

ABSTRACT

Title of Document: VIEWS OF GOD AND EVIL: A
PERSPECTIVAL APPROACH TO THE
ARGUMENT FROM EVIL

Christopher William Thomas Bernard, Ph.D.,
2008

Directed By: Professor Allen Stairs, Department of
Philosophy, University of Maryland

A view referred to as “skeptical theism” has received much attention in recent discussions of the argument from evil against the existence of God. According to skeptical theism, humans are not in an epistemic position to make the inferences necessary for the evidential argument from evil to go through. In this dissertation, I defend the importance of individual variations in epistemic position to our evaluation of the argument from evil. Skeptical theists highlight the *inadequacy* of the human epistemic position to make the relevant judgments. I underscore the importance of individual *differences* in epistemic position—perspectival differences—to our evaluation of the argument from evil.* I argue that believers and nonbelievers may be epistemically justified in drawing different conclusions about God from the same, or similar, evidence because the evidence is judged from different epistemic perspectives. In particular, my discussion focuses on two perspectival factors which

* Some would prefer the term “worldview” over “perspective.” As I use the term, an epistemic perspective includes a worldview but includes factors that are broader than our beliefs, like practical interests and social factors.

have received relatively little attention by analytic philosophers of religion: practical interests and social influences.

Discussions of the argument from evil which place epistemic considerations at the forefront are typically about the evidential argument from evil. I contend that epistemic considerations deserve to be a central focus of the logical argument from evil as well. In this dissertation, I set out the logical argument from evil and Alvin Plantinga's freewill defense to it. I consider several objections to the freewill defense. I consider objections to the freewill defense aimed at Molinism, on which Plantinga's free will defense seems to depend. I conclude that one need not settle the debate about Molinism because there are versions of the freewill defense that don't depend on it which are successful.

I then turn to an epistemic objection raised by Daniel Howard-Snyder and John Hawthorne. I defend their objection from a criticism raised by William Rowe but offer my own criticism of their objection. My discussion of this epistemic objection calls attention to the importance of individual differences in epistemic situations, which I call "perspectival factors," to our assessments of the logical argument from evil.

I turn next to a discussion of two perspectival factors that haven't received sufficient attention in the analytic philosophy of religion—practical interests and the social production and distribution of knowledge. I contend that these perspectival factors may have an important, and often ignored, bearing on the epistemic justification of the religious beliefs, or lack thereof, of both believers and non-believers.

VIEWS OF GOD AND EVIL: A PERSPECTIVAL APPROACH TO THE
ARGUMENT FROM EVIL

By

Christopher William Thomas Bernard

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2008

Advisory Committee:
Professor Allen Stairs, Chair
Professor Samuel Kerstein
Professor Jerold Levinson
Professor Charles Manekin
Professor William Stuart

© Copyright by
Christopher William Thomas Bernard
2008

Dedication

For my children:

Christopher, Abigail, Jonah, and Jacob

With Love

Acknowledgements

Many people have contributed in some way to this dissertation. I would first like to thank my mom for doing the hard work of raising me and my siblings as a single parent. I am grateful to the faculty and students of the University of Maryland's philosophy department for my graduate education.

I would like to thank Louise Gilman for her logistical support which, among other things, included finding me quiet office space during return trips to Maryland. I greatly appreciate the work of John Kohl, who read the entire manuscript twice for spelling and grammatical errors and engaged with me in many hours of philosophical conversation. Thanks to Prof. Paul Pietroski for reading an earlier section of this dissertation. I particularly appreciate the time, conversations, and comments of Prof. Allen Stairs.

I would also like to thank my friends Kevin and Lillie Warren for the help and support they gave to my family while I was on extended trips to Maryland to finish this dissertation. Lastly, and most importantly, I would like to thank my wife, Mary. Her patience (indeed longsuffering), encouragement, and support have been truly amazing.

Table of Contents

Dedication.....	ii
Acknowledgements.....	iii
Table of Contents.....	iv
Introduction.....	1
1. Skeptical Theism.....	3
2. From Skepticism to Perspectivalism.....	15
3. The Plan Forward.....	21
Chapter 1: The Logical Argument from Evil.....	25
1. The Logical Argument from Evil.....	25
2. Two Kinds of Responses: Theodicy and Defense.....	28
3. The Freewill Defense.....	29
4. Logical Objection to Transworld Depravity.....	47
Chapter 2: The Freewill Defense and Molinism.....	51
1. The FWD with Molinism.....	53
2. The FWD without Molinism.....	85
Chapter 3: Transworld Depravity and Transworld Sanctity.....	97
1. The Epistemic Amendment and Transworld Sanctity.....	98
2. William Rowe’s Objection.....	107
3. Transworld Depravity and the Ontological Argument.....	112
4. The Person-relative nature of the reasonableness of believing transworld sanctity over transworld depravity.....	121
Chapter 4: Justification, Religious Belief, and Practical Interests.....	131
1. Intuitive Examples and Rejection of Easy Dismissals.....	138
2. Contextualism.....	151
3. Subject-Sensitive Invariantism.....	166
4. Religious Belief and Practical Interests.....	182
Chapter 5: Religious Belief in a Social World.....	190
1. The Nature of Testimony.....	192
2. The Epistemology of Testimony.....	204
3. Testimony and Religious Belief.....	222
Conclusion.....	241
Bibliography.....	245

Introduction

The argument from evil is, perhaps, the chief and most enduring argument against the existence of God.¹ According to the argument from evil, evil is evidence that God doesn't exist. Although philosophers refer to "*the* argument from evil," it isn't just one argument but rather a family of arguments.² This particular family tree has three branches roughly corresponding to the three basic forms of inference—deduction, induction, and abduction.

The deductive, or logical, argument is typically considered an *a priori* argument. Perhaps it is more accurate to say that the logical argument has an important, indeed central, *a priori* component that the other versions lack. According to the logical argument, the existence of evil and the existence of God are logically incompatible. We can infer from the existence of evil, not only that God doesn't exist, but that he *can't* exist. The inductive and abductive arguments are empirical, evidential, or *a posteriori*. According to these two types of arguments, the existence of evil provides probabilistic evidence that God doesn't exist.

One response by theists to the argument from evil which has received attention in the literature in recent years is referred to by critics as "skeptical theism."³ According to skeptical theists, we are not in an epistemic position to infer that God doesn't exist from the existence of evil. We are simply not in an epistemically appropriate situation to draw the kind of inferences that advocates of

¹ Like most discussions of the argument from evil, I am using the term "evil" to mean something like "suffering."

² The locution "the problem of evil" is a synonym or near synonym of "the argument from evil."

³ So far as I can tell, the term was first employed by Paul Draper in (Draper, 1996b).

the argument from evil want us to draw. Sometimes skeptical theists claim that our epistemic impoverishment here is due to a limitation of human cognitive capacities. Other times they assert that it is due to a mere lack of knowledge and leave it as an open question why we lack the relevant knowledge.

My project in this dissertation is to highlight and sketch the legitimate role of differences in the epistemic situations of theists and non-theists in their assessments of the argument from evil. I contend that theists and atheists have different epistemic perspectives in ways that are important to making the kinds of inferences involved in the argument from evil.⁴ Unlike skeptical theism, which emphasizes the *inadequacy* of the human epistemic position, I underscore the importance of individual *differences* in epistemic position to our evaluation of the argument from evil.

What I mean by “epistemic perspective” will shortly become clearer. However, I want to stress up front, and in the strongest possible terms, that the sort of perspectivalism I am urging is *not* relativism in any way, shape, or form. I want to make it absolutely clear that I reject relativism and associated theses. Relativism is the view that what is true is relative to persons, societies, etc. I deny this. There are objective truths that describe the objective world. These objective truths are knowable. There is a person-independent reality. We can and do have genuine knowledge about the objective, person-independent world, so relativism is false.

Knowledge, however, is person-relative. It is person-relative in a rather straightforward way: different people know different things. What is known isn’t person-independent but varies from person to person. For present purposes, let’s

⁴ Even though my discussion focuses on the argument from evil, my project has broader application than the argument from evil and the philosophy of religion.

think of knowledge as justified true belief. Truth, the third component of knowledge, is not person-relative. However, the first two components, belief and justification, do vary from person to person. People both believe different propositions and possess different evidence. So, one should not confuse the rather trivial claim that knowledge is person-relative in the ways I just described with the controversial claim, one which I deny, that truth is person-relative.

The rest of this introduction is divided into three sections. In this first section, I explain the nature of skeptical theism. This discussion serves several purposes. First, it provides a discussion of the relevant literature which helps to show how my thesis is intended to advance, or at least contribute to, the dialectic. Second, it serves to introduce some of the basic terms and concepts of the debate. In the second section of this introduction, I explain the perspectival approach to the religious epistemology of the argument from evil that I sketch in this dissertation. Lastly, I briefly outline the contents of the rest of the dissertation

1. Skeptical Theism

Skeptical theism is typically made to evidential arguments from evil (Draper, 1996b, p. 175). (I take it, though, that the kind of response to the logical argument from evil that Daniel Howard-Snyder and John Hawthorne ultimately urge, which I discuss in chapter 3, is at least a weak form of skeptical theism.) We will be able to get a clearer picture of it by looking briefly at two evidential arguments for evil and the replies of skeptical theists to them.

1.1 Skeptical Theism and Rowe's Inductive Argument

1.1.1 Rowe's Inductive Argument

Consider William Rowe's inductive argument from evil and the responses of skeptical theists to it. In "The Problem of Evil and Some Varieties of Atheism," Rowe sets out an inductive version of the argument from evil. Rowe's argument has two phases. The first stage is deductive. The second is inductive. The deductive portion of Rowe's argument is:

- (1) There exist instances of intense suffering which an omnipotent, omniscient being could have prevented without thereby losing some greater good or permitting some evil equally bad or worse.
- (2) An omniscient, wholly good being would prevent the occurrences of any intense suffering it could, unless it could not do so without thereby losing some greater good or permitting some evil equally bad or worse.
- (3) There does not exist an omnipotent, omniscient, wholly good being (Rowe, 1996b, p.2).

Rowe defends (2) by pointing out that it is in accord with basic moral principles shared by theists and non-theists alike (Rowe, 1996b, p.4).

He introduces the following case to defend (1):

Suppose in some distant forest lightning strikes a dead tree, resulting in a forest fire. In the fire a fawn is trapped, horribly burned, and lies in terrible agony for several days before death relieves its suffering. So far as we can see, the fawn's intense suffering is pointless. For there does not appear to be any greater good such that the prevention of the fawn's suffering would require either the loss of that good or the occurrence of an evil equally bad or worse (Rowe, 1996b, p.4).

He uses the case to argue for (1) in two different ways. First, he argues that, given our experience and background knowledge, it is more reasonable than not to think that the suffering of the fawn is not essential for bringing about some greater good or preventing an evil equally bad or worse. In other words, it is more reasonable than not to believe that the suffering of the fawn is pointless, so it is more reasonable than

not to think that there is at least one instance of suffering that is pointless. Thus, it is more reasonable than not to believe (1).

Secondly, he argues, elsewhere, by enumerative induction that God is not justified in permitting the fawn's suffering. For an instance of suffering to not be pointless there would have to be some greater good that justifies it. Rowe argues that, upon inspection of known goods, there is no good that justifies the suffering of the fawn:

Among the goods we know of, is any such that it has J [the property a good has just in case obtaining that good would justify an omnipotent, omniscient being in permitting E1 or E2]? It is obvious that many of these goods lack J. For when we contemplate them we see that their value is not high enough (e.g. my enjoyment on smelling a good cigar) to offset the evils in question. Other goods we know of may have a good deal of intrinsic value, perhaps even more value than E1 or E2 have disvalue. But we readily see that an omnipotent, omniscient being could obtain them without permitting E1 or E2 (Rowe, 1988, pp. 123-124).

From this he inductively infers there is no good that justifies the suffering of the fawn. He infers something about the set of all good from the set of known goods, namely, none justify the fawns suffering.

...we are justified in making this inference in the same way we are justified in making the many inferences we constantly make from the known to the unknown. All of us are constantly inferring from the A's we know of to the A's we don't know of. If we observe many A's and all of them are B's we are justified in believing that the A's we haven't seen are also B's (Rowe, 1988, pp. 123-124).

Rowe's argument for (1), then, is an argument from the existence of inscrutable evil to the existence of gratuitous evil.

1.1.2 Skeptical Replies to Rowe

Stephen Wykstra points out that Rowe's argument for (1) depends on an inference of the following sort: it appears that $\sim P$, therefore $\sim P$.⁵ Rowe argues that it *appears* that there is no good reason to let the fawn suffer, therefore there *is* no good reason to let the fawn suffer. Wykstra calls this type of inference a "noseeum inference" because it moves from the fact that we "no see" (hence, no-see-um) P to the fact that there is no P (Stephen John Wykstra, 1996, p. 126). Wykstra points out that noseeum inferences are sometimes legitimate and are sometimes illegitimate. He contends that when we examine relevant cases a pattern emerges regarding when it is legitimate to make noseeum inferences. Here is one relevant case.

Suppose you walked into a room and someone told you that there could be a flea egg in the room. You do not search the room with a magnifying glass or by crawling around on your knees. You just casually look around the room from your current position. Looking around the room, you do not see a flea egg. In this case, it is illegitimate to infer that there is not a flea egg in the room because you don't see one. It would be a mistake to make a noseeum inference under such circumstances. Now suppose you walk into the same room and the same person informs you that there is a normal sized adult elephant in the room. You look around and you do not see an elephant. In this case, it is legitimate to make a noseeum inference. The fact that you cannot see the elephant is reason to infer that there is no elephant.

Wykstra contends that the difference between these two cases helps us formulate a principle for when it is legitimate to make noseeum inferences. The difference between the two cases is one of epistemic position. In the case where it is

⁵ Some theists, most notably William Hasker, have challenged (2) in Rowe's argument (Hasker, 1992). They have argued that the existence of God and the existence of gratuitous, or unjustified, evil are *not* incompatible.

legitimate to make the inference, if there were an elephant in the room we would be in an epistemic position to know it. We would see the elephant if he were there. Since, we see no elephant, we are entitled to conclude that there is no elephant in the room. In the illegitimate case, if there were a flea egg in the room we would not be in an epistemic position to see it. One cannot tell just by casually glancing around the room that it lacks flea eggs.

This observation leads Wykstra to form the Condition of Reasonable Epistemic Access (CORNEA). According to CORNEA:

On the basis of cognized situation *s*, human *H* is entitled to claim ‘it appears that *p*’ only if it is reasonable for *H* to believe that, given her cognitive faculties and the use she has made of them, if *p* were not the case, *s* would likely be different than it is in some way discernable by her (Stephan J. Wykstra, 1990, p.152).

The basic idea of CORNEA is that one is entitled to infer $\sim p$ from that fact that it appears $\sim p$, only if it is reasonable to believe that we would see *p*, if *p* were true.⁶

Wykstra then argues that CORNEA is unmet in Rowe’s fawn case and, hence, we are not entitled to make the inference. His argument for the claim that CORNEA is unmet is that we have no reason to believe that if God did have an outweighing good to permit evil, we would be in a position to know it. In fact, we have reason to think otherwise. The enormous difference between God’s knowledge and our knowledge is reason to think that it would be unlikely that we would know God’s reasons for permitting evil, if he had any. Wykstra’s argument is based upon the analogy between a parent and an infant:

⁶ In this formulation of CORNEA, Wykstra is concerned about when it is legitimate to make “appearance claims.” Elsewhere he construes CORNEA slightly differently (Wykstra 1996, pp. 135-136). On this second construal, CORNEA is a necessary condition for making legitimate noseem inferences not for making “appearance claims.” There is no fundamental difference between the two. It just seems more natural to claim that we are entitled to say it appears there is no flea but are not entitled to infer there is no flea from this appearance.

...the outweighing good at issue is of a special sort: one proposed by the Creator of all that is, whose vision and wisdom are therefore somewhat greater than ours. How much greater? A modest proposal might be that his wisdom is to ours, roughly as an adult human's is to a one-month old infant's... that we should discern most of them seems about as likely as that a one-month should discern most of his parents' purpose for those pains they allow him to suffer — which is to say, it is not likely at all (Stephan J. Wykstra, 1990, pp. 155-156).

William Alston makes a similar, but not identical, skeptical response to Rowe's argument. Alston argues we are not justified in believing there is an instance of evil that God could have eliminated without thereby losing some greater good or causing another evil that is equally bad or worse (Alston, 1991a, p. 98). To be justified in believing such a claim one must be justified in excluding all the live possibilities for what the claim denies exists" (Alston, 1991a, p.102). We are not justified in excluding those possibilities for the following reason:

Even if we were fully entitled to dismiss all the alleged reasons for [God's] permitting suffering that have been suggested [by the most plausible theodicies], we would still have to consider whether there are further possibilities that are undreamt of in our theodicies. Why should we suppose that the theodicies thus far excogitated...exhaust the field?...It is surely reasonable to suppose that God, if such there be, has more tricks up His sleeve than we can envisage. Since it is in principle impossible for us to be justified in supposing that God does not have sufficient reasons for permitting [the fawn's suffering] that are unknown to us, and perhaps unknowable by us, no one can be justified in holding that God could have no reason for permitting the Bambi and Sue cases, or any other particular cases of suffering (Alston, 1991a, p. 119).

1.2 Skeptical Theism and Draper's Abductive Argument

1.2.1 Draper's Abductive Argument

In "Pain and Pleasure: An Evidential Problem for Theists," Paul Draper offers an abductive version of the argument from evil. Draper contends that there is a serious hypothesis that: (i) is inconsistent with theism and (ii) explains the facts about pain and pleasure better than theism. He calls this hypothesis the "Hypothesis of Indifference" (HI):

HI: neither the nature nor the condition of sentient beings on earth is the result of benevolent or malevolent actions performed by nonhuman persons (Draper, 1996a, p.13).

Draper summarizes the observations "O" that HI and theism are competing to explain:

O: a statement reporting both the observations one has made of humans and animals experiencing pain and pleasure and the testimony one has encountered concerning the observations others have made of sentient beings experiencing pain or pleasure (Draper, 1996a, p. 13).

He then makes explicit that he is arguing for the following conclusion:

C: P(O/HI) is much greater than P (O/theism) (Draper, 1996a, p.14).

Draper's argument for C centers on the biological utility of pain and pleasure.

Something is biologically useful if it contributes to the biological goals of the organism (survival and reproduction) in a way that is not biologically accidental (Draper, 1996a, p.15).

Draper thinks it will help make his argument for C clearer if he breaks the facts that O reports into three exhaustive categories:

O1: Moral agents experiencing pain or pleasure that we know to be biologically useful.

O2: Sentient beings that are not moral agents experiencing pain, or pleasure that we know to be biologically useful.

O3: Moral agents and sentient beings experiencing pain or pleasure that we do not know to be biologically useful (Draper, 1996a, p.16).

O is logically equivalent to (O1 & O2 & O3). So, it follows that, for some hypothesis h:

$$P(O/h) = P(O1 \& O2 \& O3/ h)$$

From probability calculus we know that:

$$P(O/h) = P(O1 \& O2 \& O3/ h) = P(O1/h) \times P(O2/h \& O1) \times P(O3/h \& O1 \& O2)$$

So, C is true if:

$$P(O1/HI) \times P(O2/HI \& O1) \times P(O3/HI \& O1 \& O2) > P(O1/T) \times P(O2/T \& O1) \times P(O3/T \& O1 \& O2)$$

Draper argues for C by arguing that the probability of each multiplicand on the left side of the equation is greater than the probability of the corresponding multiplicand on the right side of the equation.

Draper's argument that $P(O1/HI) > P(O1/T)$ runs as follows. O1 reports facts about moral agents, or humans, experiencing pain or pleasure that we know to be biologically useful. We know independently from O1 reports that humans are goal directed systems with parts that contribute to the biological goals of these systems. So, we have reason to expect, antecedently, that pain and pleasure would also systematically contribute to those biological goals.

Pain and pleasure, though, have intrinsic moral worth. This, he thinks, shows that $P(O1/T)$ is lower than the antecedent probability of O1. God is omnipotent and could design us to achieve our biological goals without pain. If he is also omnibenevolent he would do so, unless there was some greater good which justifies his permitting pain that he could not otherwise achieve. On the other hand, $P(O1/HI)$ is unchanged from the initial antecedent probability of O1.

Draper argues that $P(O2/HI \& O1) > P(O2/T \& O1)$ by pointing out that sentient beings that are not moral agents are, biologically, very similar to moral agents. This and O1 make it antecedently likely that merely sentient beings will

experience pain and pleasure that is biologically useful. He argues the same is not the case for $P(O2/T \ \& \ O1)$ because given theism and $O1$ we have reason to think that pain plays a moral role in the lives of humans. This cannot be said for the facts that $O2$ reports since merely sentient beings are not moral agents.

Draper gives a two-part argument for $P(O3/HI \ \& \ O1 \ \& \ O2) > P(O3/T \ \& \ O1 \ \& \ O2)$. First, he argues that we have more reason to expect sentient creatures to be happy on theism $\& \ O1 \ \& \ O2$ than on $HI \ \& \ O1 \ \& \ O2$. But $O3$ reports that humans and animals experience a lot of pain and are not as happy as one would expect. Further, he contends that on theism $\& \ O1 \ \& \ O2$, one would expect to discover a close connection between moral goods, like justice and virtue, and biologically gratuitous pain and pleasure, but we make no such discovery.

Second, he argues that we have more reason to expect, on $HI \ \& \ O1 \ \& \ O2$ than on theism $\& \ O1 \ \& \ O2$, that the fundamental role of pain and pleasure would be biological and that any biologically gratuitous pain or pleasure is a biological accident due to a lack of fine-tuning. He thinks this is just what we find.

So, he concludes, the probability of each multiplicand on the left side of the equation is greater than probability of the corresponding multiplicand on the right side of the equation and so, the left side of the equation is more likely than the right side. Consequently, C is true and provides abductive evidence against theism.

1.2.2 van Inwagen's Skeptical Reply to Draper

In "The Problem of Evil, the Problem of Air, and the Problem of Silence," Peter van Inwagen argues that Draper's argument that $P(O/HI) > P(O/theism)$ fails.

It does so, he thinks, because we are in no position to assign a value, or range of values, to the epistemic probability of O on theism.⁷ He argues that we don't really know what to expect if theism is true. If we cannot assign a value to $P(O/theism)$, then we cannot legitimately claim to know that $P(O/HI) > P(O/theism)$.

van Inwagen thinks we can show that we are in no position to determine $P(O/theism)$, if we can tell some story that entails the truth of theism and O and is "true for all anyone knows," this will undermine Draper's contention that $(O/HI) > P(O/theism)$. He explains why such a story would have that effect:

Suppose that Jane wishes to defend the character of Richard III and that she must contend with evidence that has convinced many people that Richard murdered the two princes in the Tower. Suppose that she proceeds by telling a story—that she does not claim to be true or even more probable than not—that accounts for the evidence that has come down to us, a story according to which Richard did not murder the princes. If my reaction to her story is "for all I know, that's true. I shouldn't be at all surprised if that's how things happened," I shall be less willing to accept a negative evaluation of Richard's character than I might otherwise have been (Van Inwagen, 1996, p.156).

Likewise, if we can tell stories that are true for all we know and entail the truth of both theism and O, this "should lead a person in our epistemic situation to refuse to make any judgment about the relation between the probabilities of ...[O] on theism and on HI" (Van Inwagen, 1996, p.156).

Van Inwagen offers one such story. Remember, he does not claim it is true or even likely true but merely that it is true for all we know. His story consists of three propositions:

- (1) Every possible world that contains higher-level sentient creatures either contains patterns of suffering morally equivalent to those recorded by ...[O], or else are massively irregular.
- (2) Some important intrinsic or extrinsic good depends on the existence of higher-level sentient creatures; this good is of sufficient magnitude that it outweighs the patterns of suffering recorded by ...[O].

⁷ van Inwagen uses the letter "S" rather than "O." For purposes of uniformity, I use "O" when discussing van Inwagen's argument.

(3) Being massively irregular is a defect in a world, a defect at least as great as the defect of containing patterns of suffering morally equivalent to those record by...[O] (Van Inwagen, 1996, p.157).

His understanding of “higher-level sentient creatures,” “morally equivalent patterns,” and “defect in a world” is pretty straight forward. By “massively irregular world” he means a world in which the laws of nature fail in some massive way. Occasional miracles do not make a massively irregular world. A massively irregular world is a world where:

God, by means of a continuous series of ubiquitous miracles, causes a planet inhabited by the same animal life as the actual earth to be a hedonic utopia. On this planet, fawns are (like Shadrach, Meshach, and Abednego) saved by angels when they are in danger of being burnt alive. Harmful parasites and microorganisms suffer immediate supernatural dissolution if they enter a higher animal’s body. Lambs are miraculously hidden from lions, and the lions are compensated for the resulting restriction on their diets by physically impossible falls of high protein manna. On this planet, either God created every species by a separate miracle, or else, although all living things evolved from a common ancestor, a hedonic utopia has existed at every stage of the evolutionary process. (The latter alternative implies that God has, by means of a vast and intricately coordinated sequence of supernatural adjustments to the machinery of nature, guided the evolutionary process in such a way as to compensate for the fact that a hedonic utopia exerts no selection pressure.) (Van Inwagen, 1996, pp. 158-159).

If the story consisting of (1), (2), and (3) is true for all we know, then for all we know God could only obtain some important good, a good that outweighs the patterns of suffering recorded by O, only by permitting patterns of suffering morally equivalent to O or by creating a massively irregular world, which is as morally bad as O. So, for all we know, God had to permit suffering if he wanted to obtain some good.

His argument that the story is true for all we know is, in short, that we lack the cognitive abilities, or at least don’t know whether we possess them, to judge the possibility of the story being true. He contends that in order to determine the

possibility of the story, we have to be able to make highly abstract and removed modal and value judgments.

Consider (1). In order to know that it not true, for all we know, that every possible world that contains higher-level sentient creatures either contains patterns of suffering morally equivalent to those recorded by O, or else are massively irregular, we need to know whether it is possible for a world to contain higher-level sentient creatures and lack suffering while at the same time not being massively irregular. We have to be able to make highly abstract modal judgments.

The same holds for (2). Is it possible that some important intrinsic or extrinsic good depends on the existence of higher-level sentient creatures and this good is of sufficient magnitude that it outweighs the patterns of suffering recorded by O? In order to determine this, we need to be able to determine whether it is possible for some particular good G to exist independently of higher-level sentient creatures.

(3) makes a value judgment rather than a modal judgment. It tells us something about the value, or disvalue, of the defect of a world being massively irregular. It tells us that being massively irregular might have the same disvalue as the pattern of suffering described by O.

van Inwagen argues that our faculty for making modal and value judgments is suited for surviving in our immediate environment and we have no reason to think they are reliable very much beyond that context. Our modal judgments are probably reliable, for example, when they tell us whether a table could have been on the other side of the room but they should not necessarily be trusted when they tell us that there could have been transparent iron. In order to judge the possibility of (1) or (2), we

have to make highly abstract modal judgments akin to the case of transparent iron rather than the location of a table:

...assuming there are “modal facts of the matter,” why should we assume that God or evolution or social training has given us access to modal facts knowledge of which is of no interest to anyone but the metaphysician? God or evolution has provided us with a capacity for making judgments about the size and distance that is very useful in hunting mammoths and driving cars, but which is of no use at all in astronomy. It seems an analogous restriction applies to our capacity for making modal judgments. How can we be sure that an analogous restriction does not apply to our capacity to make modal judgments? My position is that we cannot be sure... (Van Inwagen, 1996, p.162).

van Inwagen thinks the same considerations apply to our ability to make the relevant sort of value judgments. In sum, van Inwagen thinks that we are not in an epistemic position to think that $P(O/HI) > P(O/theism)$.

Based on the preceding discussion, we can summarize skeptical theism thusly:

Skeptical theism (ST): Given the limitations on human cognitive faculties and/or knowledge, we are in no epistemic position to infer that God does not exist based on the existence of suffering.

2. From Skepticism to Perspectivalism

There are some merits to ST. First, if theism is true, ST no doubt reflects an important truth regarding the limits of human understanding of God. Second, it employs the more modest strategy of replying to the argument from evil without claiming to know why God actually permits suffering.⁸ But ST also has important failings.

In particular, skeptical theism and the assumptions underlying it may be too strong, even for theists. The claims of skeptical theists are so strong that it threatens to undermine the whole project of theology, in general, as well as even modest

⁸ I introduce the distinction between theodicy in defense in chapter 1.

attempts to provide positive rational support for theism.⁹ It even threatens to undermine negative defenses of theism, like Plantinga’s freewill defense. van Inwagen’s skepticism of our ability to make abstract modal judgments that are “only of interest to metaphysicians” seems to undermine the modal claim at the heart of the freewill defense, which I will discuss in chapters 1 and 2.

So, while skeptical theism has merit, I want to suggest a different approach — a perspectival approach.¹⁰ According to the perspectival approach that I am urging, whether the argument from evil justifies atheism, depends on the epistemic perspective of individual knowers. The force of the argument from evil legitimately varies from person to person and from time to time, depending on the epistemic situation of the person assessing the argument.

Consequently, the argument from evil *may* justify an atheist’s disbelief in God and at the same times it *may not* be a defeater for a theist’s belief in God.¹¹ It depends on the epistemic situation of the person evaluating the argument. If the perspectival approach is right, then the fact that the argument from evil may have different epistemic force for theists and atheists cannot always be chalked up to epistemically illegitimate differences between people, like bias or abnormal ignorance. Sometimes there are good epistemic reasons, perspectival reasons, which explain why atheists

⁹ Sometimes a distinction is made between positive apologetics and negative apologetics. Positive apologetics is the project of giving positive reasons for believing that God exists while negative apologetics is the project of responding critiques of theism. While many theists criticize the project of natural theology, many still think that arguments for God can provide *some* rational support for believing in God.

¹⁰ It is perhaps more accurate to say that what I am advocating is not a comprehensive approach called “perspectivalism” but rather that others take much more seriously the perspectival character of arguments for and against the existence of God.

¹¹ I briefly explain what epistemic defeaters are in chapter 5.

and theists are entitled to draw different, or at least weaker, conclusions from the existence of suffering.

This approach may leave philosophical partisans somewhat unsatisfied. Those committed to what Robert Nozick calls “coercive philosophy” prefer conclusions which end with one side or the other being labeled “irrational” (Nozick, 1981, pp. 4-8). While theists and atheists can be, and often are irrational, not all theological disagreement can be boiled down to either cognitive bias and deficiency or an impoverished store of knowledge. They can sometimes be explained by legitimate differences in epistemic situation.

One might object at this point that my claim is obvious, even trivial. Of course people in different epistemic situations vary in what they are justified in believing, including whether they believe in God. My reply is two-fold. First, too many philosophers don’t *really* think this is obvious, particularly when it comes to religion. When it comes to God, many philosophers are willing to grant what we might call “weak perspectivalism.”

According to weak perspectivalism about religion, people in epistemically anemic situations (e.g. like ancient Israelites or other “primitive” peoples) or with cognitive limitations (like small children) might be justified in believing in God¹² but that is only because their environment is so epistemically impoverished. Atheist philosopher Kai Nielsen conveys this idea:

...*if* people have a good scientific and philosophical education...then they should see that it is irrational to believe in God...That is, it is irrational for someone who has a good scientific and philosophical education. And I point out to you that I don’t mean to say by this that I think that I’m more rational than Professor Hick or some religious person, because I remind you that

¹² Particularly given an internalist and deontological notion of justification.

rational people can have irrational beliefs. And what I'm maintaining is that belief in God for people in the twentieth century, not people at all times and at all places, with a good scientific and philosophical education is irrational (Nielsen, 1993, pp. 48-49).¹³

Let's distinguish weak perspectivalism from moderate perspectivalism.

According to moderate perspectivalism about religion, people with normally functioning cognitive apparatuses and with no abnormal deficiency knowledge can be rational in believing in God. Contrary to Nielsen, theists can have a good scientific and philosophical education and still rationally believe in God. Atheists and theists in normal epistemic situations can legitimately evaluate the strength of the argument from evil differently. The burden of this dissertation is to explain, at least partly, how.

Second, the particular perspectival factors that I call attention to in this dissertation are both interesting and important. An exploration of some of the specific ways in which people's different epistemic situations legitimately affect their assessments of the evidence for and against the existence of God is anything but trivial. Rather, it focuses much needed attention on some of the fine-grained epistemic differences between people that should be included in sophisticated discussions about the epistemology of religion and the theory of knowledge in general.

My thesis is that differences in epistemic perspectives explain how theists and atheist can legitimately view the argument from evil differently. What is an epistemic

¹³ USC philosopher Dallas Willard's response to Nielsen seems entirely appropriate:

I begin with the plea of "not guilty" to Professor Nielsen's charge that my belief in God is an irrational belief. I do live in the twentieth century, I have been given a passably good philosophical and scientific education, and I have thought carefully about the existence of God...it seems to me, at least, that my belief in God is not irrational, though, of course, I myself may still be a wildly irrational person (Willard, 1993, p. 197).

perspective? I don't have a precise definition. I can, though, list the sort of elements that comprise what I think of as an epistemic perspective. This should give the reader sufficient idea of what I mean by the term.

Roughly, then, an epistemic perspective includes the epistemically relevant factors which vary from person to person and are such that they have a legitimate bearing on one's epistemic judgments. An epistemic perspective is comprised of things like background beliefs, experiences, and values, among others.¹⁴ How background beliefs and experiences bear on, for example, probability and modal judgments is relatively straightforward. Consider, by way of illustration, a few very simple and uncontroversial examples.

Suppose you are playing Spades with a standard fifty-two card deck and you wanted to know the probability that you would be dealt the ace of spades from a randomized deck. You conclude the odds are one in four (there are four players in Spades and the whole deck is dealt, so each gets one quarter of the deck). Suppose that you were the first person to the left of the dealer and saw the dealer mistakenly flash the top card. No one besides you saw that the card was an ace of spades. You would now judge your chances of getting the ace of spades at unity while the others would judge it to be one in four. Your probability judgment differs from that of your fellow card players because you have different experiences and background beliefs than they do.

¹⁴ This is just a partial list. I argue, for example, in chapters four and five that practical interests and one's location in social space are part of, or at least has a bearing on, one's epistemic perspective. If a worldview is the set of one's belief, then one's epistemic perspective includes but is not limited to one's worldview.

Modal judgments also depend on the knower's epistemic situation. Ancient Babylonians apparently believed that it was possible for Hesperus to outlast Phosphorus. We now know that Hesperus *is* Phosphorus and so it is not possible for one to outlast the other. As our background beliefs changed, so did our modal judgments.

The basic idea of the perspectival approach that I am advocating has relatively straightforward application to evidential arguments for evil. Evidential arguments from evil are inductive and abductive arguments. Abductive and inductive inferences are probabilistic inferences. The concept of probability involved is subjective probability, not the classical conception or the relative-frequency conception. Subjective probabilities are often thought to be a sort of credence function. They are a measure of how likely something is, given a set of background beliefs. Since theists and atheists have different background beliefs, it shouldn't be surprising that their probability judgments and hence, views of the quality of abductive inferences differ.

The probability of God's existence is judged against the backdrop of the atheist's and the theist's *total evidence*. Since, theists and atheists have different sets of background beliefs, it should not be surprising that the epistemic force of evidential argument from evil differs between them.

As Alvin Plantinga points out, one's total evidence includes not only other propositions that one believes but also non-propositional reasons for believing (Plantinga, 1996, p.88). Suppose, for the sake of argument, that some theists have had religious experiences which actually justify their belief in God.¹⁵ In that case, the

¹⁵ I understand this is a controversial assumption, to say the least, but my point is merely to illustrate how the total evidence for believers and non-believers might be different in ways other than having

believer's religious experience plays a role in his assessment of the success of the evidential argument from evil. Anyway, it is a relevant epistemic difference between some believers and non-believers.

Much of the discussion about the evidential argument from evil, then, concerns what humans are in an epistemic position to know *simpliciter*. The skeptical theist's concern over the epistemic perspective of humans in general is important and needs to be taken seriously. However, philosophers of religion should also take seriously that we humans aren't always in the same epistemic position. As I mentioned above, if, for example, religious experience is a genuine source of knowledge about God, then it isn't impossible that some people have had religious experiences may be in a different epistemic position from those who haven't in important and relevant ways. Epistemology is too complicated to merely ask what "we" are in a position to know. In this dissertation, I call attention to the relevance of individual epistemic situations and factors to the justification of one's religious belief or lack thereof. Here's how I plan to do this.

3. The Plan Forward

It is easier to see how perspectival differences between believers and non-believers have an epistemic bearing on evidential arguments from evil but not as easy to see this in the case of the logical argument from evil. It is for this reason I focus on the logical argument rather than evidential arguments. Furthermore, it is relatively straightforward to see how factors like background beliefs legitimately influence one's assessment of the argument from evil. In this dissertation, I intend to show how

different propositional evidence. For a defense of the claim that religious experience can justify belief in God see (Alston, 1991b).

epistemic perspectives bear on the logical argument from evil and how less obvious factors than background beliefs, like practical interests and location in social space, legitimately affect one's epistemic assessment of the argument from evil. So, there are two ways in which I take on a more difficult project—by focusing on the logical argument and by exploring less discussed perspectival factors, like practical interests and social considerations.

In chapter one, I set out the logical argument from evil, particular as expressed by J.L. Mackie, and Alvin Plantinga's freewill defense. I present Plantinga's negative defense, i.e. his argument that the atheist has failed to show a contradiction between the existence of God and the existence of evil. I then present his positive attempt to demonstrate that God and evil are logically consistent. At the end, I mention three kinds of objections to Plantinga's freewill defense. I consider and reject the first kind but leave the other two objections for chapters 2 and 3 respectively.

Part of the reason for discussing these objections is to make my perspectival claim more plausible. After all, if the existence of God and evil are straightforwardly inconsistent, then it is hard to see how the knowing subject in different epistemic situations can legitimately differ in their assessment of the logical argument from evil. Of course, along the way, I weigh in on the various debates surrounding the logical argument.

In chapter two, I discuss one of the most influential set of objections to the freewill defense, namely, that it assumes Molinism, the controversial claim that there are true counterfactuals of creaturely freedom. First, I set out and explain Molinism in some detail. Since counterfactuals play such a central role in the freewill defense

and in Molinism, I take the space necessary to give a reasonably full discussion of the semantics of counterfactuals. I also explain a distinction that Molinists take to be critically important—the distinction between weak actualization and strong actualization. I then set out and assess the main objections to Molinism in light of this detailed groundwork. I conclude that arguments against Molinism yield mixed results. In light of this, I then consider how the freewill defense fares if there are no true counterfactuals of creaturely freedom. I conclude that whatever we think about Molinism, a version of the freewill defense that doesn't assume it can still make it reasonable to reject Mackie's logical argument from evil

In chapter three, I discuss an epistemic objection to the freewill defense by John Hawthorne and Daniel Howard-Snyder and reply to it. First, I lay out Howard-Snyder's and Hawthorne's argument that Plantinga has failed *to demonstrate* that it is possible for God and evil to co-exist because he has failed to meet a minimum epistemic standard for such a demonstration. I then defend Howard-Snyder and Hawthorne's argument from a criticism made by William Rowe. I argue that they use the wrong epistemic standard. When we consider other more appropriate epistemic standards for judging such arguments, it becomes apparent that different people, in different epistemic situations, can legitimately make different judgments (e.g. abstract case modal judgments). Part of my argument involves making an analogy between the logical argument from evil and modal ontological arguments.

Having shown how different epistemic standards lend themselves to a more perspectival approach, even to the logical argument from evil, I shift gears somewhat. In the next two chapters I consider two perspectival factors that have generally

received less attention than factors like differences in background beliefs and the kinds of holistic and coherency considerations such differences have in justifying our making different judgments.

In chapter four I contend that differences in the practical interests of believers and non-believers have a legitimate epistemic bearing on how they decide religious questions. In this chapter, I relate the debate regarding what some have called “pragmatic encroachment in epistemology” and apply the results of this discussion to religious beliefs.

In chapter five, I discuss the relevance of social epistemology to the justification of religious belief and non-belief. This chapter contains a preliminary discussion. For reasons that I express there, I offer a less a detailed account than a programmatic approach to the relevance of social epistemology. In chapter five, I sketch and outline how social differences between believers and non-believers might also be an important influence on epistemic considerations.

Chapter 1: The Logical Argument from Evil and the Freewill Defense

Some atheist philosophers contend that the existence of evil creates a logical problem for theists. As H.J. McCloskey puts it:

Evil is a problem for the theist in that a contradiction is involved in the fact of evil on the one hand, and the belief in the omnipotence and perfection of God on the other. God cannot be both all-powerful and perfectly good if evil is real (McCloskey, 1997, p.203).

Alvin Plantinga contends that atheists have failed to demonstrate such a contradiction. He goes even further and offers what he thinks is good reason to think that the existence of God and the existence of evil are logically compatible—the freewill defense. In this chapter, I lay the foundation for the rest of the dissertation by setting out and explaining the logical argument from evil and some of the standard responses to it. In particular, I focus on Plantinga’s freewill defense.

1. The Logical Argument from Evil

The contemporary discussion of the logical argument from evil centers on J.L. Mackie's 1955 paper, "Evil and Omnipotence" (Mackie, 1955). Mackie seeks to forestall a certain move by the theist. He points out that the arguments in support of the existence of God have been effectively undermined but worries that the theist can concede this and still “retain all that is essential to his position, by holding that God's existence is known in some other, non-rational way” (Mackie, 1955, p.200). Mackie seeks to block the theist from taking this position by arguing that: “...it can be shown, not that religious beliefs lack rational support, but that they are positively irrational, that several parts of the essential theological doctrine are inconsistent with one

another" (Mackie, 1955, p.200). If Mackie's argument is correct, he thinks the theist is in the position of not merely believing something that hasn't been proved but of believing something that has been disproved (Mackie, 1955, p.200).

He advances this ambitious goal by arguing for the strong claim that it is *logically* impossible for the God of classical theism to coexist with evil:

In its simplest form the problem is this: God is omnipotent; God is wholly good; and yet evil exists. There seems to be some contradiction between these three propositions, so that if any two of them were true the third would be false. But at the same time all three are essential parts of most theological positions: the theologian, it seems, at once *must* adhere and *cannot consistently* adhere to all three (Mackie, 1955, p.200).

Mackie's claim is that it is impossible for God and evil to both exist. Since Kripke, the standard way to understand modal claims is in terms of possible worlds (Kripke, 1963b). For example, the statement "it is impossible for an object to be both a cube and a sphere" means that there is no possible world that includes an object that is both a cube and a sphere. The sentence "it is possible that unicorns exist" means that there is at least one possible world that contains unicorns. The sentence "necessarily God exists" means that God exists in every possible world. If we use possible world semantics to understand Mackie's claim, Mackie is making the very strong claim that there are no possible worlds that include both God and evil.

More formally, Mackie's argument amounts to this:

Let "G" stand for God exists and is omnipotent, omniscient, and wholly good.
Let "E" stand for evil exists.

- (1) $\sim\Diamond(G\&E)$
- (2) E
- (3) $\therefore \sim G$

The argument above is formally valid. Whether it is sound depends on the truth of the premises. (2) is uncontroversial. Evil exists in the actual world. Premise (1), then, is doing all the heavy lifting. Mackie's argument for (1) is stated a bit awkwardly:

...we need some additional premises, or perhaps some quasi-logical rules connecting the terms 'good', 'evil', and 'omnipotent'. These additional principles are that good is opposed to evil, in such a way that a good thing always eliminates evil as far as it can, and that there are no limits to what an omnipotent thing can do (Mackie, 1955, pp. 200-201).

His use of the term "quasi-logical rules" clouds his argument a bit, but what he seems to mean is that when one properly explicates the words "good," "evil," and "omnipotent" two principles or premises emerge:

- (3) A good thing always eliminates evil as far as it can.
- (4) There are no limits to what an omnipotent thing can do.

From (3) and (4) Mackie infers:

- (5) A good omnipotent thing eliminates evil completely.

After all, if there were no limits to what God could do, he *could* eliminate evil. If he were perfectly good, he *would* eliminate evil. So, if God exists, evil can't exist and if evil exists, God can't exist. And this is precisely what (1) claims.

The logical argument from evil is typically classified as an *a priori* argument because (1) is based entirely on *a priori* reasoning. Mackie argues that (1) follows entirely from the meaning of the words "good," "evil," and "omnipotent." Premise (2) is an empirical premise. We know it is true, given our experience of the world. Premise (1), though, is the contentious claim and it is supported entirely *a priori*.

2. *Two Kinds of Responses: Theodicy and Defense*

The theist has two general strategies he can employ to counter the argument from evil: theodicy and defense. A theodicy is any response to the problem of evil that purports to tell us why God permits evil. A theodicy makes a claim about why God allows evil in the actual world. A defense, on the other hand, is any response that attempts to undermine the argument from evil *without* trying to tell us why God actually permits evil. In the context of the logical argument from evil, a defense merely tries to give a logically possible reason God permits suffering. If the theist offers a theodicy, he is maintaining that God does, in fact, have a good reason to permit evil and that reason is spelled out by the theodicy.

The difference between a theodicy and a defense is critical when it comes to replying to Mackie's argument. Mackie's claim is a claim about what is logically possible. Mackie makes the very strong claim that there is no possible world where both God and evil exist. Strictly speaking, all that is necessary to undermine this strong claim is to demonstrate that there is at least one logically possible world where both God and evil exist.

