

2011

An Inquiry Into the Moral Significance of Doxastic and Epistemic States: Examining the Circumstantial Element of Moral Obligation

Gregory William Keenan
Claremont McKenna College

Recommended Citation

Keenan, Gregory William, "An Inquiry Into the Moral Significance of Doxastic and Epistemic States: Examining the Circumstantial Element of Moral Obligation" (2011). *CMC Senior Theses*. Paper 156.
http://scholarship.claremont.edu/cmc_theses/156

This Open Access Senior Thesis is brought to you by Scholarship@Claremont. It has been accepted for inclusion in this collection by an authorized administrator. For more information, please contact scholarship@cuc.claremont.edu.

CLAREMONT McKENNA COLLEGE

**AN INQUIRY INTO THE MORAL SIGNIFICANCE OF DOXASTIC AND
EPISTEMIC STATES: EXAMINING THE CIRCUMSTANTIAL ELEMENT OF
MORAL OBLIGATION**

SUBMITTED TO

PROFESSOR DUSTIN LOCKE

AND

DEAN GREGORY HESS

BY

GREGORY KEENAN

FOR

SENIOR THESIS

SPRING/2011
APRIL 25, 2011

ACKNOWLEDGMENTS

I would like to express gratitude to my parents, William and Donna Keenan, and acknowledge that without their years of sacrifice I would not have the privilege of studying and writing about the philosophical questions that inspire me. I would also like to thank my friends Rio Fischer and Matthew Wissa for generously listening (or at least convincingly pretending to listen) to me repeat the same incoherent ramblings until these ramblings developed into a focused project. Finally, I would like to thank Professor Dustin Locke for instilling in me an appreciation of epistemology and for his invaluable guidance throughout the process of writing this thesis.

TABLE OF CONTENTS

INTRODUCTION	5
CHAPTER 1	
I. A Critique of the Subjective View	23
II. A Critique of the Objective View	38
III. Motivating the Epistemic View	48
CHAPTER 2	
I. A Dilemma for the Epistemic View	52
II. The Unreasonableness Worry for the Epistemic View	56
III. The Manipulation Worry for the Epistemic View	63
IV. Concluding Thoughts	74
BIBLIOGRAPHY	81

INTRODUCTION

In this work, I seek to explore an area of intersection between ethics and epistemology. Generally speaking, ethics and epistemology can overlap in one of two ways. The first of these two ways deals with epistemological questions regarding moral issues, such as which moral theory is correct (e.g. Utilitarianism, Kantianism, etc.). Examples of epistemological questions regarding moral issues are: How can an agent know what he morally ought to do in a given situation? ; Can we ever be justified in believing that a particular moral theory is correct? ; What reasons do we have to avoid being skeptics about moral facts? The second way in which ethics and epistemology intersects deals with the ethical relevance of an agent's epistemic state. That is to say, how do an agent's justified beliefs, knowledge, and ignorance affect her moral obligations in a particular situation or her praiseworthiness or blameworthiness for her actions? It is with this latter category of epistemic-ethical analysis that I will be concerned with in this work.

Upon first glance, it may seem a peculiar thing to suggest that an agent's epistemic status has any bearing on his moral obligation. After all, many of us tend to think of moral obligation as being some sort of transcendent standard that we must live up to—for example, a set of hard-and-fast rules such as “Don't Kill, Don't Cheat, Don't Steal.” Surely what an agent does or does not know would not affect his obligation to abide by such rules, though it certainly would affect the agent's ability to live up to these

standards.¹ But this assumption that an agent's epistemic state is irrelevant to his moral obligation is not as intuitive as it might seem.

