Since context sensitivity is so useful in resolving concerns about other vague terms, it is natural to apply that apparatus to these epistemic terms. Once applied, the reward is a philosophically interesting resolution of an otherwise pernicious paradox. So why not accept a wholesale contextualism about epistemic terms?

Contextualist Sentiments in Knowledge and Its Limits

[13] Part of Williamson's objection to contextualism relies on a rejection of the universal coupling of vague terms and contextualist resolutions. More on this later.

Before addressing some of Williamson's arguments against contextualism, it is worthwhile to note that the epistemological view espoused in Knowledge and Its Limits is not obviously at odds with the view that 'knows' varies in extension depending on context. Some exegetical observations will legitimize this point. Williamson has argued, persuasively, that knowledge is the norm of assertion. The conversational activity of asserting, rather than merely saying or conjecturing, has but one constitutive rule: assert that p only if one knows that p. Generally speaking, philosophers who are opposed to this rule[14] are opposed to it on the grounds that less stringent standards govern the correctness of making an assertion. Perhaps some milder condition, such as only believing that p, or only rationally believing that p, or only believing truly that p, or only believing that one knows that p, all conditions weaker than the knowledge rule, is the appropriate rule to capture what speakers must be at least implicitly sensitive to in order to genuinely be making assertions.[15] However, opposition could come from the opposite direction. One might be so convinced of the high epistemic requirements of appropriate assertion that the knowledge rule that Williamson defends appears too weak. Williamson's response to this objection is interesting in its extensive use of contextualist reasoning:

> "After all, something is wrong even with the assertion 'A and I cannot be certain that A'. Does that not suggest that only something *more* than knowledge warrants assertion? What seems to be at work here is a reluctance to allow the contextually set standards for knowledge and certainty to diverge. Many people are not very happy to say things like 'She knew that A, but she could not be certain that A'. However, we can to some extent effect such a separation, and then assertibility goes with knowledge, not with the highest possible standards of certainty. For example, one may have warrant to assert 'A and by Descartes's standards I cannot be absolutely certain that A', where the reference to Descartes holds those standards apart from the present context. Again, it would often be inappropriate to respond to the assertion 'A' by asking 'How can you be so certain that A?'. The word 'so' flags the

[14] For example, Richard Fumerton in conversation.
[15] Williamson objects to each of these weaker rules. See chapter 11 of KAIL.

> invocation of unusually high standards of certainty. By ordinary standards you may have had warrant to assert that A even if you could not be *so* certain that A. … The present account permits such contextual variation…"[16]

If the general approach to epistemology that Williamson advocates were, at bottom, opposed to the contextualist program, then there would be little need to make room for a contextualist treatment of his rule of assertion. One of the motivations for the view that knowledge is the norm of assertion is a desire to see knowledge per se, and not some complex of other concepts, play an ineliminable explanatory role in our assessment of epistemic phenomena, of which assertion is a member. Williamson is at pains throughout Knowledge and Its Limits to avoid explanatory models which would fare equally well were the occurrences of 'knowledge' replaced by a cluster of concepts like belief, truth, justification, and whatever other element is required to overcome Gettier cases. Whenever an alternative explanation seems possible within a decompositional structure, Williamson attacks the view. So not all earlier epistemological theory is given latitude to operate inside a Williamsonian framework. The inclusion of a contextualist line of thought, without the accompanying criticisms typical of his way of arguing, signals a willingness to embrace at least some elements of contextualist thinking into the fold of his new epistemology.

A second contextualist-friendly sentiment occurs in Williamson's discussion of evidence. In most conceptions of evidence, with the possible exception of a strong coherentist position, a proposition p cannot serve as evidence for p. Vicious circularity in one's reasoning is to be avoided and the smaller the circle the more dangerous it is. But

[16] KAIL, pg. 254-255

on Williamson's account of evidence according to which one's evidence is identical with what one knows, once p is known it is part of one's evidence. Since knowing p guarantees that p is true, from the standpoint of E=K, p is outstanding evidence for p. Yet there does seem to be something inappropriate about citing p as evidence for p. One way of capturing what is wrong with citing p as evidence for p is that it violates certain conversational norms. If one is asked for a reason for p, it is usually the case that p itself is in doubt and some *other* piece of evidence is requested to support it. To cite p as a reason for p misses the point of the request for a reason, though it need not be false that p is evidence for p given a "violation of conversational norms" analysis. But a contextualist can offer a strong account of the inappropriateness of citing p as evidence for p by arguing that it is simply false that such a relationship could ever obtain.