A defense is a far more modest kind of reply than a theodicy. The advantage of a more modest reply is that it is easier to establish. The weaker a claim is, the less it needs by way of argument to establish it. There are several disadvantages to a weak reply. First, if it is very narrowly tailored to the objection, even if it is successful, there is the possibility that a small modification to the objection will suffice to undermine the reply. Second, a defense is less theoretically and psychologically satisfying.

A theodicy, if successful, doesn't just undermine the argument from evil: it answers it. A successful theodicy has the epistemic and psychological merits of an explanation. It tells us something substantial not merely formal. Because a theodicy is more ambitious, it is more difficult to establish. In fact, it is largely the perceived failure of the project of theodicy *by theists* that has shifted the debate from offering theodicies to offering defenses.

3. The Freewill Defense

Alvin Plantinga, among others, has set aside the project of theodicy in favor of offering a defense (Plantinga, 1977, pp. 7-64). Plantinga's response to the logical argument from evil has two stages. First, he argues negatively that Mackie and other atheologists have yet to *show* that God and evil are logically incompatible. They say they claim both can't exist but haven't demonstrated it. Second, and more importantly, he endeavors to actually show that they are logically compatible. Let's call the former project the negative project and the latter project the positive project.

3.1 Plantinga's Negative Project

Mackie's argument amounts to the claim that there is no possible world where both God and evil exist. How can he demonstrate this? Obviously, possible worlds are not the sort of things we can empirically examine. Even if something like David Lewis's account of possible worlds is right and they are concrete objects, possible worlds are still causally isolated (D. Lewis, 1986, p.78). So we can't empirically examine other worlds from our own. If examining possible worlds just involves

abstract reasoning, there is still the problem that Mackie's claim is a universal claim. He is telling us something about *every* possible world. Of course, Mackie doesn't need to examine each world individually to show that there is no possible world where both God and evil exist. He only needs to show that the propositions "God exists" and "evil exists" are contradictory. If so, there is no possible world where both exist.

So has Mackie demonstrated a contradiction? Plantinga doesn't think so. He argues that if there is an inconsistency here, it isn't readily apparent. He distinguishes between three kinds of contradiction: explicit, formal, and implicit. An *explicit contradiction* is just a standard form contradiction ($A \& \sim A$). *Prima facie*, (G&E) is not a standard form contradiction. Of course, Mackie probably doesn't think that the theist's inconsistency amounts to holding an explicit contradiction. However, Plantinga wants to be precise and clear about exactly what sort of contradiction the theist is supposed to be involved in.

Perhaps, he suggests, while theists don't hold to an explicit contradiction, they embrace a set of beliefs that is *formally inconsistent*. A set of sentences is formally inconsistent if one can derive an explicit contradiction from the set by making valid deductive inferences. For example, from:

- (a) *John is in New York*
- (b) *Mary is not in New York* and
- (c) *If John is in New York then Mary is in New York,*

we can derive an explicit contradiction; namely:

- (d) *Mary is in New York and Mary is not in New York.*

If we unpack G and E we get four propositions:

- (6) *God is omnipotent.*
- (7) *God is omnibenevolent*
- (8) *God exists.*
- (9) *Evil exists.*

Is there a series of deductive inferences we can make from these four sentences that results in an explicit contradiction? It doesn't appear so. This set of sentences seems to lack the necessary material to derive an explicit contradiction.

Plantinga considers a third kind of contradiction—an *implicit contradiction*. A set of sentences S is implicitly contradictory just in case there is a necessary truth p, such that the addition of p to S results in S being a formally contradictory set. Why does p have to be a necessary truth? If there were no constraints on the propositions one could add to set S, then turning S into a formally inconsistent set would be trivial. We could get a contradiction, for example, if we added:

- (10) *God is not omnipotent.*

The conjunction of (6) and (10) is an explicit contradiction. The problem is that we can't add just any proposition. To understand the relevant constraints, consider this set of two sentences:

- (e) *John is a bachelor.*
- (f) *John is married.*

Clearly, these two sentences contradict each other. But they are neither explicitly nor formally contradictory. That is to say, neither sentence is a standard form contradiction nor can a standard form contradiction be derived from them by a series of deductive inferences. They are, though, contradictory because there is a

necessary truth that when added to this set results in the set being formally contradictory, namely:

(g) *If someone is a bachelor, then he is not married.*

Thus, if John is a bachelor and all bachelors are unmarried, then John is unmarried. This is inconsistent with John being married. This sentence is necessarily true in virtue of the meaning of the words “bachelor” and “married.”

Is the relevant set of propositions (6)-(9) above, implicitly contradictory? It is if there is one or more necessary truths that when added to that set yield a formally inconsistent set. Are there such necessary truths? Mackie offers two candidates:

(11) *There are no limits to what an omnipotent being can do*

and

(12) *A good thing always eliminates evil as far as it can.*

Plantinga examines each in turn. Consider Mackie's first candidate. It's true that an omnipotent being is a being whose *power* has no limits. Mackie's claim, though, is that an omnipotent being has no limits on what he can do. There is an important difference here. Some limits on what a being can do are not due to a lack of power.

Consider a mundane example of the distinction between limits on acting due to a lack of power and limits on acting *per se*. The United States government has the power to kill Osama bin Laden but they lack the knowledge of where he is. The limit the U.S. government bumps up against here is a genuine limit but not a limit of power. After all, if they knew exactly when he would be at particular place, a bombing strike could kill him. Now, it might be argued that all this is beside the

point because God is also essentially omnipotent. While it's true that this sort of example doesn't apply to the case of God, the example still has value in demonstrating the bare conceptual difference between limits due to a lack of power and limits *per se*.

Consider an example involving an omnipotent being. Most theologians and theistic philosophers concede that even God cannot make a square circle, a married bachelor, or that $(A \& \sim A)$ is true. His lack of ability to perform these tasks is not due to a lack of power but because these tasks are incoherent—they're pseudo-tasks. So there are at least some limits on what an omnipotent being can do that do not involve a lack of power. It is not true that if God just had a little more power he could make contradictions true.

So, the first candidate doesn't seem to be a necessary truth. Perhaps Mackie can avoid this sort of response by refining his claim like this:

(11a) *There are no non-logical limits to what an omnipotent being can do.*

Is this modified candidate a necessary truth? It certainly looks to be a necessary truth given the meaning of the relevant terms. Plantinga takes a more modest view and concedes that it is at least “initially plausible to suppose” that it is a necessary truth (Plantinga, 1977, p.18). He argues, though, that even if this proposition is a necessary truth, it does not yield an implicitly inconsistent set because it is possible there is some logical limit preventing God from preventing evil.

Now consider Mackie's second candidate:

(12) *A good thing always eliminates evil as far as it can.*

Is this proposition a necessary truth? Plantinga argues that it is not. There are times when a good person doesn't eliminate an evil that he can. Consider two friends: Alvin and Paul. Suppose Paul goes for a drive on a cold winter night and runs out of gas on a deserted road. Paul is cold, miserable, and frightened. Alvin is at home warm and comfortable. Alvin has the power to alleviate Paul's pain; he has an emergency gas can full of gas at home and a car in his garage. Does the fact that Alvin isn't eliminating Paul's pain mean that Alvin is not a good person? After all, he *can* eliminate the evil but doesn't.

The obvious reply for Mackie to make is that Alvin can't help Paul because he doesn't know Paul is in need of help. In some bare sense of the predicate "is able," Alvin is able to help but it is also perfectly sensible to claim that Alvin can't help because of his lack of knowledge. Further, classical theism includes the view that God is omniscient. Rather than get bogged down in an analysis of the relevant sense of "is able," Plantinga simply alters Mackie's original candidate thusly:

(12a) *A good thing always eliminates evil it knows about as far as it can.*

Plantinga doesn't think this proposition is necessary either. Consider two counterexamples. First, suppose, as in the previous example, that Paul is stranded on a deserted road on a cold winter night. Suppose this time that Alvin knows Paul is stranded. Alvin is able to help Paul. But suppose this time that Alvin's friend Peter is in the same situation as Paul and that Alvin, Paul and Peter are far enough apart that Alvin can only help one of his friends. Alvin is able to help Paul and knows of Paul's suffering. Alvin is able to help Peter and knows of Peter's suffering. But he is not able to help both.

So, there are two evils. Individually considered, neither is such that Alvin can't eliminate it; he just can't eliminate both. Given this situation we wouldn't judge Alvin to be less morally praiseworthy because there is some evil that he can eliminate but doesn't (Plantinga, 1977, p.19). There is, though, some evil he can't eliminate, namely the conjunction of the two. The reason we hold him blameless is precisely because he can't eliminate both.

Second, consider a surgeon that can eliminate the mild suffering of a knee scrape by amputating a leg. There is pain, or evil, that she can eliminate but only by causing the greater evil of an amputated leg. We could put this point differently: there is an evil the surgeon can eliminate but only by eliminating an outweighing good state of affairs (i.e. having two functioning legs).

Plantinga defines an outweighing good state of affairs thusly: A good state of affairs G outweighs an evil state of affairs E, if the conjunctive state of affairs (G&E) is itself a good state of affairs. If G&E is a good state of affairs, then contrary to Mackie's principle, a good thing does not necessarily eliminate an evil E, because eliminating E also eliminates the good state of affairs G&E (Plantinga, 1977, p.19-20). If the surgeon refused to alleviate the pain of the scrape, we wouldn't consider her immoral because eliminating the evil would bring about a greater evil or eliminate an outweighing good state of affairs. So, (12a) isn't a necessary truth.

But what if eliminating E eliminated G&E but not G? What if G and E are not causally or physically connected? Consider the conjunction: Jones discovers a cure for cancer and Smith has a toothache. It is true that if we eliminate Smith's toothache by modern dental techniques, we would no longer have the state of affairs

that includes the cure for cancer *and* Smith's toothache but we would still have the cure for cancer. Later on, though, Plantinga says that an omnipotent being might be in a situation such that a good state of affairs G includes E such that it is impossible that G obtain and E not obtain (Plantinga, 1977, p. 22). So, while at this stage he seems to be talking about seemingly arbitrarily connected states of affairs, he later refers to goods that cannot be obtained without at least allowing for the possibility of suffering.

To help Mackie avoid this sort of counterexample, Plantinga proposes modification to (12a). He does so by introducing the notion of *properly eliminating* an evil. An evil is properly eliminated when it is eliminated without bringing about a greater evil or without eliminating a good state of affairs that outweighs the evil. So the modified candidate looks like this:

(12b) *A good thing eliminates every evil that it knows about and can properly eliminate.*

Plantinga argues, though, that even this modified form of (12) is not necessarily true. It survives the second objection to this candidate necessary truth but it still falls afoul of the first objection. By including the notion of properly eliminating an evil, (12b) can account for our intuitions regarding the moral status of a person who fails to eliminate an evil because doing so would also eliminate an outweighing good. It cannot, though, account for cases like Alvin and his stranded friends Paul and Peter, even if the example meets the added condition regarding proper elimination.

Plantinga offers the following modified example. Suppose you come across two stranded rock climbers as a violent storm approaches. If they don't get down

before the storm hits, they will fall to their deaths. Suppose that, for whatever reason, they need help getting down and you only have enough time to save one of them before the storm hits. You can *properly eliminate* either evil. You can't eliminate both. Suppose you help one and the other dies. According to (12b), you are not perfectly good because you *can* properly eliminate the death of the person who actually died. But this is counterintuitive (Plantinga, 1977, p.20).

The obvious objection to Plantinga's counterexample is that an omnipotent being should be able to save both rock climbers. If we combine (11a) and (12b) we get something like:

(13) *An omnipotent, omniscient, wholly good being can properly eliminate every evil.*

But is (13) necessarily true? Plantinga doesn't think it is obvious. But more to the point of Plantinga's negative project, he doesn't think anyone has demonstrated it to be necessarily true. Until that is shown, it can't be said that the atheologist has demonstrated that God and evil are incompatible.

3.2 Plantinga's Positive Project

Thus far Plantinga has argued negatively that the atheologist has yet to make a complete case. But just because the atheologist hasn't *demonstrated* an inconsistency, doesn't mean God and evil are compatible. Maybe they are inconsistent and atheologists just haven't proven it yet or, for whatever reason, the inconsistency will never actually be demonstrated. Plantinga isn't content to leave the matter in this

state. In the second stage of his response to the logical argument, he endeavors to show, positively, that it is logically possible for both God and evil to exist.

Plantinga's strategy for demonstrating $\diamond(G \& E)$ is to give a model where G and E are both true. He employs the following procedure: for any propositions P and Q, (P&Q) is possible if there is a third proposition R such that R is possible, P and R are consistent, and P and R together entail Q (Plantinga, 1977, p.26). So, for example, suppose someone claims both that Sally hasn't eaten any food in the last week and that Sally *has* gained weight in the last week. We might think these two claims are inconsistent. We can use the procedure Plantinga suggests to show that these two claims are consistent. So, our goal is to show that the following two propositions are consistent:

P: Sally hasn't eaten any food in the last week.

Q: Sally has gained weight in the last week.

We can do so, Plantinga says if there is some third proposition that is consistent with Sally not eating any food last week and together with this entails that Sally gained weight last week. Consider the following artificial example:

R: Sally drank enough high calorie liquids last week to cause her to gain weight.

R is such that it is consistent with P, and (P&R) entails Q. And intuitively, we know that it is possible for Sally to have not eaten last week and to have gained weight last week if she drank enough high calorie liquids.

Plantinga uses this same strategy to show that the existence of God and the existence of evil are consistent. In particular, he argues that God and evil are consistent because there is a third possibly true proposition, the freewill defense, such

that the existence of God is consistent with the freewill defense and the existence of God and the truth of the freewill defense together entail the existence of evil. So, let F stand for the freewill defense. According to Plantinga, $\diamond(G \ \& \ E)$ because $\diamond(G \ \& \ F)$ and $((G \ \& \ F) \rightarrow E)$ (Plantinga, 1977, p.29).

What is the freewill defense (F)? The freewill defense consists of three component possibility claims:

- (i) A “world containing creatures who are significantly free (and freely perform more good than evil actions) is more valuable, all else being equal, than a world containing no free creatures at all” (Plantinga, 1977, p.30).
- (ii) There are significantly free creatures (i.e. creatures who are free with respect to morally significant actions).
- (iii) Everybody suffers from transworld depravity (Plantinga, 1977, p.48).

The first component is a claim about what is axiologically possible. The second component amounts to the claim libertarian freewill is coherent and that we exercise our freewill about morally significant actions. The third component will be explained below.

Plantinga includes in the term “significantly free” the idea of a morally significant action. An “action is *morally significant*, for a given person, if it would be wrong for him to perform the action but right to refrain or *vice versa*” (Plantinga, 1977, p.30). So, for example, Maurice might be free with respect to the action of eating oatmeal for breakfast but he is not *significantly free* because whether or not he eats oatmeal is not a morally significant action (Plantinga, 1977, p.43).

By themselves, the first two components aren’t enough to do the job Plantinga needs done. In order to understand, the third component, i.e. transworld depravity, it

helps to understand why the first two components are insufficient for the task at hand. So, set the definition of transworld depravity on the back burner while we lay the dialectical foundation for its inclusion in the freewill defense.

The problem with the first two elements of the freewill defense is that they don't entail the existence of evil. As Mackie points out:

...if God has made men such that in their free choices they sometimes prefer what is good and sometimes what is evil, why could he not have made men such that they always freely choose the good? If there is no logical impossibility in a man's freely choosing the good on one, or on several, occasions, there cannot be a logical impossibility in his freely choosing the good on every occasion. God was not, then, faced with a choice between making innocent automata and making beings who, in acting freely, would sometimes go wrong: there was open to him the obviously better possibility of making beings who would act freely but always go right. Clearly, his failure to avail himself of this possibility is inconsistent with his being both omnipotent and wholly good (Mackie, 1955, p.209).

Since there are possible worlds where all free creatures only do morally good acts, the first two components of the freewill defense don't entail the existence of evil.

Plantinga points out that this sort of objection assumes that it is within God's power to create any possible world. He argues this assumption is wrong and dubs it "Leibniz's Lapse" (Plantinga, 1977, p.44) He gives an initial trivial counterexample to Leibniz's Lapse. Assuming God is a contingent being, there are possible worlds where God doesn't exist. God can't actualize one of those worlds. If God exists contingently, then any world God actualizes includes him. If a world doesn't include God, then he can't actualize it (Plantinga, 1977, pp. 39-40).

We might be tempted to doubt this by mistakenly thinking of God existing "outside" the world he actualizes. Thus, God would exist but he wouldn't exist *in* the actual world. This objection is based on a misunderstanding. It imagines God as outside all possible worlds looking at a "menu" of possible worlds and deciding

which one to actualize. Possible worlds, though, are just ways the world could have been. For Plantinga, they are abstract maximal states of affairs. The actual world is just the maximal state of affairs that actually obtains (Plantinga, 1974, p.44). So, if God contingently but actually exists, this means the actual maximally consistent state of affair includes the existence of God. Thus, it is impossible for God to actualize a maximal state of affairs, i.e. a possible world, that doesn't include him. Rather than the image of God picking from a menu, perhaps it is a better metaphor to see him as choosing from an infinite number of paths and sub-paths ahead of him.¹⁶

The atheologian can, of course, avoid this point by making the following kind of slight modification to his claim about what God's omnipotence amounts to: if God is omnipotent he can actualize any possible world in which he exists. Plantinga offers another counterexample to Leibniz's lapse, one that cannot be skirted so easily and gets to the heart of why he includes transworld depravity in the freewill defense.

Plantinga offers the following example: Suppose there is a proposal before the mayor of a town to build a bridge that would require the destruction of a historical landmark. Suppose further that the mayor, call him "Curley" opposes the proposal. The director of highways for the State, call him "Smedes," offers Curley a \$ 35,000 bribe to drop his opposition to the bridge. Imagine that Curley accepts the bribe and that Smedes finds himself wondering if Curley would have accepted the bribe if it were \$20,000 instead of \$35,000.

Given this situation, Plantinga asks us to imagine that two additional things are true. First, grant that Curley has libertarian freedom with respect to taking the

¹⁶ This metaphor isn't perfect either. It adds complications about time and what it means to "go forward," etc. The point here is to get a different basic image of the kind of choices God has when deciding which world to actualize, one that doesn't mislead us in the specific way discussed here.

bribe. This means that there is some maximal world segment S' that includes Curley being offered the \$20,000 but does not include Curley accepting or rejecting the bribe. A maximal world segment is a complete state of affairs such that if any state of affairs that is compatible with S' but not included in it is added to S' , S' becomes a possible world (Plantinga, 1977, p. 46). Plantinga stipulates that S' is a maximal world segment in order to ensure that S' includes all relevant circumstances. Second, grant that it is true that if Curley were offered \$20,000, he would have taken it. To put this more carefully, suppose the following counterfactual is true:

(14) If S' were actual, Curley would accept the \$20,000 bribe.

Given this scenario, Plantinga grants that there is some possible world W' which includes S' but where Curley rejects the bribe. But given the truth of (14), God cannot actualize that world. To do so, God would have to actualize S' and leave Curley free with respect to accepting the bribe. But assuming Curley is free, (14) tells us that if Curley were bribed, he would freely accept the bribe. If God determines that Curley accepts the bribe, then Curley doesn't have libertarian freedom. If God gives Curley libertarian freedom, then Curley will freely accept the bribe. In other words, if it is true that Curley would freely accept the bribe, not even God can make it such that Curley *freely* rejects the bribe.

The kind of counterfactual expressed by (14) is typically referred to as a "counterfactual of freedom" or a "counterfactual of creaturely freedom" (CCF). Plantinga's argument makes use of the fact that even God can't determine the truth value of CCFs. If a particular CCF is true at a particular world, God has a choice about whether or not to actualize that world but if he actualizes that world even he

can't make that CCF false. If God determines the truth value of CCFs, then God is making it true that Curley would do A, if S' were actual and that is not compatible with libertarian freedom.

Plantinga asked us to assume for the sake of argument that (14) is true. Suppose for the moment it isn't. Suppose instead that the following counterfactual is true:

(15) If S' were actual, Curley would refuse the \$20,000 bribe.

If (14) is true and (15) is false, is Plantinga's argument undermined? He doesn't think so. Plantinga contends that either (14) is true or (15) is true (Plantinga, 1977, pp. 40-41). If (14) is true, then God can't actualize a world where S' holds but Curley refuses the bribe. If (15) is true, then God can't actualize a world where S' holds but Curley accepts the bribe. Either way, there is some world that God cannot actualize. Which one it is depends on the truth values of CCFs and this is something that God doesn't have control over.

Plantinga takes the case of Curley a step further. The above case involved a narrowly defined CCF. The CCF was about what Curley would do if offered a particular bribe. Suppose that a broader CCF about Curley were true. Suppose that there is some situation, in one of the possible worlds where Curley is perfect, where he makes the moral choice but if that same situation were actualized, Curley would make the morally wrong choice. More carefully put, Plantinga asks us to imagine that there is some action A for every world W' where Curley is both free and morally perfect that has the following property:

There is a maximal world segment S' that obtains in W' such that (1) S' includes Curley being free *re* A but neither is performing A nor refraining

from A; (2) *S'* is otherwise as much as possible like *W'*; and (3) if *S'* had been actual, Curley would have gone wrong with respect to A (Plantinga, 1977, p.47).

This means that there is a possible world *W'* where Curley is both free and morally perfect but God cannot actualize that world because to do so he would also have to actualize *S'* (because *W'* includes *S'*) and if *S'* were actual, the counterfactual labeled (3) in the quote of Plantinga above tells us that Curley would go wrong with respect to A. Given the nature of freewill, even God can not determine the truth value of the CCF above. Plantinga says if this is true of Curley, then Curley suffers from what he refers to as “transworld depravity.”

We are now in a position to understand the third and final component of Plantinga’s freewill defense. The third part of the freewill defense is the claim that it is *possible*, in the broadly logical sense, that *everyone* suffers from transworld depravity. He defines transworld depravity thusly:

A Person P suffers from transworld depravity if and only if the following holds: for every world W such that P is significantly free in W and P does only what is right in W, there is an action A and a maximal world segment *S'* such that

- (1) *S'* includes A’s being morally significant for P
- (2) *S'* includes P’s being free with respect to A
- (3) *S'* is included in W and includes neither P’s performing A nor P’s refraining from performing A

and

- (4) If *S'* were actual, P would go wrong with respect to A (Plantinga, 1977, p.48).

Plantinga’s definition of transworld depravity has four numbered components, but it’s important to notice that one of these components is not like the others. Plantinga calls attention to this difference by remarking that, “(4) is to be true in fact, in the actual

world—not in that world W'' (Plantinga, 1977, p.48). (4), is importantly different than the other three. What is the difference?

(4) is a CCF. Counterfactuals are world-indexed. The truth value of a counterfactual depends on the world about which the counterfactual is asserted. Plantinga is highlighting the fact that (4) is a counterfactual truth in the actual world, not in W . In *this* world it's true that if S' were actual, P would go wrong with respect to A .

So, even though there are S' worlds where P makes the morally right choice with respect to A , transworld depravity tells us that if S' were *actual*, then P would go wrong with respect to A . Why? Because if transworld depravity is true, it's just a fact about the actual world that the counterfactual expressed in (4) is true. Plantinga's claim that it is possible that everyone suffers from transworld depravity amounts to the claim that it is possible that every person has the following property: even if there is a world where everyone freely does only morally good acts, for any person you pick there is some circumstance that calls for a morally significant choice to be made. But if that circumstance were actualized, the person would perform an immoral act.

We can better understand how transworld depravity is important to the freewill defense by reviewing the dialectical flow of the debate. Mackie contends that there are possible worlds where every creature freely performs only good actions. He further asserts that if God is omnipotent he *could* actualize one of those worlds and if God is omnibenevolent, he *would* actualize one of those worlds (i.e. there is no possible world that contains both God and evil). The actual world is not one of those worlds because it contains evil, so God doesn't actually exist.

Plantinga's reply concedes that there are possible worlds where every free creature always perform good actions but argues that it might not be possible for God to actualize those worlds. How could it be that there is a logically possible world that even God cannot actualize? (Recall Plantinga's trivial example of God not being able to actualize a possible world where he does not exist.) If everyone suffers from transworld depravity, then the following counterfactual is true: no matter which person P you pick, there is a maximal world segment S' and a morally significant action A, such that *if S' were actual*, P would go wrong with respect to A (Plantinga, 1977, p.48).

Plantinga ends his discussion of transworld depravity by adding that transworld depravity is possibly true of individual essences or natures. He adds this in order to forestall the objection that it is possible that everyone (i.e. all of us who actually exist or have existed) suffers from transworld depravity but God could just actualize a different set of persons. If it is the essence of a person that suffers from transworld depravity, then any instantiation of that person suffers from transworld depravity (Plantinga, 1977, pp. 49-53).

What is an individual essence? It is a collection of essential properties that make something what it is. An essential property is a property that an individual essence has in every world in which that individual exists (Plantinga, 1977, p.49). Take Socrates. Plantinga contends that the essential properties of Socrates include all world-indexed properties of Socrates. So, suppose that Socrates has the property of being snub-nosed in the world Kronos. Having a snubbed nose is true of Socrates in Kronos but not in other worlds. So being snub-nosed is not an essential property of

Socrates. However, the world-indexed property, “is snub-nosed in Kronos” is essential to Socrates because it is true of him in every world (Plantinga, 1977, p.50). Socrates has essential properties that are not world-indexed. Plantinga doesn’t give examples but they are just properties without which a person would no longer count as being Socrates (Plantinga, 1977, p.51).

Let’s return now to the general issue of the consistency of God and evil. Speaking merely formally, does the freewill defense meet the requirements of the procedure Plantinga laid out earlier? Consider the first requirement. Is F consistent with G and do G and F together entail E? It appears that the freewill defense is consistent with the existence of God. There is nothing that is readily apparent about the existence of God that is inconsistent with the claim about the value of worlds with morally significant actions relative to worlds without such actions. God’s existence also seems consistent with the coherence of libertarian freedom and transworld depravity. So, (G&F) looks to be consistent. Second, does (G&F) entail E? It appears so. If God exists and creates a world with significantly free creatures, then, given transworld depravity, evil will exist.

4. Logical Objection to Transworld Depravity

Several objections have been offered against Plantinga’s freewill defense. I will consider three types: a logical objection, objections to Molinism, and an epistemic objection. The first objection will be briefly considered in this section. The second and third types of objections will be explained in the next two chapters. There are several reasons to examine these objections. First it’s important to my thesis that there are no obvious surface level objections to the freewill defense.

Otherwise, it is less plausible to think that the well informed modern theist and atheist can legitimately differ in their assessment of the freewill defense. Secondly, I aim to show that the most serious sort of objections are epistemic in nature. This helps motivate and make more plausible my general thesis because knowledge and justification are, in important ways, person relative.

A common objection raised against the freewill defense targets the modal claims made by transworld depravity. Hugh Lafollette (Lafollette, 1980) and R. Zachary Manis (Manis, 2004) make these sort of objections. They both argue that if Plantinga is right when he claims that it is possible everyone is transworld deprived (TD), then it is necessary. More formally, in S5 \diamond TD implies TD. Since, it is highly implausible that necessarily everyone suffers from transworld depravity, then Plantinga's claim is highly implausible.

Of course Plantinga doesn't argue that it's possible that transworld depravity is necessary. Rather he argues merely that it is possible. Why do they think that his view of transworld depravity implies that it is necessary? Because they think it is a claim about every person in every possible world. If something is true in every possible world, then it is necessarily true.

However, Plantinga isn't saying that it is possible that every essence in every possible world is transworld deprived. He is not claiming that there are no worlds where every essence always chooses to do right all the time (Mackie-worlds). In fact, he is conceding to Mackie that such worlds exist but argues that it is possible that there are possible worlds like that that even God cannot actualize. His whole

argument assumes Mackie-worlds are possible. If Mackie-worlds are possible, then transworld depravity isn't necessary.

Rather, Plantinga is claiming that there is at least a possible world *W* where every personal essence suffers from transworld depravity. His claim is that there is a possible world *W* with a certain pattern of true counterfactuals of freedom. The patterns of true counterfactuals at *W* are such that any essence which God instantiates at that *W* will go wrong at least once in world.

As we saw earlier, a personal essence includes the world-indexed properties of that essence. Let's name our world *Z*. So, for example, the property "is a philosopher in *Z*" is an essential property of Socrates. Consider any person, in any possible world, if that person is not a philosopher in *Z*, then he is not Socrates. Likewise, Plantinga is saying that there is a possible world *W* where every essence God instantiates in *W* is transworld depraved. If this is so, then every essence has an essential world-indexed property, namely, it is transworld depraved in *W*.

So, contrary to this objection, Plantinga is *not* claiming it is possible that every personal essence will go wrong at least once in every possible world. Rather, he is saying that it is possible that there is one world where every essence God could instantiate at *that* world goes wrong at least once. Further, if this is so, then every essence God could instantiate has the essential world-indexed property "goes wrong at least once."¹⁷

As we have seen, Plantinga's freewill defense assumes that there are true CCFs. The thesis that there are true CCFs and that God knows them is Molinism.

¹⁷ For more replies along these lines see (Evans, 1983) and (Gan, 1982).

Many of the criticisms of Plantinga's freewill defense have focused on the underlying assumption of Molinism. It is to this topic I now turn.

Chapter 2: The Freewill Defense and Molinism

There are several important criticisms of Plantinga's response to the logical argument from evil. In this chapter I focus on criticisms of Molinism, on which the freewill defense appears to depend. CCFs constitute an important part of Plantinga's response to the logical argument from evil. The view that there are true CCFs and God knows them is called "Molinism." The free will defense, as developed by Plantinga, depends on the truth of Molinism. If Molinism is false, this would seem to undermine the free will defense. In this section, I will lay out and explain Molinism. I will then set out and assess the main criticisms of Molinism in order to determine whether the objections against Molinism are sufficiently strong to undermine the free will defense. Lastly, I will explore whether a different formulation of the free will defense that doesn't rely on Molinism can succeed.

Before I explain what Molinism is in detail, it might be helpful to see how the free will defense is supposed to rely on Molinism. The claim that CCFs have truth value and God knows them plays a critical role in Plantinga's account of transworld depravity and transworld depravity is, in turn, at the heart of the freewill defense. They also enter into Plantinga's general response to Leibniz's Lapse but this lays the groundwork for the role they play in characterizing transworld depravity. The function of CCFs in both of these discussions is similar: they act as a sort of legitimate limit on God's power. By "legitimate limitation on God's power," I mean a limit that doesn't undermine or diminish God's divine credentials.

What sort of limit is a legitimate limit on omnipotence? In order for it to be possible for God to be justified in permitting evil, there has to be some moral good that even God can only achieve by permitting evil. Any claim to the effect that “even God cannot achieve X without allowing for evil,” is a statement about God’s limits. In order to be omnipotent, any such limit has to be a merely logical limit. So, a legitimate limit amounts to a logical limit.

What Plantinga attempts with the freewill defense is to find a logical limit on God which provides some space to find a possible moral justification for God’s permitting evil. CCFs play the role of a *logical* limit on God. The basic idea is that even God cannot determine the truth value of CCFs. CCFs are supposed to be an important source of legitimate constraints on God. If the free will defense is right, CCFs provide God with an excuse for permitting suffering, or at least room for an excuse.

How do CCFs provide a *logical* constraint on God? They provide a logical constraint because the kind of freedom relevant to the free will defense is libertarian freedom. In order for agents to be significantly free, the truth value of the relevant CCFs have to be determined by the relevant free agents, not God. If God determines what a person would do in a particular situation, then that decision is not free. If God sets the truth value of CCFs, then they are not counterfactuals of *freedom*.

It may be worth making a general point about counterfactuals and modality here. There is a sense in which God doesn’t determine which non-CCF counterfactuals are true. God doesn’t *create* worlds, he *actualizes* them. This is an important distinction. God determines which world to actualize not which worlds are

possible. As will be explained in a moment, the truth of counterfactuals is set by the relations among possible worlds. The extrinsic relationships among possible worlds set the truth value of counterfactuals. God doesn't have control over the modal space. He doesn't determine which worlds exist in modal space and the relations between them. He determines which one is actual.

1. The FWD with Molinism

1.1 Explanation of Molinism

Molinism is a theological view of providence developed by the 16th century Jesuit theologian Luis de Molina. Molina sought to provide an account of divine providence that preserved both libertarian free will and an exhaustive view of divine providence. He held both that humans are indeterministically free and that “whatever occurs is specifically decreed by God” (Freddoso, 1998, p.463).¹⁸

Molina developed just such an account by proposing that God has three kinds of knowledge: prevolitional knowledge, postvolitional knowledge, and middle knowledge. The basic idea is that God's prevolitional knowledge is logically prior to God's act of creation and that God has no control over the truth value of propositions that he knows prevolitionally. Propositions known prevolitionally are both necessary

¹⁸ Alfred Freddoso gives the following description of the traditional Catholic doctrine of divine providence:

God freely and knowingly plans, orders and provides for all the effects that constitute the created universe with its entire history, and he executes his chosen plan by playing an active causal role that ensures its exact realization. Since God is the perfect craftsman, not even trivial details escape his providential decrees and after that description (Freddoso 1998, p.463).

Just after this description he writes that, “This much is accepted by both Molina and the Bañezians.”

and true independent of God's free will. An example of prevolitional knowledge is knowledge of necessary truths (Flint, 1998, pp. 37-38).

Molina thought that God's actualization of a particular possible world adds to his knowledge "by informing him as to which contingent propositions he may have rendered false (i.e. that Cuthbert comes into existence at *t*)" (Flint, 1998, p.38).

Postvolitional knowledge, then, is knowledge of propositions that are contingently true and true in virtue of an act of God's free will (Flint, 1998, pp. 38-39). An example of postvolitional knowledge is knowledge of the gravitational constant.

Middle knowledge is knowledge of CCFs (e.g. if Curley were offered the bribe, he would take it). It is a sort of middle category of knowledge between prevolitional and postvolitional. Middle knowledge is like prevolitional knowledge in that it is true and known independent of God's will. God doesn't determine the truth value of CCFs, free agents do. Middle knowledge is like postvolitional knowledge in that it is contingent. CCFs are contingent, not necessary truths, otherwise they wouldn't be counterfactuals of *freedom* (Flint, 1998, p.40-42).

Molina argued that God's middle knowledge allows the theologian to have his providential cake and eat it to. He contended that God's knowledge of CCFs allows God to arrange for every event to happen according to his will without violating libertarian free will. For example, if God's plan includes Curley taking the bribe and God knows the CCF "if Curley were offered the bribe, he would take it," then God need only actualize the state of affairs of Curley being offered the bribe. God gets the end result he wanted and Curley accepted the bribe of his own libertarian free will.

1.1.1 Strong and Weak Actualization

At this point it is worth taking some time to understand the Molinist's distinction between strong and weak actualization. John Laing contends that the distinction:

has proven to be a powerful tool for understanding the relationship between God's providence and human freedom. However, it must be noted that it implies that there may be some states of affairs that God cannot weakly actualize, which leads to the further conclusion that there may be some possible worlds that God cannot actualize (Laing, 2005).

This distinction is intended to clarify in what sense God actualizes the states of affairs caused by free creatures. It gets heavy use both by Plantinga in the free will defense and by Molinists like Thomas Flint to answer other various objections to Molinism (Flint, 1998, pp. 110-111).

Plantinga begins first with the definition of actualization *simpliciter*: An agent actualizes a state of affairs if he causes that state of affairs to be actual (Plantinga, 1985c, p.48). Plantinga distinguishes between strong actualization and weak actualization thusly:

Let us say that God *strongly* actualizes a state of affairs *S* if and only if he causes *S* to be actual and causes to be actual every contingent state of affairs *S** such that *S* includes *S**; and let's say that God *weakly* actualizes a state of affairs *S* if and only if he strongly actualizes a state of affairs *S** that counterfactually implies *S*. (Strong actualization is thus a special case of weak actualization) (Plantinga, 1985c, p.49).

Philosophers like Edward Wierenga and Daniel Hill have criticized Plantinga's account of the strong/weak actualization distinction, so it is worth taking the time to get the distinction clear (Hill, 2005, p.167; Wierenga, 1989, p.22). Wierenga points out that one consequence of this formulation of the distinction:

...is that any state of affairs an agent strongly actualizes is thereby one the agent weakly actualizes. It will be convenient to have a notion of weak actualization that excludes strong actualization, so let's say that whoever weakly but not strongly actualizes a state of affairs *strictly* weakly actualizes it. Having introduced this cumbersome locution, let's henceforth avoid it and use the term 'weakly actualize' to mean strictly actualize (Wierenga, 1989, p. 22).

Wierenga's complaint can be addressed more economically by adding a slight modification to Plantinga's original definition.

Rather than saying that "God *weakly* actualizes a state of affairs *S* if and only if he strongly actualizes a state of affairs *S** that counterfactually implies *S*," Plantinga would be better off saying that "there is a counterfactual of *freedom* such that the state of affairs *S** counterfactually implies *S*." The difference is this: if we say merely that *S** counterfactually implies *S*, then that is compatible with the counterfactual truth in question being true for deterministic reasons. But it's an important part of the strong/weak actualization distinction that weak actualization involves God *and* another free agent.

A formulation of the distinction that avoids the trouble Wierenga points out amounts to something like this: God *strongly* actualizes a state of affairs *S* if and only if God is the only agent involved in the actualization of *S*. God *weakly* actualizes a state of affairs *S* if and only if God's action is necessary but not sufficient to actualize *S* and the action of God together with at least one other agent are sufficient to actualize *S*.

Let me explain this in a looser but more intuitive manner. When God strongly actualizes a state of affairs *S*, he causally determines *S*. The concept of strong actualization is more straightforward than that of weak actualization. It is the kind of

actualization involved in everyday cases like: *Mike started the car by turning the key* or *Michelangelo made the statue*. When God weakly actualizes a states of affairs S^* he strongly actualizes S knowing that the following counterfactual is true: if S were to obtain, Curley would strongly actualize S^* . Weak actualization involves: (1) God strongly actualizing S and (2) God knowing the truth of some CCF like: $S \rightarrow$ Curley causes S^* .

Consider an example. Suppose God wants to weakly actualize Michelangelo's sculpture "David." To do so, God has to (1) actualize the state of affairs S which includes e.g. a large stone, Michelangelo, his tools, etc. and (2) know that the following counterfactual is true: if S were to obtain, then Michelangelo would actualize a different state of affairs S^* (i.e. make the sculpture). In this case, God actualized two states of affairs—one strongly (S) and one weakly (S^*). Michelangelo strongly actualized one state of affairs (S^*) but did not weakly actualize any state of affairs.

There are two misunderstanding we need to avoid here: First, one might think that the difference between strong and weak actualization just amounts to the difference between direct and indirect causation. God directly caused S but only indirectly caused S^* . One problem with this is there doesn't seem to be a precise and clear way to distinguish between direct and indirection causation. If a surgeon in the U.S. is manipulating buttons on a control panel in order to direct a robotic surgery arm in the Middle East, would we say he is directly performing the surgery or that he is directly manipulating the robot's controls and only indirectly performing the surgery? At the very least it is unclear how to sort out these kinds of cases. The

direct/indirect causation distinction is useful in ordinary life but it isn't sharp enough for the kind of philosophical work envisioned here.

Furthermore, there *is* a difference between strong/weak actualization and direct/indirect causation. Weak actualization *requires* the involvement of at least one other free agent. One can be the indirect cause of an event without other free agents being involved in the process. Think of example above of the surgeon remotely performing robotic surgery. Think of pilot who causes a plane crash that causes a forest fire which in turn burns down a house with someone's retirement saving stuffed in a mattress. The pilot may have caused the person to lose his retirement nest egg but he didn't do so directly.

While indirect causation is not the same as weak actualization, it may include it. If I constantly badger someone and relentlessly "push his buttons" in a way that results in his losing his cool and smashing a hole in the wall, then I am an indirect cause of the hole in the wall but another agent was essentially involved in the act. At best then, weak actualization is only a narrow kind of indirect causation because there are other applications of the direct/indirect causation distinction that don't involve more than one free agent.

The requirement that weak actualization involve at least a second free agent, may lead to a second potential misunderstanding. One might be tempted to think that the difference between strong and weak actualization is that in weak actualization there is a causal chain that starts with God and runs through a second free agent to the actualization of the final state of affairs. We have to be careful here. If we think of

these kinds of causal chains as mechanistic or deterministic, then libertarian freedom is undermined.

How should we understand the weaker sense in which God caused the sculpture? Consider the case of God and Michelangelo above. In that case, there are two cases of strong actualization and one instance of weak actualization. There are two deterministic causal chains: one originating with God and ending in the actualization of *S*, the other originating with Michelangelo and ending in the actualization of *S**. The first causal chain did not causally determine the second causal chain. There is a gap of deterministic causation between the two causal chains. The two causal chains *are* linked together but they are connected by the truth of a counterfactual of freedom about Michelangelo, not by deterministic causation. More specifically, the actualization of *S*, together with the truth of a particular CCF(s), implies *S**.

It's less misleading to think of weak actualization using the metaphor of a building rather than a chain. God builds a particular foundation *knowing* that if he did so, Michelangelo would freely build a particular superstructure on top of that foundation. While God is the cause of the foundation in the strong sense of "cause," he is the cause of the entire building only in the weak sense.

1.1.2 Robust and bare-bones Molinism

With a basic account of Molinism and the distinction between strong and weak actualization on the table, a problem seems to arise for the Molinist version of the free will defense. The robust version of providence set out by Molina seems to

make the argument from evil worse. God chose, for example, to weakly actualize the holocaust. He didn't just know that the holocaust would occur, he *intended* it to!

The Molinists who advocate the free will defense have a credible reply to this apparent problem by making a further distinction. It's important to make a distinction between robust and bare-bones Molinism.¹⁹ Robust Molinism is the view of divine providence sketched above. It has two important elements: (1) that God knows CCFs and (2) that God uses that knowledge to predetermine all events. According to robust Molinism, God not only knows what will result from creating free creatures, he intends for that result to occur. God chose to weakly actualize every state of affairs that is strongly actualized by free creatures. Bare-bones Molinism, on the other hand, has only one element: that God knows true CCFs.

Plantinga's free will defense only makes use of bare-bones Molinism. What's important for his purpose is that there are true CCFs that constrain what God can achieve. The advocate of the free will defense has no need to make any stronger claim. If Plantinga is merely asserting bare-bones Molinism, then it's merely the case that God knew which states of affairs he would end up weakly actualizing, not that he preferred those states of affairs.

Now one might object to the distinction between robust and bare-bones Molinism on the grounds that if God knew with certainty that the holocaust would occur under certain circumstances and he chose to weakly actualize those circumstances, then in *some* sense God intended the holocaust to occur. God knew the consequences of actualizing a particular world and still chose to actualize that particular world. So, God weakly actualizes the states of affairs that he wanted

¹⁹ This is my distinction. I am unaware of it being made explicit elsewhere.

actualized. If this objection is right, the distinction between robust and bare-bones Molinism collapses and Plantinga is saddled with the ugly problems of robust Molinism.

The weak Molinist has a reply here. The problem with this objection is that it mistakenly conflates “choosing X, knowing that Y would result” and “intending Y.” The principle of double effect is relevant here. Thomas Aquinas is credited with being the first philosopher to explicitly state the principle of double effect. Aquinas makes the point that:

Nothing hinders one act from having two effects, only one of which is intended, while the other is beside the intention. Now moral acts take their species according to what is intended, and not according to what is beside the intention, since this is accidental... Accordingly the act of self-defense may have two effects, one is the saving of one's life, the other is the slaying of the aggressor. Therefore this act, since one's intention is to save one's own life, is not unlawful, seeing that it is natural to everything to keep itself in "being," as far as possible (Aquinas, 1920 [1274], I-II, Qu. 64, Art.7).

The bare-bones Molinist claims that God chose a particular *pattern* of good and evil with the effect that the holocaust occurred. Following Leibniz’s lead, the bare-bones Molinist argues that God chose the best pattern of good and evil available to him given the pattern of the CCFs he was confronted with²⁰. God didn’t *want* the holocaust, or any other particular evil to occur. He intended that the best overall pattern of good and evil actualized and, unfortunately, that patterns included some pretty horrible events. So, there is no reason Plantinga needs to affirm anything stronger than what I have been calling “bare-bones Molinism” and bare-bones Molinism doesn’t imply that God wanted events like the holocaust to occur.

²⁰ Of course, there may not be a best possible world or best possible pattern of good, given the truth of CCFs God had to work with. The weak Molinist only needs it to be the case that there was no better pattern than the one God chose.

1.2 Criticism of FWD via criticism of Molinism

Plantinga's version of the free will defense assumes what I have been calling "bare-bones Molinism." True CCFs limit God's options and effectively give him a sort of excuse for not actualizing a world with moral good but no moral evil. Consequently, an important set of objections to the free will defense revolve around objections to bare-bones Molinism. According to the standard objections to Molinism, God doesn't have middle knowledge, i.e. knowledge of true CCFs, because there are no true CCFs. God doesn't have middle knowledge because there are no counterfactual facts concerning the actions of free creatures to be known.

1.2.1 The Semantics of Counterfactuals

To fully understand the objections to bare-bones Molinism, it's important to understand the semantics of counterfactuals to a certain level of detail. Additionally, much of the force of Plantinga's free will defense depends on subtleties of counterfactual and modal semantics. Consequently, I shall lay out the basics of the semantics of counterfactuals in some detail before I set out specific objections to Molinism.

What do we mean when we say, "if it were to rain, the lawn would be wet?" The goal of a semantics of counterfactuals is to answer this question in general terms. The idea is to specify the truth conditions for counterfactuals like the one above. The two most prominent accounts of the semantics of counterfactuals are offered by Robert Stalnaker and David Lewis (D. Lewis, 1973; Stalnaker, 1981b [1968])²¹.