Consider the following scenario to clarify the relationship between an agent's epistemic status and his moral obligation:

*Alejandro is driving around town in his recently purchased car. Wanting to see how fast his new vehicle can go, he accelerates down a deserted side-road, speeding up to 70mph. However, as he approaches the intersection of this deserted side-road and a heavily-traveled main road, Alejandro has a decision to make: should he slow down and cautiously approach the intersection or should he keep traveling at his current speed.*²

Ought Alejandro slow down and approach the intersection with caution? I take it to be the case that the majority of readers would think that he ought to because Alejandro should not carelessly put other drivers (and perhaps himself) in danger. But this common-sense response is not as straightforward as it appears. Imagine that you are a passenger in Alejandro's car and you say to him, as you both approach the intersection, "I think you

¹ Undoubtedly many would stop to make a nuanced concession here. Perhaps some unenlightened person who has never been taught the importance of these rules might not be held as blameworthy for violating these rules as someone who is aware of them and violates them knowingly. Nevertheless, many would still hold to the point of view that the unenlightened person ought not engage in acts such as killing, cheating, and stealing, and thus would still maintain the view that this agent's epistemic state does not affect his moral obligation to refrain from actions such as these. Rather the agent's ignorance will simply prevent him from recognizing his moral obligation and will affect his ability to meet his moral obligation.

² Ross 1939, p. 152.

ought to slow down, Alejandro. After all, you really should not to put other drivers at risk.” Alejandro, nodding in agreement, proceeds to speed up through the intersection, which fortuitously is clear of any other vehicles. Has Alejandro violated your advice? It is not clear that he has. In fact, there is a sense in which he may properly say to you that he has fully complied with your advice since he did not put any other drivers at risk by speeding through the intersection because there were, *in fact*, no other vehicles in the intersection as his car passed through it. At this point, you would most likely respond by saying to your reckless driver, “Still, you had no way of knowing that there would not be any vehicles in the intersection when we sped through it. So you should not have done that.”

This scenario highlights a very important aspect of the ethical deliberations that we each engage in on a daily basis—the decisions that we make regarding what actions to perform are made in a state of (at least partial) ignorance about the circumstances we find ourselves in, as well as, about the consequences of our actions. Further, this ignorance will undoubtedly lead to our possessing false beliefs about our circumstances and the consequences of our actions in many situations. These realizations lead us to an important ethical question: What significance, if any, does an agent’s epistemic state (i.e. what he or she has reasons for believing) bear upon what that agent ought to do in a particular situation (i.e. his or her moral obligation in that situation)?³ The hypothetical exchange

³ Throughout this paper I will use the phrase “what an agent ought to do” and that agent’s “moral obligation” interchangeably. This is not entirely correct as it is possible that an agent ought to do act X even though he does not have a moral obligation to do act X. After all, there are certainly amoral reasons (e.g. societal, pragmatic, and professional reasons) that an agent ought do something. Nevertheless, a central characteristic of moral obligation is that if an agent has a moral obligation to do something, then this

between Alejandro and you, his passenger, provides two competing answers to this question. By saying to Alejandro, “You could not have known that there would not be any vehicles in the intersection when we sped through it. So, you should not have done that...” you would be expressing the view that an agent’s epistemic status does significantly affect an agent’s moral obligation. In contrast, Alejandro’s response that there, *in fact*, was no one else in the intersection expresses the view that it is the agent’s actual circumstances that determine what he or she ought to do rather than epistemic considerations, such as what he or she could have reasonably known in this situation.

The main focus of this work is to consider three theories which attempt to explain the proper relationship between an agent’s moral obligation and that agent’s doxastic state (i.e. what that agent believes to be true), her epistemic state, and her actual circumstances. This relationship will henceforth be referred to as “the circumstantial element of moral obligation”. Roughly put, the *objective view* defines an agent’s moral obligation as a function of her actual circumstances; the *subjective view* defines an agent’s moral obligation as a function of what that agent believes to be her actual circumstances; and the *epistemic view* defines an agent’s moral obligation as a function of what that agent’s evidence gives her reason to believe are her actual circumstances.⁴

obligation will trump all other considerations. Though it does not follow that if S ought to do X, then S has a moral obligation to do X, it does follow that if S ought to do X, then it cannot be the case that S has a moral obligation to do some act other than X. Thus, if S has a moral obligation to do X, then S ought to do X. For the purposes of this investigation into ethical questions, I will set aside the scope of amoral considerations. As such, using these two phrases interchangeably should prove to be an innocuous simplification.