> "The idea would be that the question 'What is the evidence for e?', meant as a challenge, creates a context in which e falls outside the extension of 'S's evidence'. But that seems too drastic. For example, suppose that a doctor asks you, 'Do you feel a tingling sensation?' and you answer, 'No.' If you were asked 'What is your evidence for the proposition that you do not feel a tingling sensation?', you might be at a loss to answer, for the question seems to expect some *further* evidence for the proposition, and you might look in vain for such further evidence. Nevertheless, when we assess the status of your claim that you did not feel a tingling sensation on your evidence, we do not exclude that proposition from your evidence. Its presence justified your claim. This is not to deny that the extension of 'evidence' may vary slightly with context, perhaps corresponding to slight contextual variation in the extension of 'knowledge'… ."[17]

This is not exactly a ringing endorsement of contextualism, but it does reveal a willingness to accommodate contextualist sentiments inside the general epistemological program. The text appears to provide a counterexample to a contextualist resolution of the inappropriateness of citing p as evidence for p. But instead of following through and rejecting contextualism as incompatible with his proffered account of evidence,

[17] KAIL, pg. 188

Williamson gives the contextualist breathing room by suggesting that 'evidence' may indeed be context sensitive. The fact that he rejects one strong contextualist consequence while remaining agnostic about the context sensitivity of correct applications of 'evidence', and therefore of correct applications of 'knowledge' by E=K, signals again a willingness to accommodate aspects of contextualist thinking in his general epistemological framework.

It is not really surprising that Williamson does not reject contextualism when arguing for his preferred account of evidence and the norms of assertion. Williamson's epistemology upends the earlier conceptual analysis program by arguing that knowledge should be the primary tool of epistemological investigation rather than the object of conceptual analysis (in the sense of conceptual decomposition and definition). But in many ways his epistemology can extract the important insights of earlier epistemologists who were convinced of the ultimate success of analyzing the concept of knowledge into a set of more well-understood concepts.

Williamson has explicitly argued that reliability is required for knowledge.[18] But his account of what reliability amounts to is saturated by the knowledge-first attitude. Reliabilism treats reliability as a conjunct in an analysis of the concept of knowledge. Since Williamson has rejected the viability of the conceptual analysis program as applied to knowledge, he cannot concur with a standard reliabilist view on the role of reliability in understanding the nature of knowledge. But that does not imply that reliability is not an important requirement for knowledge. Unreliable belief-forming processes are surely

[18] KAIL, pg. 98-102

poor vehicles for accruing knowledge. When we encounter unreliability in someone's cognitive architecture, it is often, if not always, correct to deny that the person can come to know things by way of such hazardous means. A carnival palm-reader who forecasts wealth and good fortune for a particular person does not know that such benefits await her customer even if she is in fact correct about what will happen in the future and sincerely believes in her own prescient abilities. The unreliability of palm-reading as a way of gaining information about the future is what prevents the palm-reader from having such knowledge. But from the fact that reliability is necessary for knowledge, it does not follow that reliability is part of a non-circular set of necessary and sufficient conditions for knowledge.[19] Nevertheless, the spirit of the reliabilist account remains intact inside Williamson's epistemological view in the sense that reliability is counted as a legitimate necessary condition for knowledge. Even the difficulties in specifying the exact process which is relevant to concerns about the reliable origins of particular beliefs[20] is no impediment to utilizing reliability as an admittedly vague but powerful explanatory tool.[21] If Williamson has created a revolution in epistemology, he has done it by garnering support from earlier squabbling factions. In that sense, Williamson avoids distancing himself too much from earlier theorists, despite the radical differences in the overall explanatory structure.[22]

## An Anti-Contextualist Argument

[19] See KAIL pg 32 and chapter one of this dissertation for further discussion of this point.
[20] Conee, E. and Feldman, R. 1998 'The Generality Problem for Reliabilism' *Philosophical Studies* 89
[21] See KAIL pg 100.
[22] As I argue in chapter four, this line of thought can be extended to include defeasibility accounts as well. The absence of defeaters is arguably a necessary condition on knowledge without itself being a component of a non-circular set of necessary and sufficient conditions for knowledge.