²¹Stalnaker's essay was originally published in (Stalnaker, 1968) Harper, et. al.

1.2.1.1 Stalnaker's Semantics of Counterfactuals

In his paper "A Theory of Conditionals," Stalnaker takes as his starting point for a semantics of conditionals Frank Ramsey's test for evaluating the belief conditions of hypothetical statements (Ramsey, 1931). He suggests that the proper view of the truth conditions for conditionals corresponds, roughly, to Ramsey's account of the acceptability conditions for conditionals. Stalnaker makes two modifications to Ramsey's account which yields a first approximation of his semantic theory of conditionals. He then fleshes out his account with a model (not modal) structure and a selection function.

Stalnaker begins his discussion by considering the following, now dated, hypothetical as an example:

*If the Chinese enter the Vietnam conflict, the United States will use nuclear weapons*²².

On Stalnaker's construal of Ramsey, Ramsey advocates the following procedure to evaluate the belief conditions of the above hypothetical: add the antecedent to our stock of beliefs and then evaluate whether the consequent is true (Stalnaker, 1981b, p.43). So, adding the belief that "the Chinese enter the Vietnam conflict" to my total set of beliefs yields a new set of beliefs. Given this new set of beliefs, is it true that "the United States will use nuclear weapons?" One's belief about the conditional should be the same as one's belief about the consequent, given the addition of the antecedent to one's stock of beliefs.

²² While this conditional doesn't have the standard form of a counterfactual, it doesn't matter for our purposes. Stalnaker is giving a general theory of conditionals which includes counterfactual conditionals. Lewis, by contrast, directly offers a theory of counterfactuals and does not give us a general theory of conditionals.

Stalnaker points out that this won't work for evaluating conditionals with antecedents that one believes to be false. This is where he makes the first of two modifications to Ramsey's account that eventually results in his semantics of counterfactuals. In order to account for these cases, Stalnaker's suggestion is to add the antecedent to one's stock of beliefs by making only the minimal changes necessary to maintain consistency. So, according Stalnaker's first revision of Ramsey, in order to evaluate a conditional:

First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is then true (Stalnaker, 1981b, p. 44).

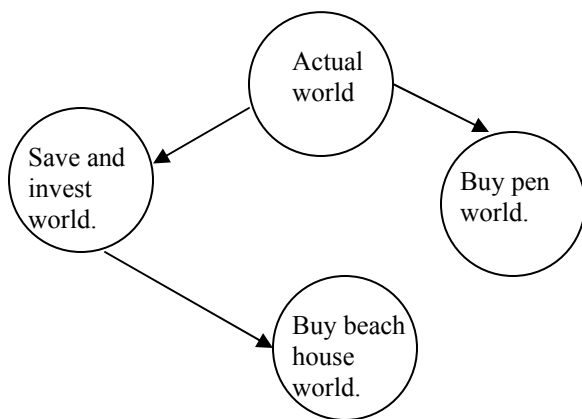
Thus far, using Ramsey's suggestion as a starting point, Stalnaker has just told us how we are to decide whether or not to believe a conditional, even if we think the antecedent is false. Stalnaker's second modification of Ramsey's test is intended to turn a theory about the *belief conditions* of conditionals into a theory about the *truth conditions* of conditionals. He bridges this doxastic-semantic divide by introducing the concept of possible worlds, "since a possible world is the ontological analogue of a stock of hypothetical beliefs" (Stalnaker, 1981b, p.45). With these two modifications, Stalnaker has a first approximation of a process for determining the truth conditions of a conditional:

Consider a possible world in which *A* is true, and which otherwise differs minimally from the actual world. "if *A*, then *B*" is true (false) just in case *B* is true (false) in that possible world (Stalnaker, 1981b, p.45).

Stalnaker fills out this account with both a model structure and a selection function. The model structure he proposes roughly follows that developed by Saul

Kripke (Kripke, 1963a). Stalnaker first defines a model structure. Let M be an ordered triple (K, R, λ) . K is the set of all possible worlds. R is an accessibility relation. So, if α and β are members of K , i.e. possible worlds, then $\alpha R \beta$ is read as “ β is possible with respect to α .” What this means is that if α is the actual world, β is a possible world accessible from α .

What is an accessibility relation? Sometimes we want to talk about what is possible in general. Other times, we want to talk about what is possible relative to a certain world. Accessibility relations help us to do this. Consider two possible worlds: one where I buy an ink pen today and one where I buy a large beach house today. Both worlds are accessible from the actual world but not in the same way. The world where I buy a pen today is *directly* accessible to the actual world because given the way the actual world is, it is financially possible for me to buy a pen today. However, the world where I purchase a beach house today is *indirectly* accessible from the actual world. Given the way the actual world is, it is financially impossible for me to buy a beach house today. There is, though, a world accessible from the actual world where I saved a lot of money and invested it wisely. The world where I purchase a beach house today is accessible from *that* world. So, the world where I buy a beach house today is accessible from the actual world but only indirectly.



As we see from this example, an accessibility relation allows us to fine-tune our modal talk. It allows us to say more than that there is a possible world where I buy a beach house. It permits us to talk about whether that world is possible relative to some other world. Accessibility relations tell us which world is possible relative to another world.

In Stalnaker's model structure, λ is the absurd world. The absurd world is the world in which everything is true. All contradictions and their consequences are true in λ . The purpose of λ is to permit conditionals with an impossible antecedents to have an interpretation. λ is an isolated element under the accessibility relation R , meaning that no world is possible with respect to it and, more importantly λ is not possible with respect to any world (Stalnaker, 1981b, p.45).

In addition to a model structure, Stalnaker includes a selection function in his account whose purpose is to tell us which particular possible world where the antecedent of a conditional is true is most like the actual world. Stalnaker's selection function, f , takes a proposition and a possible world as arguments and produces a possible world as its value (Stalnaker, 1981b, p.45). For any antecedent proposition

A and base world α the selection function picks out a world $f(A, \alpha)$. $f(A, \alpha)$ is the world most similar to the base world where the antecedent is true.

The selection function is subject to certain constraints designed to ensure that the selected world really is the most similar world where the antecedent is true (Stalnaker, 1981a). Stalnaker proposes four such conditions:

- (1) For all antecedents A and base worlds α , A must be true in $f(A, \alpha)$.
- (2) For all antecedents A and base worlds α , $f(A, \alpha) = \lambda$ only if there is no possible world with respect to α in which A is true...
- (3) For all base worlds α and antecedents A , if A is true in α , then $f(A, \alpha) = \alpha$.
- (4) For all base worlds α and antecedents B and B' , if B is true in $f(B', \alpha)$ and B' is true in $f(B, \alpha)$, then $f(B, \alpha) = f(B', \alpha)$ (Stalnaker, 1981b, p.46).

What these conditions amount to is this: First, the antecedent must be true in the selected world. Second, the absurd world can only be selected when the antecedent is impossible. Third, the base world must be selected if it is among the worlds in which the antecedent is true. The fourth condition is a technical condition that ensures a total ordering of worlds. It basically says that there are two worlds where all and only of the propositions true in one world are also true in the other, then the two worlds are the same world (Stalnaker, 1981b, pp. 46-47). Stalnaker takes these conditions on the selection function to be “sufficient to define the semantical notions of validity and consequence for conditional logic” (Stalnaker, 1981b, p.47).

With this apparatus in place, Stalnaker states his semantical rule for conditionals thusly:

$A \Box \rightarrow B$ is true (false) in α if B is true (false) in $f(A, \alpha)$ (Stalnaker, 1981b, p.45)²³.

²³ Stalnaker and Lewis use different symbols for translating conditionals/counterfactuals. Stalnaker uses the \Box symbol, called “the corner”, for the conditional/counterfactual conditional while Lewis uses

When is a conditional true? On Stalnaker's theory, a conditional is true if the consequent is true in the most similar world where the antecedent is also true. If we want to know if a conditional is true, we examine the possible world most like the base world where the antecedent is true. If the consequent is true in that world, then the conditional is true. The sentence:

If the Chinese enter the Vietnam conflict, the United States will use nuclear weapons.

is true if it is true that the United States uses nuclear weapons in the world most similar to the actual world where the Chinese enter the Vietnam conflict.

Stalnaker next sets out a formal system of conditional logic, which he names "C2." One important axiom in C2 is:

$$\diamond A \rightarrow [\sim(A \square \rightarrow B) \equiv (A \square \rightarrow \sim B)] \text{ (Harper, 1981, p.6).}$$

This axiom says that if an antecedent is possible, then the negation of a conditional is logically equivalent to that same conditional with the negation of the consequent.

Given the rest of C2, this axiom is equivalent to Conditional Excluded Middle (CEM):

$$\text{(CEM)} (A \square \rightarrow B) \vee (A \square \rightarrow \sim B) \text{ (Harper, 1981, p.6).}$$


Both of these assertions are a result of the fact that the selection function, f , selects one and only one world as the closest antecedent-world to α .

1.2.1.2. Lewis' Criticism of Stalnaker's Semantics

the symbol $\square \rightarrow$ for the "would" counterfactual and the symbol $\diamond \rightarrow$ for the "might" counterfactual. Rather than switch back and forth between two ways of symbolizing, I will follow Lewis' notation.

David Lewis has criticized Stalnaker's semantics of counterfactuals because it assumes that there is a single closest antecedent-world to the relevant base world. He raises three objections to Stalnaker's theory that arise from this assumption. Lewis offers an alternative semantics of counterfactuals that avoids the pitfalls associated with Stalnaker's theory.

Lewis begins his criticism of Stalnaker's semantics by pointing out that it rests on a particular assumption. He dubs it "Stalnaker's Assumption" (SA). According to SA: "for any world i and antecedent ϕ ...that is entertainable at i , there is a sphere around i containing exactly one ϕ -world" (D. Lewis, 1973, p.78). (Unlike Stalnaker who uses a selection function to determine relative similarity, Lewis uses the idea of spheres as a way of describing comparative similarity. I will say more about this later when I describe Lewis' positive account. For now it is enough to know that worlds located in spheres closer to a world i are more similar than worlds in sphere farther from i). SA is represented in the figure below:




Stalnaker's Assumption (SA)

The figure represents SA because there is a sphere around i that has exactly one ϕ -world. Of course there are other models that also represent SA (e.g. a model with only one ϕ -world in all the spheres, etc.).

The first problem Lewis raises about Stalnaker's semantics is that SA implies what Lewis calls the "Limit Assumption" (LA). LA is the view that there is always a closest antecedent permitting sphere to i (D. Lewis, 1973, p. 19). SA implies LA because if there is always a sphere around i that contains exactly one ϕ -world, then that sphere will always be the closest ϕ -permitting sphere to i . To see this, suppose there were a closer ϕ -world to i than the one in the drawing above. Then the sphere containing the original ϕ -world in the above representation would now contain two ϕ -world's—the original one and the closer one. It would then be the sphere around i that SA states that there must be. Further, the sphere containing the closer ϕ -world that we are imagining has been added to the above model now counts as the sphere containing exactly ϕ -world.

Lewis points out that LA doesn't imply SA. To see this consider the following model:




In this model, LA is true because there is a closest ϕ -permitting sphere but SA is false because there is no sphere with exactly one ϕ -world.

The problem Lewis has with LA is that it fails in cases where two or more worlds get indefinitely more and more similar to the base world:

Unfortunately, we have no right to assume that there always are a smallest antecedent-permitting sphere and, within it, a set of closest antecedent worlds. Suppose we entertain the counterfactual supposition that at this point


there appears a line more than an inch long. (Actually it is just under an inch.) There are worlds with a line 2" long; worlds presumably closer to ours with a line 1 1/3 inches long; worlds presumably still closer to ours with a line 1 1/8 inches long; worlds presumably still closer... But how long is the line in the *closest* worlds with a line more than an inch long? If it is $1+x$ " for any x however small, why are there not other worlds still closer to ours in which it is $1+_x$ ", a length still closer to its actual length? The shorter we make the line (above 1"), the closer we come to the actual length; so the closer we come, presumably to our actual world. Just as there is no shortest possible length above 1", so there is no closest world to ours among the worlds with lines more than an inch long, and no smallest sphere permitting the supposition that there is a line more than an inch long (D. Lewis, 1973, p.20).

Here is a model of what Lewis is saying:



The arrows represent a continuous series of worlds that get indefinitely closer to i without reaching it. In this model, there is no closest ϕ -world to i because there are two ϕ -worlds equally close “all the way down” without actually “reaching” i .

A second problem with SA is that it implies CEM. If there is always a sphere around i with exactly one ϕ -world, then (assuming bivalence) that ϕ -world is either a ψ -world or a $\sim\psi$ -world. Thus, it follows for all worlds i , either $\phi \Box \rightarrow \psi$ or $\phi \Box \rightarrow \sim\psi$. Even in cases where the comparative similarity of two or more worlds to i is definite SA rules out there being ties and, hence, affirms CEM (Stalnaker, 1981a, p.89). This problem is distinct from LA. LA does not imply CEM because there could be a closest sphere with a ϕ -world in it but that sphere could contain a tie between a $\phi\psi$ -world and a $\phi\sim\psi$:



Lewis raises the following objection against CEM. Consider the following two sentences:

- (17) If Verdi and Bizet were compatriots, Bizet would be Italian;
- (18) If Verdi and Bizet were compatriots, Bizet would not be Italian.

Is one of these conditionals true? That is, is there a single world where Verdi and Bizet are compatriots and either Bizet is Italian or Bizet is not Italian? Is there any reason to assume that one of those worlds has to be closer than the other to the actual world? At the very least it's not safe to just assume that either (17) or (18) is true. SA requires that one of those two antecedent-worlds above be closer to the actual world than the other.

Suppose that neither antecedent-world is closer to the other. On Lewis' semantics, which will be set out shortly, (17) and (18) are both false because Lewis' account only has two truth values. Since neither sentence is true, they are both false. If they are both false, then CEM is false. Lewis puts it this way:

Given Conditional Excluded Middle, we cannot truly say such things as this:

It is not the case that if Bizet and Verdi were compatriots, Bizet would be Italian; and it is not the case that if Bizet and Verdi were compatriots, Bizet would not be Italian; nevertheless, if Bizet and Verdi were compatriots, Bizet either would or would not be Italian.

That is:

$\sim(\phi \Box \rightarrow \psi) \ \& \ \sim(\phi \Box \rightarrow \sim\psi) \ \& \ (\phi \Box \rightarrow \psi \vee \sim\psi)$.

I want to say this, and think it probably true; my own theory was designed to make it true. But offhand, I must admit, it does sound like a contradiction. Stalnaker's theory does, and mine does not, respect the opinion of any ordinary language speaker who cares to insist that it is a contradiction. But the cost of respecting this offhand opinion is too much. However little there is to choose for closeness between worlds where Bizet and Verdi are compatriots by both being Italian and worlds where they are compatriots by both being French, the selection function still must choose. I do not think it can choose – not if it is based entirely on comparative similarity, anyhow. Comparative similarity permits ties, and Stalnaker's selection function does not (D. Lewis, 1973, p.80).

So, the first two problems Lewis raises against SA is that it requires us to assume both LA and CEM.

A third difficulty for views like Stalnaker's is that it can't account for the difference in meaning between “would” counterfactuals ($\phi \Box \rightarrow \psi$) and “might” counterfactuals ($\phi \Diamond \rightarrow \psi$) in English. Stalnaker gives us *one* way to determine the truth value of counterfactuals in general. Thus, might-counterfactuals and their corresponding would-counterfactuals get assigned the same truth value. Consider the following two sentences:

(19) If the Chinese were to enter the Vietnam conflict, the United States would use nuclear weapons (i.e. ($\phi \Box \rightarrow \psi$))

(20) If the Chinese were to enter the Vietnam conflict, the United States might use nuclear weapons (i.e. ($\phi \Diamond \rightarrow \psi$))

To determine the truth values for (19) Stalnaker tells us to consider the world most like our own but where it is true that the Chinese enter the Vietnam conflict. We then ask ourselves if it is true, in that world, that the United States uses nuclear weapons. Suppose it is. Then, on Stalnaker's semantics, (19) is true.

But if there is always one and only one closest antecedent world, then *every* counterfactual is a would-counterfactual. If we use the same analysis to determine the truth value of (20), we get the same results as we did for (19). Stalnaker's semantics can't account for the differences between "mights" and "woulds" because there are never ties. There is always a uniquely closest antecedent world to the actual world. As Lewis puts it:

Another manifestation of Stalnaker's Assumption is that, except in the vacuous case, the difference between 'would' and 'might' counterfactuals is lost. Given my definition of $\diamond \rightarrow$ (namely $\phi \diamond \rightarrow \psi =_{df} \sim(\phi \square \rightarrow \sim \psi)$), it is a consequence of Conditional Excluded Middle that $\phi \diamond \rightarrow \psi$ implies $\phi \square \rightarrow \psi$; and $\phi \square \rightarrow \psi$ implies $\phi \diamond \rightarrow \psi$, except in the vacuous case, both on Stalnaker's theory and on mine. Hence, $\phi \square \rightarrow \psi$ and $\phi \diamond \rightarrow \psi$ cannot differ in truth value for Stalnaker, except in the vacuous case. But surely English 'would' and 'might' counterfactuals do sometimes differ in truth value, and not only in the vacuous case (D. Lewis, 1973, p.80).

1.2.1.3 Lewis' Semantics of Counterfactuals

Lewis' semantics of counterfactuals was designed to avoid the pitfalls of Stalnaker's account. Unlike Stalnaker, Lewis doesn't assume there is one unique world closest to the actual world where the antecedent is true. Basically, Lewis' account is: $\phi \square \rightarrow \psi$ is true, if some $\phi \psi$ -world is closer to the actual world than any $\phi \sim \psi$ -worlds (D. Lewis, 1973, p.16).

More needs to be said though. For example, what does it mean for one world to be closer than another to the actual world? Stalnaker uses a selection function to

determine the relative similarity worlds bear to the actual world. Lewis uses a system of spheres to convey information about the overall comparative similarity of worlds.

The system of spheres is a centered system in that the world about which the counterfactual is asserted is at the center of the system of spheres. So, if a system of spheres $\$$ is centered on a world i , then the members of $\$_i$ are just the spheres around i . The spheres of $\$_i$ are nested. The smaller a sphere that a possible world is located in, the more similar that world is to i .

It's important that this way of representing comparative similarity allows for there to be no single nearest antecedent world to i in both of the two ways that Stalnaker's failed to. First, if there is a sequence of smaller and smaller spheres *ad infinitum* in $\$_i$, then there is no smallest sphere. Second, Lewis's system of comparative similarity is such that if there is a smallest antecedent permitting sphere in $\$_i$, it allows for two or more worlds to be in that sphere. So, Lewis' picture can handle cases of ties where the closest sphere to i , includes both a $\phi\psi$ -world and a $\phi\sim\psi$ -world (D. Lewis, 1973, pp. 19-20).

With this account of comparative similarity in place, we can now state Lewis' semantics for "would" counterfactual more formally:

$\phi\Box\rightarrow\psi$ is true at a world i (according to a system of spheres $\$$) if and only if either
(1) no ϕ -world belongs to S in $\$_i$, or
(2) some sphere S in $\$_i$ does contain at least one ϕ -world, and $\phi\rightarrow\psi$ holds at every world in S (D. Lewis, 1973, p.16).

The counterfactual $\phi\Box\rightarrow\psi$, then, is true in one of two cases. First, it is true in the vacuous case if there are no accessible, or relevantly similar, ϕ -worlds to i . Second, it is true in the non-vacuous case if the closest ϕ -worlds to the actual world

are also ψ -worlds. So $\phi \Box \rightarrow \psi$ is true even if there is no uniquely closest ϕ -world that is also a ψ -world. There may be two or more ϕ and ψ -worlds that are equidistant from the actual world. These ϕ -worlds count as different worlds because they differ from each other in some way. They count as equidistant from i if their overall similarity to i is the same.

I've already pointed out how Lewis' system for representing comparative similarity can handle both cases of indeterminate similarities involved in continuums and cases of ties. So far, Lewis hasn't told us how those cases should be accounted for semantically. He accounts for these cases by introducing "might" counterfactuals. A might counterfactual is defined in terms of the would counterfactual: $\phi \Diamond \rightarrow \psi =_{df} \sim(\phi \Box \rightarrow \sim\psi)$. $\phi \Diamond \rightarrow \psi$ is true in the cases where every sphere around a world that contains a ϕ -world also contains a world where ϕ & ψ hold. This includes cases where there is a closest $\phi\psi$ -world (i.e. when the would counterfactual is true) but it also includes cases where every sphere containing a ϕ -world also contains $\phi\psi$ -worlds and $\phi\sim\psi$ -worlds. More formally:

$\phi \Diamond \rightarrow \psi$ is true at a world i (according to a system of spheres S_i if and only if both
 (1) some ϕ -world belongs to S in S_i , and
 (2) every sphere S in S_i that contains at least one ϕ -world contains at least one world where ϕ & ψ holds (D. Lewis, 1973, p.21).

1.2.2 Undermining CEM Undermines an Important Rationale for Molinism

One common objection to God's having middle knowledge is not that CCFs are not true but rather that an important reason offered to think CCFs are true is

mistaken. Plantinga thinks that CEM implies that there are true CCFs. He asserts, for example that for counterfactual pairs like (14) and (15), one of them but not both has to be true:

It seems clear that at least one of these conditionals is true, but naturally they can't both be true; so exactly one is...[and] we are now in a position to grasp an important fact. Either...[(14) or (15) is in fact true: and either way there are possible worlds God could not have actualized (Plantinga, 1977, p.41).

Undermining CEM undermines an important reason for thinking there are true CCFs. In the section above which sets out the semantics of counterfactuals, I explained Lewis' arguments against CEM. I take it that those arguments constitute compelling reasons to think that CEM is false and, consequently, undermine an important rationale for bare-bones Molinism.

1.2.3 The Grounding Objection

Perhaps the most important criticism of Molinism is this: if CCFs are true, what are the *grounds* for that truth? The grounds being asked for here are not epistemic grounds but metaphysical grounds. The question is what makes a CCF true? The objection is that the claim that CCFs are true is implausible because it is hard to imagine what kind of thing the truth of CCFs could be grounded in. In other words, its not even clear what it would mean for a CCF to be true (R. M. Adams, 1977, p.111) (Hasker, 1989, p.29).

Initially, one might think that CCFs are grounded in the psychological make-up or character of the free agent about which they are true. Remember, though, that we are talking about freewill in a libertarian sense. If someone has libertarian freedom with respect to some decision, then psychological facts can't be a sufficient

explanation for why a free agent makes a particular free choice. If CCFs were grounded in psychological facts alone, then that choice is determined by psychological facts about the agent and this is usually regarded as inconsistent with libertarian freedom.

William Hasker points something interesting out in this regard. When Molinists offer examples of CCFs that we seem to know the truth value for, they usually are of the following sort: “If Bob Adams were to offer to take me [Plantinga] rock climbing at Tahquitz Rock the next time I come to California, I would gladly (and freely) accept” (Plantinga, 1985b, p.373).²⁴ And Robert Adams concedes that, “there does not normally seem to be any uncertainty at all about what a butcher, for example, would have done if I had asked him to sell me a pound of ground beef, although we suppose he would have free will in the matter” (R. M. Adams, 1977, p.115). Hasker makes the following observation about these kinds of examples:

In virtually every case where we seem to have plausible examples of true counterfactuals of freedom, the plausibility is grounded in...psychological facts. (Again we recall ...Plantinga’s climbing Tahquitz Rock, and Adam’s butcher selling him a pound of hamburger) [Emphasis Hasker’s] (Hasker, 1989, p. 31).

It is psychological facts about Plantinga (e.g. his strong desire to climb rocks) and the butcher (e.g. his desire to want to make money selling meat), that we think grounds the truth of those CCFs.

Plantinga doesn’t try to ground CCFs in psychological facts about free agents. Instead, he argues that, assuming the demand for grounds is clear, CCFs aren’t different than other kinds of propositions in this regard:

²⁴ Plantinga has a reputation for being an avid rock climber.

Suppose, then, that yesterday I freely performed some action *A*. What was it or is it that grounded or founded my doing so? I wasn't *caused* to do so by anything else; nothing relevant *entails* that I did so. So, what grounds the truth of the proposition in question? Perhaps you will say that the grounds of its truth is just the fact that I did *A*. But this isn't much of an answer; and at any rate the same kind of answer is available in the case of Curley. For what grounds the truth of the counterfactual, we may say, is just the fact that Curley is such that if he had been offered a \$35,000 bribe, he would have freely taken it (Plantinga, 1985b, p.374).

Plantinga's reply amounts to the claim that CCFs are just basic facts. They are not true in virtue of anything else. If it isn't problematic that other propositions aren't grounded in something else then why, Plantinga asks, is it odd to think that CCFs are not simply ultimate facts as well? They are just bare facts about what a free agent would do if placed in certain circumstances.

Critics of middle knowledge are not impressed with this reply. Although there may be ultimate facts, CCFs don't seem to be good candidates. Other types of conditionals, like the material conditional or causal conditionals are grounded in the way things are. There is something counterintuitive to the idea that CCFs are just ultimate facts grounded in nothing. William Hasker puts it this way:

The intuition is this: In order for a (contingent) conditional state of affairs to obtain, its obtaining must be grounded in some categorical state of affairs. More colloquially, truths about "what *would be the case...if*" must be grounded in truths about what *is in fact* the case. This requirement seems clearly to be satisfied for the more familiar types of conditionals. The truth of the material conditional is grounded either in the truth of its consequent, or the falsity of its antecedent, or both. More interestingly, the truth of causal conditionals, and of their associated counterfactuals, are grounded in the natures, casual powers, inherent tendencies, and the like, of the natural entities described in them (Hasker, 1989, p.30).

What sort of reply can Molinists like Plantinga make to Hasker? One possible line runs like this: true CCFs are grounded in the *agent*. Suppose, for the sake of argument, that something like agent causation is true. This means that it is true of

Curley that he would take the bribe and that this is true of Curley *simpliciter*. This fact about Curley is ultimately grounded in Curley and not in something else. This might seem might appear to be question begging or asserting the conclusion rather than arguing for it.

However, if a libertarian view of freedom is true, then the facts about what a free agent would do in a particular circumstance is just supposed to be made true by the agent himself. Think about the dialectic of the debate over freewill. Libertarians argue that determinism excludes freewill (and moral responsibility) because the agent's choice is not ultimately grounded in the agent but rather it is grounded in, or determined by something else (e.g. DNA and environment). Libertarians push the intuition that genuine freewill depends on choices being ultimately made by the *agent*.

Conversely, determinists sometimes argue (e.g. Hobart) that indeterminism undermines freewill (and moral responsibility) because it is not the *agent* per se making the choice, but rather it is indeterminate or random conditions that determine the outcome of the choice. Both sides seem to agree that freewill (and moral responsibility) require a certain kind of control by the agent. They agree freewill requires that it really be the *agent* making the choice. They disagree about which view of freewill gives the agent this kind of control over his actions. They accuse each other, in other words, of locating what ultimately determines the choice as something other than the agent.

Here's one way to think about this reply: Hasker rightly points out that CCFs can't be grounded in psychological facts, physical facts, etc. But isn't this just

another way of saying that CCFs just have to be ultimate facts about agents? If something like this characterizes a portion of the debate over freewill, then it is not so implausible to think that CCFs are just grounded in the agent *simpliciter*. This sort of reply seems to rely on some kind of agent causation. Critics of Molnism might reply that to the extent that the claim that CCFs are supposed to be basic facts about agents depends on something like agent causation, it depends on a thesis as obscure as the original claim. I leave it for the reader to decide.²⁵

However, one might make a more general point. The problem with Molinism isn't that it makes CCFs bare facts about agents. The mistake Molinists make is that they fail to capture our intuitions. What we are really saying when we assert that a person is free with respect to an action *A* is that the person *might* do *A* or might not do *A*. If Lewis gets the semantics of "might" counterfactuals right, then the truth of the above two might-counterfactuals rules out the truth of both corresponding would-counterfactuals.

1.2.4 The Circularity Objection

Robert Adams and Anthony Kenny argue that God's having middle knowledge involves a certain sort of circularity. This objection assumes that there are true CCFs and that God knows them. It denies, though, that God's knowledge of them is logically prior to God's knowledge of which world is the actual world. This means that God's middle knowledge is of no use to him when trying to actualize a

²⁵ Critics will, of course, want to know what an agent *simpliciter* is. When we remove psychological and physical facts and physical, facts, etc. from an agent, what precisely is it that is left over? They will want to know what exactly is it that we are being asked to imagine here. Many contend that the concept of agent causation is itself obscure and mysterious. Defenders of agent causation are not without replies but I will leave the details of the discussion over freewill for the reader to pursue.

world with moral good but no moral evil. According to the standard semantics of counterfactuals, the truth value of a counterfactual depends on whether the closest antecedent-worlds to the actual world are also consequent-worlds. Whether a CCF is true, then, depends on which world is the actual world. However, which world is the actual world depends on whether the antecedent is actualized. Adams puts it this way:

Consider a deliberative conditional,

(9) If I did x , y would happen.

Is (9) true? According to the possible worlds explanations, that depends on whether the actual world is more similar to some world in which I do x and y happens than to any world in which I do x and y does not happen. That in turn seems to depend on which world is actual. And which world is the actual world? That depends in part on whether I do x . Thus, the truth of (9) seems to depend on the truth or falsity of its antecedent (R. M. Adams, 1977, p.113-114).

Anthony Kenney contends that:

Prior to God's decision to actualize a particular world those counterfactuals cannot yet be known: for their truth-value depends...on which world is the actual world. It is not simply that God's knowledge of these counterfactuals cannot be based on a decision which has to be taken subsequent to knowledge of them...the problem is that what makes the counterfactuals true is not yet there at any stage at which it is undecided which world is the actual world...the difficulty is simply that if it is to be possible for God to know which world he is actualizing, then his middle knowledge must be logically prior to his decision to actualize; whereas, if middle knowledge is to have an object, the actualization must already have taken place. As long as it is undetermined which action an individual human being will take it is undetermined which possible world is the actual world—undetermined not just epistemically, but metaphysically (Kenny, 1979, pp. 68-71).

Edward Wierenga has called this the “not true soon enough” objection (Wierenga, 1989, pp. 148-149).

Plantinga's reply to this objection is that the sense of “depends on” that

Adams and others are using is not transitive:

This argument, I fear, does not warrant the trust Adams apparently places on it. It is true that:

(1) the truth of [(9)] depends on which world is actual

In the sense that [(9)] is true in some worlds and false in others; it is also true that

(2) which world is actual depends on whether the antecedent of [(9)] is true

Again, in the sense that the antecedent of [(9)] is true in some worlds and false in others. It doesn't follow, however, that the truth of [(9)] depends on the truth of the antecedent. Consider the following analogue:

(1*) the truth of *The Allies won the Second World War* depends on which world is actual;

(2*) which world is actual depends on whether I mow my lawn this afternoon;

Therefore

(3*) the truth of *The Allies won the Second World War* depends on whether I mow my lawn this afternoon.

Clearly the relation expressed by the relevant sense of 'depends' isn't transitive (Plantinga, 1985a, p.376).

Plantinga is right that the circularity objection as expressed by Adams and Kenny is mistaken. But I am not sure he has properly diagnosed the problem with it. The problem is not that the relevant sense of "depends on" is not transitive. Rather, the problem is that the sense of "depends on" is importantly different in the two cases. A kind of equivocation is taking place here.

The truth of a counterfactual depends on the layout of modal space. It depends on the external relations between possible worlds. Different counterfactuals are true at different possible worlds. If Lewis is right and the term "actual" is an indexical term, then whether an antecedent of particular counterfactual is true (i.e. actual) depends on "where" we are in modal space. The dependency relation in the case of counterfactuals is extrinsic in that whether a counterfactual is true at a

particular possible world is independent of which world is actual. However, whether an antecedent of a counterfactual is actual isn't independent of which world is actual.

We have considered three objections to bare-bones Molinism, so far. The first objection we considered was directed at positive reason for thinking there are true CCFs. If this objection to CEM is right, at best it takes away a reason for Molinism but it isn't, strictly speaking, a reason for thinking Molinism is false. The grounding objection, if right, undermines the truth of Molinism. The grounding objection is a serious objection but the Molinist has a reasonable move to make by arguing that CCFs are bare facts and that the nature of libertarian freedom makes this somewhat plausible. The circularity objection rests on a conflating two kinds of dependency. Bare-bones Molinism has taken some serious blows but has not been clearly shown to be false.

2. The FWD without Molinism

At this point in the dialectic, we should pose the question: if Molinism is false, does this undermine the freewill defense? At first glance, it might seem so because CCFs function as a sort of legitimate limitation on God in the freewill defense. It is a legitimate limitation in the sense that it doesn't undermine or diminish God's divine credentials. If CCFs are removed from the equation, then an important source of legitimate constraints on God disappear. Without true CCFs, God has no excuse.

On the other hand, if God doesn't know what a free agent would do in various counterfactual circumstances how can he be responsible for the results of his creatures' free agency? He could very well have created the best world he was

capable of and didn't know it would result in so much suffering. So, God's lack of knowledge of CCFs seems to provide him a moral excuse for the results of his creatures' free agency. Plantinga seems to think this latter situation is the case:

Without the assumption of middle knowledge it is much harder to formulate a plausible deductive argument from evil; and it is correspondingly much easier, I should think, to formulate the free will defense on the assumption that middle knowledge is impossible (Plantinga, 1985a, p.379).

But is he right?

Molinism can fail for two reasons. Either all CCFs are false or some are true but God doesn't know them. If the latter is the case, God isn't omniscient because divine omniscience requires that God knows everything that is knowable. So, let's lay this option aside and focus on the former possibility. If CCFs are false, then God can strongly actualize states of affair with free agents but because he doesn't know what they would freely chose, he doesn't weakly actualize anything. So, he is not morally responsible for what his creatures strongly actualize.

Another way to put it is this: Mackie assumes that it is within God's power to weakly actualize any possible world. Thus, he assumes it is possible for God to weakly actualize a world with free agents that always freely chose to perform good acts. But if God doesn't have middle knowledge he can't weakly actualize that possible world, or any other possible world for that matter. Since Mackie's argument depends on this assumption, his argument is undermined if God lacks middle knowledge.

So, Plantinga's version of the free will defense depends on God having middle knowledge. If God lacks middle knowledge Plantinga's free will defense fails. Interestingly enough, Mackie's argument from evil also depends on God having

middle knowledge. Mackie's original objection and Plantinga's free will defense share a common assumption: Molinism. If this assumption is false, both Mackie's objection and Plantinga's response fail.

On the one hand neither side can make their case. On the other hand, Mackie was taking a positive burden on himself to *demonstrate* an incompatibility here. Remember: Mackie's argument is that "...it can be shown, not that religious beliefs lacks rational support, but that they are positively irrational, that several parts of the essential theological doctrine are inconsistent with one another" (Mackie, 1955, p.200). If his demonstration fails to get off the ground because a crucial assumption is false, does it really matter that Plantinga's free will defense against that demonstration also fails for the same reason?

2.1 The Risk-taking Objection

Mackie's argument, though, can be modified to avoid this result if it is reformulated so that it doesn't depend on the truth of Molinism. The reformulated logical argument from evil centers on the claim that a being that is perfectly good would not take the kind of risks God would have to take in creating the universe without middle knowledge. Call this the "risk-taking objection."

Thomas Flint, a defender of Molinism, argues that:

...all things considered, God's lacking middle knowledge would make the problem of evil even more difficult for the Christian to handle. For if God knows only probabilities, then he takes enormous risks in creating significantly free beings: he risks creating a world in which many, or most, or even all of his free creatures consistently reject him, a world in which they use their freedom to degrade others and themselves. It seems to me that one can reasonably argue that a good and loving God would not take such a risk (Flint, 1998, p.107).

Kenneth Perszyk agrees that if Molinism is false, then a different version of Mackie's argument from evil centered on the risk-taking objection succeeds. Perszyk appeals to the same kind of intuition as Flint:

...creation is a serious business, and...caution (not luck) is a moral category, especially in light of cognitive limitations. If God didn't ('antecedently') know what any (libertarian) free creatures would do, would he create them? If God knew that free creatures might never produce moral good, or might produce more evil overall than good, or might even produce only moral evil, would he create them? If God's primary end or plan in creation is for free creatures to enter into loving personal relationships with him, but he didn't know that any would do so, would he create them? It's not at all obvious that he would, or that doing so would be compatible with his perfect goodness (Perszyk, 1998, p.49).

So Flint and Perszyk both see the potential failure of Molinism as a net plus for the logical argument from evil.

So, what are we to make of the risk-taking objection? To begin with, there are obvious cases of risk-taking that are perfectly moral. People risk death every time they drive to work. Surely, it is not immoral, or even less than perfectly moral, to drive to work because it involves some risk of death. This means that risk-taking *per se* doesn't diminish the moral status of the risk-taker. Many proponents of the risk-taking objection speak of risk-taking in such general terms that they sometime seem to be asserting that if God took any risk, it would be immoral. This is clearly false, if God's risk-taking is of the same sort we take when we drive to work.

There are, however, also cases of risk-taking that are immoral, say betting your family's house and car on the roulette wheel at a casino. If the kind of risk God takes in creating free creatures is akin to "betting the house," then the risk-taking objection has merit. So there are morally acceptable, even obligatory, risks and

immoral risks. The risk-taking objection depends on which category of risk-taking God's creating free creatures fall under.

When is risk-taking immoral? Thomas Flint offers a two rules of thumb concerning immoral risk-taking. First, he argues risk-taking is less moral the higher the stakes are:

First, the greater the stakes, the less we generally want to engage in risky endeavors. Buying a lottery ticket engages me in a relatively risky enterprise, one where my odds of success are quite small. Provided that I am wagering only a dollar, though, most people would find nothing objectionable in such a gamble. But if the bet is substantially higher, so high I risk losing my house, car, job, and the like, most of us would insist that we have gone beyond the realm of innocent amusement (Flint, 1998, pp.104-105).

Second, Flint argues that taking risks with other people is more morally perilous:

...where others are involved, especially others to whom we have some special relationship, we are especially concerned to lessen or eliminate the risks our actions may pose to them. For example, it is arguably immoral for a father to drive without wearing his seat belt. But it seems clearly worse for him to drive and not provide a proper child-restraint seat for his two-year-old daughter. If he could completely eliminate the risks which driving poses for her, surely he would do it in an instant (Flint, 1998, p.105).

What should the defender of the free will defense without Molinism say about Flint's principles of the morality of risk-taking? First, both principles have some merit. They are intuitive. But just as importantly, they are incomplete. Flint isn't giving us a complete account of the morality of risk-taking. This point isn't meant as an objection to Flint because he never intended to give such an account. But even the kinds of cases he cites are more complicated than he lets on.

He is right that how high the stakes are has bearing on the moral status of a particular act of risk-taking but more needs to be said. Suppose, for example, Bill Gates is worth fifty billion dollars and he bets forty-nine billion dollars on a lottery

worth forty-nine billion dollars. Suppose further that with forty-nine billion dollars he can buy enough lottery tickets to make it an even bet. The stakes are enormously high but most of us would not say he acted immorally. If there is any sense of moral perverseness in this kind of case it involves “playing” with that amount of money rather than using it for more productive means. The high stakes by themselves don’t seem to make a bet immoral.

Flint might argue that the stakes aren’t really high because if Gates loses his bet, he still has a cool billion in the bank. It’s odd at the very least to argue that a bet with someone standing to lose forty-nine billion dollars is not a high stakes bet. However, at this point we should take care not to get into a needless verbal dispute over the meaning of “high stakes.”

What Flint’s example of gambling one’s house, car, and job away and my example of Bill Gates gambling forty-nine of his fifty billion dollars show us is that it is not high stakes *per se* that make a bet immoral but making a bet that puts a minimal standard of living at risk. In Flint’s example, one stands to lose the basic ingredients in a modern Western life. In the Bill Gates example, the stakes are high but a minimal standard of living is not at risk. This suggests that there is a floor to risk-taking. If we engage in risky behavior that threatens a minimum standard of living, then that behavior may be immoral. However, it *is* sometimes moral to take risks that threaten important and basic goods (e.g. our life). I’ll say more about this in discussing Flint’s second principle.

Now let’s consider Flint’s second principle regarding the morality of taking risks with others. People take risks with their children every day. We take a risk

when we drive our children to school. The fact that people put their children in car seats doesn't eliminate the risk. It decreases it but the risk could be decreased even more by not bringing one's child to school and home school instead. We take risks with others everyday and much of the time this is morally acceptable.

What's important is *why* we are taking a risk and *how likely* the risky outcome is to occur. Flint and other Molinists tend to overstate the value of minimizing risk. Caution is morally important but so is risk-taking. Risk-taking is not only morally acceptable in many cases, like driving a child to school it is sometimes morally *required*. It is immoral to not let your child get an education, or visit grandparents, play outside, etc. in order to minimize risk.

Flint's remarks about risk-taking tend to center around risk-taking for the sake of entertainment and fun. For example, Flint consistently discusses risk-taking in the context of game-playing and entertainment. Commenting on the father's insistence that his daughter wear a safety belt, Flint writes:

Openist cries of "How dull!"; "Take a chance!"; "Live a little!" and the like would presumably have minimal effect [on the father]...though we have used the image of playing games extensively in this chapter, it is worth reminding ourselves that, for the Christian at least, our relationship with God is not most properly as a matter of fun and games; it is a matter of life and death, in the truest and fullest sense of those terms (Flint, 1998, p.105).

If the valuable and important risks we take in life are mainly about gambling, game-playing, and entertainment, then Flint might have a point. But his examples are misleading. We take risks for much more important and valuable reasons than fun and games. We do so to achieve enriching elements of the good life.

To see this imagine a person that minimizes risk as far as is feasibly possible. This person doesn't have a job because traveling to work involves risk of death and

because it is easier to get a communicable disease in concentrations of humans (e.g. schools, offices, etc). Instead, he lives a meager life from maybe a modest inheritance, or collects government and/or private charity instead. This person might not go on romantic dates or partake in community life all to minimize risk. Perhaps he agrees with Flint that we need to minimize risk when it involves others, so he decides to play it safe and not get married or have children. Most of us would see that this kind of safe life is deeply unfulfilling not because the fun and games are missing from it but because it lacks, or minimizes, important goods like family and other relationships, job satisfaction and productivity, etc. Taking a certain amount of risk to achieve these goals is perfectly moral.

Returning to the free will defense, what we should take away from this discussion of Flint's two principles about the morality of risk-taking is this: whether it is immoral for God to take the risks involved with creating free agents depends on (1) why he is taking the risk and (2) how likely he is to achieve his goal. The free will defense simply requires that it is *possible* that such risk-taking is morally justified.

If it is possible that God's reason for assuming the risks of creating free agents is to achieve the kinds of deeply valuable goods mentioned above rather than for mere entertainment value, then God's taking such a risk might be justified. This is precisely the claim the free will defense makes: it is possible that a world with significantly free creatures but suffering is more valuable than a world without suffering but also lacking free agents. This seems possible in the broadly logical sense of possibility.

What about the second principle: how likely is God to achieve his goal?

Presumably if God has knowledge of the conditional probabilities of achieving his goal, given his creation of the world, he would minimize the risk of failure. Failing to achieve a world with enough moral good to outweigh any moral evil is still possible on this account of God's knowledge but it is minimized if God has sufficient knowledge of conditional probabilities.

But does God have such knowledge? Ordinary humans are often familiar with the behavioral tendencies of people we know. We know how our spouses, children, friends, or colleagues are likely to act in certain scenarios. If we mere "cognitive mortals" have such knowledge it is hard to see how God doesn't. Put alternatively, if conditional probabilities are knowable (because we humans know some about the people closest to us), then an omniscient God knows them.²⁶

A worry arises here. It's possible that there are no objective probabilities regarding the behavior of free agents but that we think there are because we assign a subjective probability to a free agent's behavior. We give a certain amount of credence to the belief that our spouse will behave a certain way in a certain situation, even though there is no objective fact regarding the probability of her so behaving. The worry is that we confuse subjective and objective possibility such that we take our subjective assigning of credence to a belief as an objective probability about the belief.

²⁶ Some theists, like Richard Swinburne argues that there are knowable truths that even God can't know. For example, the statement "it is now $t1$...can only be known at $t1$." This also includes *de se* truths regarding others. Thus, Swinburne defines omniscience as God knowing everything it is logically possible to know (Swinburne, 1995)

There is, however, strong *prima facie* reason to think that there are objective patterns of behavior that free agents tend to follow. Various free agents have a tendency to behave certain ways in certain circumstances. Assuming libertarianism is true, much of the work in the social sciences confirms that free agencies have behavioral tendencies. Also there is strong anecdotal evidence that people really can predict to at least to a certain extent the behavior of people they know well.

Of course our knowledge of conditional probabilities is limited, so God's might also be limited. Although free agents have tendencies to behave in certain ways, they are just that: tendencies. Perhaps God knows every conditional probability about free agents that is knowable and that amount of knowledge isn't enough to make reasonable assessments of risk. Maybe people aren't predictable enough to take the kinds of risks involved in creating free agents.

It is important to keep in mind, though, that all the free will defense needs is for it to be logically possible for God to have sufficient knowledge of conditional probabilities. Is it logically possible that God has sufficient knowledge of the behavior tendencies of free agents to risk creating them? It is hard to see how such a possibility is *logically* impossible. In other words, it's hard to see a *contradiction* involved in making such a claim.