⁴ N.B. This restricts the question of the circumstantial element of an agent’s moral obligation to an agent’s beliefs and evidence about her circumstances, not her beliefs and evidence about moral facts (i.e. the correct moral theory). This will be elaborated upon shortly.

This question regarding the circumstantial element of moral obligation stems from the fact that an agent's beliefs about his circumstances often pull apart from what his circumstances actually are. There are two ways in which this can happen: 1) the agent is mistaken about fact (i.e. he believes a falsehood) 2) the agent is ignorant of fact (i.e. he fails to believe the truth). For example, John might glance at a passing car and mistakenly believe that he sees a stranger driving the vehicle, when it is in fact his close friend, because he is mistaken about the circumstances. Or, John might fail to glance at a passing car altogether, and therefore fail to believe that his friend is in the car due to his ignorance of the circumstances. In both instances John's beliefs about his circumstances differ from what his circumstances actually are.

Such cases can best be explained in terms of the objective and subjective elements of a situation, as W.D. Ross does in his discussion of our central question.⁵ The objective element of a situation consists of an agent's actual circumstances. In the example above, the objective element of the situation would be that John's friend is driving the car which passes by John. On the other hand, the subjective element of a situation consists of what an agent believes to be his actual circumstances.⁶ In the example above, in which John mistakenly believes that he sees a stranger driving the passing car, the subjective element of the situation would be that a stranger is driving the passing car. In the example above, in which John fails to believe that his friend is driving the car because John fails to notice the car altogether, the subjective element of the situation would be that there is no passing

⁵ Ross 1939, p. 146.

⁶ It is important not to describe the distinction between the objective element and the subjective element of the situation as being the distinction between real and not real. After all, each element has a legitimate realness to it.

car, and thus, no driver of the passing car.⁷ And, if it were the case that John correctly believed that his friend was driving the passing car, then the subjective element and the objective element of the situation would both be that John's friend was driving the passing car—though they would maintain their status as independent elements of the same situation.

At this point, I will explore each of our three views in greater detail and extrapolate what each view entails about the circumstantial element of an agent's moral obligation. Consider the following example, from which I demonstrate each of the three views:

Proper moral theory dictates that one should always keep one's promises, and Bob has promised to arrive at his uncle's dinner party on time. Bob knows from numerous previous experiences that it takes only twenty-five minutes to get from his house to his uncle's house, where the dinner party is being held. However, Bob has just watched a traffic report on television, which has informed viewers that Clarkson Street (the only street that Bob can take to get to his uncle's house) is experiencing delays that last up to thirty-five minutes. Nevertheless, Bob disregards this information because he has an unfounded, albeit deep-seeded,

⁷ It is important that we be careful to define the subjective elements of an agent's situation as that agent's beliefs about what his circumstances **are** rather than merely saying that it is an agent's beliefs about his circumstances. After all, John might believe the claim that his red shirt is hanging in the closet and he might also believe the claim that red is an ugly color. Only the first belief will constitute a subjective element of the situation. The relevant difference between the two believed claims is that the first claim is a factual claim, whereas the latter is a value claim. We must be careful to limit the phrases objective and subjective elements of a situation to factual claims.

distrust of traffic reports.⁸ Thus, Bob believes that the twenty-five minutes it would, under ordinary circumstances, take to get to the dinner party are in fact sufficient for his arriving to the party on time. As it turns out, Bob loses track of time and leaves only fifteen minutes prior to the party begins. Fortunately, he still makes the party on time.

The first view that I will discuss is the objective view. According to the *objective view*, an agent, S, ought to perform act X if and only if X is actually the best option that S has. By *best option*, I mean that act which optimally captures the moral good as determined by the proper moral theory.⁹ An agent's best option is the product of two factors, namely, the moral facts (i.e. facts about what is morally good) and circumstantial facts (i.e. the agent's actual circumstances).¹⁰ For example, if the proper moral theory is Utilitarianism, then X is S's best option if and only if X, of all the acts metaphysically available to S, is the act that maximizes utility given S's circumstances.¹¹ From this definition of the objective view, we can extrapolate that it is morally right for S to perform X if and only if S has no better option than X, and it is morally wrong for X to

⁸ This distrust is not based upon any past experiences, statistics, or any other type of evidence but rather just a deeply held belief that such reporters are not to be trusted. If, from a psychological point, the reader argues that his belief must be based off of some experience, assume this belief to be as unfounded and unjustified as possible.