In "Knowledge, Context, and the Agent's Point of View", Williamson offers an argument against contextualism that begins by noticing the apparent lack of context-sensitivity of 'wrong' and then draws a noteworthy parallel between judgments about the use of 'wrong' and judgments about the use of 'knows'. If 'wrong' is indeed context-insensitive and if correct applications of 'knows' are parasitic on correct applications of 'wrong', then 'knows' (and its cognates) will turn out to be context-insensitive as well. But close attention should be paid to whether the analysis of the use of 'wrong' that Williamson offers is really the correct one. Additionally, there are concerns over whether the argument is compatible with the spirit of Williamson's epistemology as articulated in Knowledge and Its Limits.

He begins by presenting a case in which an agent, Clare, is deliberating over whether to quit her job. It is a difficult decision for her because there are reasons both for and against her leaving the job. If she is to quit her job, then she wishes to do so only as a matter of principle. If she is to retain her job, then, again, she wishes to do so only as a matter of principle. But the reasons both pro and con are reasons of principle, and she vacillates between judging positively and negatively whether on balance the reasons favor one or the other option.

As we saw earlier with respect to statements (1), (2), and (3), contextualist resolutions of paradoxical sets of statements attempt to preserve the correctness of apparently competing judgments by relativizing crucial terms to contexts. Since different

propositions are expressed by the same sentence type in different contexts, allegedly competing claims can be maintained simultaneously. However, in Clare's situation, Williamson argues, a contextualist resolution of her difficult choice seems hollow. Clare's vacillations between the judgments that she should resign and that she should not resign are caused by changes in the salience of principled reasons. A contextualist solution is that when the reasons in favor of resigning are more salient, that creates a context in which it is not wrong to resign. When the reasons against resigning are more salient, that creates a context in which it is wrong to resign. She need not reject either position. Each is correct given its appropriate context. We can also consider the third-person case of two evaluators taking opposite views about whether Clare should resign. Again, differences in the context of the two evaluators can render different verdicts about the truth of the sentence 'Clare should resign' as uttered by each evaluator. Suppose there is a time limit on Clare's deliberations. If she is to resign at all, she must do so by time t. At t, Clare must make a decision. She cannot simultaneously resign and not resign.[23] By taking one course of action, she leaves the motivating principles of the contrary course of action behind. She cannot maintain both positions forever since they differ with respect to what she must do at t.

Here is Williamson's assessment of what is wrong with the contextualist position:

> "In cases of decision-making, one context is distinguished above all others: that of the agent at the moment of action. The primary question is whether the sentence 'It would be wrong for me to resign' expresses a truth as uttered in the context in which the speaker is Clare and the time is that for

[23] In passing it is worth noting the strong parallels between Clare's situation and what William James calls a "genuine option". Resigning and not resigning are both live options, the decision is forced, and the consequences are momentous. On James' view, it is permissible for Clare to let her "passional nature" decide the matter.

> resigning if she is to do so at all. Call that context the *agent's context*, and the proposition which the relevant first person present tense sentence (such as 'It would be wrong for me not to resign') expresses in it the *agent's proposition*. If at some point in her agonizing Clare uses the sentence to express a proposition other than the agent's proposition, which might fail to match the agent's proposition in truth-value, she is no longer concentrating on the relevant practical problem. Similarly, an external commentator who uses the sentence 'It would be wrong for Clare not to resign' to express a proposition other than the agent's proposition is no longer concentrating on the relevant practical problem for Clare."[24]

Since a hallmark of contextualism is that the speaker's context (sometimes also called 'the attributor's context') is the correct one for evaluating context-sensitive elements in sentences, Williamson's analysis, if correct, reveals that contextualism fails for 'wrong'. If the only relevant question is the one framed relative to Clare's context when Clare is the speaker, then there is no room for a contextualist resolution of Clare's dilemma. One of the competing views must be false.