What are we to make of the risk-taking objection? First, that risk-taking *per se* is not immoral. The mere fact that God takes risks isn't sufficient reason to think that God is less than perfectly moral. Molinists and proponents of the form of the logical argument from evil that lacks Molinism need to go further and show that the kind of risk God takes in creating free agents is immoral. At best, Flint as provided

us with the first premise of a two-premise argument. As it stands, his argument is incomplete.

2.2 *The Reckless Risk-taking objection*

What the Molinist would need to show is that God is guilty of *reckless* risk-taking. There are two ways to argue that God's creating free agents involves reckless risk-taking. First, if God has some but not enough knowledge of the conditional probability that the world would have a favorable pattern of moral good to moral evil. Second, if God creates free agents without *any* knowledge of conditional probabilities.

If God exists, what reason would one have to think he didn't have sufficient knowledge of probabilities to risk creating free agents? Presumably the nature of free will is such that either God can't know at all what a free agent would probably do in certain circumstance or he can know to a sufficiently limited extent.

The free will defense only requires that it be logically possible that God has sufficient knowledge of the relevant conditional probabilities. So the proponent of the reckless risk-taking objection is asserting the very strong claim that libertarianism *requires* either that free agents have no behavioral tendencies or that they not have enough for God to make reasonable predictions. *Prima facie* there is no reason to think that libertarian freedom by its very nature precludes free agents from having behavioral tendencies. And so far as I can tell, no prominent advocate of libertarianism denies that free agents behave according to certain tendencies. The

defender of the reckless-risk objection bears the burden of proof here, given the strong nature of the claim and the lack of any obvious reasons to hold it.

Further, it is possible that God can limit the risk involved in creating free agents by miraculously intervening in the world. Initially one might think that God's intervening and overriding an agent's freewill undermines the free will defense by removing free will. But nothing about the free will defense requires that *every* act be free. If someone acts freely 90% of the time and is determined 10% of the time, surely we would still count that person as having significant freedom. To be sure, there is a limit to how much God can intervene in order to avoid undermining the free will defense but the fact that God can intervene to prevent *some* evil means he can reduce the risk he is taking to *some* extent.

So how does the free will defense fare without Molinism? It seems to be better off than it would be if Molinism were true. With less knowledge of consequences comes less responsibility for results. However, with too little knowledge of consequences the charge of reckless risk-taking arises. To withstand the charge of reckless risk-taking the free will defender only need it to be *logically* possible that God has enough knowledge of conditional probabilities to avoid the charge of reckless risk-taking. Arguing that something is logically possible is harder than showing it is logically impossible but generally, it is reasonable to believe that some claim is logically possible if there is no apparent contradiction in holding it

Chapter 3: Transworld Depravity and Transworld Sanctity

In “Transworld Sanctity and Plantinga’s Free Will Defense,” Daniel Howard-Snyder and John Hawthorne aim to show that Plantinga’s freewill defense fails (Howard-Snyder & Hawthorne, 1999, pp.41-42). Howard-Snyder and Hawthorne are quick to add that their “aims are not ultimately unfriendly; indeed, by understanding why Plantinga’s defense fails, one can see what he should have said, which by our lights is every bit as decisive as his defense has been thought to be” (Howard-Snyder & Hawthorne, 1999, pp.1-2).

The authors argue that Plantinga’s freewill defense fails because its goal is too ambitious. Plantinga’s goal is to *demonstrate* that the existence of evil is logically compatible with the existence of God. Demonstration inherently has an epistemic element. Plantinga’s project, they argue, needs to meet a minimum epistemic threshold in order to be a successful demonstration. Howard-Snyder and Hawthorne contend that the freewill defense fails to meet such a minimum epistemic standard and, thus, fails as a demonstration of the compatibility of the existence of God and the existence of evil.

The author’s friendly gesture is their argument that the freewill defense succeeds at establishing the more modest claim that it is *not reasonable to believe* that the existence of God and the existence of evil are incompatible. The upshot of their article is this: Plantinga should have made the more modest claim that it is not reasonable to believe that the existence of God and the existence of evil are

incompatible because that claim can be successfully defended (Howard-Snyder & Hawthorne, 1999, pp. 14-18).

In this chapter I will argue for a thesis that is in some ways stronger than Howard-Snyder's and Hawthorne's and in other ways weaker. I will argue that the essential role the epistemic element of the freewill defense plays introduces a person-relative element to the freewill defense. Beliefs, justification, and knowledge are person-relative in the sense that different people have different beliefs, justification (or lack thereof) for their beliefs, and knowledge. Therefore, there is something epistemically ham-handed about discussing what is reasonable or unreasonable to believe (or refrain from believing) *simpliciter*. Rather, discussions about God and evil, the freewill defense, and transworld depravity should be epistemologically nuanced in a way that acknowledges that what is reasonable or unreasonable to believe varies in important ways from person to person. This is particularly the case in arguments that depend on very abstract (and somewhat fragile) modal judgments and intuitions.

1. The Epistemic Amendment and Transworld Sanctity

Recall that Plantinga's strategy for demonstrating $\diamond(G \& E)$ is to give a model where G and E are both true. He employs the following procedure: for any propositions P and Q, (P&Q) is possible if there is a third proposition R such that R is possible, P and R are consistent, and P and R together entail Q (Plantinga, 1977, p.26). In particular, he argues that God and evil are consistent because there is a third possibly true set of propositions, which make up the freewill defense, such that the

existence of God is consistent with the freewill defense and the existence of God and the truth of the freewill defense together entail the existence of evil. So, let F stand for the freewill defense. According to Plantinga, $\diamond(G \ \& \ E)$ because $\diamond(G\&F)$ and $((G\&F)\rightarrow E)$ (Plantinga, 1977, p.29).

Howard-Snyder and Hawthorne contend that even if Plantinga's strategy succeeds, he hasn't necessarily *demonstrated* or *shown* that $\diamond(G \ \& \ E)$: "So what's our worry? This: to *show* that G is compatible with E is in part an epistemological task; thus, one succeeds at it only if the claims that constitute one's defense meet certain epistemic standards" (Howard-Snyder & Hawthorne, 1999, p.2).²⁷ They propose that the following minimal epistemic standard, which they call the "epistemic amendment," be added to the project of giving a defense:

²⁷ Howard-Snyder and Hawthorne credit Nelson Pike and Marilyn Adams with making a similar point about the importance of epistemic considerations in the debate over God and evil. However, what Pike and Adams say in the works cited by Howard-Snyder and Hawthorne are general enough and scant enough to make it unclear whether Pike is making a similar point.

While discussing the portions of Hume's *Dialogues Concerning Natural Religion* which are about the problem of evil, Pike points out that a certain strategy is not open to Philo. Philo can't *show* that God is not justified in permitting suffering by listing potential reasons for his permitting suffering and then showing that each one in turn is inadequate. Pike points out that Cleanthes could always argue that God's reasons for permitting evil might be one of the reasons Philo has yet to discuss. Cleanthes, Pike argues, could always retreat to unexamined reasons. Pike's point is apparently that Philo hasn't *shown* or *demonstrated* that God has no good reason for permitting evil because he hasn't examined every possible reason and shown them wanting (Pike, 1963, pp.41-42). Adams refers to Pike's argument as an "epistemic defense" and that is all she says in the work cited by the authors (M. M. Adams, 1990, p. 4).

Two points are worth mentioning in this regard. First, although Pike points out an epistemic move that Cleanthes could use to counter Philo's argument, it doesn't look to be the same point that Howard-Snyder and Hawthorne are making. In order for the point to be similar, it would have to be *reasonable* for Cleanthes to refrain from believing that God has no sufficient reason to permit suffering. And *that*, in turn, depends on whether it is *reasonable* to believe that the set of unexamined reasons might contain just such a morally sufficient reason. The epistemic amendment requires a proof to meet the minimal epistemic standard of reasonable to refrain from believing some relevant proposition. One might think that Pike's example requires a higher epistemic standard because refraining from believing X until possible reason against X has been examined and found wanting may strike some as unreasonable. Second, it's odd that the authors cite Marilyn Adams' discussion of Pike's article in the "Introduction" of her anthology on the problem of evil as evidence that she agrees with Pike's point. As far as I can tell, she was just explaining Pike's position without endorsing or rejecting it (M. M. Adams, 1990, p.4).

One shows that G is compatible with E by deploying a Plantinga-style defense only if it is not reasonable to refrain from believing those claims that constitute it (Howard-Snyder & Hawthorne, 1999, p.2).

The addition of the epistemic amendment to the project of defense means that for any defense D, it must not be reasonable to refrain from believing:

1. D specifies a *logically possible* reason for God to permit evil,
2. D is logically compatible with G, and
3. the conjunction of G and D entails E (Howard-Snyder & Hawthorne, 1999, p. 2)²⁸.

They also add the following corollary to the epistemic amendment, which they take to be a natural extension of it:

If one believes that D is possible solely on the basis of a particular argument and it is reasonable for one to refrain from believing at least one premise of that argument, then it is reasonable for one to refrain from believing that D is possible (Howard-Snyder & Hawthorne, 1999, p.3).

How does Plantinga's freewill defense measure up to the epistemic amendment? Let's consider the three elements of a defense which it must not be reasonable to refrain from believing, enumerated above in reverse order. First, consider 3. Is it reasonable to refrain from believing that the conjunction of "God exists" (G) and "the freewill defense is true" (D) entail "evil exists" (E)? It is not. If God exists and creates a world with significantly free creatures, then, given transworld depravity, evil must exist. That G and D entail E is obvious enough to make it unreasonable to refrain from believing it.

²⁸ First, this is not a direct quote, though the reference located at the end of the sentence might imply that. Howard-Snyder and Hawthorne later add a fourth proposition that it must not be unreasonable to refrain from believing, namely, "that D is possible." For simplicity's sake I built this fourth proposition into 1. Second, there is a slight difference in nomenclature here. I have been using "F" to refer to the relevant defense. They use "D." There is no substantive difference.

Second, consider 2. Is it reasonable to refrain from believing that the existence of God (G) is logically compatible with the claims made by Plantinga's freewill defense (D)? There is no apparent inconsistency between those claims. It seems unreasonable to refrain from believing that G and D are logically consistent without there being any apparent contradiction or some other reason to take them to be inconsistent. Of course, two claims might seem to be consistent and at the same time it is reasonable to refrain from believing they are consistent. For example, we might have empirical reason to think that we are bad at making those kinds of judgments (e.g. as our tendency to accept the Gambler's Fallacy shows empirically that we are inclined to make certain sorts of mistakes in probability reasoning).

Of course, whether they are consistent also depends on whether each claim, by itself, is coherent. To be consistent with each other, it has to be possible for the conjunction (G&D) to be true. If one conjunct is impossible, though, then the conjunction as a whole can't be true. We will consider that consideration next. So, it's not clear whether it is reasonable to refrain from believing 2.

Lastly, consider 1. This is where Howard-Snyder and Hawthorne think the problems lies. Is it reasonable to refrain from thinking that the freewill defense is possible? That depends, in part, on whether it is reasonable to refrain from thinking that transworld depravity is possible, because transworld depravity is an essential element of Plantinga's freewill defense.²⁹ This also follows, they assert, from their

²⁹ Other versions of the freewill defense do not depend on Molinism and thus do not depend on transworld depravity being possible. So, not all versions of the freewill defense depend on the possibility of transworld depravity being true of everyone.

corollary to the epistemic amendment.³⁰ So, if it is reasonable to refrain from believing that transworld depravity is possible, then it's reasonable to refrain from thinking that Plantinga's freewill defense is possible. And, in fact, the authors argue that it *is* reasonable to refrain from believing that it is possible that everyone suffers from transworld depravity, hence Plantinga's freewill defense fails the test set out by the epistemic amendment. Plantinga hasn't *shown* or *demonstrated* anything.

The heart of their argument involves the concept of transworld sanctity. What is transworld sanctity? Howard-Snyder and Hawthorne define it this way:

An essence *E* is blessed with transworld sanctity if and only if for every world *W* such that *E* contains the properties *is significantly free in W* and *always does what is right in W*, for **no** action *A* and for **no** maximal world segment *S* such that

- (1) *S* includes *E's being instantiated* and *E's instantiation being free with respect to A* and *A's being morally significant for E's instantiation*, and (2) *S* is included in *W* but includes neither *E's instantiation's performing A* nor *E's instantiation's refraining from A*, it is the case that
- (3) if *S* were actual, then the instantiation of *E* would have gone wrong with respect to *A*.

(The bold type indicates changes from Plantinga's definition of transworld depravity.) (Howard-Snyder & Hawthorne, 1999, pp.5-6)³¹

It won't have escaped the careful reader's attention that transworld sanctity looks very similar to transworld depravity:

An essence *E* suffers from transworld depravity if and only if for every world *W* such that *E* contains the properties *is significantly free in W* and *always does what is right in W*, there is an action *A* and a maximal world segment *S'* such that

- (1) *S'* includes *E's being instantiated* and *E's instantiation being free with respect to A* and *A's being morally significant for E's instantiation*,

³⁰ I am not sure how much it adds to explicitly define this corollary to the epistemic amendment. It seems clear enough that if it is reasonable to refrain from believing the premise of an argument, then it is reasonable to refrain from accepting the whole argument. Further, if that argument is the only reason to believe the conclusion, then it is reasonable to refrain from believing the conclusion. It just doesn't seem to add much.

³¹ Again, it is critical to notice that the third component is a counterfactual.

- (2) *S'* is included in *W* but includes neither *E's* instantiation's performing *A* nor *E's* instantiation's refraining from *A*
and
(3) if *S'* were actual, then the instantiation of *E* would have gone wrong with respect to *A* (Plantinga, 1977, pp. 52-53).

The only differences between Howard-Snyder and Hawthorne's definition of transworld sanctity and Plantinga's definition of transworld depravity are in bold font in the definition of transworld sanctity above.

The difference amounts to this. According to transworld depravity, if an essence *E* suffers from transworld depravity, then if *E* is actualized³², *E* will go wrong with respect to at least one action. According to transworld sanctity, if an essence *E* is blessed with transworld sanctity and *E* is actualized, *E* will *not* go wrong with respect to any action. Transworld depravity says that an actual essence that is transworld depraved will do at least one wrong thing, while transworld sanctity says that an essence that has it will never perform an immoral action.

With the definition of transworld sanctity in place, Howard-Snyder and Hawthorne ask us to consider that it is possible (epistemically) that³³:

TS. Necessarily, some essence or other is blessed with transworld sanctity (Howard-Snyder & Hawthorne, 1999, p.6).

They point out that they are not asking us to consider that some particular individual essence is transworld sanctified. Rather they are asking us to consider that, necessarily, some essence or other is transworld sanctified. Recall that Plantinga's assertion that transworld depravity means that it is possible that there is a world where the truth value of the counterfactuals of creaturely freedom are such that no

³² That this is a counterfactual statement is important to bear in mind.

³³ The possibility here is epistemic because their whole argument is epistemic. They contend that the epistemic amendment has not been met and hence Plantinga has not demonstrated consistency.

matter which essence is actualized it will go wrong at least once. The claim that it is possible (epistemically) that necessarily some essence or other is transworld sanctified is a claim about all possible worlds. If it is true, then there is at least one person in every possible world who does no wrong.

It's important to notice for clarity's sake that there is a difference between TS and transworld sanctity. Transworld sanctity is a property. TS is the claim that necessarily some essence or other has that property. For consistency, let TD stand for the claim that:

TD. It is possible that every essence suffers from transworld depravity. "TD" also does not stand for "transworld depravity," which is a property. TD is a proposition, which claims that it is possible that every essence is transworld depraved.

Howard-Snyder and Hawthorne point out that the possibility that everyone is transworld depraved and the necessity that at least one essence is transworld sanctified can't both be true. If TS is true, then TD is false. If TD is true, then TS is false. Why? TD claims that there is at least one possible world where, no matter which personal essence God instantiates, that essence will go wrong at least once. TS claims that in every possible world there is at least one essence that doesn't perform even one wrong act. If TS is right, then every world has one morally perfect saint. If every world has one moral saint, then there is no possible world where every essence goes wrong at least once. Conversely, if there is a possible world where every essence goes wrong at least once, then it is false that every world contains one person that never goes wrong.

Next Howard-Snyder and Hawthorne return to their main epistemological point. They argue that “for any propositions p and q , if we know that p entails q and it is reasonable to refrain from believing q , then it is reasonable to refrain from believing p ” (Howard-Snyder & Hawthorne, 1999, p.8). We know that TD entails the denial of TS. So, TS and TD can’t both be true. The fact that TS and TD have this logical relationship and we know it, has important epistemic implications.³⁴

They then make the following epistemic claim: if it is reasonable to refrain from believing not-TS, then it is reasonable to refrain from believing that TD. Given this epistemic claim, the important question is whether it is reasonable to refrain from believing that TS is false. Whether it is reasonable to so refrain depends on what we have to go on, epistemically. So, what *are* the relevant reasons in favor or against both TS and TD?

Let’s start with TD. What is there by way of epistemic support for TD? Surprisingly, Plantinga doesn’t offer much by way of argument or reason for his possibility claim. What he says in support of his claim that it is possible that everyone suffers from transworld depravity is that it is “interesting fact” and is “clearly true” (Plantinga, 1974, p.186, 188) (Plantinga, 1977, p.53). Howard-Snyder and Hawthorne point out that this doesn’t constitute a demonstration. Plantinga hasn’t *shown* it is possible, he has merely asserted it (Howard-Snyder & Hawthorne, 1999, p.5). The fact that Plantinga offers no reasons to accept TD is part of what makes it is reasonable to refrain from believing TD.

³⁴ It’s important to include that fact that we *know* the logical relation in order to avoid complications about whether knowledge is closed under entailment. It’s controversial whether the fact that we know p , and p entails q means that we also know q . The case above, however, is one where we know p , and we *know* that p entails q .

This is where TS enters the epistemic picture. The authors contend that we have as much reason to affirm TS as we do TD. There is a sort of epistemic stalemate here. If TS and TD are logically incompatible and it is as reasonable to believe TS as it is TD, then it is reasonable to refrain from believing TD. Further, it seems that if it's reasonable to entertain TS as something that *might* be a metaphysical possibility, then to that same extent, it's reasonable to *refrain from believing* that TD is a metaphysical possibility. If it is reasonable to refrain from believing TD, then it is reasonable to refrain from affirming the freewill defense. If it is reasonable to refrain from believing the freewill defense, then it is reasonable to refrain from believing that God and evil are compatible.

Howard-Snyder and Hawthorne are quick to point out that this works both ways. If it is reasonable to refrain from believing TS, then it's reasonable to refrain from thinking that God and evil are inconsistent. Keep the dialectic in mind here. To show that Mackie is wrong, all Plantinga has to do is show that there is one possible world where transworld depravity holds of everyone. This is why TS claims that it is *necessary* that at least one person is transworld sanctified. The point of making it necessary is to rule out that *any* world is a world where transworld depravity is true of everyone.

For Mackie to be right, TS has to be right (i.e. there can't be any worlds where at least one person never goes wrong). So, their point is, ultimately, still friendly to Plantinga's larger aims: it is not unreasonable for the atheist to refrain from thinking God and evil are compatible *but neither* is it unreasonable for the theist to refrain from thinking that God and evil are incompatible. Plantinga hasn't shown

consistency but neither has Mackie succeeded at his goal of showing an inconsistency. The epistemic stalemate means that neither has *demonstrated* their conclusion.

2. *William Rowe's Objection*

In "In Defense of the Free Will Defense" William Rowe defends TD against Howard-Snyder and Hawthorne's objection (Rowe, 1998). He contends that we do, in fact, have reason to think that TS is false. Rowe lays the foundation for his objection with a rather detailed example. He asks us to consider two scenarios.

Scenario #1: Rowe asks us to consider an urn containing two balls numbered "1" and "2." Suppose that one necessarily³⁵ randomly selects one of the two balls from the urn but does not return the selected ball to the urn. One then selects the remaining ball. Now Rowe calls on us to make a few modal judgments about this scenario. Is it possible that the first ball selected was not ball number 2? Yes it is. Is it possible that the second ball selected was not ball number 2? Yes again. Is it possible that neither of the selections is ball number 2? No. It is not possible that neither selection is ball number 2.

Scenario #2: Now Rowe changes the example a little. Suppose this time that we return the selected ball to the urn once it has been removed. Rowe asks us to make the same modal judgments as we did above. Is it possible that the first ball selected was not ball number 2? Yes it is. Is it possible that the second ball selected was not ball number 2? Yes again. Is it possible that neither of the selections is ball

³⁵ Rowe's example includes the use of modal terms, like "necessarily" in ways that sound like unnatural English. But because Howard-Snyder and Hawthorne's objection involves modal judgments, Rowe is trying to recreate the modal situation in his analogy.

number 2? Unlike the example above, it is possible that neither selection is ball number 2. If the first ball selected is ball number 1 and it is returned to the urn, then the second ball selected could be ball number 1 as again. In which case, neither of the two selections involves removing ball number 2.

The difference between the two scenarios amounts to this: in the first case, we've imposed a modal constraint—it's impossible to get ball #2 on each drawing, though it is possible to get it on either drawing. In the second situation, there is no such constraint. We can draw ball #2 on both drawings and either drawing. Rowe wants to use these scenarios to call our attention to a few logical relations.

In the first situation, this proposition:

◇ (the first ball selected is not number 2 and the second ball selected is not number 2)

does not follow from this one:

◇ (the first ball selected is not number 2) and ◇ (the second ball selected is not number 2).

It doesn't follow because, given the setup,

(if the first ball selected is not number 2, then the second ball selected is number 2).

In the second situation, the proposition:

◇ (the first ball selected is not number 2 and the second ball selected is not number 2)

is consistent with

◇ (the first ball selected is not number 2) and ◇ (the second ball selected is not number 2).

The lesson we are to take away from these two scenarios is this: In the first scenario where the selected ball is not returned to the urn, if the first ball selected is not ball number 2, then it follows *logically* that the second ball selected is ball number 2. In the second scenario, where the selected ball is returned to the urn, nothing logically follows from the fact that the first ball selected was not ball number 2.

How does this apply to TS? Rowe asks us to imagine that, for simplicity's sake, there are only two essences, say Al Capone's and Gandhi's. According to TS, one of them has to be transworld sanctified, i.e. morally perfect. Suppose that Capone is not transworld sanctified. That means that Gandhi is transworld sanctified. Rowe asserts that if TS is true, then the following proposition is true: necessarily if Capone is not transworld sanctified, then Gandhi is transworld sanctified. But Capone isn't transworld sanctified, so Gandhi is.

Rowe points out that this means that whether Gandhi is transworld sanctified is determined by a fact external to Gandhi. Whether Gandhi is morally perfect depends, not on Gandhi's nature alone but on Capone's. But how can facts about Capone determine the moral quality of choices made by Gandhi? But Rowe thinks the problem is even worse than this.

He contends that if TS is right, then a fact about Capone *necessitates* that Gandhi is transworld sanctified. Facts about Capone don't merely determine that Gandhi is transworld sanctified. They determine whether he is necessarily transworld sanctified. As Rowe sees it, if TS is right, then whether Gandhi is *necessarily* transworld sanctified depends on Capone. As Rowe says, "some fact *external* to

Gandhi himself would necessitate what he would do with his freedom if he were created in certain circumstances” (Rowe, 1998, p.119). Ultimately, some fact about Capone *necessitates* how Gandhi uses his freewill.

So, Rowe seems to be raising two related objections to TS. First, it is not Gandhi’s nature or character alone that determines whether he acts morally perfect but something external to it, namely Capone’s nature. Second, Capone’s nature *necessitates* how Gandhi uses his freewill and there is something incoherent about saying that someone is freely necessitated to act a certain way.

What are we to make of Rowe’s objection? First, Rowe’s argument seems to involve a modal fallacy. From the fact that it is necessarily true that either Capone or Gandhi is transworld sanctified and the fact that Capone is not transworld sanctified, it doesn’t follow that Gandhi is *necessarily* transworld sanctified. In other words, from $(C \vee G)$ and $\sim C$ it doesn’t follow that G . We can correct for Rowe’s mistake by weakening the claim by removing the necessity operator in the conclusion. So, rather than conclude $\Box G$, Rowe should simply conclude G . What should we think about this weaker conclusion?

This brings us to a second consideration. It is important to keep the epistemological issues separate from the metaphysical and logical. Consider the following example. Suppose Larry, Curley, and Moe are sealed in a room. The lights go out for 30 seconds and when they come back on Curley is dead on the floor, murdered by someone in the room (assume it was murder not suicide or accidental death). Assuming the room is sealed and no one has been able to come and go from

the room, then either Larry or Moe is a murderer. Suppose we can rule Larry out as the murderer. This means that Moe is the guilty party.

If we know that either Moe or Larry is the murderer and we also know that Larry is not the murderer, then Moe is the murderer. Does the fact that Larry is not the murderer *make* Moe the murderer? It seems confused to think so. Let's add the modal quantifier to our example just to make the parallel with Rowe's example clearer. Suppose that *necessarily* Larry or Moe is the murderer and that Larry is not the murderer. Then, Moe is the murderer. Rowe's point that no fact about Larry can necessitate whether Moe is a murderer, seems to conflate metaphysical and epistemological considerations.

Facts about Larry combined with the general claim that someone in the room is a murderer gives us *reason to believe* that Moe is the murderer. In this context, we need to take care that we don't confuse epistemic with metaphysical or causal interpretations with locutions like "means that" or "makes it the case that." If all Rowe is saying is that, according to TS either Capone or Gandhi is blessed with transworld sanctity, we know that Capone isn't so Gandhi must be (epistemic "must").

Third, let's give Rowe the benefit of the doubt and give a non-epistemic interpretation of his objection. Does Capone not being transworld sanctified *make* Gandhi transworld sanctified? Rowe's objection seems to assign a causal or metaphysical meaning to the word *make*. But Howard-Snyder and Hawthorne need not claim that Capone's being transworld depraved causes Gandhi to be transworld sanctified. There clearly is no causal connection here.

Consider again the case of Larry and Moe. The fact that either Larry or Moe is the murderer and it is not Larry doesn't mean that facts about Larry *make* Moe the murderer. What makes Moe a murderer is that he murdered Curley. What makes Gandhi transworld sanctified is that he doesn't do anything morally wrong in the relevant world.

Howard-Snyder and Hawthorne seem to only need to claim that if either Gandhi or Capone is transworld sanctified and Capone is not transworld sanctified, then it follows logically Gandhi is. We are stipulating a situation in advance: either Capone or Gandhi is transworld sanctified. If we assume that is true and also stipulate another fact, i.e. that it isn't Capone, then it logically has to be Gandhi who is transworld sanctified.

So, as long we avoid giving a causal, ontological, or metaphysical interpretation of Howard-Snyder and Hawthorne's claim, Rowe's objection seems to lose much of its force. If what Capone's being transworld depraved "means" or "makes" Gandhi transworld sanctified in either an epistemic sense or logical sense, there doesn't seem to be much of a problem.

3. Transworld Depravity and the Ontological Argument

There is an interesting and illuminating parallel between the debate about transworld depravity and Plantinga's ontological argument. Understanding this parallel helps not only to clarify the authors' point but also helps us understand what Plantinga might have intended to say in epistemic defense of TD. It also puts us in a better position to assess the epistemic lay of the land. So it is worth taking some time

to briefly set out Plantinga's modal ontological argument and one of the main objections to it.

The basic strategy behind modal ontological arguments is to take advantage of a fact about the modal logic system S5. One of the axioms of S5 is: $\Diamond\Box P \rightarrow P$.³⁶ As a matter of substitution, if it is possible that God exists necessarily, then God exists. Modal ontological arguments tend to include a premise equivalent to $\Diamond\Box G \rightarrow G$ (where G is "God exists"). If the antecedent of preceding conditional is true, then it is not hard to get a sound argument for this existence of God. Not surprisingly, much of the debate over modal ontological arguments is precisely over whether it is possible that God exists necessarily.

Consider Plantinga's version of the ontological argument:

1. There is a possible world W in which maximal greatness is instantiated (premise).
2. Necessarily, a being is maximally great only if it has maximal excellence in every world (definition).
3. Necessarily, a being has maximal excellence in a world if it has omniscience, omnipotence, and moral perfection in that world (definition).
4. In W the proposition *there is no omnipotent, omniscient, and morally perfect being* is necessarily false (from 1, 2, and 3).
5. If a proposition is necessarily false in a possible world W , then it is false in all possible worlds (definition of logical necessity).
6. So, the proposition *there is no omnipotent, omniscient, and morally perfect being* is false in the actual world (from 4 and 5).
7. Therefore: it is true that God exists in the actual world (from 6)³⁷.

³⁶ Its more accurate to say that S5 *can* be formulated equivalently by adding the axiom typically referred to as "B" to S4:

$$(B) P \rightarrow \Box\Diamond P$$

Additionally, $\Diamond\Box P \rightarrow P$ is provable from (B) ((Garson, 2007).

³⁷ This specific statement of Plantinga's modal ontological argument is a composite from various versions mentioned in *God, Freedom, and Evil*. Because Plantinga lays out his argument in a dialectical form, changing versions as he goes along to account for objections, there is no one place where the argument is stated in its complete form.

If S5 is right, then the argument is valid. If the premises are true, it is sound. The only premise that is questionable is 1. So, if it is possible that God necessarily exists, Plantinga's modal argument is a sound argument for the existence of God.

Not surprisingly, critics of this ontological argument are unwilling to grant the first premise. His first premise is pregnant. It is stuffed with modal claims. Given S5, granting the first premise amounts to granting that God exists. Richard Gale puts it this way:

The explicit S5 formulation of Plantinga's argument brings out just what is involved in the Fool's acceptance of the possibility premise...It would be unreasonable to ask him to assent to a proposition until he fully understands what it means. Once the Fool understands what unsurpassable greatness involves—in particular that it involves necessary existence—and that for an entity to have necessary existence is for it to be necessary that it exists, and moreover, accept the S5 theorem that what is possibly necessary is necessary, he would have to be quite a schmuck to grant [it]... the principle of informed consent applies to disputation and not just within the area of biomedical ethics (Gale, 1991, p.227).

Because granting the possibility premise isn't as cheap as one might initially think, one might expect a good reason to think it is true. But Plantinga doesn't offer one. He concedes that he cannot demonstrate the possibility premise is true.

He argues instead that it might be rational to accept the possibility premise even though it hasn't been demonstrated to be true. He offers Leibniz's Law as an example. According to Leibniz's Law:

For any objects x and y and property P , if $x = y$, then x has P if and only if y has P .

According to Plantinga, Leibniz's Law and the possibility premise of his modal ontological argument are in a similar epistemic boat. Despite the fact that neither can be established by argument, both may be rational to accept:

Some philosophers reject [Leibniz's Law]...; various counterexamples have been alleged; various restrictions have been proposed. None of these 'counterexamples' are genuine in my view; but there seems to be no compelling argument for [Leibniz's Law]...that does not at some point invoke that very principle. Must we conclude that it is improper to accept it, or to employ it as a premiss? No indeed. The same goes for any number of philosophical claims and ideas. Indeed, philosophy contains little else. Were we to believe only what is uncontested or for which there are incontestable arguments from uncontested premises, we should find ourselves with pretty slim and pretty dull philosophy...So if we carefully ponder Leibniz's Law and the alleged objections, if we consider its connections with other propositions we accept or reject and still find it compelling, we are within our rights in accepting it—and this whether or not we can convince others. But then the same goes for [the possibility premise of the ontological argument]...And perhaps that is all that can be expected of any such argument (Plantinga, 1974, pp. 220-221).

What Plantinga seems to be saying is that some propositions may have a sort of basic intuitive appeal. If we consider all the epistemic pros and cons of such a proposition and carefully consider how that proposition relates to everything else we believe and still find it epistemically appealing, then we are within our epistemic rights to accept it.

Now let's return to the discussion of TD and TS. Consider the analogy between Plantinga's ontological argument and transworld depravity. First, both cases involve apparently valid arguments that rest on a crucial modal claim about what is possible. Plantinga's critics don't typically charge either the freewill defense or his modal ontological argument with being invalid. If TD is true, then it follows that God and evil are logically compatible. Likewise with his modal ontological argument: if we agree to the possibility of maximal greatness, it follows that God exists. The critics rightly focus on the truth of the possibility premises, not what Plantinga claims follows from them.

Second, in both cases critics charge that the possibility premises are less benign than they initially appear. The possibility premises are “stuffed with modal concepts” and so granting the possibility premises involves granting far more than it might appear at first blush. In addition to these relatively less important parallels, there are similarities that are quite instructive.

For example, a third parallel is that in neither case does Plantinga offer an argument for the truth of the possibility premise. He asserts that the two possibility claims are true. In both cases, Plantinga makes a general appeal to the reader’s modal intuition and points out that there is no obvious or apparent inconsistency in the possibilities premise. It is on this parallel where bringing up the ontological argument might start to pay off.

Plantinga offers no reason to think that TD is true. He does, however, have something general to say in favor of the possibility premise in his ontological argument. In his appeal to the epistemic status of Leibniz’s law, he attempts to provide a defense of the rational permissibility of one’s accepting the possibility premise of his ontological argument. The parallel between the freewill defense and the ontological argument suggests that what Plantinga says in defense of his modal claim in the ontological argument might be what he has in mind regarding the rational permissibility of accepting TD.

Plantinga would presumably say that the epistemic status of TD is like Leibniz’s Law and the modal premise of his ontological argument. Suppose we carefully ponder TD and the alleged objections to it. We consider its connections with other propositions that we accept or reject. If at the end of all of this reflection

we still find it compelling, we are within our rights in accepting it—whether or not we can convince others. If we consider everything we know about TD and its relations to the rest of what we believe and find it has a certain epistemic appeal, then it is rationally permissible to believe TD.

What are we to make of this general kind of defense of accepting TD? Howard-Snyder and Hawthorne point out that the same could be said in defense of TS. Suppose we carefully ponder TS and the alleged objections to it. We consider its connections with other propositions that we accept or reject. If at the end of all of this reflection we still find it compelling, we are within our rights in accepting it—whether or not we can convince others.

Howard-Snyder and Hawthorne concede that we are sometimes within our rights in believing a proposition that we find compelling upon careful reflection, even if we don't have an argument for it. However, they point out that just because we find a proposition for which we have no argument compelling upon reflection does not mean we are within our epistemic rights in believing it. They offer the following mathematical example.³⁸

Let's call a real number "septiquatenary" if "7777" occur in its decimal expansion. Let's call a number "perimetric" if it measures a circumference of a circle whose diameter is 1. With these two definitions in place consider the following modal claim:

SP: Possibly, no number is both septiquatenary and perimetric.

³⁸ Interestingly, the example they cite is originally from Peter van Inwagen in an article he wrote on the ontological argument. It is in a section discussing the epistemology of the possibility premise of Plantinga's ontological argument (Van Inwagen, 1977, p.391).

The authors ask us to imagine a mathematician who takes Plantinga's line with respect to these propositions. "I have carefully pondered SP and the alleged objections to it. I considered its connections with other propositions that I accept or reject. At the end of all of this reflection I still find it compelling, so I am within my epistemic rights in accepting it—whether or not I can convince others." Howard-Snyder and Hawthorne point out that such a person is *not* within his epistemic rights to believe SP. They think this case is really no different than Plantinga's defense of believing the possibility premise of the TD (Howard-Snyder & Hawthorne, 1999, p. 11).

Whatever we make of this type of criticism against the rationality of believing the possibility premise, the main thrust of Howard-Snyder and Hawthorne's criticism is to go a step farther and show that it is not rational to believe the possibility premise. This brings us to a fourth and very important similarity between Plantinga's freewill defense and his ontological argument. They are both susceptible to particular kind of criticism.

Howard-Snyder and Hawthorne's argument involves finding a proposition, namely TS, that is inconsistent with TD and is at least as intuitively plausible as TD. The same type of argument is used to counter Plantinga and others' modal ontological arguments. Richard Gale, William Rowe, and others urge a similar strategy against the possibility premise in Plantinga's modal ontological argument (Gale, 1991; Rowe, 1996a). Gale contends that a successful response involves finding:

...some property that (i) intuitively seems more likely to admit of the possibility of instantiation than does having unsurpassable greatness and (ii) is *strongly incompatible* with it in that if either property is instantiated in any possible world, the other is instantiated in none (Gale, 1991, p. 227).

Plantinga, himself aware of this objection, mentions two such candidates: near-maximality and non-maximality. The property of near-maximality, for example, is “enjoyed by a being if and only if it does not exist in every possible world but has a degree of greatness not exceeded by that of any being in any world” (Plantinga, 1974, p. 218). William Rowe offers the more intuitive property of *being in less than perfect company*. A being has this property at a world *W* if and only if every being in *W* has at least one imperfection, no matter how minor. As Rowe points out:

It may be that we enjoy (or are burdened with) this property in the actual world. But even if we are not, surely, one would think, it is *possible* that this property be instantiated. Surely there is some possible world in which every person has some imperfection, however, slight... (Rowe, 1996a, p. 199)

Of course, if it is possible that this property is instantiated, that means it is instantiated in at least one possible world. If it is instantiated in one possible world, then God does not exist in that world because Anselm’s God is without imperfection. And if God doesn’t exist in one possible world, then he doesn’t exist in every possible world and, hence, is not a necessary being.

The type of reply made by Gale, Rowe, and others is not only that we can conceive of properties such that if it is possible for them to be instantiated it is not possible that maximal greatness is instantiated (i.e. God does not exist necessarily) and that it is more plausible that those properties are possible than Plantinga’s property of maximal greatness. We can feel the pull of this argument when we consider Rowe’s property of *being in less than perfect company*.

If maximal greatness and *being in less than perfect company* are mutually exclusive, which of the following two claims is epistemically more plausible: that every possible world is such that one being instantiates maximal excellence (note:

don't confuse maximal greatness and maximal excellence) or that there is *one* possible world that instantiates *being in less than perfect company*? The possibility that *being in less than perfect company* is instantiated in at least one but maybe only one possible world is a more modest claim than the claim that possibly every world instantiates maximal excellence.

This common line of attack on modal ontological arguments has interesting consequences for Howard-Snyder and Hawthorne's argument against TD. Their argument is that TS is inconsistent with TD and it is as reasonable to think TS is true as it is TD is true. Thus, there is a sort of epistemic stalemate which means it is reasonable to refrain from believing TD. However, if the general line of attack on modal ontological arguments explained above is reasonable, then it raises the possibility of an analogous reply on behalf of the freewill defense.

What makes Gale, Rowe, etc. response to Plantinga's modal ontological argument appealing is that we think it is reasonable to prefer one of the two competing modal claims over the other. It is rational and reasonable to prefer the more modest claim that it is possible that the property of *being in less than perfect company* is instantiated in some possible world than that the property of maximal excellence is instantiated in *every* possible world.

In the jostling for epistemic advantage between TD and TS, something similar holds true. In this case, though, epistemic situation is reversed. Plantinga's claim is the more modest claim while Howard-Snyder and Hawthorne's is the more ambitious. In the case of the ontological argument Plantinga is making a strong claim because it is a claim about *every* possible world. The critic of the ontological argument is

making the more modest claim because it is a claim about one possible world. In the case of the freewill defense, Plantinga is making a more modest claim about one possible world while Howard-Snyder and Hawthorne are making a more ambitious claim about what is true in every possible world.

So, if the critics of the modal ontological argument are right and it is reasonable to prefer Rowe's more modest claim than Plantinga's more ambitious claim, then the same would seem to be true in the debate over the freewill defense. We might say that what is sauce for the ontological argument's goose is sauce for the freewill defense's gander. It seems reasonable to prefer TD to TS because TD is a more modest claim. If this is right, it blunts the force of the authors' objection by decreasing the claimed epistemic parity between the two claims and raising doubt about whether TD and TS are in the same epistemic boat.

This doesn't mean the freewill defense doesn't have any problems. It's just to say that Howard-Snyder and Hawthorne's objection involves some overreaching. Why is their argument over reaching? What diagnosis should be made of their argument? This is the topic of the next section.

4. The Person-relative nature of the reasonableness of believing transworld sanctity over transworld depravity.

In the previous section, I argued that it might be rational to prefer one modal claim over another that is inconsistent with it, even though neither proposition has sufficient independent evidence to justify believing it. I took as an example, the debate regarding the modal ontological argument. In particular, I argued that when it comes to the incompatible propositions *it is possible maximal greatness is*

instantiated and it is possible that being in less than perfect company is instantiated, it is reasonable to prefer the latter over the former on the grounds of modesty. The analogy with the possibility premise of the modal argument shows, I take it, that Howard-Snyder and Hawthorne have the epistemology wrong. Their epistemic framework seems mistaken in a few related ways.

First, there might not be *one* relevant epistemic standard. It might be, as seems plausible, that there are different epistemic standards depending on the purpose of an argument. Epistemic standards might be like rulers in a tool box. Which ruler one picks as a standard depends on the purpose for measuring. If one needs a fine scale, one might pick a smaller ruler. This doesn't mean that just any epistemic standard will do. Just because there is a flexible, or even indefinite, set of standards rather than one general standard doesn't mean that just any standard can legitimately be included in the canonical set of standards. Which standard is relevant depends, in part, on the purpose. If, for example, one's purpose is to demonstrate that it is rationally required to believe that God and evil are compatible, then one has a higher epistemic standard to meet. If the purpose is to demonstrate the more modest claim that it is rationally permissible to believe that God and evil are compatible, then one has a lower standard to meet.

Second, for many purposes, their standard is too high. Plantinga acknowledges in other contexts, such as the arguments for the existence of God, that arguments are rarely so strong as to be rationally coercive, i.e. their conclusions can be avoided only on pain of irrationality. Notice that Howard-Snyder and Hawthorne's standard is that it must not be "reasonable to refrain from believing."

An argument that meets their standard is such that it is unreasonable, and one presumes irrational, to not believe the conclusion of the argument.³⁹ Presumably that means that for Plantinga to succeed he would have to produce premises that are rationally coercive. One can only avoid the force of this kind of argument on pain of being unreasonable. Surely this standard is too high.

Consider examples of modal judgments from other domains of philosophy. Is it possible for the wet substance XYZ on Twin Earth, which looks and functions identically to water, to be water (Kripke, 1980; Putnam, 1975)? Is it possible for David Chalmers to have a phenomenal zombie twin? This zombie twin is molecule for molecule identical to Chalmers. It is physically and functionally identical to Chalmers but it lacks conscious experience. It is psychologically conscious because it can perform certain functions. It can introspect. It can report the results of its introspection. But it lacks phenomenal experience (Chalmers, 1996).⁴⁰ In each case, it seems reasonable to refrain from believing either claim. In other words, those who make different modal judgments on these highly abstract claims which are somewhat removed from ordinary life, aren't necessarily being unreasonable.

A third mistake is that they consider whether it is reasonable to refrain from believing \sim TS independently of the reasonableness of believing TD. One might very well consider TD by itself and conclude with Plantinga that it seems possible.

However, Howard-Snyder and Hawthorne point out that when one considers TS by

³⁹ To meet their standard, it must be unreasonable to be agnostic about the matter. One might think this is reasonable. After all, if the quality of evidence is such that one can still reasonably be agnostic about p , has one really *shown p*? My point is that the evidence for p can be sufficient to justify the belief that p without being so strong as to make it unreasonable to be agnostic about p .

⁴⁰ Consider also the modal judgments involved in Max Black's example of the spheres in the debate surrounding the identity of indiscernibles (Black, 1952).

itself, one gets the same result. Since TD and TS are incompatible, they conclude it is reasonable to refrain from believing not-TS. Considering each claim independently, one notices that it is reasonable to refrain from believing not-TS. But, as the debate about Plantinga's modal ontological argument, and the examples above, shows us, we are often called on to make *comparative* judgments.

Consider their standard as applied to the debate about the ontological argument. Is it reasonable to refrain from believing *it is possible maximal greatness is instantiated*? It is. Is it reasonable to refrain from believing *it is possible that being in less than perfect company is instantiated*? It is. What Howard-Snyder and Hawthorne don't seem to consider is that it is also reasonable to prefer Rowe's modal claim over Plantinga's and, hence, it is reasonable to reject Plantinga's modal ontological argument. As I pointed out earlier, we have similar reasons to prefer TD over TS. TD is in a similar boat compared to TS as Rowe's modal claim is to Plantinga's modal claim. It is more modest and has slightly more intuitive appeal.

A fourth mistake, and one closely related to the previous three, is that they intentionally remove the person-relative nature of knowledge from consideration. Which epistemic standard is the relevant standard is not just a matter one's purpose in making the argument but also at *whom* the argument is targeted. I pointed out above that if Plantinga's goal were to show that it is irrational to believe that God and evil are incompatible, he will have to meet a higher standard than if his goal were to demonstrate that it is rational to believe that God and evil are compatible. Likewise, if Plantinga's purpose is to demonstrate *to everyone*⁴¹ that it is possible that God and

⁴¹ Of course "everyone" should be suitably modified (e.g. to exclude people with low I.Q.s, etc.).

evil both exist, then he has to meet a higher standard than if he is attempting to demonstrate it to some smaller subset of rational people.

Suppose I am right in taking away from the debate over the modal ontological argument that a lower epistemic standard sometimes applies to these sorts of arguments. Rather than the standard “it is unreasonable to refrain from believing P” we should sometimes prefer “it is reasonable to believe P” or the even weaker standard “it is reasonable to prefer P to Q” (where P and Q are incompatible). These lower standards seem to highlight the person-relative nature of knowledge more so than their standard.

Here’s why. Suppose both that TS and TD fail to meet the minimum epistemic standard set out by the epistemic amendment and that there is an ever so slight epistemic difference between TD and TS in TD’s favor. The epistemic advantage enjoyed by TD might be reason to prefer it to TS while at the same time the reasons for TD are not so strong as to make it unreasonable to refrain from believing it.