⁹ The use of the phrase "best option" to explain each view is adopted from Michael J. Zimmerman in Zimmerman 2008, pp. 2-3.

¹⁰ The schematic division between moral facts and circumstantial facts is inspired by Terrance McConnell in McConnell 1988.

¹¹ An act is "metaphysically available" to S if and only if it is an act that S can actually perform. Contrast this with "doxastically available" acts, which are acts that S believes he can actually perform. Thus, for the purposes at hand, for X to be S's best option, must be metaphysically available to S insofar as S can do X. The term metaphysically available is borrowed from John Martin Fischer in Fischer 2006, p. 17.

perform X if and only if S has a better option than X. Under the objective view, merely engaging in the act which S believes is his best option (even if S believes with very good reason) is not sufficient for S's satisfying his moral obligation. Nor is it morally wrong for S to engage in the act which he believes is not his best option. In the above case, Bob ought to leave the house at least X minutes prior to the party, where X represents the minimum number of minutes necessary to actually make the party on time—in this particular case, fifteen minutes. Under this view, Bob would have violated his moral obligation *if* fifteen minutes had not been enough time to make the party on time. Since, however, the fifteen minutes turned out to be enough time to make the party Bob has satisfied his moral obligation.

According to the *subjective view*, an agent, S, ought to perform act X if and only if S believes that X is the best option that he has.¹² From this definition, we can extrapolate that it is morally right for S to perform X if and only if S believes that he has no better option than X, and it is morally wrong for S to perform X if and only if S believes that he does have a better option than X. Merely engaging in the act that is S's best option is not sufficient for S's satisfying his moral obligation. Nor is S's failing to perform his best option sufficient for S's action being morally wrong. In the case above, Bob has violated his moral obligation and is wrong because he left with only fifteen minutes to spare when he believed that it would take more time to get to his uncle's house. But it is not because his evidence indicated that he would need more time to arrive

¹² As previously noted, throughout this discussion moral facts are being held fixed. Thus, when we say that the subjective view defines moral obligation as a function of what the agent believes to be his best option, this is tantamount to saying that this view defines moral obligation as a function of what his best option would be if his beliefs about what his circumstances are were correct.

at his uncle's house (as the epistemic view asserts) that his action is wrong. Under this view, Bob ought to leave at least twenty-five minutes prior to the party because he believes that it will take him twenty-five minutes to arrive at the party on time. The fact that Bob's belief is unfounded and unjustified given his evidence (i.e. the traffic report) has no bearing on his moral obligation. Thus, it is important to notice that under the subjective view, neither the actual circumstances of an agent's situation nor that agent's evidence will affect his moral obligation in any way.¹³

According to the *epistemic view*, an agent, S, ought to perform act X if and only if S's evidence indicates that X is the best option that S has. From this definition, we can extrapolate that it is morally right for S to perform X if only if S's evidence indicates that S has no better option than X, and it is morally wrong for S to perform X if and only if S's evidence indicates that S has a better option than X. Neither performing the act that actually is his best option nor performing the act that he believes is his best option is sufficient for S's satisfying his moral obligation. Further, neither engaging in the act which he believes is not his best option nor engaging in the act which actually is not his best option is sufficient for S's action being morally wrong. In the case above, Bob has violated his moral obligation and is wrong because he left with only fifteen minutes to spare when his evidence indicated more time was needed. But it is not because he believed that he needed more time to get to the party (as the subjective view would assert) that his action is wrong. Under the epistemic view, Bob ought to leave 60 minutes

¹³ That is to say that neither the actual facts about an agent's situation nor his evidence directly affect an agent's moral obligation. If the facts or his evidence influence the agent's beliefs, then they do indirectly affect an agent's moral obligation. But even in such a case his obligation is such as it is only in virtue of his beliefs and not as a direct result of the facts or evidence that influence these beliefs.