The conclusion that 'wrong' is context invariant is crucial for Williamson's argument that 'knows' is context invariant. The next step in the argument is to offer an account of belief according to which one is wrong to believe that p if one does not know that p. In short, knowledge is the norm of belief.[25] If the thesis sounds unbelievably extreme, bear in mind that it relies on a sharp distinction between outright belief and degrees of belief. Put more carefully, the thesis is that it is wrong to have an outright belief that p if one does not know that p. The sense of 'wrong' here is not moral, though it does carry normative force. The proposed standard for non-wrong belief is a parallel of his account of assertion. One is satisfying a certain norm of belief only if one believes what one knows. One can be corrected for believing what one does not know, precisely on those

[24] Williamson, T. 'Knowledge, Context, and the Agent's Point of View' in G. Preyer and G. Peter, *Contextualism in Philosophy: Knowledge, Meaning, and Truth,* Oxford: Clarendon Press 2005
[25] This view is suggested early in KAIL and developed in the chapter on evidence.

grounds, in the same way that one can be corrected for asserting what one does not know. One has overstepped one's evidential authority, in Williamson's terminology. A body of evidence might be strong enough to legitimize a high degree of belief that p is true without legitimizing the outright belief that p is true. It is a strict standard, but one which fits the model of a knowledge-first theoretical structure and accounts for the intuition that when people believe what they do not know they have missed the cognitive mark.

Another way Williamson argues for the wrongness of believing what one does not know is by way of our judgments about excuses.[26] Suppose Victoria unwittingly wanders into a region of phony barns. She knows what barns look like, though she has never before seen a phony barn. Her eyes fix on the nearest phony barn in her vicinity. She makes a judgment about her experience, namely that she is seeing a barn. Since it is not a barn, she does not know that it is a barn. But she confidently believes that it is a barn. Given her visual evidence, it is highly probable that she is seeing a barn. It is also, arguably, highly probable on her evidence that she knows that she is seeing a barn. Since it is so highly probable that it is a barn, and since Victoria is in no position to know that she is not seeing a barn, she can be excused for believing that it is a barn when in fact it is not a barn. But excuses are only applicable to wrong activity. Right activity needs no excuses. In particular, when one knows that p one needs no excuse for believing that p. One knows that p if and only if it is not wrong to believe that p.

The final step of this anti-contextualist argument examines a contextualist resolution of apparently competing views about whether some person S knows that p, and then

[26] 'Knowledge, Context, and the Agent's Point of View' pg. 17

evaluates that resolution in light of the connection between wrong belief and the lack of knowledge. Consider a case in which a person S believes truly on the basis of memory that p. Since the belief is true and S's memory is generally reliable, there is a prima facie case in favor of the claim that S knows that p. An evaluator A who is not a philosopher but is otherwise a competent speaker of English says 'S knows that p'.[27] An evaluator B who is a graduate student in an epistemology seminar (with comparable linguistic competence to A) says 'S does not know that p'. The contextualist treatment of the situation is to point out differences in the contexts of A and B which yield differing contents for 'know' in each of their ascriptions. In this way, an apparent disagreement is defused. Next, each of A and B deduce from their claims the corresponding judgment about whether S is wrong to believe that p. A says 'It is not wrong for S to believe that p' and B says 'It is wrong for S to believe that p'. On a contextualist reading of 'wrong', each of them is correct to say what she said in her proper context. Here is Williamson's line of reasoning at this point:

> "Since epistemic standards vary so widely between the ordinary context and the epistemological context, they vary markedly between at least one of those contexts and the agent's context, [S's] context. Yet the primary question seems to be 'Is it wrong for me to believe that [p]?' as uttered by [S]. The epistemic standards relevant to answering that question are those operative in [S's] context."[28]

As a consequence, a contextualist treatment of 'know' leads to a mistake about the acceptability of the epistemic standards in place in at least one of either A's context or B's context.

[27] It is customary to point out in situations like this that there are innocuous use/mention confusions here.
[28] 'Knowledge, Context, and the Agent's Point of View' pg. 18

Responding to the Anti-Contextualist Argument

A problem[29] with this argument is that it runs afoul of the spirit of certain aspects of Knowledge and Its Limits. Williamson often argues there that knowledge per se is an important philosophical concept because of the role it plays in the best explanations of a wide array of human phenomena. Knowledge is the norm of assertion[30], one's evidence is identical with what one knows[31], and knowledge as a better explanation of successful action than even (mere) justified, true belief[32] are each examples typical of his view of knowledge. Let us focus on the role of knowledge in the explanation of successful action.