However, because there is *some* epistemic difference between the two, it might be reasonable to *prefer* TD to TS. So, presumably one could meet the epistemic standard laid out in the epistemic amendment *and still* have reason to prefer TD over TS just by having some reason to prefer TD to TS. But that reason still might not be strong enough to make it unreasonable to refrain from believing the denial of TS. For people in this kind of epistemic situation, Plantinga might have demonstrated that TD is preferable to TS just as Rowe demonstrated to some that his possibility premise is preferable to Plantinga’s. Hence Plantinga’s project would not

be a failure. Their standard isn't flexible enough to accommodate itself to our diverse purposes.

How does a less rigid standard, or set of standards, suggest that we not lay aside the person-relative nature of proof as the authors did at the outset of their paper? If we allow enough flexibility in choosing epistemic standards to accommodate our diverse epistemic goals, some of the relevant epistemic standards might be more sensitive to individual differences.

Think of the physical weight of different people. Suppose that for some job requires a certain minimum weight standard. Suppose, for example, one must weigh 120 pounds for this job. That weight standard will be sensitive to one fact about people— whether they weigh at least as 120 lbs. It will sort them into only two categories. If, however, a job requires a weight standard that is more complex and fine grained, results might be more sensitive to individual differences. Suppose, for example, that, for whatever reason, this job needs pairs of people whose difference in weight is within five pounds. The same people that might have met the previous weight standard might not meet this one and *visa versa*. Because the jobs are different, the standards are different. Because the standards are different, individual differences, contexts, and purposes cannot simply be dismissed as failing to meet *the* weight standard.

Perhaps Rowe's modal claim by itself doesn't meet some general epistemic standard. That fact, just by itself, is not enough to tell us whether it fails to meet some other relevant epistemic standard. In fact, it does seem to meet some relevant

epistemic standard, namely, being epistemically preferable to Plantinga's modal claim (and incompatible with it, etc.).

It seems that we have to take more seriously than the authors admit the person-relative nature of the epistemic. This less rigid epistemic standard, I am suggesting means that Plantinga's freewill defense may have cogency for more people than others and, hence, not be a failure just because it doesn't meet some very general epistemic standard.

It's interesting that, while Howard-Snyder and Hawthorne admit at the outset that proofs are person-relative they drop this point right away. They make an important observation upfront but then ignore it. It's not that they make the point and then set it aside. Rather they make the point and then proceed to contradict it. Their conclusion, after all, is that Plantinga's freewill defense is a failure. But given the point they make at the outset, whether something is a successful demonstration or not depends on, in part, *to whom* the demonstration is being made. A conclusion more appropriate to their observation is that it is a failure for people in a certain epistemic situation. For all they have said, most people are not in that epistemic situation and, hence, it is a successful demonstration for most people.

But suppose that Plantinga's freewill defense ends up being a demonstration (given my more flexible set of standards) to theist but not atheists. Would it be a failure? Suppose it successfully defends the rationality of theists without showing that atheists are irrational to believe God and evil can't co-exist. Given a "defensive" rather than "offensive" purpose, it might not be a failure.⁴²

⁴² It's interesting that the person-relative nature of proof might end up strengthening Howard-Snyder and Hawthorne's conclusion relative to the background beliefs of traditional Christians. The kind of

In conclusion, I have argued in this chapter that Howard-Snyder and Hawthorne have mistaken criteria regarding the success of Plantinga's freewill defense. In order for Plantinga to be successful, they argue that a certain minimum epistemic standard must be met. The standard they propose is that it not be reasonable to refrain from believing \sim TS. I criticized this standard on several grounds and argued that arguments like Plantinga's can be epistemically successful by meeting weaker standards. Other epistemic standards which make it reasonable to make very abstract and fragile comparative judgments. These judgments depend often depend on very general holistic considerations as well as other perspectival factors. I now turn to a general discussion of two of these sorts factors.

Taking Stock and Looking Forward

Let's pause to take stock of where we have been and to look forward to where we are going. The first three chapters of this dissertation have discussed the logical argument from evil. In this first chapter, I set out the logical argument from evil as expressed by J.L. Mackie and Plantinga's reply to it. I then addressed three objections to Plantinga's freewill defense. At the end of the first chapter, I addressed a modal objection to Plantinga's position. In chapters two and three, I addressed two

highly abstract and removed modal judgments that we are called on to make to assess the freewill defense or the ontological argument seem to be influenced by the content of one's belief structure. In cases like abstract modal judgments, when we don't have many closely related specific beliefs to base our judgment on, we seem to rely on the "big picture" of our belief structure. We seem to tend to make holistic type judgments because there are no narrow set of beliefs to go on. In these cases, the general content of one's beliefs seems to play a larger role. This seems to have the consequence that a Christian, for example, who believes that Christ is essentially sinless, might assign different relative weights to TS and TD than a member of a different religion or group that has a more pessimistic view of human nature. So, ironically, the consideration of the person-relative nature of knowledge might lead a Christian to give slightly more weight to TS than other people. This doesn't mean that such reasons are enough to prefer TS to TD but it means that the Christian might have less reason to prefer TD over TS.

other objections to the freewill defense. In chapter two I addressed objections aimed at the Molinism on which the freewill defense appears to depend. In chapter three I addressed an epistemic objection to Plantinga's project.

This discussion of the logical argument from evil and the freewill defense has the following aims. First, explain to the reader, in some detail, the nature of the logical argument from evil and the substance of the freewill defense. The second goal is to show that none of the major objections to the freewill defense are so serious as to be insurmountable. The modal objection discussed in chapter one is based on a misunderstanding. The objections to Molinism discussed in chapter two end up being less important than one might have thought because there are versions of the freewill defense that do not depend on Molinism. The epistemic objection discussed in chapter three depends on an epistemic standard that is set too high.

Besides showing that a particular epistemic objection to the freewill defense is not successful, chapter three makes an important contribution to the dialectic of the dissertation. At the end of chapter three, I argued that the standard Howard-Snyder and Hawthorne advocate is too high and that a more appropriate standard is a comparative one. I made the more general point that there is not always one appropriate epistemic standard. Which standard is relevant depends on one's goals.

This discussion of epistemic standards serves to highlight the person-relative nature of proof. To use Howard-Snyder's and Hawthorne's language, what it is reasonable to believe or refrain from believing partly depends on factors which vary from person to person like background beliefs, experiences, and values. In the next two chapters, I change gears and discuss two perspectival factors that receive less

attention than epistemic factors like background beliefs and experiences. In chapters four and five I discuss the epistemic relevance to religious epistemology of practical interests and social connections.

Chapter 4: Justification, Religious Belief, and Practical Interests

Blaise Pascal (Pascal, 1995 (1670)) and William James (James, 2003 (1896)) are prominent defenders of the importance of the role of pragmatic considerations in the philosophy of religion. They contend that pragmatic considerations bear on the rationality of religious belief. Pascal, for example, argues that if there is an epistemic stalemate between belief and non-belief, then practical consideration may be brought to bear in deciding whether or not to believe. Pascal doesn't think that practical reasons have any epistemic bearing on the situation. Rather, he contends they may legitimately come into play when the evidence can't settle the issue one way or the other.

In this chapter I am going to explore a different model of the relationship between pragmatic considerations and religious belief. According to this new model, pragmatic consideration can and do affect the epistemic status of a belief. If this is right, it might help explain how some religious believers and non-believers may differ in their conclusion about the argument from evil and other religious questions without one side or the other suffering from some epistemic defect. Let me briefly summarize the basic idea before I delve into specifics.

The central project of epistemology involves thinking about the ingredients of knowledge. And in particular, epistemologists have focused on the factor or factors that are thought to distinguish knowledge from true belief. Philosophers typically refer to that element of knowledge which, when added to true belief, yields

knowledge as “justification.”⁴³ It’s something of a received truth in epistemology that the factors which make true belief count as knowledge are purely epistemic.

Richard Feldman and Earl Conee refer to this view as “evidentialism:”

What we call “evidentialism” is the view that the epistemic justification of a belief is [entirely] determined by the quality of the believer’s evidence for that belief (Feldman & Conee, 1985, p. 15).

David Owens describes it this way: “evidentialism is the doctrine that epistemic norms invoke only evidential considerations” (Owens, 2000, p.24). Jeremy Fantl and Matthew McGrath define evidentialism more formally:

Evidentialism. For any two subjects S and S’, necessarily, if S and S’ have the same evidence for/against p, the S is justified in believing that p iff S’ is, too (Fantl & McGrath, 2002, p.68).

Other names for this epistemological view are “epistemic purism” (Fantl & McGrath, 2007, p. 558), and “intellectualism”(Stanley, 2005, p.6).⁴⁴ Although they all name the same position, I will use the term “epistemic purism.”⁴⁵

The most common theories of justification proposed by epistemologists are “epistemically pure” theories.⁴⁶ They are theories where the factors which convey justification are what Jason Stanley calls *truth-conducive* (Stanley, 2005, p.1). An epistemic factor is truth-conducive if its existence makes the truth of the belief more

⁴³ In a famous 1963 article, Edmond Gettier called into question the traditional view of knowledge as justified true belief (JTB) (Gettier, 1963). Some early responses to Gettier proposed adding a fourth condition to the standard account of knowledge. In some ways, whether one adds a fourth condition to the JTB theory of knowledge or simply modifies one’s theory of justification to be able to handle Gettier-style counterexamples is a matter of accounting. For simplicity’s sake, I will use “justification” in the more traditional way as that which when added to true belief yields knowledge.

⁴⁴ Stanley credits Earl Conee for suggesting this term.

⁴⁵ Does it really matter what we call this position? In this case, it seems to me that it is. By denying “evidentialism” or “intellectualism” there is a danger of being misunderstood as denying the importance of evidence or the intellect or as advocating some form of relativism or postmodernism. The term “epistemic purism” seems more apt to avoid such misunderstandings.

⁴⁶ By “epistemologists” I mean “Analytic epistemologists.” I am not sure the extent to which the views of epistemology in either Continental philosophy or Eastern philosophy are epistemically pure.

likely. So, for example, if one has the perceptual experience of seeing a tree in a yard, under normal visual conditions, that experience makes it more likely that there actually is a tree in the yard. The perceptual experience is a truth-conducive factor.

Broadly speaking the theories of justification proposed by epistemologists can be classified as either externalist or internalist. Roughly, externalists contend that the truth-conducive factors which convey justification are external to the knower.

Externalists typically contend that a belief is justified if it was formed by reliable belief-forming processes. Alvin Goldman, for instance, has proposed a causal theory of justification, according to which, a belief p is justified if and only if it was caused by the fact that p (Goldman, 1967). If there was a certain kind of causal chain from the fact that p to the belief that p , then the belief that p is justified.

According to internalist theories of justification the factors which confer justification are those factors to which the knower has introspective access. According to externalism, the knower need not have introspective access to the factors which confer justification. It's enough that the process which formed the belief be reliable. So, while there are important differences between the two major categories of theories of justification both kinds view the factors which convey justification as purely truth-conducive and, hence, are instances of epistemic purism.

As of late, this epistemological orthodoxy has been questioned by philosophers like John Hawthorne (Hawthorne, 2004), Jason Stanley (Stanley, 2005), Jeremy Fantl and Matthew McGrath (Fantl & McGrath, 2002, 2007), David Owens

(Owens, 2000), and Christopher Hookway (Hookway, 1990), among others.⁴⁷

According to these philosophers, at least part of the difference between true belief and knowledge is pragmatic.⁴⁸ Stanley, for example, contrasts his thesis with the epistemic purism common to both internalist and externalist theories of justification:

It is no surprise that epistemologists have widely shared the assumption that the additional facts that make a true belief into knowledge are uniformly truth conducive (either objectively or from the point of view of the subject). The differences between true belief and knowledge are matters that fall within the purview of theoretical rationality, which many philosophers hold to be guided solely by the normative purpose of discovering the truth. My purpose in this book is to challenge this conception of knowledge. I will argue that the factors that make true belief into knowledge include elements from practical rationality. One consequence of my arguments is that the distinction between practical and theoretical rationality is less clear than one might wish (Stanley, 2005pp. 1-2).

Stanley argues that one's *practical interests* have epistemic import. Typically epistemologists consider the practical investment one has in the truth or falsity of one's belief as irrelevant to the epistemic status of that belief. Stanley, by contrast, argues that whether a true belief is also knowledge is not completely independent of the cost of being wrong (Stanley, 2005, p. 2).

Stanley is *not* claiming that none of the factors which confer justification on a belief are truth-conducive. They are the most important factors. His thesis is weaker. It is that, contrary to epistemic purism, the factors which render true belief into

⁴⁷ Looked at from a more practical and applied point of view, one might reasonably include contextualist philosophers in this list. Technically, contextualist are epistemic purists, but the practical application of their ideas might not be so different from proponents of SSI.

⁴⁸ There are times in this discussion where it might be easy to confuse two uses of the word "pragmatic." "Pragmatic" can be used generally to refer to practical considerations or a practical viewpoint. It can also refer to the field of linguistics that studies how meaning, over and above merely semantic meaning, is conveyed (i.e. speaker meaning). Unless otherwise specified, when I use "pragmatic" or one of its lexical forms, I mean it in the more general sense of pertaining to practical considerations not in the technical sense used in linguistics.

knowledge are mainly truth-conducive *but not only* truth conducive. Practical facts also play a role.

Hawthorne, Fantle and McGrath, and others subscribe to views similar to Stanley's. The specific views of these philosophers vary but they all agree on the basic contention that a subject S's practical situation plays *some* role in determining whether S knows that *p*. But what kind of role?

More will be said later but a temporary rough and ready answer is this: the facts about a putative knower's practical situation help determine how much evidence, or which standard of evidence, is required for a belief to be justified. David Owens gets at the general idea this way:

[My case against epistemic purism] can be made simply by asking: how are you going to tell us, in purely evidential terms, what level of evidence is needed to justify belief? Unless this question can be answered, evidentialism (internalist and externalist) must be abandoned (Owens, 2000, p. 26).

Owens argues that something else, besides evidence, is needed to "complete" justification by helping fix the amount of evidence required for justification. "How," he asks, "could this be anything other than the (perceived) needs and interests of the believer" (Owens, 2000, p.26)?

Owens indicates that setting the standard of evidence of a particular belief to be justified by appealing *only* to epistemic norms results in an arbitrary standard of evidence. For Owens, practical facts enter at this point to help set the standard of evidence in a non-arbitrary way. When the practical stakes are high, you need a different amount of evidence for your belief to be justified than when the stakes are low. It's important to note that this doesn't mean that the *only* factors relevant to how

much evidence we need to be justified are practical. It just means that practical factors at least play *some* contributing role.

By what name should we call this kind of view? Stanley refers to it as “Interest-Relative Invariantism”(Stanley, 2005, p. 85). John Greco has called it “Interest-Dependent-Invariantism” (Greco, 2006). Hawthorne calls it “Moderate-Sensitive Invariantism” (Hawthorne, 2004, p. 157). I will use the term “Subject-Sensitive Invariantism,” used by Keith DeRose and others (DeRose, 2004). You’ll notice that every variation of the name above includes the word “Invariantism.” Why this is so will become apparent later.

With this rough approximation of how, according to Subject-Sensitive Invariantism (SSI), practical facts help determine whether a true belief counts as knowledge, I need to make two important points up front—one terminological and the other theoretical. First, the advocates of SSI use different terminology at different points in the discussions. Sometimes they speak as if the pragmatic element they are urging is an element of knowledge. Other times they speak as if it is a condition on justification. Put in terms made popular post-Gettier, their language is not always clear as to whether the pragmatic condition is a fourth addition to the JTB theory of knowledge or whether it is a condition on justification. Either way, the result is the same.

Suppose there is a fourth condition on knowledge, a practical condition. Suppose this condition functions by setting the minimum threshold for the justification condition on knowledge. This fourth condition regulates, in a way, the justification condition. Now suppose the pragmatic condition is directly dependent

on justification. It is one of the “internal requirements” of justification. Either way, the putative pragmatic functions to set the relevant standard of evidence. Whether it does so from “outside” or from “inside” the justification element doesn’t seem important. I will generally speak of the pragmatic condition on justification but keep in mind that if there is such a condition it is also a condition on knowledge, just indirectly via the justification requirement.

The important theoretical point is that SSI is not meant to be a theory of either justification or knowledge. It doesn’t tell us what justification or knowledge amount to. SSI is intended to be a component of any complete theory of knowledge. Within certain limits, SSI is intended to be added to one’s preferred theory of justification or knowledge.

So, in this chapter I will argue that there is a reasonable case to be made for SSI. The blueprint for this chapter is as follows. First, I will set out a few examples where practical matters appear to bear on a subject’s epistemic position. I will briefly consider a few responses intended to easily dismiss the intuitions we have about such examples. Second, I will layout the contextualist position. Contextualism is an important competitor to SSI. They embrace the kind of intuitions elicited by the forthcoming examples while denying that pragmatic factors are a constitutive element of justification or knowledge. I will also explain some of the objections to contextualism.

Third, I will then turn to positive accounts of SSI. While they differ about the details, most philosophers who argue for some form of SSI base their argument, at least partly, on the connection between knowledge and action. I describe the account

of SSI developed by Jeremy Fantl and Mathew McGrath and their argument for it because I take it to be representative of the broad kind of reasoning for this position. I buttress the case for SSI by briefly mentioning other explanatory payoffs of SSI. I contend that the initial examples which motivate SSI and contextualism, the argument for SSI based on the connection between knowledge and action, and the broader theoretical payoffs of SSI constitute a reasonable case for SSI.

Lastly, I argue that if SSI is correct, it seems to apply to religious beliefs and it helps provide an explanation for how some religious believers and non-believers in similar epistemic situations might legitimately, in an epistemic sense, come to different conclusions on questions like the strength of argument from evil.

1. Intuitive Examples and Rejection of Easy Dismissals

1.1 Intuitive Examples

Part of what motivates SSI, and contextualism for that matter, is a certain class of examples where knowledge seems to depend on practical matters about a putative knower's context. Contextualists have formulated numerous examples of the following sort:

Train Case 1. You're at Back Bay Station in Boston preparing to take the commuter rail to Providence. You're going to see friends. It will be a relaxing vacation. You've been in a rather boring conversation with a guy standing beside you. He, too, is going to visit friends in Providence. As the train rolls into the station, you continue the conversation by asking, "Does this train make all those little stops, in Foxboro, Attleboro, etc?" It doesn't matter to you whether the train is the "Express" or not, though you'd mildly prefer it was. He answer, "Yeah, this one makes all those little stops. They told me when I bought the ticket." Nothing about him seems particularly untrustworthy. You believe what he says (Fantl & McGrath, 2002, p.67).

In these circumstances, do you *know* that the train will stop in Foxboro (Suppose the train actually will stop in Foxboro. So your belief is true). Intuitively, you have good enough evidence to think that the train stops in Foxboro. Your belief is justified. In these circumstances, you do seem to know that the train will stop in Foxboro. Now consider a different case:

Train Case 2. You absolutely need to be in Foxboro, the sooner the better. Your career depends on it. You've got tickets for the south-bound train that leaves in two hours and gets into Foxboro in the nick of time. You overhear a conversation like that in Train Case 1 concerning the train that just rolled into the station and leaves in 15 minutes. You think, "That guy's information might be wrong. What's it to him whether the train stops in Foxboro? Maybe the ticket-seller misunderstood the question. Maybe he misunderstood the answer. Who knows when he bought the ticket? I don't want to be wrong about this. I'd better go check it out myself (Fantl & McGrath, 2002, pp. 67-68).⁴⁹

In the case, do you *know* that the train stops in Foxboro? Do you have sufficient reason to believe that the train will stop in Foxboro? Intuitively, you don't. You are not justified in believing that the train stops in Foxboro, and so you don't know it.

But in both case, the evidence for believing the train will stop in Foxboro is the same. What varied wasn't the evidence but the practical situation of the putative knower. If the justification of a true belief is merely a matter of evidence, then you should be in exactly the same epistemic situation in both of the train cases above.

Now your intuitions about these cases might differ from contextualists and advocates of SSI. You might think that in both train cases you don't have knowledge. Alternatively, you might think that in both cases you do have knowledge. It's important to notice, though, that you can disagree with whether the beliefs in the

⁴⁹ These examples are modeled on Stewart Cohen's airport example, well known among Contextualists (Cohen, 1999).

above examples count as knowledge and still have the same basic *kind* of intuition they have.

If your intuitions differ regarding the strength of evidence in the two cases *at all*, then you have a roughly similar intuition. You might, after all, think that both cases constitute knowledge but still see that despite this, the strength of the justification in the two cases varies. You might just think the amount of variation doesn't make a difference to the bottom line of whether the belief counts as knowledge.

The intuitive appeal of this kind of example might be strengthened by looking at a few more cases. Contextualists in epistemology have built up a stock of examples. I will cite a series of examples offered by Jason Stanley.⁵⁰ Stanley gives five variations of a scenario. Each of the five situations changes a relevant variable. By comparing and contrasting our intuitions about these five scenarios, Stanley thinks we end up with reason to think that, contrary to epistemic purism, practical interests play an important role in whether a true belief counts as knowledge. Considering each of these five variations will also help us later when we consider attempts at “easy dismissal” of SSI.

Prior discussions of these kinds of examples in the debate about epistemological contextualism have resulted in the making of a few distinctions that call our attention to important features of the examples. The first distinction is between cases where there is a high cost to being wrong and a low cost to being wrong. Stanley refers to high stakes cases and low stakes cases. A second distinction

⁵⁰ Stanley points out that many of these examples owe themselves to the thinking of Stewart Cohen and Keith DeRose.

is between the putative knowing subject and a third party attributor asked to make a judgment about whether the putative knowing subject has knowledge. The following five examples have a different combination of the traits above. The five permutations are: low stakes, high stakes, low attributor-high subject stakes, ignorant high stakes, and high attributor-low subject stakes. Let's look at each in turn:

Bank Case 1: Low Stakes.

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks, though it is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning" (Stanley, 2005, pp. 3-4).

In this low stakes case, Hannah is right. She does know that the bank will be open tomorrow. It seems hard to argue against this. The evidence Hannah has is exactly the kind of evidence we have in numerous everyday circumstances where we ordinarily take people to have knowledge. If she was at the bank on a Saturday two weeks ago, then, *ceteris paribus* (e.g. has no reason to think that it was a special promotion by the bank to be open on just that particular Saturday, etc), she knows the bank will be open on Saturday (assuming it is true it will be open on Saturday, etc.).

Bank Case 2: High Stakes.

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their checking account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. I don't know that the bank will be open tomorrow" (Stanley, 2005, p.4).

Intuitively, it looks like Hannah is also right in this case when she claims that she does not know the bank will be open on Saturday. This intuition might be hard for some to see. It might be helpful to “crank up” the stakes here. Suppose *a lot* counts on Hannah and Sarah depositing the check. Suppose that if they fail to deposit the check, then on Monday they will lose their house and Hannah won’t continue to get her cancer treatments. If Hannah fails to get her cancer medicine, even for a short period of time, she will die. Surely, in a case where the stakes are *this* high, we can imagine Sarah saying to Hannah, “look, it is absolutely critical that we deposit that check. Think about what you know. Do you *know* that the bank will be open tomorrow?” In such a case, it seems that Hannah is right when she asserts that she doesn’t know that the bank will be open tomorrow.

If our intuitions in these first two cases are as I have suggested, we can already see evidence begin to emerge for the important role practical interests play in justification. The evidential situation in both cases is identical. If whether a true belief is also knowledge *only* depends on purely epistemic, truth-conducive factors, then our intuitions about these cases should not differ. What did change between the two cases is the practical cost of being wrong. When the cost went up, the quality and quantity of evidence that justified Hannah’s belief in the first case isn’t enough to do so in the second case. A change in the practical cost of being wrong caused a shift in the epistemic standard sufficient for a true belief to count as knowledge.

Bank Case 3: Low Attributor-High Subject Stakes.

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their checking account, it is very important that they deposit their paychecks by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw

her. Sarah points out to Hannah that banks do change their hours. Hannah utters, "That's a good point. I guess I don't really know that the bank will be open on Saturday." Coincidentally, Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill, and she knows nothing of Hannah's situation. Wondering whether Hannah will be there, Jill utters to a friend, "Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday" (Stanley, 2005, p.4).

This is the first example where we have a third party attributor, Jill. The stakes for Jill are low. Jill knows the bank will be open because, given her low stakes situation, she has enough evidence to justify that belief. She is in a similar situation to Hannah in the Low Stakes case. In this situation, though, Hannah is in relatively high stakes situation. If Hannah is wrong, it'll cost her. Hannah is in the same situation she was in High Stakes situation. So, Hannah doesn't know the bank will be open. Jill's assertion that Hannah knows that bank will be open on Saturday is false.

Bank Case 4: Ignorant High Stakes.

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their checking account, it is very important that they deposit their paychecks by Saturday. But neither Hannah nor Sarah is aware of the impending bill, nor of the paucity of available funds. Looking at the lines, Hannah says to Sarah, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning" (Stanley, 2005, p.5).

Stanley contends that in this scenario, our intuitions are that Hannah does *not* know that the bank will be open on Saturday. Hannah is in a situation similar to the High Stakes case, even though she is not aware that the cost to her of being wrong is high.

Bank Case 5: High Attributor-Low Subject Stakes.

Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their checking

account, it is very important that they deposit their paychecks by Saturday. Hannah calls up Bill on her cell-phone, and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, "Well, I was there two weeks ago on a Saturday, and it was open." After reporting the discussion to Sarah, Hannah concludes that since banks do occasionally change their hours, "Bill doesn't really know that the bank will be open on Saturday" (Stanley, 2005, p.5).

In this situation, Stanley takes it that our intuitions are that Hannah is right. Bill does not know the bank will be open on Saturday. The cost of being wrong is low for Bill. He is in a similar situation as Hannah was in the Low Stakes scenario. But the stakes are high for Hannah in this case. She is reluctant to attribute knowledge to Bill based on his being at the bank on a Saturday two weeks ago for the same reason she didn't attribute such knowledge to herself in the High Stakes example—the cost of being wrong is too high to permit the level of evidence Bill has to count as sufficient for knowledge.

All five cases involve people who have the same evidence for thinking the bank will be open in Saturday. Additionally, whether a person was willing to attribute knowledge to the relevant subject varied from situation to situation. The facts about whether an attributor was willing to ascribe knowledge to the relevant subject appeared to vary with changes in the cost of being wrong. Practical facts, the facts about the cost of being wrong, co-varied with the facts about knowledge attribution. If these intuitions are right, particularly in the first two cases, this seems to provide *prima facie* reason to think that the factors which make true belief into knowledge are not purely epistemic. Facts about the practical situation can also play a role.

1.2 Rejection of Easy Dismissals

Epistemic purists have suggested a few attempts to easily dismiss the kind of intuitions we tend to have in the examples above. One attempt at an “easy dismissal” Stanley attributes to Jonathan Kvanvig (Stanley, 2005, p.6). According to this objection, the standards of evidence required to make true belief into knowledge don’t change when the stakes are high. Rather, our confidence is shaken by the highness of the stakes which reduces our degree of belief below the minimum threshold required for knowledge. The practical facts cause a change in the belief element of knowledge not the justificatory elements. The practical situation doesn’t influence the epistemic factors but the doxastic ones.

Stanley points out that Kvanvig’s response might seem like a plausible explanation of the High Stakes case but it doesn’t plausibly explain the Ignorant High Stakes case. In that case, Hannah didn’t know that the stakes were high, so her confidence couldn’t be shaken by the highness of the stakes. In both cases she doesn’t know the bank will be open but a weakening of her degree of belief can’t explain the latter of the two cases.

Lay this point aside, though. Stanley points out that Kvanvig’s criticism has an odd consequence: increasing our ignorance about our situation sometimes increases our knowledge. How so? The difference between Hannah in the High Stakes case and in the Ignorant High Stakes case is that she knows the stakes are high in the former but not the latter. In the Ignorant High Stakes case there is nothing to shake her confidence. The fact that she is ignorant of the high stakes means her degree of belief hasn’t been shaken enough to be lowered below the minimum threshold of belief required for knowledge. If Kvanvig’s analysis is right, Hannah

has knowledge in the Ignorant High Stakes case. Even if your intuition about the Ignorant High Stakes case is that Hannah does have knowledge, it is important to see that Kvanvig's explanation of the difference in the two cases is counterintuitive. As Stanley puts it, "it does not seem correct that adding a little ignorance increases knowledge" (Stanley, 2005, p.7).

Further, Kvanvig's explaining our intuitions as changes in the doxastic strength, doesn't account for our intuitions about the High Attributor-Low Stakes case. Our intuition in that case, according to Stanley, is that Hannah and Sarah are right to deny that Bill knows that the bank will be open. But how do Hannah and Sarah's high stakes shake *Bill's* confidence that the bank will be open (Stanley, 2005, p.7).

Another criticism offered by epistemic purists to quickly explain away the examples is to argue that our intuitions in these kinds of cases are mistakenly picking up on whether a subject *knows that he knows p* rather than just whether a subject knows *p*. Some philosophers claim that in order to know *p*, one has to know that one knows *p*. This position is called the KK thesis. We mistakenly think that Hannah doesn't know that the bank will be open in the High Stakes and Ignorant High Stakes cases and that Bill doesn't know the bank will be open in the High Attributor-Low Subject Stakes scenario because we mistakenly apply the KK thesis in those cases.

According to this second criticism by epistemic purists:

our awareness of the raised stakes for Hannah in High Stakes leads us to think that she needs to know that she knows that the bank will be open, and not merely know that the bank will be open" (Stanley, 2005, p.8).

It's not that Hannah and Bill don't know the bank will be open. It's that they don't know that they know the bank will be open. According to this objection, our

intuitions mistakenly invoke the KK thesis. Since our intuitions invoke the false KK thesis, they should be dismissed.

Stanley provides a two-fold response here. First, he contends that the advocate of this objection needs to provide an independent reason for thinking that what is going on here is that we are inclined to deny an agent knows that p because that agent doesn't know that he knows that p . I take Stanley's point to be that advocates of this response can't merely *claim* that is what is going on here, they must provide some reason to think their diagnosis is correct. His second response is that those who advocate this diagnosis of our intuitions must explain why we don't seem to require that the KK thesis be met in the Low Stakes scenario. If knowledge that p requires knowing that we know that p , then why doesn't it require it in the Low Stakes case?

It seems to me that Stanley has missed a potential third response to this objection. Suppose we grant that our awareness of the raised stakes leads us to think that Hannah needs to know that she knows that bank will be open, and not merely know it. This, it seems to me, implicitly concedes Stanley's point. If the raised stakes leads us to believe that the higher standard of evidence required by the KK thesis needs to be met, that is an admission that our intuitions are such that higher stakes requires a high standard of evidence. If we mistakenly pick the wrong higher standard, i.e. the KK thesis, this doesn't diminish the point that our intuition is to require *some* higher standard.

A third attempt at an easy dismissal of our intuitions in the above cases is by explaining away our responses as resulting from a psychological framing effect.

Stanley states the objection thusly:

It has been established that our judgments about the rationality of various inferences are highly dependent upon idiosyncratic facts about how the background situation is described. It would be unwise to put very much weight upon this evidence in claims about the nature of rationality. Similarly, one might think that the intuitions we have in the above cases are also due to psychological framing effects. If so, they are unlikely to be helpful in inquiry into the nature of the knowledge relation (Stanley, 2005, p.9).

Stanley responds by arguing that there are important differences between the framing effect we see in ordinary speaker's judgments about rationality and our judgments in the five cases above. The framing effect we see in ordinary speakers' judgments about rationality don't follow a discernable pattern. The judgments we make in the kinds of cases that Stanley brings to our attention above do follow a pattern. Stanley contends that they follow the kind of pattern we would expect, given a certain view of the connection between practical reasoning and knowledge (Stanley, 2005, p. 9).

As we shall see later, philosophers who advocate a form of SSI a (e.g. Fantl & McGrath, 2002; Hawthorne, 2004) think that our intuitions follow something like the following kind of pattern: if S knows that p, then it is rational for S to act as if p. We use knowledge attributions to criticize and justify actions. When asked why I went to one store rather than another, I might respond by saying that I knew what I wanted was at the store I went to. If it turns out that I was wrong and the store I went to didn't have what I was looking for, a third party (e.g. boss, spouse, friend) might criticize my action by pointing out that I only believed that the item was at the store:

To say that an action is only based on a belief is to criticize that action for not living up to an expected norm; to say an action is based on knowledge is to declare that the action has met the expected norm” (Stanley, 2005, p.10).

It is this connection between knowledge and action that explains our intuitions in the cases he describes above. In the High Stakes case, we think it is wrong for Hannah to act on her belief that the bank will be open on Saturday. We think this because of the aforementioned connection between knowledge and action. Since we don’t judge that she knows, in the High Stakes case, we don’t think she should perform a particular act. Our intuitions about the kinds of cases that Stanley raises, then, are dissimilar from our faulty intuitions in the case of the psychological framing effect mentioned above. They follow a recognizable pattern, the pattern one might expect given a particular antecedently plausible view of the connection between knowledge and action (Stanley, 2005, pp. 10-11).

None of the attempts at “easy dismissals” of our intuitions in the above cases, seem very compelling. Where does this leave epistemic purism? It seems to leave the epistemic purist with no easy ways to explain away the intuitions elicited by the above examples. Stanley thinks this leaves them with three options.

First, they can just flat deny that our intuitions about these cases are correct. Two responses are appropriate here. First, contextualist philosophers and advocates of SSI both have the intuitions described above. Even many opponents of contextualism and SSI share the intuitions but offer other explanations of them, as we saw above. So, there is a wide enough set of people among proponents and opponents of contextualism and SSI to at least make it reasonable for those who have those intuitions, to be reasonable in holding them. There are, of course, people who

don't share these intuitions and for those people, we should acknowledge that contextualism and SSI will have less appeal. When it comes to appeals to intuitions, we sometimes just have to "let the reader decide." That's not uncommon in philosophy.

Second, as later sections in this chapter will make clear, these intuitions alone are not the basis of the case for SSI. As I mentioned at the beginning of this chapter, the case is three-fold. It includes the idea that SSI best explains these intuitions but also includes an argument based on the connection between knowledge and action and the ability of SSI to explain other puzzles in epistemology.

As Stanley sees it, a second option open to epistemic purists is to account for apparently semantic judgments by offering a pragmatic explanation. This is not an uncommon move in philosophy. But Stanley suggests a conservative approach to deploying this kind of strategy. He takes the view that running to a pragmatic explanation of semantic data too quickly seriously undermines the entire "semantic project." Keith DeRose makes the same point:

It's an instance of a general scheme that, if allowed, could be used far too easily to explain away counterexamples marshaled against any theory about the truth-conditions of sentence forms in natural language. Whenever you face an apparent counterexample—where your theory says that what seems false is true, or when it says that what seems true is false—you can very easily just ascribe the apparent truth (falsehood) to the warranted (unwarranted) assertability of the sentence in the circumstances problematic to your theory. If we allow such maneuvers, we'll completely lose our ability to profitably test theories against examples (DeRose, 1999a, p. 198).

There is a legitimate place for pragmatic accounts of seemingly semantic intuitions. Such moves should be made sparingly and generally be limited to cases

where the pragmatic account provides an explanation from clear *general* conversational principles. Stanley points out that he isn't aware of any such accounts.

A third move open to epistemic purists is to deny that the claim about knowledge expressed by Hannah in the Low Stakes case is not the denial of the claim she makes in the High Stakes case (i.e. that they are different senses of the word "knows") (Stanley, 2005). This third option is contextualism. It is to this position that I now turn.

2. Contextualism

Contextualism is the view that knowledge ascriptions are context-sensitive in a distinctively epistemological way (Stanley, 2005, p.16). Strictly speaking, contextualism is not an epistemic thesis but rather a semantic one. According to contextualism, knowledge ascriptions are such that the particular relation picked out by "knows" varies from context to context.

There are various ways to cash out the details contextualism. In order to get a grip on the basic thesis, it helps to compare contextualism to other contextually sensitive expressions. One familiar kind of context-sensitive expression is an indexical. In the sentence "I was born in the month of May," the extension "I" varies from context to context. In particular, the indexical "I" is sensitive to which speaker is uttering the word. Another kind of context-sensitive expression is gradable adjectives.

Gradable adjectives are adjectives that come in degrees or gradations. The adjective "tall" is an example. Suppose Jones is 6'3" tall. A basketball coach may

look at Jones and utter “Jones is not tall” (Hawthorne, 2004, pp. 53-54). On the other hand, a director casting actors to play the role of Munckins in *The Wizard of Oz*, might say “Jones is tall.” The exact relation picked out by “is tall” in the sentence “Jones is tall” is partly determined by the context in which it is expressed.

Let’s consider another example in a bit more detail. A perennial favorite of contextualists is the predicate “is flat.” Flat is gradable. It comes in degrees. If we look at a typical parking lot, whether we say that the parking lot is flat depends on the context. If we are discussing riding a bike on it, we might say it is flat. If we are discussing playing billiards on it, we might say that it is not. On one view, flatness isn’t a single property but a family of properties each with differing degrees of flatness. If something is perfectly flat, let’s say that it is F_1 . If it is slightly less flat we might refer to it as F_2 . F_2 is slightly less flat than F_3 and so on (i.e. $F_1, F_2, F_3 \dots F_n$, etc.).⁵¹

When we assert “ x is flat” in relevantly different contexts, the particular property of the “is flat” family that is referred to varies. In the context of riding bikes on the parking lot, “is flat” might pick out the property F_{20} , while in the context of playing billiards it might pick out F_{10} . Which particular extension of “is flat” that is being picked out by the sentence “the parking lot is flat” depends on the context in which it is asserted.

There are other kinds of context-sensitive expressions than the two I mentioned above. Stanley mentions eight kinds, including:

⁵¹ It might seem odd that I begin my scale of flatness by assigning perfect flatness as F_1 . I do this because absolute flatness is a definite value while lesser degrees of flatness could continue on into infinity

... (a) 'core' indexicals such as 'I', 'here', and 'now', (b) demonstratives such as 'this' and 'that', (c) gradable adjectives such as 'tall', 'flat', and 'large', (d) determiners such as 'many', (e) modals auxiliaries such as 'can', (f) relational expressions such as 'local' and 'enemy', (g) noun phrases such as 'every student' and 'most professors', and (h) adverbial quantifiers such as 'usually' and 'always' (Stanley, 2005, p. 32).

Contextualists typically advance their semantic thesis by way of the analogy to the two kinds of context-sensitivity I used to introduce the concept of context-sensitive expressions—core indexicals and gradable adjectives.

Consider the analogy with gradable adjectives. Contextualists who think that “know” is a gradable expression like “flat,” hold to the following sort of account: When we assert “*x* knows the *p*” the word “knows” is referring to or picking out a relation. (Roughly, the knowledge is a particular kind of relation that holds between a putative knower and a true proposition.) Contextualists contend that “knows,” like “is flat,” is a family of related and gradable properties (i.e. $K_1, K_2, K_3, \dots K_n$). In different contexts it picks out a different particular knowledge relation. In the Low Stakes case it might pick out a less stringent kind of knowing, say K_{20} , while in the High Stakes case it might pick out a more stringent knowledge relation, say K_5 .

Alternatively, consider the contextualist analogy with core indexicals. Assuming a view of indexicality roughly along the lines sketched out by David Kaplan (Kaplan, 1989), two sentences that might appear contradictory on the surface really aren't:

- (1) I am currently located in Maryland.
- (2) I am currently located in Washington State.

These two sentences are contradictory if uttered by the same person at the same time. There is no contradiction, though, if different people utter the sentences. Despite the

fact that both instances of “I” have what Kaplan calls the same character, they have different content. Stewart Cohen thinks “know” is analogous to “I” in this way:

Thus the theory I wish to defend construes ‘knowledge’ as an indexical. As such, one speaker may attribute knowledge to a subject while another speaker denies knowledge to that subject, without contradiction (Cohen, 1988, p.97).

Both the analogy with gradable adjective and the analogy with core indexicals amount to a similar kind of claim: the content of “know” varies in different contexts.

In order to get a better grip on contextualism, it will be helpful to understand what motivates it. As is common in epistemology, contextualism developed with the problem of epistemic skepticism “in mind.” One argument from skepticism points out that in order to know that I have hands, I have to be able to rule out that I am a brain in a vat. Since I cannot rule that out, I don’t know that I have hands. This kind of reasoning generalizes, for almost everything I believe there is some competing skeptical thesis that explains the data but which I cannot rule out. DeRose calls this skeptical argument “the argument from ignorance” and schematizes it thusly. Let H be a skeptical hypothesis and O be a proposition about the world that one thinks he ordinarily knows:

1. I don’t know that not-H
2. If I don’t know that not-H, then I don’t know that O.
- C. So, I don’t know that O (DeRose, 1999b, p.183).

One attempt to address this kind of problem which spurred the development of contextualism was Fred Dretske’s Relevant Alternatives theory of knowledge (RA) (Dretske, 1970). According to RA, “*x* knows that *p*” means that *x* believes *p*, *p* is true, and *x* is able to rule out relevant alternatives at time *t* (DeRose, 1999b, p.190;

Stanley, 2005, p.17). The stock example from the discussion surrounding RA is that when we are at a zoo looking at a zebra, we know that it is a zebra. However, there is an alternative to my view that I am looking at a zebra. Perhaps I am looking at a mule cleverly painted to look like a zebra. According to RA, in the context of seeing a zebra during an ordinary visit to the zoo, I don't need to rule out that what I see is a cleverly painted mule in order to know it is a zebra. I don't have to rule out that possibility is because it is not a relevant alternative in those circumstances.

Now consider whether I know a different, but related, thesis: I know that what am looking at is not a cleverly painted mule. According to Dretske, I don't know this because I can't rule out a relevant alternative, namely, that what I am looking at is a cleverly painted mule.⁵² The idea is that in the ordinary context of visiting a zoo and looking at the zebra display, I know what I see is a zebra because there are no relevant alternatives that can't be ruled out. However, I don't know that what I am seeing is not a cleverly painted mule because there is a relevant alternative that I can't rule out, namely, that what I am looking at is a cleverly painted mule.

There are several obvious problems for RA. Some of the responses to these problems initiated the development of contextualism. The first problem for RA is that it's not clear which alternatives are the relevant ones. Second, even if the distinction between relevant and non-relevant alternatives was clear, it is not obvious

⁵² It's easy to get mixed up here. Don't confuse the thesis being considered in this paragraph with the one above. In the above example, the question was whether I know I am seeing a zebra. In the current example, the question is whether I know I am not seeing a cleverly painted mule. These two questions are related but not identical. It's natural to think that if I know I am seeing a Zebra, then I know that I am not seeing a cleverly painted mule. Part of the reason it is easy for us to read over the difference is because Dretske's examples involve a violation of single-premise epistemic closure. More will be said about this shortly.

what it means to “rule out” an alternative. Third, the truth of RA results in the failure of single-premise epistemic closure (Stanley, 2005, p.17).

Roughly, epistemic closure is the idea that if a subject S knows: (1) p and (2) If p , then q , then S also knows q . RA runs afoul of this in the following way:

According to RA, I know that what I am looking at is a zebra. I *also* know that if what I am looking at is a zebra, then it is not a cleverly painted mule. Given single-premise epistemic closure⁵³, I also know that what I am looking at is not a cleverly painted mule. But Dretske denies this. He claims that I know that I am looking at a Zebra *and* that I don’t know that I am not looking at a cleverly painted mule. RA allows for a sentence of the following type to be true: I know that I have hands, but I don’t know that I am not a brain in a vat. This is a high intuitive cost of RA.

Contextualist-style moves were suggested to buttress RA. First, some contextualists make use of the concept of saliency in their accounts of what makes various alternatives relevant. Stanley contends that the concept of salience used by contextualist is an improvement over the idea of relevant alternatives because it is a psychological notion (Stanley, 2005, pp. 19-20).⁵⁴ Accounts of saliency have been developed independently in other disciplines that are concerned with context-sensitivity

⁵³ Multiple-premise epistemic closure is importantly different. According to multiple-premise epistemic closure if a subject S knows: (1) $p_1...p_n$ and (2) if $p_1...p_n$, then q , then S also knows q . As Stanley points out multiple-premise closure has one thing going against it that single-premise closure does not. Generally, a conjunction is less probable than its conjuncts. Inferring a conjunction from multiple premises is more likely to decrease the probability of the conclusion than making it from less premises (Stanley, 2005, p. 18).

⁵⁴ I take it that when we talk about psychological salience that we are talking about the options that are psychologically “before us.” Presumably this provides a way to limit possibilities in a way that epistemological accounts of relevance can’t. After all, there are practical psychological limits at any time to my thinking whereas the bare epistemic possibilities are much larger.

(Stanley, 2005, p.20).⁵⁵ Stanley also points to the work of Gail Stine which he says can account for the intuitions that originally gave rise RA and at the same time preserves single-premise epistemic closure (Stine, 1976).

Contextualism, then, developed in the context of responding to epistemic skepticism. It took seriously Dretske's attempt to reply to skepticism by limiting the scope of possibilities we need to consider to count as knowing that *p*. It improved on the general approach to skepticism taken by RA by addressing some of the weaknesses in it. Much of the interest in contextualism is in its ability to provide a reply to skepticism.