prior to the party because his evidence indicates that it will take him 60 minutes to arrive at the party on time. Thus, it is important to notice that (similarly to the subjective view), under the epistemic view, the actual circumstances of an agent's situation do not affect his moral obligation in any way; obligation under this view is a result of what an agent's evidence indicates to be the case, regardless of the validity of what the evidence indicates.¹⁴

At this point, two very important questions arise regarding the epistemic view: "What does it mean for something, E, to be S's evidence?" and "What exactly does it mean for one's evidence, E, to indicate some proposition, P?" I will explore, in greater depth, an answer to the first question in my criticism of the epistemic view (Chapter 2) and much has been written elsewhere by others on the second question.¹⁵ For the time being, I think a quick overview of an intuitive understanding of evidence will suffice. Consider the above example, in which Bob knows from numerous experiences that it takes twenty-five minutes to get from his own house to his uncle's house, where the dinner party is being held. These consistent experiences are properly considered a piece of Bob's evidence. Additionally, however, Bob has just seen a news report indicating delays up to 30 minutes. This news report is also a piece of Bob's evidence. Taken

¹⁴ Of course, the claim that the epistemic view is indifferent to the validity of what the evidence indicates is contingent upon the claim that defective evidence (i.e. evidence which indicates a falsehood) can exist. I think, however, it is safe to assume that defective evidence does exist. For example, if I backed into your car and left a note on your windshield that reads, "Sorry, I ran into your car. You can reach me at 516-633-2633" this would be evidence that (you could reach the person who hit your car at the above phone number) (call this X). Nevertheless, if I fabricated that phone number and you could not reach me at that number, then the note would still be evidence for X, but in this case it would be defective evidence since X is not true.

¹⁵ A particularly insightful analysis is provided by Branden Fitelson in Fitelson 2001.

together, these two pieces of evidence (along with any other relevant pieces of evidence) serve as Bob's (total) evidence.

Further, it is obvious that this evidence indicates something that has a relevant bearing on Bob's decision of when to leave for his uncle's house. Bob's evidence indicates that it will take him 60 minutes to travel to his uncle's house. If he decides to leave 25 minutes before the party is set to begin, then Bob's evidence indicates that this will not be enough time, and thus it indicates that this is not Bob's best option. While this example does not explore the complexities of what it takes to qualify as a piece of evidence nor does it explore in exhaustive detail the relationship of evidence's indicating something, hopefully it provides a concrete example of an agent having a set of evidence and elucidates the role that this evidence plays by indicating certain things.

At this point, it is important to defend the claim that our question (i.e. what is the circumstantial element of moral obligation?) is not a question about which moral theory is correct.¹⁶ Kantians, Contractualists, Consequentialists, etc. could all answer the question regarding the circumstantial element of moral obligation in the same way. Further, there can, and in fact do, exist divisive answers regarding the circumstantial element of moral obligation within any particular camp of moral theory.¹⁷ For example, even if moral theorists agreed upon utilitarianism as the proper moral theory, the distinct question remains as to whether agents will have a moral obligation to do what *in fact* will maximize utility, to do what they *believe* will maximize utility, or to do what their *evidence indicates* will maximize utility. Nevertheless, the above example about Bob and

¹⁶ McConnell 1988, pp. 79-80.

¹⁷ In Moore 1912, G.E. Moore proposes an objective view of consequentialism. In Jackson 1991, Frank Jackson espouses a non-objective view of consequentialism.

the supplemental analysis may lead some to believe that the objective view is tied to a consequentialist moral theory. After all, it appears that Bob only satisfies his obligation, under the objective view, in virtue of the consequences of his actions.