When it comes to explanations of immediate human action, belief-desire psychology is sufficient without mentioning knowledge to adequately explain the success of the action. Why did Bob begin running? He did so because he believed the bus was going to run him over if he did not move. That explanation is no worse than that he knew the bus was going to crush him if he did not move. But when it comes to sustaining true belief over time, knowledge fares better than (mere) justified true belief insofar as the former is more robust against counterevidence than the latter. When maintaining a true belief about one's circumstances is essential to the increased likelihood of success of an action, the robustness of knowledge is an explanatory asset. Consider the case of a burglar who

[29] There are doubtlessly other problems with Williamson's argument here. For example, the reasons that Williamson adduces in favor of the view that the agent's context is the only important one for setting the epistemic question are less than conclusive.
[30] KAIL Chapter 11. See also chapter four of this dissertation.
[31] KAIL Chapter 9. See also chapter two of this dissertation.
[32] KAIL pgs. 62, 75, 80, and 101.

desires to steal a diamond from a house.[33] If the diamond is well-hidden and his desire for it remains unchecked, the likelihood that he will succeed in thieving it from the house is better if he knows that the diamond is there than if he merely has a justified true belief that it is there. Since in the latter case, by stipulation, he does not know that the diamond is in the house, there must be some factor which is complicating his epistemic position. Whether it be the belief's essential reliance on false premises, the existence of undefeated defeaters, or some other negative epistemic feature, such elements represent avenues by which the burglar can lose his belief that the diamond is in the house. Counterevidence against p which reveals such epistemic inadequacies can undermine belief in ways that originally knowing that p can resist.

This account of the ineliminable role of knowledge in our best explanations of human action is decidedly third-person. It is a way of evaluating whole classes of cases from an objective explanatory perspective rather than just looking at what the agent has access to. These are explanations of actions which, plausibly, take the evaluator's context as primary, not the agent's. Of course, Williamson nowhere in Knowledge and Its Limits argues explicitly in favor of contextualism, but the apparatus he deploys to argue for the explanatory importance of knowledge is first and foremost a device for external evaluators rather than agents. If the earlier argument leveled against contextualism were really a function of Williamson's general epistemological program, then one would expect there to be some way of making the agent's context essentially primary in evaluations of the likelihood of the agent's success (such as in the burglar example).

[33] This case is articulated in KAIL, pg. 62

Instead, contextualist insistence on the importance of the attributor's context is far more natural.

Compatibility without Consistency

Contextualism as such may turn out to be a defunct theory if it is inconsistent. If so, then nothing can be consistent with it. But there is an important sense in which a theory can be compatible with an inconsistent view in virtue of that theory not *itself* entailing that the inconsistent view is false. Of course, if contextualism is inconsistent, then every theory entails that it is false in the sense that, necessarily, if the theory is true, then contextualism is false. For comparison, consider the case of the freedom of the will and the debate between those who favor libertarianism and those who advocate some version of compatibilism . One reason for the perennial entrenchment of each side is the competing views' compatibility with the observational data[34]. Libertarianism and compatibilism have differing views about what free will in fact amounts to, so, for either of them, if they are false, then they are necessarily false. Strictly speaking, any theory, whether true or false, entails that a necessarily false theory is false. But there remains an important sense of 'compatible', even in the face of these facts, according to which the observational data is compatible with both theories. There is nothing specific to the one theory which makes it incompatible with the second. This is the same sense in which at least some aspects of contextualism are compatible with Williamson's general epistemological program.

[34] Not all versions of libertarianism and compatibilism suggest this, though enough do.

# Jeffrey Glick

Rutgers University, Philosophy Department
26 Nichol Ave.
New Brunswick, NJ 08901
jeffglick@gmail.com

Education:

B.A. 1999, Philosophy, California State University, Northridge
PhD (Filing July 2007), Philosophy, Rutgers University

Dissertation Title: The Importance of Knowledge Per Se

Dissertation Committee: Peter Klein (supervisor), John Hawthorne, Ernest Sosa, and Igal Kvart

Employment:

University of Rochester: 2007-2008 Visiting Assistant Professor

Rutgers University: 2002-2007 Teaching Assistant