How is contextualism supposed to solve the problem of skepticism? Recall the argument from ignorance set out above. According to the contextualist, the first premise of the argument is mistaken in ordinary contexts. Remember that contextualists think that the extension of "knows" varies in different contexts. In ordinary contexts, the extension of "knows" is a knowledge relation with a sufficiently low standard of epistemic justification, say K_{20} . This doesn't mean the standards required for K_{20} to be satisfied are low. It just means that aren't so high as to rule out the kind of everyday knowledge with take ourselves to have. In skeptical

⁵⁵ Presumably he is referring to two things: (1) linguists who study the context-sensitivity of less controversial cases than "knows," (e.g. indexicals like "I" and gradable adjectives like "is tall"). (2) Cognitive psychologists who study phenomena like how the mind is able to make decisions that potentially involve astronomically high amounts of information processing in a relatively short period of time by deciding which information is salient and processing it (i.e. fast and frugal heuristics). The entry for "salience" in *A Glossary of Semantics and Pragmatics* is informative here:

Something in a perceptual field is salient to the extent that it readily becomes the focus of attention. This may be because it actively attracts attention more than its neighbors or surroundings, or because it requires less cognitive effort to bring it to the centre of attention. Linguistic expressions make certain parts of a domain salient by profiling. Of two profiled elements, one may be more salient than the other...mere mention may enhance salience, hence, the maleness of a father is rendered more salient by describing him as a *male parent* than by describing him as a *father* (Cruse, 2006, p.154).

contexts, on the other hand, the extension of “knows” is a knowledge relation that requires very high standards of evidence, say K_5 .

The contextualist reply to skepticism, then, is to grant the skeptic’s argument in contexts where skeptical scenarios are salient but deny that the skeptic’s argument holds in ordinary contexts. All the skeptic’s argument shows is that we don’t have knowledge, K_5 , the knowledge relation expressed in high standard scenarios. It doesn’t show that we lack knowledge, K_{20} , the extension of “knows” expressed in ordinary contexts.

Remember that some contextualist see “knowledge” as an indexical. If this is right, it has an important consequences, the same person with the same evidence can know and fail to know the same proposition in different contexts. Both of the following sentences can be true:

- (3) John knows he has hands (uttered in an ordinary context).
- (4) John doesn’t know he has hands (uttered in an epistemology class).

We don’t know that we have hands in the epistemology classroom but we do know it while we are cutting onions. The contextualist response to skepticism concedes (4) to the skeptic but denies him (3). Thus, contextualists take themselves to have defended ordinary knowledge from skeptical arguments.

The skeptic takes himself to be undermining ordinary knowledge, not just knowledge in an epistemology class. The contextualist diagnosis here is that the skeptic fails to see that the amount of evidence required to count as knowledge in skeptical contexts is higher than in ordinary contexts. The skeptic is, in John Hawthorne’s words “semantically blind” (Hawthorne, 2004, p. 107). The skeptic unconsciously exploits a fact about many users of the word “knows,” namely, that we

are semantically blind to the context sensitivity of knowledge ascriptions (Stanley, 2005, p.29).

With this account of bare-bones contextualism in place, we are in a position to place it in the context of the broader dialectic of this chapter. Contextualists embrace the kinds of intuitions elicited by the examples at the beginning of the chapter. In fact, contextualists were the first to really focus on them and they have developed a large pool of them. The contextualist view of these intuitions is that they show us that knowledge ascriptions are context-sensitive.

Despite the fact that contextualism and SSI are in basic agreement about the intuitions in the cases at the beginning of the chapter, they are competing theories. Contextualism is a competitor of SSI because it contends that the practical features of the environment determine, not the standard of evidence necessary for a belief to be justified but rather the relevant semantic content of “knows.”

In low stakes cases, the semantic content of “knows” is a knowledge relation with a weaker standard of evidence, say K_{20} . In high stakes cases, the semantic content of “knows” is a stronger knowledge relation, say K_{10} . Thus, the claim about knowledge expressed by Hannah in the Low Stakes case is not the denial of the claim she makes in the High Stakes case (i.e. that they are different senses of the word “knows”).

Contextualism is a relatively popular view. As such, the relative strength of SSI can be increased by casting doubt on contextualism. The literature on contextualism is vast; so it should come as no surprise that there are many objections to contextualism. I will set out enough of the case against contextualism to raise

doubts about it, but the case I raise is by no means conclusive. For convenience, I will focus on the objections offered by Jason Stanley. His objections can be organized into two categories: objections to the claim that knowledge ascriptions are gradable and objections to the more general claim that knowledge ascriptions are context-sensitive. Let's consider a few from each category.

Stanley begins his case against contextualism by criticizing the analogy contextualists make between knowledge ascriptions and expressions like "flat" and "tall." He points out that there are important differences between these uncontroversial examples of graded expressions and knowledge ascriptions. I am not going to recite all of Stanley's points but will summarize a few. Stanley claims that there are two linguistic tests for gradability. First, gradable expressions allow for modifiers. Consider "flat." Something can be "very flat," "really flat," "extremely flat," etc.

"Know" isn't a gradable expression, Stanley argues, because it doesn't have these properties. Knowledge ascriptions just don't have the defining characteristics of gradable expressions. One might object by pointing to expressions like "I don't really know that the bank is open on Saturday." "Know" is modified by "really" and, hence, capable of modification like other gradable adjectives. Stanley counters that "really" in these cases is not a degree modifier. Compare the following two sentences:

- (5) The parking lot is flat, but not really flat.
- (6) Hannah knows the bank is open, but doesn't really know the bank is open.

(6) is somehow inappropriate or infelicitous while (5) isn't. (5) says that the parking lot is flat but not very flat. It makes an assertion about the extent to which the parking lot is flat. Stanley contends that in (6), when "know" is grammatically modified by "really," it is not a degree modifier but a "hedge." A linguistic hedge functions as a comment on the appropriateness of the assertion. An example of a linguistic hedge is "strictly speaking" as in "strictly speaking the water is clean." In (6) when "really" modifies "know" it amounts to a comment on the appropriateness of making the knowledge assertion. The part of (6) after says "I don't really know the bank will be open" which contradicts the part of the sentence before the conjunction. This explains why (6) doesn't sound right. So, this example fails to show that knowledge ascriptions have the characteristic feature of being able to take degree modifiers that uncontroversial instances of gradable expressions are (Stanley, 2005, pp. 36-38).

A second attribute that is characteristic of gradable expressions, according to Stanley, is that they can be used in comparative constructions.⁵⁶ "Flat" is capable of the comparative form "flatter than." We can say "John is taller than Fred." Stanley contends that knowledge ascriptions are not capable of being in these sorts of constructions (Stanley, 2005, p.36). One potential counterexample to this claim are sentences like "John knows Bill better than Fred." Stanley points out that "knows" in these cases is not propositional knowledge, but knowledge of acquaintance. Here, "knows" doesn't pick out a relation between a knowing subject and a proposition but rather a relations of acquaintance between two people (Stanley, 2005, p. 38).

⁵⁶ What he actually says is they can be "conceptually related to a natural comparative construction" (Stanley, 2005, p.36).

Stanley goes through many other potential linguistic counterexamples to his claim that knowledge ascriptions fail to meet two of the defining characteristics of gradable expressions. The two examples above should suffice to give the reader a flavor of the kind of argumentation that Stanley engages in. Stanley's point is that knowledge ascriptions don't have the typical characteristics of gradable expressions that either means they are not gradable or there is some idiosyncratic reason to think they are gradable which the contextualist has yet to offer. The burden is on the contextualist to offer linguistic evidence *independent* of the semantic ascriptions at the beginning of this chapter (Stanley, 2005pp. 33,45).

The objections we just considered were targeted at the claim that knowledge ascriptions are gradable. Stanley's next objection is against the broader claim that knowledge ascription are context-sensitive. He offers two arguments. Stanley first argues that knowledge ascriptions don't act like other context-sensitive expressions. This doesn't show that knowledge ascriptions *can't* be context-sensitive, but it seems to make it less likely. Second he argues that expressions that are uncontroversially context-sensitive have a property that if knowledge ascriptions had it would undermine contextualism. Let's consider each in turn.

First, he offers empirical argument that knowledge ascriptions fail two common tests for context-sensitivity involving certain kinds of speech-act reports and propositional anaphora. He doesn't claim that passing these tests is a necessary condition for being context-sensitive. In fact, he conceded that there are expressions that are uncontroversially context-sensitive which fail these tests. He argues that the

failure to pass the tests gives some probabilistic reason to doubt the expressions are context-sensitive.⁵⁷ Let's briefly examine the two tests.

Consider the following minimal pair of speech-act reports. Suppose two people, a non-philosopher (A) and a philosophy professor (B) are talking about technology:

TECHNOLOGY

A. It's possible to fly from London to New York City in 30 minutes.

B. That's absurd! No flights available to the public today would allow you to do that. It's not possible to fly from London to New York in 30 minutes.

A. I didn't say it was. I wasn't talking about what's possible given what is available to the public, but rather what is possible given all existing technology (Stanley, 2005, p.53).

The last response by A makes perfect sense. It does so because "possible" is a context-sensitive term that can take a different semantic value in the same discourse when the context shifts. Now consider an analogous conversation with "know" instead of "possible:

ZOO

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

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

A. I guess I can't rule this out.

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

A. I didn't say I did. I wasn't considering the possibility that it could be a cleverly painted mule (Stanley, 2005, p.52).

In this case, the last response of A doesn't make sense, unless A is lying. "Know" doesn't behave like other context-sensitive terms in these kinds of cases, so it fails this particular test of context-sensitivity.

⁵⁷ A piece of evidence might be neither necessary nor sufficient to justify a belief but that doesn't mean it doesn't make the belief more or less probable.

A second test for context-sensitivity is the way they handle propositional anaphora. An anaphora is an instance of an expression whose referent is provided by an antecedent expression. For example, in the pair of sentences “the tall man was reading the book. He put it down when he was finished,” “he” is an anaphora because it gets its reference the antecedent expression “the tall man.” Pay attention to the anaphoric expression “what I said” in this discourse:

It is raining here. Had I been inside, what I said still would have been true.
But now that I am in fact inside, it is not raining here (Stanley, 2005, p. 55).

Contrast that with:

If I have hands, then I know that I have hands. But come to think of it, I might be a brain in a vat, in which case I would believe I have hands, but wouldn't. Now that I'm considering such a skeptical possibility seriously, even if I have hands, I don't know that I do. But what I said earlier is still true (Stanley, 2005, p.54).

Both examples contain the anaphoric expression “what I said” but the first case makes sense while the second sounds infelicitous. The reason is that the first case involves the context-sensitive term “here.” The second case involves the term “know” which doesn't handle anaphora the way other uncontroversially context-sensitive expressions do.

Lastly, context-sensitive expressions typically have the following feature: if the same context-sensitive term occurs twice in the same expression, it can have two different values. An example of this is the sentence: “That field is flat and that billiard table is flat.” The two occurrences of the term “flat” in the sentence express different semantic values (Stanley, 2005, p.58). But knowledge ascriptions not only don't have this property, they can't have it. Consider:

KNOWEDGE SHIFT

A. If that is a zebra, I know it is a zebra.

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

A. No I can't.

B. So you admit you didn't know it was a zebra?

A. If that is a zebra, then I knew it was a zebra. But now, after your question, even if it is a zebra, I don't know it is a zebra (Stanley, 2005, p.63).

Speaker A is not saying that now that B raises the skeptical possibility, he realizes he *didn't* know. He is saying that he *did* know that was a zebra at the beginning of the discourse but doesn't know it at the end. He didn't say he thought he knew at the beginning but it turns out he doesn't. He is saying he actually did know at the beginning *and* doesn't know at the end. This oddity is referred to as the “now you know it, now you don't” objections to contextualism.

So the contextualist is in a dilemma of sorts. Either knowledge expressions are like other context-sensitive ascriptions in this way or they are not. If they are, the “now you know it, now you don't” objection arises. On the other hand, if they aren't like other context-sensitive expressions this provides another reason to doubt whether they really are context-sensitive.

The contextualist might concede that knowledge ascriptions aren't like other context-sensitive expressions in all the ways the critic mentions. Among the class of context-sensitive expressions, “know” is *sui generis*. But if knowledge attributions don't behave linguistically like other context-sensitive expressions it raises doubt as to whether they really are context-sensitive. The contextualist might claim it is *sui generis* but he would have to provide some reason, independent of the examples that led off this chapter. Otherwise, it seems to amount to a mere assertion that contextualism is true.

3. *Subject-Sensitive Invariantism*

3.1 *What is SSI?*

By way of contrast with contextualism, invariantists contend that the semantic value of “know” does not vary from context to context, i.e. it is invariant. “Knows” picks out only one knowledge relation. There are several versions of invariantism. According to *skeptical invariantism*,⁵⁸ the knowledge relation picked out by “know” is such that all, or most, of what we claim to know, we actually don’t. The knowledge relation requires certainty or near certainty for a true belief to count as knowledge. Thus, most of our knowledge ascriptions are false. One prominent proponent of skeptical invariantism is Peter Unger (Unger, 1975).

Moderate invariantists agree with skeptical invariantists that the semantic value of “know” does not vary from context to context. However, they contend that the knowledge relation picked out by “know” is such that very many, if not most, of the knowledge ascriptions we make in ordinary life are true. The knowledge relation is satisfied in cases where our evidence is lower than certainty.

SSI is a form of moderate invariantism. Proponents of SSI agree with skeptical and moderate invariantists that the semantic value of “know” does not vary from context to context. They also agree with moderate invariantists that the knowledge relation doesn’t require certainty to be satisfied and that very many of our everyday knowledge ascriptions are true. What the advocate of SSI adds to moderate invariantism is that whether a putative knowledge relation really is a knowledge relation is sensitive to facts about the practical environment. As Hawthorne puts it:

⁵⁸ I follow the names for the various forms of invariantism assigned by John Hawthorne (Hawthorne, 2004).

Restricting ourselves to extensional matters, the verb ‘know’ picks out the same ordered triples of subject, time, and proposition in the mouths of any ascriber. However, whether a particular subject-time-proposition triple is included in the extension of ‘know’ depends not merely upon the kinds of factors traditionally adverted to in accord with knowledge—whether the subject believes the proposition, whether the proposition is true, whether the subject has good evidence, whether the subject is using a reliable method, and so on—but also upon the kinds of factors that in the contextualist’s hands make for ascriber-dependence. These factors will thus include (some or all of) the attention, interests, and stakes of that subject at that time (Hawthorne, 2004, p. 158).

3.2 Fantl and McGrath’s Argument for and Account of SSI

As I mentioned earlier, there are various versions of SSI. The version I will present is that of Jeremy Fantl and Matthew McGrath (Fantl & McGrath, 2002).

Fantl and McGrath offer the following account of the pragmatic condition on knowledge:

(KA) S knows that p only if S is rational to act as if p (Fantl & McGrath, 2002, p. 559).⁵⁹

This principle is based on the important connection between knowledge and action (hence the principle’s label “KA”). Fantl and McGrath begin building their case by noting that we often defend and criticize actions by citing knowledge. Suppose I don’t take the pie from the oven and it burns. If my spouse scolds me for not taking the pie out, I might reply that I didn’t know the pie was in the oven. A campaign worker of Dennis Kucinich who comes to see that Rep. Kucinich is going to lose might say to him: “Come on Dennis, you know you are going to lose. It’s time to quit” (Fantl & McGrath, 2002, p. 561). Fantl and McGrath argue for (KA) on the

⁵⁹ KA is a condition on knowledge but the authors intend it to be only an indirect condition on knowledge. It is more directly a condition on justification and via justification on knowledge. Fantl and McGrath have written two papers on this topic in which they use slightly different terminology reflecting the different argumentative purposes of each paper (Fantl & McGrath, 2002, 2007). It is quite clear, though, that their thesis about the pragmatic condition on knowledge is the same in both papers (Fantl & McGrath, 2002, p. 79; 2007, p. 565). My discussion uses the different terms from both papers while at the same time preserving the unity of their position.

basis of this kind of connection between knowledge and rational action, “the fact that knowledge-citations play the role they do in defending and criticizing actions is evidence that KA is true” (Fantl & McGrath, 2002, p.564).

Suppose I know that there is a pie baking in the oven and suppose I also know that the pie will burn if I don’t take it out now. If I have these two pieces of knowledge then I am rational to take the pie out of the oven. It doesn’t follow that if it is rational to take the pie out of the oven, then I know it is in there and will burn. The claim here is not a biconditional. Very roughly, the form of the claim is $(K \rightarrow R)$, where K is “S knows that p” and R is “is rational to act as if p.” This doesn’t mean $(R \rightarrow K)$. It does mean, though, $(\sim R \rightarrow \sim K)$, since this is the contrapositive of $(K \rightarrow R)$. So, if you know something, it is rational to act as if it is true. Logically, this means that if it is not rational to act like p is true, you don’t know p. But this doesn’t mean that if it is rational to act as if p, then you know p.

Fantl and McGrath generalize from the kind of connections we make between knowledge and actions:

- (1) S knows that p.
- (2) S knows that if p, then A is the best thing she can do.
- (3) S is rational to do A (Fantl & McGrath, 2002, p.73).

Two things should be noted about this argument. First, it yields the following approximation of KA:

(pre-KA) If S knows that p, then if S knows that if p, then A is the best thing she can do, then S is rational to do A (Fantl & McGrath, 2007, p. 565).

Second, Fantl and McGrath see the above argument as a composite of the following two more immediately compelling arguments:

Argument #1:

- (1) S knows that p.
- (2) S knows that if p, then A is the thing to do.
- Therefore,
- (3) S knows that A is the thing to do (Fantl & McGrath, 2002, p.72).

The locution “thing to do” here means something like “the best thing one can do in light of all of one’s goals.” This first argument is just a general single premise closure argument. The second part of their generalized argument is:

- Argument #2:
- (3) S knows that A is the best thing she can do (in light of all her goals).
 - Therefore,
 - (4) S is rational to do A (Fantl & McGrath, 2002pp. 72-73).

Combined, these two arguments yield the above principled argument for (pre-KA).

Fantl and McGrath propose strengthening (pre-KA) in four ways. This strengthening of (pre-KA) results in KA. First, they point out there is no need to limit the original closure argument just to acts. It can be extended to include preferences. After all, if you know both *p* and *if p, then A is better for you than B*, then you are rational to prefer A over B (Fantl & McGrath, 2002, p.73). Strengthened in light of this consideration, (pre-KA) becomes:

- S knows that p only if, for any states of affairs A and B, if S knows that if p, then A is better for her than B, then she is rational to prefer A to B (Fantl & McGrath, 2002, p.74).

A second way they propose to strengthen their principle is by reformulating it to include cases where it is rational to prefer A to B without actually knowing that A is better than B. We can rationally prefer A to B based on having good reasons to think A is better than B, even when those good reasons aren’t strong enough to count as knowledge or even justified belief.

Given there is a burglar in my house, it is rational for me to prefer the burglar being unarmed to his being armed. Their point is, I take it, that I am rational to have that particular preference even though I don't actually know that if burglar breaks in it is better for me that he is unarmed. For all I know, the unarmed burglar breaking into my house is more jumpy without the psychological security of having a gun. Perhaps he feels threatened by me when he otherwise wouldn't and ends up attacking me. So, I don't *know* it is better for me that the burglar is unarmed but I can still rationally prefer it.

If I am rational, given that a burglar breaks into my house, to prefer that he be unarmed, even though I don't really know it will be better for me that he is unarmed, then if I *know* a burglar is breaking into my house, I am rational to prefer he is unarmed in fact. What this second modification amounts to is the idea that if it is rational for S to prefer A to B, given p and S knows p, then it is rational for S to prefer A to B in fact. With this second modification, Fantl and McGrath are basically proposing to strengthen the second premise of Argument #1 above. To see how, compare these two premises:

(2) S knows that if p, then A is the thing to do (the consequent ends up being "is rational to prefer A").

(2') S is rational to prefer A to B, given p.

(2) explicitly states S knows *if p, then A*. (2') lacks the knowledge claim. It just claims that S is rational to prefer A to B, given p. S can rationally prefer *A to B, given P* without knowing p. Substituting this notion as expressed by (2'), in the place of (2), yields:

(1') S knows that p.

(2') S is rational to prefer A to B, given p.

Therefore,
(3') S is rational to prefer A to B in fact (Fantl & McGrath, 2002, p.74).

S no longer has to *know* A is better than B for S to be rational to prefer A to B.

Notice this new form of their argument is no longer an instance of epistemic closure, at least explicitly.⁶⁰ The moving away from epistemic closure is, I take it, why they think this is a strengthening of their conclusion. This modified argument yields the following form of (pre-KA):

S knows that p only if, for any states of affairs A and B, if S is rational to prefer A to B, given p, then S is rational to prefer A to B, in fact (Fantl & McGrath, 2007, p.76).

A third strengthening of their principle is to make it a requirement on justification not just knowledge. Assume that a subject S is justified in believing p but doesn't know p. Suppose also that, given p, S is rational to prefer A to B. It is rational, then, for S to prefer A to B. Fantl and McGrath ask us imagine there is another subject S'. Suppose S' knows p and is rational to prefer A to B, given p. Suppose further S and S' have the same evidence for p and the same fundamental preferences. Since what is rational to prefer is determined by one's evidence and fundamental preferences, S' is rational to prefer A to B in fact, if and only if S is.⁶¹

⁶⁰ Their argument does rely on the weaker epistemic notion of "good reasons to believe."

⁶¹ This is a bit mysterious. S and S' have the same evidence and fundamental preferences. Yet, S' knows p while S is merely justified in believing p. It is not completely clear how they can differ epistemically. Assuming knowledge is something like justified true belief, then what is the difference in their epistemic situations that makes one count as knowledge and the other not? It can't be the truth requirement because if S' knows p, then p is true. It can't be the justification requirement because that is identical by stipulation. Perhaps it is the belief requirement. Presumably one can be justified in believing p and yet not actually believe it. You have enough evidence for p but you don't "pull the doxastic trigger." Perhaps it is a difference in their practical situations. Maybe S and S' have the same evidence but because evidentialism is false, one can have knowledge while the other doesn't. However, this is just to use what they are arguing for. Further, S and S' both have the same fundamental preferences. Presumably "fundamental preferences" isn't understood so broadly as to include all of one's practical interests.

The argument, modified to be about justification rather than knowledge *per se*, results in:

- (1'') S is justified in believing that p.
- (2'') S is rational to prefer A to B, given p.
- Therefore,
- (3'') S is rational to prefer A to B in fact (Fantl & McGrath, 2002, p.76).

The corresponding form of (pre-KA) is:

S is justified in believing that p only if, for any states of affairs A and B, if S is rational to prefer A to B, given p, then S is rational to prefer A to B in fact (Fantl & McGrath, 2002, p.77).⁶²

The fourth, and final, strengthening of the principle is to make the consequent of the conditional into the following biconditional: S is rational to prefer A to B, given p, iff S is rational to prefer A to B in fact. Recall the discussion of the second modification to the original argument. The idea was that it can be rational to prefer A to B, given p without knowing A is better than B, given p. Fantl and McGrath concluded from this, that if S is rational to prefer A to B, given p and S knows that p, then, S is rational to prefer A to B in fact.

They now argue that this works in the “opposite direction” as well. Putting it in argument form as they have with previous modifications, they are arguing:

- (1'') S is justified in believing that p.
- (2'') S is rational to prefer A to B, in fact.
- Therefore,
- (3'') S is rational to believe A to B, given p.

Suppose that S is, in fact, rational to prefer A to B, and also that S knows that p, then, they argue, S is rational to prefer A to B, given p. Suppose, for example, that you're justified in believing that the bank will be open on Saturday (p). Assume also that

⁶² We might symbolize this as follows: $J_{sp} \rightarrow (A >_p B \rightarrow A > B)$.

you are, in fact, rational preferring going to the bank on Saturday (A) than Friday (B). Could it turn out that it is not rational to prefer going to the bank on Saturday over Friday, given you are justified in believing the bank will be open Saturday? Fantl and McGrath contend it could not.

Just in terms of the logical space, there are two ways it might be rational in this case to not prefer going Saturday to going Friday. But neither possibility works. First, you could be rational to prefer going Friday (B) to going Saturday (A), given that the bank will be open Saturday (p). But by hypothesis you know p and are still rational to prefer Saturday to Friday. Second, you could be rational to be indifferent between Saturday (A) and Friday (B), given that the bank is open Saturday (p). But you know p, and are in fact rational to prefer Saturday to Friday. So, it's hard to see how you could prefer A to B, in fact, know that p, and be rational to have not preference for A to B, given p (Fantl & McGrath, 2002, p. 77). With this modification, their proposed condition on justification becomes:

S is justified in believing p only if, for any states of affairs A and B, S is rational to prefer A to B, given p, iff S is rational to prefer A to B in fact.

Thus, on their account, one is epistemically justified in believing p, only if a certain biconditional is true. The biconditional is fulfilled if one has the same rational preference both given p and in fact. Although their account relates epistemic justification to rational preferences more generally, their final thesis connects justification to rational acts. So their final strengthening, when the kind of preferences is narrowed to preferences to act, is:

S is justified in believing that p only if, for all acts A, S is rational to do A, given p, iff S is rational to do A, in fact (Fantl & McGrath, 2002, p.78).

They condense this principle to:

(JA) S is justified in believing that p only if S is rational to act as if p (Fantl & McGrath, 2002, p.78).⁶³

In this more compact form “rational to act as if p” means the same as “rational to do act A, given p, iff S is rational to do act A in fact.” If it is rational to do act A, given p and it is rational to do act A in fact, then it is rational to act as if p.

How, then, do we get from (JA) to (KA)? Fantl and McGrath state that their condition is “on justification and therefore knowledge” (Fantl & McGrath, 2002, p.79). On a standard account of knowledge, knowledge is justified true belief. If justification is a necessary condition for a belief counting as knowledge and JA is a necessary condition on justification, the JA is a necessary condition on knowledge, hence, KA holds if JA does.

So, by way of summary, their argument for a pragmatic condition on knowledge (and justification) is based on two more immediately compelling arguments: the first is an instance of single premise epistemic closure. The second is the following principle of pragmatic rationality: if S knows that A is the best thing she can do, in the light of all of her goals, then S is rational to do A. Their argument is not based on an appeal to the intuitions elicited in the kinds of cases mentioned at the beginning of this chapter. They make their case on independent grounds but their condition does help explain our intuitions of the earlier cases. This explanatory power adds to the appeal of their thesis.

⁶³ What I call “JA”, they call “PCA.” I changed it to JA to more closely match their terminology in their 2007 paper (see footnote 4 in this chapter).

Let's see how this kind of pragmatic principle handles the cases described at the beginning of this chapter. In the Low Stakes scenario Hannah is rational to act as if the bank is open Saturday. She is rational to do so because she has a little evidence to think it will be open and not much is lost if she is wrong. More specifically, she has reason to prefer the bank's being open on Saturday, given that it is open on Saturday. (She is rational to prefer this because she sees a very long line at the bank on Friday as she drives by and is rational to prefer going on Saturday to waiting in a long line). And she is also rational in fact, to prefer, that the bank is open Saturday (for the same reason). What she is rational to prefer given *p* and in fact are the same. Thus, Fantl and McGrath's account, the pragmatic constraint on knowledge has been met.

In the High Stakes scenario, Hannah is not rational to act as if the bank is open on Saturday. She is not rational to do so because she has a little evidence it will be open but if she is wrong, the costs are very high. More specifically, Hannah is rational to prefer going to the bank on Saturday, given that the bank is open on Saturday. But she is not, in fact, rational to prefer going to the bank on Saturday because it is not rational to take such a high risk compared to the payoff. Thus, Fantl and McGrath's pragmatic constraint on knowledge has *not* been met.

Fantl and McGrath next set out various test cases intended to confirm the kind of epistemic role they give to pragmatic factors. Consider one:

Case of the Time-Consuming Call. You're at home. You're taking the train to New York tomorrow. You have a distinct recollection from making your reservation by phone that your train leaves at 5pm. You are not worried about being wrong. You have good evidence, and in any case, if you miss the 5pm, you'd take the 7pm. Calling up Amtrak again would mean waiting on the phone for 35 minutes or so (Fantl & McGrath, 2002, p.81).

You are justified in believing the train leaves at 5pm. You don't need to inquire further to be justified. Fantl and McGrath's principle confirms this. What you are rational to prefer given that the train leaves at 5pm is the same as what you are rational to prefer in fact. Given that the train is leaving at 5pm, you are rational to act as if the train leaves at 5pm because you have some evidence it leaves at 5pm and the cost of being wrong is pretty low. In fact, you are rational to prefer not to inquire further over inquiring further because it doesn't really matter to you if you are wrong and there is a high cost to checking further.⁶⁴

So Fantl and McGrath argue for a pragmatic condition on knowledge. The account of that condition is able to make sense of the connection between knowledge and action and also account for our intuitions about the cases at the beginning of this chapter.

Fantl and McGrath's are not alone in defending SSI. John Hawthorne, for example, argues for a principle similar to KA. Donald Smith refers to Hawthorne's principle as: "The Practical Reasoning Constraint:"

The Practical Reasoning Constraint (PRC): Suppose that p is practically relevant for S at some time t . Then, if it is unacceptable for S to use the premise that p in a piece of practical reasoning at t , then S does not know that p at t (Smith, 2005, p. 127).

And Hawthorne's reasoning for PCR is similar to that of Fantl and McGrath:

Insofar as it is unacceptable—and not merely because the content of the belief is irrelevant to the issues at hand—to use a belief that p as a premise in practical reasoning on a certain occasion, the belief is not a piece of knowledge at that time (Hawthorne, 2004, p. 176).

⁶⁴ They note and respond to the following objection: any time it is rational for S to inquire further whether p , then S is not justified in believing p . They point out that it can be rational to inquire further on something you are justified in believing because you are offered a lot of money just to check further. But the costs of being right and wrong haven't changed only the cost of the process of inquiry.

Jason Stanley takes pains to point out that SSI is not a theory of justification but should be a component of one's favored theory of knowledge. For that reason, when he sets out his positive account he does so by imbedding it in a "toy" theory of knowledge. His goal is not to propose the particular theory of knowledge he sets out, but only to use it as an illustration of how a theory of knowledge can be married to SSI. Stanley's positive account is that knowledge is an ordered quadruple with variables for names of persons (x), worlds (w), times (t), and propositions (p):

Knows $\langle x, w, t, p \rangle$ if and only if (1) p is true at w (2) $\sim p$ is not a serious epistemic possibility for x at w and t (3) If p is a serious practical question for x at t , then $\sim p$ has sufficiently low epistemic probability, given x 's total evidence (4) x believes at t that p on the basis of non-inferential evidence, or believes that p on the basis of a competent inference from propositions that are known to x at t (Stanley, 2005, pp. 89-90).

So, while there are other versions of SSI than Fantl and McGrath's, their version and their argument for it are fairly convincing.

3.3 Additional Explanatory Payoffs of SSI

In addition to explaining our intuitions about the cases laid out at the beginning of the chapter and the argument based on the connection between knowledge and action layed out in the last section, another consideration in SSI's favor is that it can explain other epistemic phenomena. In particular, John Hawthorne argues that it offers a reasonably good solution to the lottery paradox. This theoretical benefit of SSI contributes, at least to some extent, to the reasons in its favor.

The lottery paradox, as discussed by John Hawthorne and many others is this: Suppose a person of modest means, John, declares that he knows he will not be going

on a six-month tour of Europe this year. He just won't have enough money. Ordinarily, we are inclined to think that what John said is true, barring any farfetched possibilities such as John unexpectedly inheriting a large amount of wealth or winning the lottery. We are aware that people of modest means, like John, sometimes unexpectedly inherit wealth or win a lottery. Nevertheless, under normal circumstances, it's not uncommon for people of modest means, like John to judge that they don't have enough money to go on a six-month tour of Europe. Further, their conversational partners are inclined to agree with their claim (Hawthorne, 2004, pp. 1-2).

Suppose a person of modest means, like John, buys a lottery ticket. If John wins, the payout is worth several million dollars. We know that people of modest means occasionally win the lottery. Now suppose John announced in advance that he *knew* he would not win the lottery. We are normally not inclined to accept his judgment as true. We don't normally think that people know in advance whether they will win or lose a lottery that they have entered. So, we don't tend to think that John knows he will lose the lottery (Hawthorne, 2004, p. 2).

Now the proposition that the John will not have enough money to go on a 6 month tour of Europe this year, entails that John will not win the lottery. If John knows he will not have the money to go on the European tour next year, then doesn't he also know (or be in a position to know) that he will not win the lottery? But how can he know in advance that he won't win the lottery? The form of the argument is:

(5) John knows the he won't have enough money to go on a six-month tour of Europe this year.

(6) If John knows the he won't have enough money to go on a six-month tour of Europe this year, then John knows (or is in a position to know) he won't win the lottery he entered.

(7) Therefore, John knows (or is in a position to know) he will not win the lottery he entered.

Jonathan Vogel and others have pointed out that this problem generalizes (Vogel, 1990). For example, I am inclined to think that I know where my car is parked right now. I don't think, though, that I know whether my car is one of the few that has been stolen today. You claim to know you will have dinner with your wife next week. But you don't think you know that you will not have a heart attack and die between know and then. So the term "lottery paradox" is a term that covers a relatively common phenomenon.

How are we to solve this kind of apparent paradox? In *Knowledge and Lotteries*, Hawthorne argues that his version of SSI is the best candidate for solving the lottery paradox. His argument is basically two-fold. First, he shows how he thinks SSI explains the knowledge judgments we make in cases like the lottery paradox:

We now have the beginnings of a diagnosis of our epistemological puzzlement. We underestimate the contribution of practical environment to the truth of knowledge ascriptions. The picture just given is compatible with the idea that most ordinary knowledge claims come out true. But when we reflect as philosophers, it does not occur to us that issues about practical environment may be relevant to the truth of those ordinary ascriptions. We are insensitive, and attempt to evaluate knowledge ascriptions out of context. In particular, we fail to consider the deliberative context of the subject. No wonder we get confused (Hawthorne, 2004, p. 180).

Second, Hawthorne examines several competing responses and how they come up short. The main ones we are already familiar with: contextualism, skeptical

invariantism, moderate invariantism, and sensitive moderate invariantism (his form of what I have been calling “SSI”). He judges the success of the responses against a set of desiderata. These desiderata include:

1. The Moorean Constraint: many, if not most, of our ordinary knowledge attributions are true (Hawthorne, 2004, p.111; McGrath, 2004).
2. The Assertion Constraint: one should not assert what one doesn't know (i.e. knowledge is the norm of assertion) (Hawthorne, 2004; McGrath, 2004, p. 111).
3. The Practical Reasoning Constraint: one ought only to use that which one knows as a premise in one's deliberations (Hawthorne, 2004, p. 111, 30; McGrath, 2004).
4. Single-premise Closure (SPC): Necessarily, if S knows that p , and S competently deduces that q while retaining knowledge that p , then S knows that q (Hawthorne, 2004, pp. 111, 34; McGrath, 2004).
5. Multi-premise Closure (MPC): If S knows p_1, \dots, p_n , and S competently deduces q from p_1, \dots, p_n while retaining knowledge of p_1, \dots, p_n throughout, then S knows that q (Hawthorne, 2004, pp. 111, 34; McGrath, 2004).
6. The Epistemic Possibility Constraint: It is possible that p for S at t (i.e there is a chance that p for S at t) iff p is consistent with what S knows at t (Hawthorne, 2004, pp. 111, 26; McGrath, 2004).
7. The Objective Chance Principle: If at t , S knows that there is a nonzero objective chance that p at t (where p supervenes on the intrinsic facts about the future relative to t), then, at t , there is a nonzero epistemic probability for S that p (Hawthorne, 2004, p. 92, 112).
8. The Disquotational Schema for “knows” (DSK): If an English speaker E sincerely utters a sentence s of the form ‘A knows that p ’, and the sentence in the that-clause means that p and ‘A’ is a name or indexical that refers to a , then E believes of a that a knows that p , and expresses that belief by s (Hawthorne, 2004, p. 101, 112).

Hawthorne argues that his version of SSI explains the lottery paradox better than competitors when measured against the above desiderata. He doesn't claim that

SSI meets all of these desiderata in an ideal manner. He merely claims that compared to competing solutions, it fares the best overall. The details of this argument take up his entire book. For brevity's sake I will provide just a few examples.

Skeptical invariantism, for instance, fails the Moorean Constraint. According to skeptical invariantism, we don't know most of what we claim to know in ordinary life. It also fares poorly by the Assertion Constraint and the Practical Reason Constraint. Moderate invariantism fares well by the Moorean Constraint, for example, but poorly by the Assertion Constraint (i.e. if I know my lottery ticket will lose, why can't I assert it will lose?). In addition to citing common objections to contextualism, Hawthorne contends that it fares poorly, for instance, by way of the Assertion Constraint and the Practical Reasoning Constraint (Hawthorne, 2004pp. 85-91).

Hawthorne's argument is sophisticated and philosophically charitable. He doesn't disparage the competitors to SSI. In fact, he admits that each putative solution has something going for it. But when judging these responses to the lottery paradox in the broad context of other things we care about in epistemology, he finds his version of SSI the best overall response. If Hawthorne is right, and his case appears impressive, SSI helps explain that two ordinary patterns of thinking aren't as incongruous as they first appear.

I have outlined three streams of reasoning in favor of SSI. First, SSI accounts for the kind of intuitions elicited by the examples at the beginning of this chapter, better than competitors. Second, SSI takes seriously the connection between

knowledge and action. Thirdly, SSI explains better than competitors the kinds of problems raised by the lottery paradox.

4. Religious Belief and Practical Interests

So far in this chapter I have presented what I take to be a reasonably strong case for the truth of some form of SSI. If something like SSI is right, then epistemic purism is false: two people with the same evidence for p might differ as to whether they are justified in believing p . I will now argue that if this is right, it seems to also apply in cases regarding the justification of religious beliefs. Differences in the practical situations of a religious believer and a non-believer may sometimes have the result that a believer and a non-believer with the same reasons for and against a religious proposition vary with respect to whether they are justified in believing that proposition.

More specifically, if, as advocates of SSI suggest, pragmatic facts help determine which normative epistemic standards must be met for a belief to be a justification, then differences in the pragmatic situation of a religious believer and non-believer may result in them having to meet different standards of evidence. So, they could have the same reason to think p true but one be justified and the other not because the minimum evidential-threshold for their belief to be justified is different.

4.1 Religious Examples and JA

Let me address a preliminary issue up front: there may be either an ideological or psychological impediment for some readers when thinking about my application of

SSI to religious beliefs. They might think that religious beliefs are obviously false and this might prevent them from making the kind of judgments I am asking them to make. Consider, for example, the Hannah and the bank case. In the Low Stakes case many, if not most of us, have the intuition that Hannah has knowledge. When it comes to religion, many philosophers think that there is no religious knowledge because there are no religious truths. Thus, when I present a low stakes case analogous to the case involving Hannah and the bank, some may not have the intuition that Hannah has religious knowledge.

There are two ways to address this potential issue. First, we are talking about justification. Surely at least some people are *justified* in having religious beliefs. Remember a belief can be justified, even if it is false. The fact that one thinks that religious beliefs are false isn't by itself a reason to think that religious beliefs are never justified.

This point can be made even using the stronger concept of knowledge. The question at hand is: what are the factors that make true belief into knowledge? So, we might suppose just for the sake of argument that God exists. The question we are interested in, then, is: what are the factors which would make a religious believer S's belief that God exists count as knowledge? So we can assume true belief about religious propositions just as we do about non-religious propositions and then as the analogous questions about knowledge and justification.

Second, I am about to give examples of the justification of religious belief in different practical situations. I will give a low stakes and a high stakes case. Suppose, for the reasons mentioned above, you don't have the intuition in the low

stakes case that the religious person is justified in hold his religious belief. Suppose you just can't get past the fact that you think religious beliefs are false. You can still feel the pull of the argument if you can see *some* epistemic difference between the low stakes and high stakes case. You might think that in neither case is one justified in holding the religious belief in question but see that the believer is even less justified in one case. So long as you can see some epistemic difference based on the practical situation, that is enough to see the basic point.

With this preliminary point out of the way, let's consider some examples involving religious belief that are analogous to the kind of cases contextualists and proponents of SSI appeal to in defense of their respective positions:

Low Stakes Religious Case. Hannah believes in God but is not a particularly religious person. Her family and friends are pretty much areligious. She attended Sunday school as a child and is familiar with basic bible stories and basic religious teachings. As an adult she doesn't attend church much. Every few years she might go to church on a Christmas or an Easter. While working on her college degree, Hannah decides to take a philosophy course to fulfill a humanities requirement. Hannah admires and respects the professor and would like the professor to think highly of her. The professor has already made it clear he doesn't think much of theists. During the philosophy course, the professor presents to the class the argument from evil against the existence of God. After careful consideration Hannah thinks the argument provides at least some evidence to think that God doesn't exist. On that basis, Hannah forms the belief that God does not exist.

In this case, I take it that our intuition is that Hannah is justified in believing that God does not exist based on the argument from evil.

High Stakes Religious Case. Hannah believes in God. In fact, she is a pastor. She has shepherded her congregation for fifteen years. All of her immediate family and most of her friends are religious. Hannah

decides to take a philosophy course. Hannah admires and respects the professor and would like the professor to think highly of her. The professor has already made it clear he doesn't think much of theists. During the philosophy course, the professor presents to the class the argument from evil against the existence of God. Hannah thinks the argument provides at least some evidence to think that God doesn't exist. On that basis, Hannah forms the belief that God does not exist.

In this high stakes case, is Hannah justified in believing that God does not exist?

What our intuitions are here is complicated by our general views of religion, as I noted above. It also depends on how much reason the argument from evil provides for Hannah to think that God doesn't exist. Let me address these concerns in reverse order. First, let's stipulate that in both cases Hannah's total evidence is the same. Suppose there are no evidential differences between the two cases.

Second, as I addressed in my preliminary comments in this section, if we can see at least some difference in the justification between the two cases, then the basic point is made. So let's consider the two cases and determine whether there is such a defense. In the first case there is not much cost to Hannah giving up her belief in God. It is reasonable to judge that the argument from evil provides sufficient evidence for her to believe that God doesn't exist.

In the second case, the cost to Hannah of becoming an atheist is high. She will lose her job. She may lose at least some of her close friends. We can even imagine cases where a very religious spouse, especially one who married a pastor, is inclined to divorce his wife for becoming an atheist. Suppose Hannah stands to be divorced, lose her job, lose her friends, and get in a nasty custody dispute over her kids. We could imagine a friend advising Hannah: "this is going to change your life a lot. You'd better be damned sure of what you believe."

Whatever our intuition is about whether Hannah is actually justified in either case. It seems clear to me that our intuitions in the two cases are not identical. The Hannah of the second case needs to have a better reason to think that God doesn't exist than the Hannah of first case. Differences in the practical environment in two cases helps determine the amount of evidence required to believe that God does not exist.

Let's apply Fantl and McGrath's reasoning to these two cases. Remember that as Fantl and McGrath see it, the pragmatic condition on justification is:

(JA) S is justified in believing that *p* only if S is rational to act as if *p* (Fantl & McGrath, 2002, p.78).⁶⁵

Also remember that "S is rational to act as if *p*" gets cashed out to mean "S is rational to do A, given *p*, iff S is rational to do A, in fact." In essence, one is rational to act as if *p* if what one is rational to prefer given *p* is what one is rational to prefer in fact.

In the Low Stakes scenario it is rational for Hannah to act as if God doesn't exist. She is rational to do so because she has a little evidence to think that God does not exist and not much is lost if she is wrong. More specifically, she has reason to prefer being an atheist to being a theist, given that God doesn't exist. (She is rational to prefer this because she knows the professor will think less of her if she is a theist and nobody else in her life will care at all). And she is also rational in fact, to prefer, being an atheist to being a theist (for the same reason). What she is rational to prefer given *p* and in fact are the same. Thus, Fantl and McGrath's account, the pragmatic constraint on knowledge has been met.

⁶⁵ What I call "JA", they call "PCA." I changed it to JA to more closely match their terminology in their 2007 paper.

In the High Stakes scenario, Hannah is not rational to act as if God does not exist. She is not rational to do so because she has a little evidence that God does not exist but if she is wrong, the costs are very high. More specifically, Hannah might be rational to prefer being an atheist, given that God does not exist. But she is not, in fact, rational to prefer being an atheist to being a theist because the cost of being an atheist is, in fact, very high. Thus, Fantl and McGrath's pragmatic constraint on knowledge has *not* been met in this case. There is a difference between what Hannah is rational to prefer given that God does not exist and what she is rational to prefer in fact.

4.2 Contextualism, SSI, and Religious Belief

One question worth asking is this: suppose the basic intuitions elicited by the kinds of examples at the beginning of the chapter are correct. Does it matter for my purposes whether SSI or contextualism offers the better explanation? Both views take the pragmatic environment to bear on knowledge claims. The chief difference is that contextualism understands the practical facts to be influencing the semantics of "know" while SSI see the practical facts as affecting whether a particular belief is justified.

SSI seems to be the stronger claim. If SSI is right and I am right that it applies to religious beliefs just as it does non-religious beliefs, then if a justified religious belief is true, it counts as knowledge in the same sense of everything else we know. If contextualism is correct and applies in cases of religious belief, then if a justified religious belief is true, it is knowledge but the knowledge relation picked out

by “know” might be a very weak relation. Assuming God exists and at least some religious people are justified in believing that, religious people might just barely know that God exists. They might know God exists in the same way we very weakly know other things.

However weak the relevant knowledge relation expressed by “know” is here it still properly counts as a piece of knowledge. If contextualism constitutes a decent reply to the epistemic skeptic, I see no reason why it doesn’t constitute an acceptable reply to the religious skeptic. If the only point is that religious knowledge, if there is any, is more like ordinary everyday knowledge rather than the kind of knowledge that a skeptic would accept, so what? Contextualism hopes to safeguard the knowledge of ordinary life from the epistemic skeptic. If it is good enough for the knowledge of everyday life, it is good enough for everyday religious life. So, while SSI might be preferable, the application of contextualism to religious knowledge seems to have similar enough results to make it worthwhile.