Recall, however, that according to the objective view an agent S ought to perform act X if and only if X is actually the best option that S has, where best option means the act which actually maximizes the moral good, as determined by the proper moral theory, given S's circumstances. In short, there exist two questions. First, what is the proper moral theory? The answer to this, in conjunction with the agent's actual circumstances, will determine the agent's best option. The second question is ought the agent perform what is actually his best option, what he believes to be his best option, or what his evidence indicates is his best option? The point is that moral theory determines what an agent's best option is. The appropriate view of the circumstantial element of moral obligation, on the other hand, determines what the relationship between the agent's best option and the agent's moral obligation is. For example, if it is the subjective view, then the agent ought to do what he believes to be his best option. But the fact remains that the moral facts are fixed and are independent of what the agent believes to be the case.¹⁸ Thus, the objective view is perfectly compatible with deontological moral theories. For example, if it the proper moral theory deems lying immoral, then moral theory precludes

¹⁸ There is nothing any more consequentialist about the objective view's stating that Bob has fulfilled his obligation because he engage in the action that *actually* satisfied his best option (i.e. getting to the party on time) than there is, for example, about the subjective view's stating that Bob would have fulfilled his obligation if he had engaged in the action that he *believed* would satisfy his best option (i.e. getting to the party on time). Both views judge the moral rightness or wrongness of the agent's actions based upon that actions relation to his best option. Each simply employs a different relationship to his best option as the standard by which to judge the moral worth of his actions.

any act of lying from being the agent's best option to lie (no matter how beneficial or profitable the consequences!) As such, the objective view could never deem an act of lying as being an agent's moral obligation if this were the case. As we can see, the question of what moral theory is correct is in fact distinct from the question of the circumstantial element of moral obligation.

Now that we have demonstrated the distinction between questions regarding the circumstantial element of moral obligation and those regarding moral theory, it is important to establish a clearer understanding of what is meant by the phrase *moral obligation*, since a clear understanding of the concept of moral obligation is necessary to ensure that there is no confusion over what exactly it is of that each of these views attempts to provide an account. There are two essential characteristics of any satisfactory account of moral obligation. First, there exists the fundamental characteristic of moral obligation according to which if an agent has a moral obligation to do X, then it would be wrong for that agent to do other than X.¹⁹ In short, we always have a duty to satisfy our moral obligations and whatever those moral obligations are, it would be wrong not to satisfy them.²⁰ For example, if an agent has a moral obligation to keep his promises, then it would be wrong for him not to keep his promises regardless of any other consideration.

¹⁹ In Zimmerman 2008, p. 2, Zimmerman expresses this condition in terms of a moral ought is a "contrary of 'wrong'". In Pink 2007, p. 1, Thomas Pink expresses this condition in the following way: "moral obligation seems to involve a reason that is demanding." As such, if X is S's moral obligation then it is not the case that doing X is just good or advisable or prudent, but rather that doing X is something that S "must do."

²⁰ It is important to note here that an agent's performing a morally wrong act need not entail that agent's being morally blameworthy for that act. The relationship between an act's wrongness and an agent's blameworthiness for that act's wrongness are closely tied together, but the one need not entail the other. It will be important to keep this distinction between moral wrongness and moral blameworthiness in mind, for the purposes of the discussion at hand; an analysis of this distinction will be provided later in this chapter.

And if there were any other considerations that made it such that the agent should in fact neglect to keep his promise, then it would be incoherent to say that this agent had a moral obligation to keep his promise. Further, there exist derivative obligations resulting from such primary moral obligations. That is to say, just as an agent might have a moral obligation to keep his promises, he would, as a result, have obligations to perform and refrain from specific acts that satisfy this primary moral obligation. For example, he would have derivative obligations such as not making promises which he could not keep, returning a book that he promised that he would return, etc.

This leads us to the second essential characteristic of moral obligation, namely that moral obligation is prescriptive. Moral obligation pertains to an agent's actions, and, as such, we are under moral obligation "to do things or refrain from doing" things.²¹ In other words, moral obligation is directly tied to how an agent should use his or her agency by prescribing a certain course of action in a given situation. Since moral obligation trumps all other consideration and it prescribes a given course of action, an agent in a situation inquiring into what he or she ought to do should primarily be concerned with his/her moral obligation. Thus, an agent's moral obligation should provide the answer to the question "what ought I do?" when an agent is attempting to choose between a set of actions to engage in, in a given situation.²² For example, if an agent S has a moral obligation to do what will actually maximize utility (an objective view of utilitarianism), then this obligation *will* provide the answer to the question "ought the agent to engage in act X?" Of course this does not necessarily mean that the answer to

²¹ Pink 2007 p. 1.

²² Zimmerman 2008, pp. 1-2.

this question need be obvious to the agent, nor even accessible to the agent²³. This is a matter of debate, and one's position on this issue will be contingent upon one's answer to our central question concerning the circumstantial element of moral obligation. Objective view holders will be committed to the position that agents can have moral obligations of which they are unaware.²⁴ Many subjective view proponents, on the other hand, motivate their view of moral obligation by appealing to the argument that an agent's moral obligation should make the answer to this question accessible.²⁵ Still, the relationship between an agent's moral obligation and the answer to the question "what ought I do?" is clear and is neutral to the debate about the circumstantial element of moral obligation. The agent's moral obligation will answer the question what action that agent ought to perform in a given situation.

Now that our central question is clear, it is imperative to ask why this question's answer is a worthwhile pursuit. The reason that this investigation is an important one is that subjective elements of a situation hold much more importance than we might at first glance give them. Often when we think about subjective elements of a situation that are in contrast with the objective elements of that situation, we think of the former as being somehow defective or peripheral, and thereby unimportant. As such, this might serve as

²³ Whether or not moral obligation is such that moral agents will always have access to the answer of the question "what ought I do?" is a matter of debate. And this is an issue which will be explored later in this paper. The point I am making here is simply to show that regardless of how one answers our central question, proponents of any of the three views should agree that an agent's moral obligation answers the question what action ought that agent perform in any given circumstance. This is an essential characteristic of moral obligation and one which does not incline one to any of the three views under consideration.

²⁴ Moore 1912, pp. 192-195.

²⁵ Ross 1939, pp. 163-165.

one motivation to be quick to dismiss alternatives to the objective view. But we should not be so quick to dismiss subjective elements as having importance only in virtue of their convergence with objective elements nor underestimate just how commonplace situations in which objective elements and subjective elements pull apart really are. Both elements have their own *independent* importance, and the realization of this fact will help highlight the importance of the question at hand, insofar as it will force us to more seriously consider alternatives to the objective view.

In fact, even those who subscribe to the objective view of moral obligation will permit subjective elements of the situation to possess a substantial independent importance as a determinant of an agent's blameworthiness for a particular act. Consider the following example involving two doctors who give their respective patients a pill that actually kills them rather than cures them. Though both doctors carry out extensive research and inquiry, the first doctor gives her patient the pill with the belief that the pill will cure the patient, whereas the second doctor gives her patient the pill with the belief that the pill will kill the patient. While both doctors have, under the objective view, done wrong by killing their patients, many would find it appropriate to hold the first doctor less, if at all, culpable than the second doctor for his wrongdoing.²⁶ Thus, even objective view holders will affirm the independent value and importance of the subjective elements of a situation by permitting the degree of blameworthiness of an act to be a function of

²⁶ This understanding of the independent importance of subjective elements of a situation manifests itself in our legal tradition under the concept of *Mens Rea* (*guilty mind*)

the subjective element.²⁷ In light of the fact that we often permit subjective elements to hold significance independently of their conformance to objective elements, we realize that the objective view cannot just be the default view, but one which must be argued for.

Further, situations in which the subjective elements of that situation and the objective elements of that situation pull apart are quite common, and the normalcy of such situations drives the importance of our central question. When we realize that on a daily basis we are aware of only a small portion of the objective elements of our situation, we realize just how foreign to us many of the objective elements of our situations in fact are. Further, the subjective elements of a situation are always near and dear to us. That is, they are in fact what we take ourselves to experience when we have experiences. For example, if I believe that I just saw my friend Steve drive by my house, when in fact some stranger named James was driving the car that I just saw, the experience I take myself to have just had—that is, the subjective element—is that Steve just drove by my house. We may say sure, but the experience that you *really* had is that James drove by your house. This is true, but the experience that I have had is that of Steve driving by. We might even say that I should not have believed that Steve just drove by given a certain set of evidence. For example, perhaps Steve had earlier that week told me that he would be out of the country for the next month and it simply slipped my mind. Again, while both of these contentions are true (i.e. that the subjective elements of a situation can be either wrong or unjustified or both), neither changes the fact that I, as the agent in this situation, both take my experience to be the subjective elements of the

²⁷In Moore 1912, pp. 192-193, G.E. Moore relies upon such a distinction between moral obligation and moral blameworthiness in response to criticisms leveled against the objective view.

situation and am ignorant of the objective elements of situation. This being the case, we should not dismiss the importance of subjective elements of a situation in regards to moral obligation without good reason and serious reflection, for the resultant inaccessibility of the objective elements will certainly carry with some considerable costs (more on this later).