Conclusion

If the kind of argument I am urging in this chapter is right, then we have good reasons to think that religious believers and non-believers in a similar epistemic situation can come to different conclusions. They can end up with a different set of justified religious beliefs because they are in different pragmatic situations. Thus, religious believers and non-believers with similar information can have different views on religion without it being the case that one side or the other is suffering from a serious epistemic defect like ignorance or bias.

Chapter 5: Religious Belief in a Social World

One fact about knowledge that has been neglected by analytic epistemologists, until lately, is that knowledge is socially distributed. It might be helpful to think of the accumulated knowledge of humanity as being spread across a network of sorts. The network in this case is a social network. Think of people as individual nodes on this social network. Different individual people and groups of people know different things (they also know many of the same things).

Information can be, and is, passed from one node of the social network to another. The transfer of information occurs, for example, by word of mouth, books, newspapers, computers, radio, and television, etc. It's nearly impossible to overstate our epistemic dependence of the social production, distribution, and transmission of knowledge. Much, perhaps most, of the interesting things we know are based on the testimony of others. Indeed, much of what we know, we could only know on the basis of testimony. Elizabeth Fricker describes this observation well with the phrase "the testimony-soaked nature of all of our knowledge" (Fricker, 1995b, p. 394).

We know that Alexander the Great conquered the western world before he was 32 years of age and that Plato was a philosopher only because of testimonial knowledge. In fact, we only know who our mothers are because someone else told us. It's not just in the area of history that we depend on others for our knowledge. Most of our scientific knowledge is also socially produced and disseminated. We know about that the Big Bang occurred about 11-20 billion years ago and that eating a diet high in fat content is bad for our health only because others have told us.

The social transfer of information can't legitimately be done in any old way. There are epistemically appropriate and inappropriate ways for one or more nodes on the network to access information on other nodes of the network. For example, we shouldn't trust the testimony (understood broadly as "tellings") of someone we *know* to be either untrustworthy or ignorant about the issue at hand.⁶⁶

Given our broad-based reliance on the social production and distribution of knowledge, it should not be surprising that much of what we know about religion also has an important social element. In this chapter, I am going offer a sketch of how religious beliefs might be socially justified. My discussion is intended to offer another example of how the religious believer and non-believer can *legitimately* view the evidence for and against the existence of God, for example, differently. In chapter 4, I argued that differences in practical interests can have an important epistemic bearing on religious questions. In this chapter I want to point out the importance of social relations to the justification of religious belief and non-belief.

I want to state upfront that my aims in this chapter a very modest indeed. There are several reasons for this. As mentioned earlier, analytic epistemologists haven't really paid much attention to social epistemology until recently. Even now it receives nowhere near the attention it deserves. If the amount of literature on social epistemology is relatively sparse, the amount on the social epistemology of religion is downright scanty. In important ways, this is undeveloped territory. Therefore, my purposes here should be seen as merely programmatic. I merely intend to outline

⁶⁶ I italicized the word "know" to emphasize the relatively uncontroversial role that internalistic defeaters play in undermining testimony. Matters become more controversial when it comes to other conditions that need to be met or avoided in order to be justified in believing testimony. More will be said about this in the second section of this chapter.

how social location might be relevant to the justification of both religious belief and religious non-belief.

Much of the analytic literature on social epistemology has centered on elucidating the concept of testimony and on the epistemology of testimony. In this first section of this chapter, I set out and explain the concept of testimony. In the second section, I offer a brief description of the debate regarding the epistemology of testimony between reductionists and non-reductionists. Although I don't intend to offer a defense of non-reductionism, my discussion should implicitly make clear that I prefer non-reductionism and my reasons for doing so. In the last section, I suggest how social location might partly account for how some religious believers and non-believers are justified in their respective positions.

1. The Nature of Testimony

There has been much discussion in the literature about what precisely counts as testimony. For present purposes, all that is required is a truncated presentation of that discussion which concludes with what I take to be the best current elucidation of the concept of testimony.

In his influential book *Testimony*, C.A.J. Coady, distinguished between “formal testimony” and “natural testimony” (Coady, 1992, p. 26). Roughly speaking, formal testimony is the kind of testimony given in a court of law. It is legal or quasi-legal testimony. Natural testimony is a broader concept referring to the kinds of non-legal testimony we experience in ordinary life: providing accounts of current events, giving of sports scores, driving directions, and reporting an event witnessed at work

(Coady, 1992, p. 38). I will mainly be concerned with natural testimony. So unless otherwise mentioned, when I use “testimony” I mean “natural testimony.” Anyway, much of what is said about natural testimony will also apply to formal testimony.

Coady’s account of testimony is influential enough to warrant using it as a first approximation a theory of testimony:

A speaker *S* testifies by making some statement *p* if and only if:

- (1) His stating that *p* is evidence that *p* and is offered as evidence that *p*.
- (2) *S* has the relevant competence, authority, or credentials to state truly that *p*.
- (3) *S*’s statement that *p* is relevant to some disputed or unresolved question (which may, or may not be, *p*?) and is directed to those who are in need of evidence on the matter (Coady, 1992, p.42).

Before examining the specific conditions Coady imposes on testimony, it is important to get clear on the way he uses the term “evidence” in his definition. It should be apparent that the concept of evidence plays a prominent role in Coady’s account of testimony. The notion of evidence he employs has consequences for his account of testimony that are not insignificant. What does Coady mean by “evidence?” Coady helps himself to a notion of evidence sketched by Peter Achinstein (Achinstein, 1978).

Achinstein uses the following example to distinguish between three notions of evidence. Suppose over the weekend Alan notices that his skin is yellow. On Monday he goes to the doctor. The doctor examines his skin and tells Alan that he has jaundice, which is the external manifestation of increased concentration of bilirubin (i.c.b) in the blood. The doctor takes some tests and makes a follow up appointment with Alan on Friday. When Friday comes, the doctor tells Alan the test results indicate Alan does not have increased concentration of bilirubin in his blood.

His skin is yellow as a result of a dye with which Alan was working (Achinstein, 1978, p.22). On Friday what should the doctor think of Andy's yellow skin?

- (i) Alan's yellow skin was evidence of an i.c.b. and still is.
 - (ii) Alan's yellow skin was but no longer is evidence for an i.c.b.
 - (iii) Alan's yellow skin is not and never was evidence of an i.c.b.
- (Achinstein, 1978, p. 22).

The concept of evidence expressed by (i) Achinstein calls "potential evidence." Roughly, for some bit of evidence e to count as potential evidence for a hypothesis h requires that there be an "objective connection, association or regularity," between e and h such that e makes h antecedently more likely to true (Graham, 1997, p. 228). Further, potential evidence is such that e can be potential evidence for h , even if h is false. However, e cannot entail h . Achinstein contrast the concept of potential evidence with the concepts of "veridical evidence" and "X's evidence" The concept of veridical evidence is expressed in (iii). Veridical evidence is potential evidence where h is required to be true. X's evidence is what a subject takes to be evidence. It is a subjective concept and is expressed by (ii) (Achinstein, 1978, pp. 22-27).

The concept of evidence that Coady uses, then, is Achinstein's concept of potential evidence. When Coady speaks of testimony as a statement of evidence, he means that there is an objective connection, association, or regularity between e and h in cases where h is true. But a statement can be evidence for some hypothesis even if that hypothesis ends up being false. This is important for reasons I will explore shortly.

Coady's view has received its share of criticism. Examining the criticism will help us better understand Coady's account and help set the stage for understanding

competing theories of testimony. Consider the first clause of Coady's account. It is a conjunction that requires two things: that S's stating that p is evidence that p and is offered as evidence for p . Contrary to Coady, though, there are legitimate cases of testimony that don't meet either conjunct. Let's consider each conjunct in turn.

Consider the following kinds of cases. John tells Mary that Osama bin Laden is in Iran (p). John claims to know this on the basis of his having ESP. He genuinely claims to have ESP, i.e. actually believes what he is saying. John states p and intends his statement of p to be evidence for Mary that p . However, to count as evidence, according to Achinstein's concept of potential evidence, there must be an objective connection, association, or regularity between ESP and finding hidden people. There is no such objective connection, so John's statement isn't evidence for p and, hence, not testimony.

So, according to Coady's view, John's statement of p isn't testimony. But *prima facie* John is offering testimony that p . To see this, suppose that Mary came to believe that p on the basis of what John has told her. The source of Mary's belief that p is John's assertion that p . In such a case, it is John's testimony that caused her to believe p . Her belief is false and formed by way of an epistemically unreliable process but the epistemic status of Mary's belief is not necessarily relevant to what activity John is engaging in, namely, testifying.

To see what is wrong with Coady's definition on this specific point, it will help to also notice that John's statement fails to meet the second clause of Coady's definition because John lacks the relevant competence, authority, or credentials to state truly that p . These examples expose a problem with Coady's account: it

conflates the nature of testimony with the nature of *good* testimony. To be testimony an expression of information must be potential evidence and potential evidence *e* involves a kind of objective connection, association, and regularity between *e* and *h*. To count as testimony, an informational expression must express at least *prima facie* good evidence. His account doesn't allow *prima facie* bad testimony as testimony.

In many ways, the mistake here is analogous to the kind of mistake some neophytes make when considering the nature of art. They form a definition that fails to count bad art as art. John may not be offering good testimony that *p* but he *is* offering testimony that *p*. Consider an example involving formal testimony. If a witness lies in a court of law and testifies that *p*, he *is* giving formal testimony. It's just bad testimony because it is not evidence that *p*. Coady's account conflates this distinction (Lackey, 2006b, pp. 179-180). This is no small flaw in Coady's account.

Let's look at a few other problems. Consider the second conjunct of the first clause of Coady's account: that to count as testimony *S* must intend his statement to be evidence that *p*. Suppose that you happen to run into your neighbor, Smith, at the doctor's office. Smith is in the room adjacent to yours and the walls are thin. You hear the doctor say, "Smith, as you already know, we diagnosed you with colitis last week. Today we need to decide on a course of treatment." Suppose you form the belief that Smith has colitis based on the doctor's statement to Smith.

Surely, in the broad sense, your knowledge is based on testimony. The source of your knowledge is the testimony of the doctor, *even though he did not intend his testimony as evidence*. The doctor and Smith already both knew and discussed that

Smith had colitis last week. The doctor meant his statement to set the stage for discussing treatment not to give Smith evidence that he had colitis.

Think of another kind of case: posthumously published works. In particular, consider personal diaries never intended to be read by others that were discovered and published after the testifier's death. Or suppose I mistakenly think you already know a secret and, so, I ask you a question that reveals information you didn't already know. I revealed a secret by asking a question that assumes you already knew the secret. I didn't intend to offer my words as testimony that *p*. These kinds of cases run afoul of Coady's first clause because it requires that S's testimony that *p* be "offered as evidence that *p*" (Coady, 1992, p.42). So, the second conjunct of the first clause of Coady's account is also problematic.

Lastly, consider the third clause of Coady's theory. It requires that the statement of *p* be relevant to some dispute or unresolved question in order to count as testimony. Suppose John and Mary are having a discussion about whether the earth is flat. Mary is a scientist and *knows* that the earth is a sphere. Suppose John tells Mary that his ESP gives him reason to think that the earth is flat. The question of whether the earth is flat is not an open question for Mary. To her, the matter is not in dispute. She is a scientist and, by stipulation, *knows* the earth isn't flat. It just isn't an open question for her.

Coady would not count John's act of stating his ESP gives him reason to think that earth is flat as testimony. (Again, it is bad testimony but that doesn't mean it isn't testimony *per se*.) But, contrary to Coady's account, it does seem as if John's utterance is testimony. Mary, the hearer, isn't open to John's information but that

doesn't change the fact that John is trying to provide it. He is offering testimony to Mary, but she just isn't buying it. Think again of the private diaries never intended to be read by others that were discovered and published posthumously. Suppose they make statements about which the reader has no need for evidence. Suppose a wife is chronically depressed her whole life, receives lifelong treatment for depression, and frequently tells her husband she is depressed. Her husband is not "in need of evidence on the matter" (Coady, 1992, p.42). Such diaries were not offered as evidence that *p* nor were they directed at those in need of evidence but seem to count as written testimony anyway (Lackey, 2006b, pp. 180-181).

So, there are three general sorts of problems for Coady's approach each of which results in an overly restrictive view of testimony. First, it conflates an account of testimony with an account of good testimony. Second, it doesn't take account of cases where testimony is a source of knowledge even when the testifier never intended to be an epistemic source. And third, his view can't accommodate cases of a speaker testifying when the receivers of that testimony either aren't in any epistemic need of it or are not open to it more generally (Lackey, 2006b, pp. 179-181).

Jennifer Lackey classifies Coady's overly restrictive account of testimony as a *Narrow View of Testimony* (NVT). NVTs tend to exclude cases we typically count as instances of testimony. The requirements of Coady's view are simply too demanding. Partially in response to the kind of problems that plague NVTs, some epistemologists, like Ernest Sosa, Robert Audi and Elizabeth Fricker offer less restrictive accounts (Audi, 1997; Fricker, 1995a; Sosa, 1991).

Audi, for example, thinks of testimony broadly as "...people's telling us things" (Audi, 1997, p. 406). Fricker says that testimony should be understood as "tellings generally" and should have "no restrictions either on subject matter, or on the speaker's epistemic relation to it" (Fricker, 1995a, pp. 396-397). Lackey calls an account of testimony that have these much looser requirements, like Audi's, Fricker's, and Sosa's, a *Broad View of Testimony* (BVT). A barebones BVT looks something like this:

BVT: S testifies that p if and only if S's statement that p is an expression of S's thought that p (Lackey, 2006b, p. 182).

A BVT like the barebones one above is broad enough that, unlike Coady's NVT, it doesn't rule out seemingly uncontroversial cases of testimony. However, BVTs threaten to count expressions as testimony when they really aren't. In particular, they tend to allow non-informational expressions of thought to count as testimony. As Peter Graham points out:

...mere statements are not testimony. Saying 'It is a nice day' is not usually taken as testimony about the weather (though it is when said by the weatherman). Repeating what you have already said over and over does not count as testimony either, unless you have forgotten each previous utterance (Graham, 1997, p. 231).

Lackey provides another example. Suppose two people are taking a stroll and one says "boy, it's a nice day today." This expression is not necessarily a testimony. Suppose it is not meant to convey information but is a conversational filler, a linguistic sigh of contentedness (Lackey, 2006b, p. 183). It expresses a thought but that thought doesn't have any informational content. It is just a contented sigh of "ah" but with more words. It's hard to see this kind of expression as testimony.

One account that is broader than Coady's NVT but narrower than BVTs is Peter Graham's view. Graham's view is what Lackey calls a *Moderate View of Testimony* (MVT). Graham's view is broad enough to avoid excluding clear cases of testimony, like Coady's does, but narrow enough not to include non-informational expressions, like BVTs. According to Graham:

...a speaker *S* *testifies* by making some statement *p* if and only if:
G1. *S*'s stating that *p* is offered as evidence that *p*
G2. *S* intends that his audience believe that he has the relevant competence, authority or credentials to state truly that *p*
G3. *S*'s statement that *p* is believed by *S* to be relevant to some question that he believes is disputed or unresolved (which may or may not be *p*) and is directed at those whom he believes to be in need of evidence on the matter (Graham, 1997, p. 227).

Graham states his account in a similar form to Coady's. The important difference is that the conditions of Graham's account are less objective and more subjective than Coady's. For example, while Coady requires that *S* has the relevant competence, authority or credentials to state truly that *p*, Graham merely requires that *S* *intends* that his audience *believe* that he is qualified. Thus, on Graham's view, a false witness in a court of law counts giving testimony even though the witness is not qualified (e.g. didn't really see what he claims to or isn't really a ballistics expert) if he intends his audience to believe he is qualified. Graham also doesn't require the evidence to be potential evidence in Achinstein's sense of the term. This allows bad testimony, i.e. testimony where there is no objective connection, association, or regularity, to still count as testimony *per se*. The cumulative effect of these two changes is to allow for bad testimony, broadly construed, to count as testimony. Hence, Graham's account avoids Coady's conflation of an account of testimony with an account of good testimony.

Graham also loosens the conditions on testimony by only requiring S to *believe* his statement is relevant to some dispute rather than Coady's stricter requirement that S's statement actually be relevant to some dispute. This relaxing of that requirement allows cases of testimony even when the hearer neither needs nor is interested in the speaker's statement.

So Graham's account is an improvement on BVTs because, like NVTs, it doesn't count non-informational expressions as testimony. It is also an improvement on NVTs like Coady's because it doesn't exclude bad testimony. However, Graham's broadening of Coady's NVT is not enough to include cases where the hearer takes a statement as testimony even when the speaker doesn't intend it as such (e.g. posthumously discovered and published diaries or revealing a secret you thought the hearer already knew).

Jennifer Lackey offers an account of testimony that seems to avoid the problems of Coady's NVT, Audi, Fricker, and Sosa's BVT, and Graham's MVT. Her diagnosis of the problems with the previous accounts is two-fold. BVTs fails for the reason noted above: they include cases of non-informational expressions. Her diagnosis of the problems with NVTs, and by extension Graham's MVT is insightful.

She points out that our concept of testimony has two important elements: First, we think that testimony may be a source of knowledge for hearers even when the speaker doesn't intend his testimony to be an epistemic source. Second we think that a speaker's intent to communicate information to others may count as testimony even if the hearers don't have an epistemic need or interest in that information. Lackey notes that Coady's account, "requires the *conjunction* of these features (i.e.

the speaker's intents *and* the hearer's needs) rather than their *disjunction*" (Lackey, 2006b, p. 181).

Our concept of testimony, she notes, is not univocal. It has two distinct and sometimes independent aspects: the speaker's intent to be a source of knowledge and the hearer's reasonably taking the speaker's expression as a source of knowledge.

Lackey captures this essential fact about testimony by providing a disjunctive account. According to Lackey:

T: S testifies that *p* by making an act of communication *a* if and only if (in part) in virtue of *a*'s communicable content, (1) S reasonably intends to convey the information that *p*, or (2) *a* is reasonably taken as conveying the information that *p* (Lackey, 2006b, p. 193).

Lackey's account is expressed in a form that is different than Coady's and Graham's. The differences in terminology might confuse some readers. So, before I show how Lackey's account succeeds where her predecessors fail, I will briefly explain a few of the less crucial differences which help explain the differences in form and terminology than the accounts set out above.

The previously mentioned proposals tend to define testimony in terms of statements. Both Coady's and Graham's definition of testimony include something like: "S testified by making some statement *p* if and only S's stating that *p*..." But we can, of course, offer testimony by nodding our heads in answer to a question. We can point in a direction when asked for the location of some building. Both these kinds of responses contain information. They may be a statement in some loose sense of the word. Lackey proposes to remove any vagueness by defining testimony in terms of "acts of communications" rather than statements (Lackey, 2006b, p. 187). This allows for testimony to consist of nods and pointing.

Using “act of communication” is also intended to address the following shortcoming in other accounts. Competing theories of testimony tend to make “the assumption that there is always a one-to-one correspondence between the explicit content of one’s testimony and that which is being testified to” (Lackey, 2006b, p. 186). She points out, however, that there are cases of testimony where the meaning conveyed is more than the semantic content of the speaker’s statement. Thus, if you ask me if it is raining outside and I reply “there is an umbrella in the closet,” my testimony conveys both that it is raining and that there is an umbrella in the closet (Lackey, 2006b, p. 186). “Act of communication” also allows for there to be more to the content of testimony than the semantic content expressed by the speaker. It allows testimony to convey speaker meaning not merely semantic meaning.

Lackey’s view of testimony is preferable to BVTs because it doesn’t count non-informational expressions as testimony because such expressions lack “communicable content.” Unlike NVTs, her account is broad enough not to exclude either cases of bad testimony or cases where a hearer lacks an epistemic need or interest in a speaker’s act of communication. Unlike Graham’s MVT (and also NVTs), Lackey’s account can handle cases where a speaker didn’t intend his act of communication to be testimony even though the hearer took it as such (e.g. cases involving private, posthumously read diaries and unintentional and unknowing revealing of secrets). Whether Lackey’s stands up to further philosophical scrutiny, it is sufficiently accurate and illuminating for present purposes.

2. *The Epistemology of Testimony*

The antagonists in contemporary debates about the epistemology of testimony are grouped into two general camps—reductionists and non-reductionists. Very roughly, reductionism about testimony is the view that testimony is not a basic source of justification. The justification of testimonial beliefs depend on, and can be reduced to, more basic sources of justification like sense perception and induction (Coady, 1992, p.79; Lackey, 2006a, p.5).⁶⁷ Just as roughly, non-reductionism about testimony is the view that testimony is just as basic a source of justification as sense perception, memory, etc. (Lackey, 2006a, p.4). The exact nature of reductionist and non-reductionist theories of testimony tends to be unclear and somewhat muddled. I hope in this section to clarify the central issues under dispute.

It is not uncommon for the various positions in a contemporary philosophical debate to have historical representatives. For example, although there are numerous differences, the debate surrounding Chomsky's nativist account of grammar has antecedents in the debate between rationalists and empiricists. Rationalism, considered very generally, can be traced back to Descartes and Plato, among others. Empiricism can be traced back to Locke and Aristotle. The same is true of reductionism and non-reductionism about testimony.

Both the reductive and non-reductive views of testimony also have impressive historical pedigrees. David Hume is widely considered to be the father of reductionism about testimony. In fact, Hume is thought to be the first philosopher to

⁶⁷ Contemporary philosophers tend to distinguish between abduction and induction. Abduction is inference to the best explanation while induction typically refers to enumerative induction. However, historically the kind of inferences that philosophers now tend to classify as abduction was referred to as induction. Unless otherwise specified, *in this chapter* I will follow Hume's usage and lump both kinds of inferences under the heading "induction."

offer anything like a sustained account of testimony (Coady, 1992, p.79; Lackey, 2006a, p.5). The historical roots of non-reductionism are typically traced to Thomas Reid (Coady, 1992, pp. 54-55, 120-121; Lackey, 2006a, p.4). I will set out reductionism and non-reductionism by reference to the views of Hume and Reid, respectively. My goal here is not to do Hume or Reid scholarship. Rather, I am using Hume's and Reid's discussions as a convenient way of conveying the general debate about the epistemology of testimony.

The *locus classicus* of Hume's views of testimony occurs in the context of his discussion of miracles. Like other's who are concerned with Hume's view of testimony, I will ignore Hume's arguments against miracles and focus on his view of testimony which comes out in his discussion of miracles. It's important to note at the outset that Hume fully acknowledges the central role that testimony plays in human knowledge:

...there is no species of reasoning more common, more useful, and even necessary to human life, than that which is derived from the testimony of men, and the reports of eye-witnesses and spectators (Hume, 2000 [1748], p. 84).⁶⁸

Though Hume understands the importance of testimony to knowledge, he doesn't see it as an independent source of justification. According to Hume, our reliance on testimony, in general, is justified by sense perception and induction:

This species of reasoning, perhaps, one may deny to be founded on the relation of cause and effect. I shall not dispute about a word. It will be sufficient to observe, that our assurance in any argument of this kind is derived from no other principle than our observation of human testimony, and of the usual conformity of facts to the reports of witnesses. It being a general maxim, that no objects have any discoverable connexion together, and that all the inferences, which we can draw from one to another, are

⁶⁸ All quotations of Hume in this chapter are from the Oxford's critical edition of the *Enquiry* edited by Tom Beauchamp.

founded merely on our experience of their constant and regular conjunction; it is evident, that we ought not to make an exception to this maxim in favor of human testimony, whose connexion with any event seems, in itself, as little necessary as any other (Hume, 2000 [1748], pp. 84-85).

As Hume sees it, our reliance on testimony is justified by our prior observation of a connection between testimony and the events testified to. We see people testify to facts over and over again and inductively infer from these observations that testimony is often a source of knowledge. The justification of testimonial beliefs, then, depends on whether these beliefs can be justified using sense perception and induction inference (as well as other sources of justification like memory).

Thus, Hume's view of testimony is an empiricist view in the sense that it requires there be empirical evidence that testimony is reliable. Testimony must be justified by *both* sense perception *and* inductive inference:

The reason why we place any credit in witnesses and historians, is not because of any *connexion*, which we perceive *a priori*, between testimony and reality, but because we are accustomed to find a conformity between them (Hume, 2000 [1748], pp. 85-86).

If, for example, we have observed regular connections between George's claims about the world and the world itself, we are justified in forming beliefs based on George's testimony. Reductionism, then, requires we have good empirical reasons to believe that testimony is reliable before we believe something on the basis of testimony. The hearer of testimony bears a certain epistemic burden. He has positive epistemic work to do in order to be justified in believing something on the basis of testimony.

Given the way Hume states his position it might seem that the positive reasons required for believing testimony amount to repeated observations of a similar kind of connection from inference via enumerative induction is then made. However, whatever Hume's particular views turn out to be on this matter, reductionists in general are not committed to such a narrow kind of evidence. What reductionists generally require is inferential evidence that the witness is both trustworthy and in a position to know the claim he is testifying to. Testimonial credibility has two dimensions—a sincerity dimension and a competence dimension (Audi, 2003, p. 138). The evidence for these two aspects of witness credibility can be broadly inferential and include both abduction and enumerative induction. Furthermore, any independent evidence we have from the truth or falsity of the claim being testified to will also have some bearing on the epistemic appropriateness of believing the testimony.

There is a common confusion made at this point. At first glance, Hume's requirement for the empirical justification of testimony might seem obvious. After all, it is clear that we come to know testimony via sense perception. It is important, though, to make a distinction between the justification of a claim and the way we are made aware of content of that claim. Hume's view is that the *justification* of testimony is based on other cognitive faculties—i.e. sense perception and inductive inference. We are not entitled to believe some bit of testimony unless we can provide positive empirical reasons for thinking the testifier has credibility in the two ways mentioned above.

One can acknowledge that we are made aware of the contents of testimony by way of sense perception without thinking that what *justifies* us in believing those contents are empirical arguments for the reliability of the testimony. I will discuss the details of this sort of view shortly when I discuss the Reidian tradition. My point here is just to note that Hume's view is not *merely* that testimony in *some way* depends on sense perception. All the parties to the dispute agree that testimony depends on sense perception in *some way*. Rather reductionists like Hume require an empirical justification before any testimony is accepted.

An interesting question of interpretation arises about Hume's view of testimony. Hume says that we know that testimony can be a reliable source of knowledge because we have observed a regular connection between testimony and events testified to. To whom exactly, though, is Hume referring to when he says that *we* have observed such a connection? Does he mean that we as individuals have observed the regular connection or that we, collectively, have observed the connection? Hume doesn't seem to clearly distinguish between personal observations of the reliability of testimony and what Coady refers to as "communal" observations (Coady, 1992, p. 80). Hume writes:

It is no miracle that a man, seemingly in good health, should die on a sudden; because such a kind of death, though more unusual than any other, has yet been frequently observed to happen. But it is a miracle, that a dead man should come to life; because that has never been observed, in any age or country. There must, therefore, be a uniform experience against every miraculous event, otherwise the event would not merit appellation (Hume, 2000 [1748], p.87).

Here Hume seems to be referring to communal observations. Testimony about resurrections, he says, lack credibility because *no one* has observed a

resurrection. Again, set aside the specific discussion about miracles.⁶⁹ There seems to be a certain amount of circularity here. In order to count the observations of people “in any age or country,” we have to rely on testimonial reports. If testimony, as a source of belief, is justified based on inductive inferences from *other’s* observations, then we are using testimony to justify testimony.⁷⁰

However, the reductionist can make provision for this kind of objection, even if Hume doesn’t do so. The reductionist can argue that one should not count the observations of others *until* one has first established the other’s reliability via personal observations. J.L. Mackie makes exactly this move:

Knowledge that one acquires through testimony, that is, by being told by other people, by reading, and so on, can indeed be brought under the heading of this authoritative knowledge, but only if the knower somehow checks, for himself, the credibility of the witnesses. And since, if it is a fact that a certain witness is credible, it is an external fact, checking this in turn will need to be based on observations that the knower makes himself—or else on further testimony, but, if an infinite regress is to be avoided, we must come back at some stage to what the knower observes for himself (Mackie, 1970, p. 254).

Another small question of interpretation is whether Hume is telling us why we are justified in relying on testimony as a general species of knowledge or when we are justified in believing specific testimonial reports. Hume seems focused on the former but the two need not be disconnected. As mentioned above, reductionists like Hume and Mackie can urge us to only believe specific acts of testimony based on our

⁶⁹ C.S. Lewis makes the following comment about Hume’s point:

Now of course we must agree with Hume that if there is absolutely “uniform experience” against miracles, if in other words they have never happened, why then they never have. Unfortunately we know the experience against them to be uniform only if we know that all the reports of them are false. And we can know all the reports to be false only if we know already that miracles have never occurred. In fact, we are arguing in a circle (C. S. Lewis, 1947, p. 123).

For a recent criticism of Hume’s argument against miracles, see (Earman, 2000). For a recent defense see (Fogelin, 2003).

⁷⁰ Coady offers more textual evidence for Hume’s conflation of personal and communal observations and the importance of this mistake (Coady, 1992, p. 81). I will not pursue the details here.

personal observations. But when we collect a sufficient number of observations we are entitled to inductively infer that relying on testimony as a general species of knowledge is justified.

So, according to reductionist views, testimony is not a basic source of justification. Testimonial beliefs depend for their justification on there being positive empirical reasons to believe the relevant testimony is accurate. So, testimony doesn't just depend on other more basic sources of information, sense perception and memory, to transfer informational content from one person to another but it also depends on other basic sources of knowledge for its justification. Specifically, one is justified in believing testimonial reports only if one has established by empirical observation the reliability of the relevant source of testimony.

As I mentioned previously, the historical roots of non-reductivism are typically traced back to Thomas Reid. According to Reid's account of testimony, we aren't required to empirically justify our reliance on testimony in order to be justified in believing testimonial reports. Reid contends that we have an innate cognitive tendency or disposition to believe testimony:

The wise Author of nature hath planted in the human mind a propensity to rely on evidence before we can give a reason for doing so. This, indeed, puts our judgment almost entirely in the power of those about us, in the first period of life; but this is necessary both to our preservation and to our improvement. If children were so framed, as to pay no regard to testimony or authority, they must, in the literal sense, perish for lack of knowledge (Reid, 2002 [1785], pp. 487-488).

This disposition to believe the testimony of others amounts to a kind of instinct:

In all this, I deal with the Author of my being, no otherwise than I thought it reasonable to deal with my parents and tutors. I believed by instinct whatever they told me, long before I had the idea of a lie, or thought of the possibility of their deceiving me. Afterwards, upon reflection, I found they

acted like fair and honest people who wished me well. I found, that if I had not believed what they told me, before I could give a reason for my belief, I had to this day be little better than a changeling. And although this natural credulity hath sometimes occasioned my being imposed upon by deceivers, yet it hath been of infinite advantage to me upon the whole; therefore I consider it as another gift of Nature (Reid, 1997 [1764], pp. 170-171).

According to non-reductionists, like Reid, one may be justified in believing testimony without first establishing that the source of the testimony is reliable.

Contrary to Hume, he doesn't think we have to be able to make an empirical argument based on prior observations of testimony in order to be justified in believing testimony. This doesn't mean that an inductive argument for the reliability of a source of testimony is worthless. It is just not necessary.

In order to clarify what non-reductionists have in mind, let me briefly sketch one possible picture of non-reductionism. The point of this sketch is not to advance a specific account but rather to provide a general picture of non-reductionism which will help clarify the basic claim of non-reductionism. As we can see in the quotes of Reid above, non-reductionists think that our cognitive faculties have been forged by nature in such a way that testimonial beliefs can be justified without empirical reasons to believe the testimony is reliable. But how might this work?

One way it might work is like this: human cognitive faculties are such that when given certain kinds of informational inputs, they churn out certain kinds of belief outputs without inferring the output beliefs from the informational inputs. Rather than *inferring* conclusions from these inputs, the inputs *trigger* the output of beliefs. Consider an analogy with sense perception. When I have the sense experience that a tree is in front of me, I form the belief that a tree is in front of me. In the typical case, I don't form the belief that there is a tree in front of me by running

through a quick argument. Usually, I don't infer the existence of a tree based on the experience of a tree. I see a tree and this causes, triggers, or in some other non-inferential way results in my believing there is a tree in front of me. It's implausible to think that in-between our experience of a tree and our forming the belief that a tree is in front of us, we run through the following kind of argument:

- (1) I am having the experience of a tree.
- (2) If I am having the experience of a tree, there is probably a tree in front of me.
- (3) Therefore, there is probably a tree in front of me.

It's more plausible to think that the experience itself acts like a sort of trigger for our cognitive faculties which then output a belief that there is a tree in front of us.

The analogy between sense perception and testimony is merely meant to illustrate, *very* generally, the kind of view the non-reductionist about testimony has in mind. It helps us to see, very roughly, how our beliefs *could* result non-inferentially from sense perception. So, to that extent, the case of sense perception is illuminating. But there are important dissimilarities between sense perception and testimony.

Earlier I noted that there is a distinction between source of knowledge conveying certain content and its justifying that content. It is uncontroversial that testimony depends on sense perception. The debate is over *how* it depends on it. The reductionist thinks that sense perception provides material for inductive reasoning in support of some instance (or instances) of testimony. The reductionist argues that any testimonial belief that lacks such support is unjustified. Non-reductionists, on the other hand, think that sense experience provides the content of the testimonial belief and acts as a sort of trigger for the relevant belief-forming processes.

Perhaps, though, memory, not sense perception, is a better analogy. Memory also depends on sense perception (and other faculties) for information. Memory stores information that originated with other more basic sources of information. Testimony also acts as a sort of social storage system for information. (It is also a social conduit for the transfers of information).

Memory seems to be similar to sense perception in precisely the way we have been discussing. When we form a belief based on memory, we don't typically believe what that memory conveys because we have an argument for believing that memory. Usually, we have a memory and that memory triggers the output a belief. For example, I believe that I had breakfast this morning, not because I have a memory that functions as a premise in an inductive argument that I quickly run through in my head. Rather, I have a memory and this memory seems to non-inferentially and directly result in my belief about breakfast. So, memory is like sense perception in the relevant way. But it is also like testimony in the relevant way. Like testimony, it is not usually an original source of knowledge.

Does memory, then, in some sense, reduce to sense perception?⁷¹ It does in the sense that it depends on sense perception for its inputs. But it doesn't in the sense that the justification for our memory beliefs depends on our having empirical reasons to think that our memories are accurate. The non-reductionist view of testimony, then, is similar to the rough account of memory that I just sketched. There is a sense in which testimonial beliefs depend on sense perception without necessarily depending on empirical observations to justify them.

⁷¹ Of course we remember things besides sense perception (e.g. bits of reasoning, etc). I am limiting my comments to sense perception for simplicity's sake.

The rough sketch I just outlined provides a model that explains how our testimonial beliefs (as well as other kinds of beliefs) might be produced by empirical inputs without being justified by empirical observation. But it is one thing to say how such beliefs might be produced non-inferentially and quite another to say how they are justified non-inferentially. How could beliefs formed non-inferentially by the triggering of belief-forming processes be justified?

The kind of thing that non-reductionists have in mind is this: whether beliefs formed in this way are justified depends on whether the process that formed them is reliable. In other words, one move, among others, that non-reductionists could make is to offer a reliabilist account of testimony (Goldman, 1979).⁷² If Reid is right and our cognitive system include a faculty or system which produces beliefs based on testimonial inputs, whether those beliefs are justified or not depends on whether the process of belief formation is reliable.⁷³

Consider a few non-testimonial examples. If I see something at a great distance in the middle of a dark night, that process is less reliable at producing true beliefs based than forming beliefs based on seeing things close up in the day time. The process of forming beliefs based on unclear and vague memories is less reliable than forming beliefs based on clear and distinct memories. Likewise, forming beliefs based on the testimony of experts is more reliable than forming beliefs based on ill-informed people.

⁷² There are other forms of externalism as well. For example, Alvin Plantinga and Michael Bergmann defend externalist accounts of justification/warrant which center on the proper functioning of cognitive faculties (Bergmann, 2006; Plantinga, 1993).

⁷³ Of course there are all sorts of modifications and qualifications to be added here. Advocates of externalist theories of knowledge like reliabilism and “proper functionalism” are aware that certain details have to be added to their theories.

It is critical to notice that reliabilists don't think that one has to have empirical reasons to think a process is reliable in order for beliefs formed from said process to be justified. In many ways the debate between reductionists and non-reductionists about testimony really amounts to the debate between epistemic internalists and epistemic externalists. Reductionists advocate an internalist account of testimony while non-reductionists urge an externalist account, of which reliabilism is the most prominent. This observation may help to clarify the nature of the debate between reductionist and non-reductionists.⁷⁴ Let me explain.

Epistemic internalism is the view that whether or not a belief is justified depends on whether or not the knower has introspective access to the justifications for that belief. As Laurence Bonjour puts it, whether a belief is justified depends "on elements that are internal to the believer's conscious states of mind in a way that makes them accessible to his conscious reflection (at least in principle)" (Bonjour, 2003, p.7). To be justified in holding a belief, one needs to have good reasons for holding the belief in question *and* know (or be able to know) what those reasons are.⁷⁵

Michael Bergmann's description of internalism is helpful:

A crucial ingredient of internalist accounts of justification is the thesis that it is not enough for a belief's justification that it *has* some good-making feature; in addition, the person holding the belief must be *aware* of that good making feature (Bergmann, 2006, p. vii).

⁷⁴ Although philosophers of testimony don't always make it explicit, it seems to me that a significant portion of the literature on the epistemology of testimony makes the underlying assumption that the debate here amounts to externalist versus internalist accounts of the epistemology of testimony. The language, concepts, and arguments used by reductionists and non-reductions makes this clear, even if the philosopher's involved don't always make it explicit. However, I am not entirely convinced that one could not be both a non-reductionist and have a classical internalist and deontological epistemology of testimony. After all, one might hold both that one has to be aware that one has fulfilled one's epistemic obligations in order to be justified but just deny that one always has an epistemic obligation to establish the trustworthiness and competence of a witness before believing his testimony.

⁷⁵ Again, there are all sorts of relevant clarifications here that are beyond the scope of this chapter.

Reductionist views of testimony are internalistic because they require the knower to “have” reasons for thinking the testimony is reliable. By “have reasons” reductionists mean that the putative knower must have introspective access to the reasons they have for holding the belief. If a person believes the testimony of another without having good reasons for thinking the testifier is reliable, then the reductionist would say that belief is not justified.

Epistemic externalism is the view that one need not have introspective access to the factors which confer justification. Consider the belief-forming process constituted by my forming beliefs based on seeing things in good light and at a close distance. Suppose the belief forming process is reliable in the sense that it produces true beliefs at least 51% of the time.⁷⁶ According to reliabilism, which is the most prominent form of externalism, a belief so formed is justified *even if* one doesn’t know that the belief-forming process is reliable. It is enough that the belief is formed by a reliable process regardless of whether the person holding the belief knows it is reliable.

Non-reductionist views of testimony are externalistic because they don’t require the person relying on testimony to have positive epistemic reasons for thinking the testifier is reliable. If a belief based on testimony is formed by a belief-forming process, what Reid calls “instinct,” which is reliable, then it is justified. The person holding the testimony-based belief doesn’t have to first have good reasons to think the testifier is reliable.

⁷⁶ In discussions of reliabilism it is not uncommon to define a process as reliable if it produces true belief at least 51% of the time. One might dispute whether producing a majority of true beliefs is enough to count a process as reliable but once again, the details of this dispute are beyond the scope of this chapter.

That is not to say there is no internalistic component to externalist theories of knowledge, in general, and non-reductionism about testimony, in particular.

Externalist theories of knowledge typically include an internalistic component but that component is a negative component. Internalists require one to be aware of positive reasons for believing some proposition p in order to be justified. Externalists require that one *not* be aware of sufficiently strong reasons for disbelieving p . If one is aware of sufficiently good reasons for disbelieving p , then one is not justified in believing p , even if the belief that p was formed by a reliable process.

In such cases, we have reasons that *defeat* the external justification for believing p . A belief d is a defeater for S's belief that p if and only if d is a belief or experience that (i) makes it epistemically appropriate for S to retract her belief that p and (ii) if S were to persist in believing p , she would take her belief that p to be irrational (given d) (Bergmann, 1997, p. 93). Externalist theories of knowledge usually require that a belief be undefeated in order to count as being justified. Let's call this internal element of externalist theories of knowledge the "No-Defeater Condition:"

No-Defeater Condition (NDC): S's belief that p is justified only if S doesn't have a defeater for p .⁷⁷

The reason externalists include something like (NDC) in their accounts of knowledge is rather straightforward. It is based on the strong intuition that it is irrational to hold a belief that one consciously thinks is false, whether or not it was produced by a reliable belief-forming process.⁷⁸

⁷⁷ This is my summary of Bergmann's NDC in (Bergmann, 2006, p. 163).

⁷⁸ Not all externalists think that a subject has to be consciously aware of a defeater in order for a belief to be defeated. I will shortly introduce the distinction between conscious and reflective defeaters.

Externalists have labored to spell out the nature of defeaters in some detail.⁷⁹

Let me briefly make a few observations and distinctions regarding defeaters. First, John Pollock was the first to distinguish between undercutting defeaters and rebutting defeaters (Pollock, 1986, pp. 37-39). Rebutting defeaters give you reason to think that your belief is false. Undercutting defeaters give one reason to question the grounds one has for holding a belief. As Plantinga puts it:

...the paradigm case of rebutting defeat occurs when you first have evidence for a certain proposition, and then get evidence for its denial. But there are also (following Pollock) undercutting defeaters. You visit a factory: the items coming down the assembly line look red and you form the belief that they are indeed red, a belief that has warrant by virtue of the way you are appeared to. You are then told by the local authority that this part of the assembly line is a quality control module, where the items are irradiated by red light in order to make it easier to detect a certain kind of flaw. You then no longer believe that the items you are looking at are red—not because you have reason to believe that they are some other color, but because your belief that they are red has been undermined by what you were told (Plantinga, 1993, p. 41).

Defeaters are themselves open to defeat. One can be justified in believing *p*, even if one has a defeater for *p* so long as the defeater for *p* also has a defeater. Let's continue with Plantinga's analogy here. Suppose the president of the firm tells us that for some reason the shop superintendent regularly lies about the widgets being irradiated with red light. They aren't irradiated and never have been.⁸⁰ So, more precisely, externalist theories of justification require that a justified belief has no undefeated-defeaters.

Another distinction between defeaters is between conscious and reflective defeaters. A conscious defeater is a belief or experience that the knowing subject is

⁷⁹ In particular see (Bergmann, 1997, 2006).

⁸⁰ This particular example by Plantinga of a defeater later being defeated itself is noted by Bergmann (Bergmann, 1997, p.104). Bergmann's reference is to an unpublished manuscript of Plantinga's (Plantinga, Unpublished).

consciously aware of. A reflective defeater is a belief or experience that the knowing subject would become aware of, *given reflection* (Bergmann, 1997, pp. 111-112).⁸¹

Externalists disagree about whether (NDC) should specify conscious or reflective defeaters. Bergmann thinks NDC should only require the lack of conscious defeaters while Plantinga seems to think that warrant requires that there be no reflective defeaters for the relevant belief (Beilby, 2005, p. 172; Bergmann, 1997, p. 116).

Let's return now to non-reductionist theories of testimony. Non-reductionist theories of testimony, like most or all externalist theories, have an internalistic requirement. They require that testimony-based belief also meet (NDC). Robert Audi explains it well:

There is a different, and I think more plausible, account [than the reductionist account] that can also explain the psychological role of background beliefs. On this account, beliefs about the credibility of the attester and beliefs pertinent to the attested proposition play a mainly filtering role: they prevent our believing testimony that does not "pass," for instance because it was insincere; but if no such difficulty strikes us, we "just believe" (non-inferentially) what is attested. These filtering beliefs are like a trapdoor that shuts only if triggered. Its normal position is open, but it stays in readiness to block what should not enter. The open position is a kind of *trust*. The absence or laxity of filtering beliefs yields *credulity*; excessively rigorous ones yield *skepticism* (Audi, 2003, p. 134).

Jennifer Lackey summarizes the non-reductionist view of testimony thusly:

Accordingly, so long as there are no relevant defeaters, hearers can justifiedly accept assertions of speakers *merely* on the basis of a speaker's testimony. Otherwise put, so long as there is no available evidence *against* accepting a speaker's report, the hearer has no positive epistemic work to do in order to justifiedly accept the testimony in question (Lackey, 2006a, p.4).

⁸¹ A third kind of defeater, i.e. external defeater, is usually contrasted with conscious and reflective defeaters. An external defeater is a belief or experience that is neither a conscious nor a reflective defeater and but would defeat the relevant belief in an "appropriately specified counterfactual situation" (Bergmann, 1997, p. 112). I left it out because it is an external condition on knowledge and not an internal one.

It might initially seem that the non-reductionist picture of testimony, if true, would justify an excessive credulousness. However, as we mature we acquire more defeaters. Remember that defeaters are just beliefs and experiences that undermine justification. As we accrue more beliefs and experience we are, in effect, acquiring more “filters” for believing testimony. As we add these epistemic “red flags” to our body of knowledge we are less and less justified in taking a witnesses’ word for it. Surely by the time we are adults, if not much earlier, we have acquired enough beliefs and experiences so as to be nowhere near excessive credulousness.

Its also important to understand that non-reductionists are not committed to the claim that one *never* first needs positive reasons for believing a testifier is both reliable and in a position to know. If one has a defeater for a testimony-based belief, that defeater itself might be defeated with positive evidence for the credibility of a witness. To use Audi’s metaphor, once the “trapdoor” is closed, one might require positive evidence akin to the kind reductionist demand all of the time in order to be justified in believing some particular bit of testimony. If we know a witness to be generally untrustworthy, we have a defeater for his testimony. In those circumstances, one might only be justified in believing his testimony when one has sufficient positive reasons for doing so. What is more, one might often require such positive epistemic reasons, depending on the nature of the claim being testified to and the nature of the rest of one’s beliefs and experiences.

My purpose here is not to provide a thoroughgoing defense of non-reductionism over reductionism. It is to clarify the nature and epistemology of testimony. However, I want to briefly explain why I tend to prefer non-reductionist

accounts of testimony. Some of these reasons are implicit in the prior discussion. First, non-reductionism, as Audi points out, seems more psychologically plausible. We don't seem to run through quick inferences and arguments when people tell us things unless there is a reason to do so. This is most obvious, as Reid points out, in the case of young children. We take the word of our parents when they teach us about the world without first establishing their credibility as testifiers. More to the point, we actually acquire knowledge, and *a fortiori*, justified belief in this way.⁸²

Second, non-reductionism strikes me as more epistemically plausible. One of the general motivations for externalist theories of epistemology is the failure of internalism. It is nigh well impossible to provide empirical arguments for all of the things we know. Put differently, the history of epistemology should makes us less than sanguine about the possibility of providing internalistic justification for a non-trivial amount of what we actually know. One might take the kind of arguments David Hume makes to be representative of the history I have in mind. It seems pretty clear that we can't provide independent justification for sense perception or memory. In fact, this is part of the motivation behind naturalized epistemology.⁸³ Testimony doesn't seem to be any different in this regard. It is hard to see how we can justify, by way of internalist grounds, much of the testimony we are actually justified in accepting.

⁸² Not surprisingly, other sources of knowledge seem to be similar to testimony in this way. Young children learn a lot about the world from their senses and memory, for example, but they obviously can't provide an independent epistemic justification for doing so.

⁸³ Michael Devitt has said in the class room and in personal conversation: the best reply one can make to the epistemic skeptic is "I'm not going to play that game." Devitt contends that we should acknowledge that we can't provide the kind of grounds the skeptic demands. Rather, we should start with the premise that we possess knowledge and get on with the scientific work of describing how we have it.

It's important to reiterate here that even on non-reductionist accounts of testimony we are not necessarily justified in believing the report of some testimony just because we lack a defeater for it. The idea is that we are justified in believing it if it was caused or triggered by a reliable process *and* it lacks a defeater. Consider an analogy with sense perception. Suppose one forms a belief based on a visual experience in an epistemically unfriendly environment (e.g. in the dark and at a long distance). One is not justified in holding that belief even if he lacks a defeater for it (suppose the knowing subject is a small child who has not yet formed beliefs about when perception is reliable). Such a belief can fail to be justified on either externalist ground (e.g. unreliable process) or internalist grounds (believes it was formed by an unreliable process).

3. Testimony and Religious Belief

In the last two sections I gave a brief accounting of the concept and epistemology of testimony. In this section, I will offer a brief sketch of how the different social positions and relationships of religious believers and non-believers might result in epistemically legitimate differences of belief. The believer and the non-believer are justified in holding different religious views given their different locations in the social space and the different relational connections they have with others.

I mentioned previously, but it is worth repeating here, that my goal in this chapter is rather modest. I merely intend to offer a preliminary sketch of how the social connections of believers and non-believers may rightly result in their holding

different religious beliefs. My project in this chapter is to outline some interesting issues along these lines, not to development them very fully.

The previous section of this chapter explained the debate between reductionist and non-reductionists about testimony. I explained that the debate between the two seems to just amount to the debate between epistemic internalists and externalists applied to testimony. The debate between internalists and externalists and their analogous debate in the epistemology of testimony is enormously complex. Despite this, I intended the discussion to make non-reductionism not only appear plausible but hopefully at least somewhat more plausible than reductionism.

If this is the case, it seems that there is at least one straightforward way for the social positions of believers and non-believers to legitimately affect what they are justified in believing. Recall that on the non-reductionist model, we are justified believing the testimony of others when the process that brings about that belief is a reliable process and we lack a defeater for it. We reject testimony, or at least refrain from believing it, when we have defeaters for the beliefs being testified to. These defeaters, themselves beliefs and experiences, act like filters. They filter out the testimonial attestations which we think we have reasons not to believe.

If something like this picture is correct, then when it comes to these defeaters, at birth we are a *tabula rasa* when it comes to testimony filters. We have very few, if any, beliefs and experiences. It is in our very first environment that we begin to acquire those beliefs and experiences. Many of them are provided by our parents, caregivers, and teachers. Often, though not always, forming beliefs on the basis of

the testimony of our caregivers is a reliable process. (Again, just as often, though not always, forming our beliefs on the basis of sense perception is a reliable process).

This means that a child raised in a religious home, with religious caregivers, and religious teachers (either from attending a religious private school or perhaps Sunday school), will have a different set of filters than a child raised in a secular environment. Over time, as these filters accumulate, the religious believer and non-believer will, not surprisingly, have a different set of filters regarding the testimony of others. They will make different judgments and do so legitimately.

Consider a simple example. Consider a child is raised in a home where the parents are atheists. Suppose these parents consciously teach their children that the Bible is a seriously flawed document. Not only does it contain a legion of factual errors, it also teaches morally repugnant values such as, for instance, a smaller set of rights for women than men. Now suppose this child attends church someday after a sleepover with his religious neighbor. Suppose the Sunday school teacher on that Sunday teaches some religious or moral principle based on the Bible. It's reasonable to think that child raised by atheist parents might *justifiedly* reject the testimony (again broadly "tellings") of the Sunday school teacher because he has a set of filters which act as a defeaters for the Sunday school teaching.

One might think that the difference here is merely that the two children have different information. First, this may be partly the case. Recall that one thesis I am defending is that two people can have different epistemic perspectives because they have acquired a different set of beliefs *but* those different beliefs weren't accumulated in epistemically impoverished, dysfunctional, or deficient environments.

Second, and more importantly, it's not just that the children have different antecedent beliefs. Rather they have different *justified* beliefs in virtue of occupying different social space. In particular, they are justified in relying on different sources of epistemic authority. It's not just that they are exposed to different information or beliefs but that the epistemic legitimacy of the sources of testimony they rely on vary *because of their social connection to the relevant source*. The different social environment includes different authoritative testifiers.

Throughout this chapter, I have been using a rough and ready form of reliabilism as a general model for non-reductionist accounts of testimony. The following objection might be raised. Given a non-reductionist account of testimony along reliabilist lines, whether the two children in the example above are justified depends on whether they formed their beliefs in a reliable way. Discussions of reliabilism tend to take for granted that a reliable process is one that produces true beliefs at least 51% of the time. Then is it really plausible to think that children who form their religious beliefs based on what their parents say are relying on reliable belief-forming process? After all, what reason do we have to think that parents are experts, let alone competent about religion? Furthermore, the theist parents and the atheist parents can't both be right 51% of the time since they have seemingly inconsistent views.

A few observations here might be helpful. First, we need to distinguish between global and local experts. My local family practice doctor may not be an expert in dermatology in the same way that a research dermatologist is. The research dermatologist is an expert on a larger, more global, scale. However, when I ask my

local doctor to examine a rash, I am justified in relying on his opinion. Parents aren't usually global experts in biology but when child asks his parents what frogs eat, the child is justified in believing his parents claim that frogs eat insects. In the child's local environment, his parents are epistemic authorities on what frogs eat. Likewise, most pastors aren't bible scholars or theologians but they count as local experts on the bible and theology. In a similar way, it's not always implausible to regard parents as local experts of sorts on religious matters.

John Greco provides a helpful discussion that reinforces this point. He discusses the epistemic role of one's caregivers as socially approved sources of information (Greco, 2008).⁸⁴ Greco begins his discussion by calling our attention to a problem raised by Sanford Goldberg (Goldberg, 2008). The problem Goldberg raises is how young children can learn from testimony when they are indiscriminating consumers of testimony. Greco summarizes the problem Goldberg raises by phrasing it as an inconsistency in our pre-theoretical commitments about testimonial knowledge:

Specifically, we want to say the following:

- a. that young children can learn from the testimony of their caregivers; i.e. that they can come to *know* through such testimony,
- b. that testimonial knowledge requires a reliable consumer of testimony; i.e. that the hearer can reliably discriminate between reliable and unreliable sources of testimony⁸⁵, and
- c. that young children are not reliable consumers of testimony (Greco, 2008, p. 335).

⁸⁴ Again, "socially approved" doesn't necessarily indicate that the subject is aware that the source is a reliable source.

⁸⁵ It is important to note that one can reliably discriminate between reliable and unreliable sources of testimony without being *consciously aware* of the discrimination. I add this note because some might have a tendency to read this through the eyes of epistemic internalism.

Sanford and Greco discuss various cases to illustrate the intuitions behind these commitments. For example, the following case is supposed to confirm (c):

Room Full of Liars. Sid is gullible in the extreme: he accepts anything anyone says merely in virtue of the fact that someone said so. Sid is in a room full of inveterate liars. He immediately and uncritically believes everything each of them says. At one point he happens to bump into Nancy, the only reliable person around. Nancy reliably tells Sid that *p*, and (as a matter of course) he believes her (Greco, 2008p. 336).

Sanford and Greco both agree that in this case Sid doesn't have knowledge even though Nancy is a reliable source of testimony. This is because Sid is an unreliable consumer of testimony. The process of belief-formation here is not, on the whole, reliable. The transmitter of the testimony is reliable but the receiver is not.

Sanford and Greco offer numerous other cases designed to illustrate various points. For our purposes the most relevant is the following sort of case:

Preschool. Gabe is gullible in the extreme: like other preschool children, he accepts almost anything anyone says merely in virtue of the fact that someone said so. Gabe is told by his preschool teacher that frogs eat bugs, and this testimony is both true and reliable. However, Gabe is surrounded by his preschool classmates, who are notoriously unreliable reporters. One of them, who happens to be confused on the issue, could easily have reported that frogs never eat bugs, and Gabe would have believed this (Greco, 2008pp. 336-337).

Greco and Sanford agree that in this case Gabe has knowledge because his teacher told him so, even though he would have believed otherwise if his classmates had told him otherwise.

Greco and Sanford both deny the pre-theoretical view of testimony specified as (c) above (but for different reasons).⁸⁶ Greco argues that small children are

⁸⁶ Goldberg favors a view where the social environment involves epistemic caretakers and guides, like parents and teachers. Parents, for example, monitor and give guidance about epistemic sources. So, while a child may not by himself be a reliable consumer of testimony, he becomes so with the epistemic restrictions placed on him by his parents, teachers, and caregivers. Goldberg uses the

reliable consumer of testimony when they form their beliefs on the testimony of their caregivers:

How is it that small children can come to know from their caregiver's testimony, even if they are indiscriminating vis-à-vis unreliable testifiers in their vicinity? Answer: By forming their beliefs on the basis of testimony from their caregiver specifically, which is a reliable process. In effect, we deny (c) above: that young children are not reliable consumers of testimony. They are, so long as they are consuming the testimony of their caregivers. Moreover, the solution generalizes. How is it that people can learn from the testimony of "socially approved" sources in general (parents, teachers, experts, and definitive sources) even when the hearers themselves are locally unreliable consumers of testimony? Answer: By forming their beliefs on the testimony from approved sources specifically, which is a reliable process (Greco, 2008, p. 346).

Returning now to the religious example I gave above. The child reared in the religious home and the child reared in the atheist home have different parents, caregivers, and teachers. Thus, they are entitled to give more epistemic weight to their own parents and caregivers than to other sources of testimony. In fact, our doing so often helps us reliably form beliefs in epistemically hostile environments (e.g. the Room Full of Liars case) or in cases where we are locally poor discriminators of testimony (e.g. the Preschool case). So, at the very least we can see how occupying different social spaces (i.e., in this case, being in a different family) results in different epistemic entitlements in a way that doesn't *merely* amount to having a different set of beliefs.

analogy of a new teenager driver. The teenage driver might not be a reliable driver by herself. But she is when she follows the driving guidance and restrictions placed on her by her caregivers: driving, e.g. only in the day, in the company of an experienced driver, in good weather, etc. Under these restrictions, she is a reliable driver. Likewise, young children are reliable consumers of testimony under the epistemic restrictions and guidance imposed by their epistemic caregivers but not independently of these. Greco's punch line in this article is that social facts contribute to a solution of a very narrow kind of generality problem. He doesn't think that social factors solve all generality problems for reliabilism but help in the narrow cases of young children being reliable consumers of testimony mentioned above.

A second reply to the objection that children who form their religious beliefs on the basis of their parents and teachers are not engaging in a reliable belief-forming process is that it depends on which particular belief-forming process one is relying on. We can conceive of three kinds of cases here: cases where it is plausible that to think of parents as local experts, cases where it is not, and unclear cases. It doesn't seem far-fetched to think that parents are reliable local experts on the basic teaching of *their* religion. Maybe many parents aren't competent to teach the contents of their religious tradition's teachings but some are, so we can see in these cases how relying on their testimony might be a reliable process. On the other hand, it's hard to see how the testimony of parents is a reliable indicator of whether God exists or not. An example of an unclear case is whether the New Testament teaches that Jesus is divine. In some typical cases, parents in a religious family might count as a local expert of sorts on these kinds of religious beliefs.

Third, when reliabilists use the rough figure of 51% in discussion of reliability it is often meant as a convenient place holder and not necessarily a principled description of reliability. Further, the kinds of cases that originally motivated reliabilism lend themselves to thinking of a reliable process as one that results in true beliefs more often than not. An early paradigm case cited by reliabilists was the thermometer (Armstrong, 1973). We generally wouldn't think that a thermometer was a reliable indicator of the temperature if it was sufficiently off more often than not.

In the case of our reliance on expert testimony however, it is not at all obvious that experts must be reliable indicators of the truth most of the time. Are the world's

experts in physics, for instance, reliable indicators of the truth about the physical world 51% of the time or more? Our limited knowledge of many fields may be such that we are justified in trusting experts even though in the end it turns out that they are not right more often than not.

Fourth, we can certainly imagine cases where two sets of parents are reliable indicators of the truth of religious beliefs, even though they hold radically different views— like theists and atheists. Imagine an atheist and a theist who each have ten religious beliefs. Of course, they disagree on one religious belief— whether God exists. But suppose they agree on the other nine and that the other nine were all true. They agree that every New Testament manuscript contains at least one textual error. They agree that a person named Jesus lived who claimed to be God. They agree that Paul wrote the letter to the Romans, etc. Suppose further that when it comes to religious questions outside these ten beliefs, the parents just admitted they didn't know as refused to make any pronouncements.

In this case, even though the parents disagree on an important question, they are both reliable sources of religious knowledge in that they are reliable indicators of religious knowledge 51% of the time or more. One might object that the parents are agnostic on so many religious questions that they are not really reliable. After all, would we count a thermometer as reliable if it was right 90% of the time it indicated the temperature if it only gave a reading very seldom. In fact, I think we would consider it reliable indicator because it only gives misleading information 10% of the time. One might doubt whether the process of children forming their religious beliefs on the basis of their parents' testimony is actually reliable. However, we can at the

very least imagine cases where someone forming their beliefs in that way is using a reliable belief-forming process.

A second worry involves cases where a child acquires his defeaters (i.e. beliefs and experiences) in an environment that is systematically wrong or mistaken. Suppose, to take an extreme case that a child's parents set out to systematically deceive him. If he acquires his defeaters in such an environment, then his "testimony filter" will be systematically warped. If we count the acquiring and functioning of defeaters as an element in the belief-formation process, then in such cases the child's beliefs are not formed reliably. The "defeater module" of the process is radically unreliable. My intention is to outline some plausible cases where different locations in the greater human social network might epistemically justify a believer and a non-believer in having different religious views. My claim is not that social differences always, or even frequently, play such a role. Cases like this one fall into the category of social relations that hinder or undermine the acquisition of knowledge.

A second way that social location seems to bear on epistemic justification involves our choosing between experts. One's location in social space might legitimately lead one to defer to one set of experts over another. Let's consider three different types of cases from science where laymen rely upon experts.

First, consider cases where there is unanimity among experts. Someone is taught that the earth is a sphere. It might look flat but we take the word of the experts that it is not flat. The unanimity of opinion among experts is part of one's total evidence. Further, given what it means to be an expert, one should think that the experts are better situated epistemically than any particular layman whether the earth

is flat or not. When one considers all of one's evidence—the fact that the earth appears flat *and* the fact that all the experts say it isn't— it is more rational to believe that the earth is not flat. If someone were to persist in their belief that the earth is flat, knowing full well the unanimity of experts that it is a sphere, that person would be irrational. It is sometimes rational to defer to experts, even when one has a certain sort of privileged access:

For example, when I think, based on my own direct experiences, that I am having a heart attack, but then my doctor checks it out externally and tells me no, that I am only suffering from something called “acid reflux,” then I should probably just abandon my own prior opinion and adopt his in its place (Everett, 2001, p. 25).

Now consider a case where there is a majority but not unanimous opinion among experts. For example, someone is told that there is man-made global warming. There is a clear consensus among experts that this is so. There is not, however, unanimity. There are a small minority of scientists that dissent.⁸⁷ In this case, it still seems rational to believe the consensus opinion. Again, given one's total evidence, which includes the consensus among experts, it is rational to defer and to believe that there is man-made global warming.

Lastly, consider an example where there is no consensus among scientific experts. Consider controversial claims from the field of nutrition. There is a non-trivial amount of both changing claims by nutritionists and lack of consensus. Rather than using an actual example, it is more expedient for present purposes to use a

⁸⁷ For example: Richard Lindzen who is the Alfred P. Sloan Professor of Atmospheric Science at MIT and William Gray of Colorado State University. I mention a few contemporary examples merely to show that those who deny global warming, though a very small minority, are not in the same epistemic situation as flat-earthers. I can provide no names of contemporary tenured research scientists who are flat-earthers. It is this kind of distinction—the distinction between a very small but credible set of experts and no experts at all—that I am attempting to describe here.

fictional example.⁸⁸ Even though the example I am about to offer is fictional, it is the kind of example that is not uncommon in the actual world.

Suppose *roughly* half of experts on nutrition claimed that consuming a particular preservative, common in processed food, is unhealthy and the other half of the relevant experts said it was not unhealthy to consume. What is it rational for the layman to believe? In the first two cases, matters were much clearer because there was a consensus among the relevant experts.

One option here is to suspend judgment. This is a rationally *permissible* option but it hardly seems rationally *obligatory*. Suppose one's customary daily menu consisted of many products which included this ingredient. On the one hand, you might be eating enough of the product that it would harm your health. There is a cost associate with changing one's eating habits. If the ingredient is harmless, one needlessly bore the costs of changing one's diet. My point is that practical matters might force one to make a decision, either way. It is hard to see how one would be obligated to suspend judgment, if practical matters force one to make a decision. So, suspending judgment in these cases is rationally permissible but not obligatory.

One might consider the epistemic situation a stalemate and make a decision on purely practical grounds but I want to consider a different option. The suggestion I want to consider is whether one's location in social space can help one rationally pick which set of experts to defer to. Let's consider a few examples of the social, relational, and communal considerations which might have a legitimate bearing on matters.

⁸⁸ It is unnecessary and unwise to complicate this chapter by attempting to document specific claims from the technical literature in an unfamiliar field.

One possibility is that it is rationally permissible for members of a particular community to defer to the experts of *their* community in cases where there is no consensus among experts. Suppose, for the sake of argument, that there is no consensus among economists on the following question: do lower marginal taxes rates increase tax revenue, all other things being equal?

Suppose two laymen, one a Democrat and the other a Republican are thinking about the matter. Both laymen understand that economists are divided on the matter. The Democrat notices that left-leaning economists, like Paul Krugman, contend that tax cuts don't increase tax revenues. The Republican observes that right-leaning economists, like Milton Friedman, contend that they do increase tax revenues. If the experts are divided is it rational for the Democrat to defer to experts of "his community" and the Republican to defer to experts of "his community?" It strikes me that it is. I am not claiming that the strength of one's belief in these cases should be all that great. Perhaps the deference to one set of experts over another, in cases like this should be slight.

The first thing to note about my suggestion is that it is a widespread practice. Deferring to experts of one's own social community, loosely understood, is a common practice. It occurs as much among the intellectually sophisticated as it does among the ignorant. It is probably as common among those well-schooled in epistemology as among those who aren't.⁸⁹ Deferring to the experts of one's social community in disputed matters may be a common practice, but is it ever justified?

It seems to me that it is. Consider the example above. Suppose a voter is trying to determine whom to vote for. This voter wants a tax cut but doesn't want to

⁸⁹ Allen Stairs helped me see how widespread the practice is.

increase the Federal budget deficit. The one candidate says he will oppose tax cuts because they will “bust the budget.” If we decrease rates, we will decrease the amount of money collected. The other says that taxes cuts don’t decrease tax revenues but often increase them by increasing overall economic activity. The voter understands both sides of the argument but doesn’t understand economics enough to decide the issue. So, he turns to the experts and finds they are divided. Suppose the voter in question is a Democrat and notices that experts who tend to be associated with Democrats, like Paul Krugman, contend that lowering tax rates lowers tax revenues. It seems rational for him to defer to economists like Krugman rather than Friedman.

One objection here might be that the voter should suspend judgment. As I mentioned above, this is not always an option. Practical matters sometimes force us to choose among experts. If the voter really cares about both taxes cuts and budget deficits enough that the answer will determine his vote, it doesn’t seem like he is epistemically obligated to suspend judgment and not defer to any experts.

Suppose my intuition here is correct. What explains it? Why is it rationally permissible to defer to the experts of one’s own religious community in cases of parity among the experts? One possibility is this: one might be deferring to the experts of one’s community because of the fact that they hold similar beliefs suggests that general similarity of thinking and values. In a way, when one defers to the opinions of experts, one is giving the experts an epistemic power of attorney. It is not uncommon in cases involving the law to prepare for times when we are cognitively incapacitated by giving legal power of attorney to people who share our values and

way of thinking. We pick people who we think are likely to make decisions similar to the ones we would make in similar circumstances.

There is a sense in which experts sometimes act as a sort of epistemic stand-in for us. We don't have the reservoir of knowledge or refined judgment that the expert has so we defer to his judgments. It strikes me as perfectly rational when choosing an expert to act as an epistemic substitute for me that, all things being equal, I chose one that shares the same general beliefs, values, and outlook that I do. In deferring to experts of our community in cases where there is a lack of consensus among experts, we are basically picking someone who we think would make similar judgments to the ones we would make if we had the relevant expertise.

So, there is this widespread practice among both the intellectually sophisticated and the less so to sometimes defer to experts of one's community when there is no general consensus among experts. I offered an economic example where this seems rational. I also offered what I take to be a plausible explanation of why it is rational: when we defer to experts we give them epistemic power of attorney on a limit question. It is rationally permissible to select an epistemic stand-in based upon agreement on a broader set of beliefs, values, and shared judgments.

One objection to my claim is that by choosing to defer to the experts of our community because we share a broader set of beliefs and values, we are just engaging our biases. Sometimes, maybe even oftentimes, this is the case. But it need not be the case. It depends on the reason for deferring to the particular expert or group of experts. If it is just to get the answer one wants, then it is illegitimate. If, however, out of epistemic good faith one defers to the experts of one's community because they

have similar general beliefs and values, this doesn't strike me as mere bias. Krugman and Friedman, are partisans and some, even many, partisan Democrats or Republicans will defer to "their guy" for crass partisan reasons. But one can defer to experts for the reasons I suggested and if one does so, he isn't necessarily being biased.

So far, I have been making my argument using an economic example. Let me now give a religious example. Suppose there is a rough numerical parity between bible scholars on the dating of, say, the Gospel of Matthew. Suppose biblical scholars with a more traditional bent contend it was written prior to 70 A.D. while biblical scholars of a more progressive or secular bent contend it was written much later. My intuition here is that it may be rational for the average man in the pew to defer to the scholars of his community if he does so because they share a larger set of beliefs and values. So, there are cases where it appears perfectly reasonable, all things being equal, for a traditional Christian to defer to the judgment of the scholars of their community rather than secular ones and vice versa.

Consider cases that are more based on relationship and less based on community. Here I am even more tentative than above. Suppose you go to the doctor and are diagnosed with a certain disease and two options are laid out before you. Your doctor recommends the first course of treatment. Being very cautious, you get the opinion of three more doctors. (Assume that they are all specialists in the relevant area of medicine.) In the ends, two doctors think the first course of treatment will probably result in a better prognosis while the other two think the second will. (Assume also these opinions are representative of the field at large). Now suppose that one of the four doctors is your sibling. Is it rational, in the case of a rough

consensus among experts to defer to an expert with whom one has a close relationship?

Again, it seems to me that it might be rational to do so. Why? I am less sure in this case. One reason might be that it is rational to defer to an expert in whom one has a deep level of personal trust— someone with whom one has a very close, long-time, personal relationship. Notice that this doesn't provide you with more evidence that one course of treatment is preferable to another. You have no reason to think that your sibling is a better doctor than the others. Assume, for the sake of argument that you know that all the experts are of equal quality. It still seems that the relational proximity—the personal trust— makes it rational to defer to the expert opinions of those to whom one is socially closest, in cases of parity among experts. If I am right, then it is sometimes rational for social and relational proximity to play some role when deferring to experts.

Now consider a religious example. Suppose you are a traditional Christian and you believe in, say, the doctrine of the Trinity or the dual nature of Christ. Suppose a clever Jehovah's Witness shows up at your door and tells you that neither the Trinity nor the dual nature of Christ are taught in the Bible. Both, he says, were invented at the Council of Nicaea by the recently converted Constantine for political purposes. The Jehovah's Witness clearly knows the Bible well.

Naturally enough, you ask your pastor about this. Your pastor assures you that the Bible teaches these doctrines and that they were not invented at the Council of Nicaea. He explains some of the verses the Jehovah's Witness cited. You understand both positions and see that both men seem to understand the Bible well. It seems

rational in such a case to defer to the opinion of your local expert, your pastor. There may be many reasons to do so but one reason seems to be a matter of personal trust and closeness. Again, this doesn't provide any *evidence* that your pastor's position is right. Both men seem to be local experts of a sort when it comes to Bible knowledge but the close personal relationship one has with a pastor and the personal trust that comes with it may be one reason among others make to defer to his epistemic judgment and expertise.

In conclusion, in this chapter I have offered a tentative sketch explaining how certain kinds of social proximity might play an epistemically legitimate role in deferring to certain experts. One way social location seems relevant is in the slow acquiring of defeaters in a different social environment. Location in the social space also seems to bear on who has more epistemic authority (i.e. parents, caregivers, and teachers) and which experts we defer to in cases where there is no clear consensus. This last role is *not* an evidential one but it nevertheless seems to bear on the epistemic situation via shared values and thinking and trust. If I am right, part of the explanation of the difference in perspective between religious believers and non-believers is social and legitimately so.

In conclusion, the fact that knowledge is socially produced and distributed is an important fact that has not been sufficiently developed by analytic philosophers. In this chapter I have only offered a preliminary outline of a few ways in which differences in the social environments of religious believers and non-believers may affect justification of their religious beliefs. I laid the groundwork for this brief sketch by first discussing the philosophical literature on the concept of testimony and

then by explaining and weighing in on the debate surrounding the epistemology of testimony.

Conclusion

In this dissertation, I have argued that philosophers of religion, particularly those thinking about the argument from evil, should give far more weight to perspectival factors than they typically do. Much of the recent discussion of the evidential argument from evil has revolved around what we humans are in an epistemic position to know *simpliciter*. As I pointed out in the introduction, skeptical theists argue that we are in no epistemic position to infer that God does not exist from the existence of evil. While the kind of concerns raised by skeptical theists are important and need to be taken seriously, analytic philosophers of religion should place more emphasis on the fact that we humans aren't always in the same epistemic position.

This point can be seen when we consider religious experience and the evidential argument from evil. Assume, for the sake of argument, that religious experience is a genuine source of knowledge about God. Some people have had religious experiences. If, in at least some cases, religious experiences provide genuine knowledge of God, then the epistemic situation of those who have had religious experiences and those who haven't differ in important ways.

Consider a religious believer who has had religious experiences that justify his belief in God. When such a person considers the evidential argument from evil, part of what he will have to take into account is his total evidence. He might concede that the existence of evil makes the existence of God less likely but also see that his religious experiences make the existence of God more likely. When he considers all of the evidence he has for and against the existence of God, he might be justified in

believing that God exists despite the contrary evidence provided by the existence of evil. A non-believer, who has had no such experiences, might very well have a different pattern of total evidence for and against the existence of God and justifiedly come to a different conclusion. It should be apparent, then, how differences in epistemic perspective might justify believers and non-believers to come to different conclusions about the evidential argument from evil.

The argument I have made in this dissertation differs from the preceding point in two ways. First, I have argued that differences in epistemic perspective affect what one is justified in believing regarding the logical argument from evil not just the evidential argument. I did this in the first three chapters by laying out the logical argument from evil, Plantinga's freewill defense, and three objections to the freewill defense. I argued that the objections I consider in chapters one and two are not sufficient to undermine the freewill defense.

But it is the discussion of the objection discussed in chapter three that makes it apparent how perspectival factors play an important role in the logical argument from evil. As Howard-Snyder and Hawthorne point out, the logical argument from evil and Plantinga's freewill defense both involve an essential epistemic component. Mackie takes himself to *show* that God doesn't exist. Plantinga's project involves *demonstrating* the logical compatibility of God and evil. To show or demonstrate either that the existence of God and evil are compatible or that they are incompatible is partly an epistemic task. Because both Mackie's and Plantinga's tasks have an epistemic component, perspectival considerations come into play.

As the example involving religious experience above shows, differences in perspectival factors like experiences, background beliefs, and values partly determine whether a demonstration is epistemically successful for a particular person.

Bayesians call attention to the importance of background beliefs to epistemic justification. Others have made much of the role of worldviews in our reasoning. In this dissertation, I chose to focus on perspectival factors that are more neglected and less obvious than the three mentioned above.

In the second half of the dissertation I argued that in addition to the more obvious ways that people can be in different epistemic positions (e.g. by having different experiences and background beliefs), we should also consider the epistemic role of less apparent perspectival factors like practical interests and location in social space.

In chapter four I argued that practical interests have a legitimate bearing on epistemic justification. Knowledge is usually thought of as being something like justified true belief. I argued that philosophers like Jason Stanley and John Hawthorne are on to something when they contend that there is also a practical condition on knowledge. I concluded that differences in the practical interests of believers and non-believers have a legitimate epistemic bearing on how they decide religious questions.

In chapter five I emphasized the social production and distribution of knowledge. Knowledge is spread across a social network. Different people know different things. We share the knowledge we have with other nodes on the social network. But the successful transfer of knowledge from one node to another must be

done in certain appropriate ways. I argued that one's location in social space plays an important epistemic role by, among other things, helping to determine which set of experts or epistemic authorities one defers to.

If the perspectival considerations I discuss are important in the ways I suggest, then it is possible for believers and non-believers to look at similar evidence for and against the existence of God, like the argument from evil, and legitimately draw different conclusions. The importance of epistemic position is not only important to the success or failure of the evidential argument from evil but also the logical argument for the reasons I set out in the dissertation. However, the importance of perspectival considerations has broader applications in the philosophy of religion and philosophy in general.

Bibliography

- Achinstein, P. (1978). Concepts of Evidence. *Mind*, 87, 22-45.
- Adams, M. M. (1990). Introduction. In M. M. Adams & R. M. Adams (Eds.), *The Problem of Evil* (pp. 1-24). New York: Oxford University Press.
- Adams, R. M. (1977). Middle Knowledge and the Problem of Evil. *American Philosophical Quarterly*, 109-117.
- Alston, W. (1991a). The Inductive Argument from Evil and the Human Cognitive Condition. *Philosophical Perspectives*, 5, 29-67.
- Alston, W. (1991b). *Perceiving God: The Epistemology of Religious Experience*. Ithaca: Cornell University Press.
- Aquinas, T. (1920 [1274]). *The Summa Theologica of St. Thomas Aquinas* (F. o. t. E. D. Province, Trans. 2nd and Revised ed.). London: Burns Oates and Washbourne.
- Armstrong, D. (1973). *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.
- Audi, R. (1997). The Place of Testimony in the Fabric of Knowledge and Justification. *American Philosophical Quarterly*, 34, 405-422.
- Audi, R. (2003). *Epistemology: A Contemporary Introduction to the Theory of Knowledge* (2nd ed.). New York: Routledge.
- Beilby, J. (2005). *Epistemology as Theology*. Aldershot: Ashgate Publishing.
- Bergmann, M. (1997). *Internalism, Externalism, and Epistemic Defeat*. Unpublished Ph.D. dissertation, University of Notre Dame, South Bend.
- Bergmann, M. (2006). *Justification without Awareness*. Oxford: Oxford University Press.
- Black, M. (1952). The Identity of Indiscernibles. *Mind*, 61(242), 153-164.
- BonJour, L. (2003). *Epistemic Justification*. Malden: Blackwell.
- Chalmers, D. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press.
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford: Clarendon Press.
- Cohen, S. (1988). How to Be a Fallibilist. In J. E. Tomberlin (Ed.), *Epistemology*. Atascadero: Ridgeview.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57-89.
- Cruse, A. (2006). *A Glossary of Semantics and Pragmatics*. Edinburgh: Edinburgh University Press.
- DeRose, K. (1999a). Contextualism: An Explanation and Defense. In J. Greco & E. Sosa (Eds.), *The Blackwell Guide to Epistemology* (pp. 187-205). Oxford: Blackwell.
- DeRose, K. (1999b). Solving the Skeptical Problem. In K. DeRose & T. A. Warfield (Eds.), *Skepticism: A Contemporary Reader* (pp. 182-219). New York: Oxford University Press.
- DeRose, K. (2004). The Problem with Subject-Sensitive Invariantism. *Philosophy and Phenomenological Research*, LXVIII(2), 346-350.

- Draper, P. (1996a). Pain and Pleasure: An Evidential Problem for Theists. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 12-29). Bloomington: Indiana University Press.
- Draper, P. (1996b). The Skeptical Theist. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 175-192). Indianapolis: Indiana University Press.
- Dretske, F. (1970). Epistemic Operators. *Journal of Philosophy*, 67, 1007-1023.
- Earman, J. (2000). *Hume's Abject Failure: The Argument Against Miracles*. New York: Oxford University Press.
- Evans, J. (1983). Lafollette on Plantinga's Free Will Defense. *International Journal for Philosophy of Religion*, 14, 117-121.
- Everett, T. (2001). The Rationality of Science and the Rationality of Faith. *The Journal of Philosophy*, 98(1), 19-42.
- Fantl, J., & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1), 67-94.
- Fantl, J., & McGrath, M. (2007). On Pragmatic Encroachment in Epistemology. *Philosophy and Phenomenological Research*, LXXV(3), 558-559.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48, 15-34.
- Flint, T. (1998). *Divine Providence*. Ithaca: Cornell University Press.
- Fogelin, R. J. (2003). *A Defense of Hume on Miracles*. Princeton: Princeton University Press.
- Freddoso, A. J. (1998). Luis De Molina. In E. Craig (Ed.), *Routledge Encyclopedia of Philosophy* (Vol. 6, pp. 461-465). London: Routledge.
- Fricker, E. (1995a). Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony. *Mind*, 104, 393-411.
- Fricker, E. (1995b). Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony. *Mind*, 104(414), 393-411.
- Gale, R. (1991). *On the Nature and Existence of God*. New York: Cambridge University Press.
- Gan, B. (1982). Plantinga's Transworld Depravity: Its Got Possibilities. *International Journal for Philosophy of Religion*, 13, 169-177.
- Garson, J. (2007). *Modal Logic*. Retrieved 12-14-2007, 2007
- Gettier, E. (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121-121.
- Goldberg, S. (2008). Testimonial Knowledge in Early Childhood Revisited. *Philosophy and Phenomenological Research*, 76(1), 1-36.
- Goldman, A. (1967). A Causal Theory of Knowing. *The Journal of Philosophy*, LXIV(12).
- Goldman, A. (1979). What is Justified Belief. In G. Pappas (Ed.), *Justification and Knowledge*. Dordrecht: D. Reidel.
- Graham, P. (1997). What is testimony. *The Philosophical Quarterly*, 47(187), 228-232.
- Greco, J. (2006). *Contextualism vs. Interest-Dependent Invariantism*. Retrieved 2-05-08, 2008, from <http://knowability.blogspot.com/2006/11/contextualism-vs-interest-dependent.html>
- Greco, J. (2008). Discrimination and Testimonial Knowledge. *Episteme*, 4(3), 335-351.

- Harper, W. (1981). A Sketch of Some Recent Developments in the Theory of Conditionals. In W. Harper, R. Stalnaker & G. Pearce (Eds.), *Ifs: Conditions, Belief, Decision, Chance, and Time* (pp. 3-38). Dordrecht: D. Reidel Publishing Co.
- Hasker, W. (1989). *God, Time, and Knowledge*. Ithaca: Cornell University Press.
- Hasker, W. (1992). The Necessity of Gratuitous Evil. *Faith and Philosophy*, 9, 23-44.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hill, D. J. (2005). *Divinity and Maximal Greatness*. New York: Routledge.
- Hookway, C. (1990). *Scepticism*. London: Routledge.
- Howard-Snyder, D., & Hawthorne, J. (1999). Transworld Sanctity and Plantinga's Free Will Defense. *International Journal for Philosophy of Religion*, 44(1), 1-21.
- Hume, D. (2000 [1748]). *An Enquiry Concerning Human Understanding*. Oxford: Oxford University Press.
- James, W. (2003 [1896]). The Will to Believe. In M. Raymond & C. Bernard (Eds.), *God Matters: Readings in the Philosophy of Religion* (pp. 199-202). New York: Longman Press.
- Kaplan, D. (1989). Demonstratives. In J. Almong, J. Perry & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481-463). Oxford: Oxford University Press.
- Kenny, A. (1979). *The God of the Philosophers*. Oxford: Oxford University Press.
- Kripke, S. (1963a). Semantical Analysis of Modal Logics. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9, 67-96.
- Kripke, S. (1963b). Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16, 83-94.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge: Harvard University Press.
- Lackey, J. (2006a). Introduction. In J. Lackey & E. Sosa (Eds.), *The Epistemology of Testimony* (pp. 1-21). New York: Oxford University Press.
- Lackey, J. (2006b). The Nature of Testimony. *Pacific Philosophical Quarterly*, 87, 177-197.
- Lafollette, H. (1980). Plantinga on the Free Will Defense. *International Journal for Philosophy of Religion*, 11(2), 123-132.
- Laing, J. D. (2005). Middle Knowledge. In *Internet Encyclopedia of Philosophy*.
- Lewis, C. S. (1947). *Miracles: A Preliminary Study*. New York: The McMillan Company.
- Lewis, D. (1973). *Counterfactuals*. Cambridge: Harvard University Press.
- Lewis, D. (1986). *On the Plurality of Worlds*. New York: Blackwell.
- Mackie, J. L. (1955). Evil and Omnipotence. *Mind*, 64(254), 200-212.
- Mackie, J. L. (1970). The Possibility of Innate Knowledge. *Proceedings of the Aristotelian Society*(70).
- Manis, R. Z. (2004). The Possibility of Transworld Depravity. *Southwest Philosophy Review*, 20(1), 201-210.
- McCloskey, H. J. (1997). God and Evil. In P. Angeles (Ed.), *Critiques of God: Making the Case Against Belief in God* (pp. 203-224). Amherst: Prometheus Books.
- McGrath, M. (2004). *Review of Knowledge and Lotteries*. Retrieved 2-1-08, 2008, from <http://ndpr.nd.edu/review.cfm?id=1467>

- Nielsen, K. (1993). No! A Defense of Atheism. In J. P. Moreland & K. Nielsen (Eds.), *Does God Exist: The Debate between Theists and Atheists* (pp. 48-54). Buffalo: Prometheus Books.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge: Harvard University Press.
- Owens, D. (2000). *Reason without Freedom: The Problem of Epistemic Normativity*. London: Routledge.
- Pascal, B. (1995 [1670]). *Pensees* (A. J. Krailsheimer, Trans.). London: Penguin Books.
- Perszyk, K. J. (1998). Free will defence with and without Molinism. *International Journal for Philosophy of Religion*, 43(1), 29-64.
- Pike, N. (1963). Hume on Evil. *Philosophical Review*, 72, 180-197.
- Plantinga, A. (1974). *The Nature of Necessity*. Oxford: Clarendon Press.
- Plantinga, A. (1977). *God, Freedom, and Evil*. Grand Rapids: William B. Eerdmans Publishing Co.
- Plantinga, A. (1985a). Replies. In J. E. Tomberlin (Ed.), *Alvin Plantinga* (pp. 313-399). Dordrecht: D. Reidel.
- Plantinga, A. (1985b). Reply to Robert M. Adams. In J. E. Tomberlin & P. Van Inwagen (Eds.), *Alvin Plantinga*. Dordrecht: D. Reidel.
- Plantinga, A. (1985c). Self-profile. In J. E. Tomberlin (Ed.), *Alvin Plantinga* (pp. 3-93). Dordrecht: D. Reidel.
- Plantinga, A. (1993). *Warrant and Proper Function*. New York: Oxford University Press.
- Plantinga, A. (1996). Epistemic Probability and Evil. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 69-97). Indianapolis: Indiana University Press.
- Plantinga, A. (Unpublished). *Naturalism Defeated*. Unpublished manuscript.
- Pollock. (1986). *Contemporary Theories of Knowledge*. Savage: Rowman and Littlefield.
- Putnam, H. (1975). The Meaning of 'Meaning'. In H. Putnam (Ed.), *Philosophical Papers, Vol. II : Mind, Language, and Reality* (pp. 215-271). Cambridge: Cambridge University Press.
- Ramsey, F. (1931). General Propositions and Causality. In R. B. Braithwaite (Ed.), *Foundations of Mathematics and Other Logical Essays* (pp. 237-257). New York: Harcourt, Brace and Company.
- Reid, T. (1997 [1764]). *An Inquiry into the Human Mind on the Principles of Common Sense*. University Park: The Pennsylvania State University Press.
- Reid, T. (2002 [1785]). *Essays on the Intellectual Powers of Man: A Critical Edition*. University Park: Pennsylvania State University Press.
- Rowe, W. (1988). Evil and Theodicy. *Philosophical Topics*, XVI(2), 119-132.
- Rowe, W. (1996a). Modal Versions of the Ontological Argument. In M. Y. Stewart (Ed.), *Philosophy of Religion*. Sudbury: Jones and Bartlette.
- Rowe, W. (1996b). The Problem of Evil and Some Varieties of Atheism. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 1-11). Bloomington: Indiana University Press.

- Rowe, W. (1998). In Defense of 'The Free Will Defense'. *International Journal for Philosophy of Religion*, 44, 115-120.
- Smith, D. (2005). Critical Notice of Hawthorne's Knowledge and Lotteries. *Philosophical Books*, 46, 123-131.
- Sosa, E. (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Stalnaker, R. (1968). A Theory of Conditionals. In N. Rescher (Ed.), *Studies in Logical Theory* (pp. 98-112). Oxford: Basil Blackwell.
- Stalnaker, R. (1981a). A Defense of Conditional Excluded Middle. In W. Harper, R. Stalnaker & G. Pearce (Eds.), *Ifs: Conditions, Belief, Decision, Chance, and Time* (pp. 87-104). Dordrecht: D. Reidel Publishing Co.
- Stalnaker, R. (1981b). A Theory of Conditionals. In W. Harper, R. Stalnaker & G. Pearce (Eds.), *Ifs: Conditions, Belief, Decision, Chance, and Time* (pp. 41-55). Dordrecht: D. Reidel Publishing Co.
- Stanley, J. (2005). *Knowledge and Practical Interests*. New York: Oxford University Press.
- Stine, G. (1976). Skepticism, Relevant Alternatives, and Deductive Closure. *Philosophical Studies*, 29, 249-261.
- Swinburne, R. (1995). *The Coherence of Theism* (Revised Edition ed.). Oxford: Clarendon.
- Unger, P. (1975). *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press.
- Van Inwagen, P. (1977). Ontological Arguments. *Nous*, 11, 375-395.
- Van Inwagen, P. (1996). The Problem of Evil, the Problem of Air, and the Problem of Silence. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 151-174). Bloomington: Indiana University press.
- Vogel, J. (1990). Are There Counterexamples to the Closure Principle? In *Doubting: Contemporary Perspectives on Skepticism* (pp. 13-25). Dordrecht: Kluwer.
- Wierenga, E. R. (1989). *The Nature of God: An Inquiry Into Divine Attributes*. Ithaca: Cornell University Press.
- Willard, D. (1993). Language, Being, God, and the Three Stages of Theistic Evidence. In J. P. Moreland & K. Nielsen (Eds.), *Does God Exist: The Debate between Theists and Atheists* (pp. 197-217). Buffalo: Prometheus Books.
- Wykstra, S. J. (1990). The Human Obstacle to Evidential Arguments from Suffering: On Avoiding the Evils of 'Appearance'. In M. M. Adams & R. M. Adams (Eds.), *The Problem of Evil* (pp. 138-160). New York: Oxford University Press.
- Wykstra, S. J. (1996). Rowe's Noseeum Argument from Evil. In D. Howard-Snyder (Ed.), *The Evidential Argument from Evil* (pp. 126-150). Bloomington: Indiana University Press.