In the first chapter I will explore the dialogue between objective view advocates and subjective view advocates, highlighting the important criticisms that each camp levels against the other. In the process, I will attempt to argue that these criticisms are the result of two worries—one faced by the subjective view and one faced by the objective view—which the ideal view of the circumstantial element of moral obligation will avoid. Then, I will attempt to build a strong case for the epistemic view on these grounds, by demonstrating the ways in which the epistemic view seems to avoid these worries which plague the other two views. In chapter two, I will propose a dilemma for the epistemic view, whereby it appears necessary for the epistemic view to become vulnerable to either the worry plaguing the subjective view or the worry plaguing the objective view. Ultimately, I will argue that this dilemma for the epistemic view reveals an interesting and troubling relationship between the two worries in question.

CHAPTER 1

Section One: *A Critique of the Subjective View*

In this section, I will attempt to establish good reason to reject the subjective view. In particular, I will explore two criticisms of the subjective view: that the subjective view leads to agent moral infallibility and that the subjective view defeats the need for genuine moral deliberation. Then, I will attempt to show that each of these criticisms result from a larger problem for the subjective view, which I will call the “manipulation worry.” Ultimately, I aim to demonstrate that the subjective view’s susceptibility to the manipulation worry provides good reason to reject the subjective view.

Let us now compare the objective and subjective views to get a better understanding of the reasons to doubt the subjective view. Upon first glance, the objective view seems obviously correct. After all, it seems to make intuitive sense that an agent ought to engage in the act that is *actually* his best option. For example, consider the following case, which depicts a practical application of these two theories²⁸:

Dr. Anderson has a patient who suffers from a serious but not life threatening disease. She must choose between either giving the patient drug A or drug B with

²⁸ I have borrowed this case from Zimmerman 2008, Chap 1.

the knowledge that one of these drugs kills the patient and the other completely cures the patient, though she does not know which drug has which effect. Dr. Anderson, fully believing that drug A is the drug that will cure the patient, gives her patient drug A. Unfortunately for all persons involved, drug A actually is the drug which kills the patient.

Should we say (as the subjective view entails) that Dr. Anderson did what she ought to have done by prescribing the drug which she believed would actually cure the patient and that she would have been morally wrong if she had given her patient the drug which would have actually cured him? I imagine that most readers would find such an analysis unconvincing and would be inclined to say that it would have been morally right of her to give drug B instead. Obviously more detail is necessary to assign the appropriate amount of blame for her wrongdoing. If, for example, Dr. Anderson had asked five other doctors, all of whom agreed with her judgment, and she had done additional research to confirm her belief, then we might say that she exhausted all reasonable resources and thereby should be absolved from blame for her wrongdoing.²⁹ Nevertheless, it would seem odd to say that by simply believing that giving drug A is her best option, Dr. Anderson would render her giving the drug which actually cures the patient (drug B) to be the morally wrong act and her giving the drug which actually kills the patient (drug A)

²⁹ As we can see from this type of reasoning, the objective view proponent might consider the agent's epistemic state and doxastic state in assigning blame for that agent's act. That is, an agent's ignorance or defective evidence could alleviate an agent from blame. Still, ignorance or defective evidence is not a factor that could change the agent's moral obligation, under the objective view. Thus, it is still the right action to give the patient drug B in the above case, though we can understand why Dr. Anderson did not.

to be the morally right act. By permitting an agent's beliefs to play an important role in assigning praise or blame for an agent's actions without permitting those beliefs to affect that agent's moral obligation, the objective view captures our *prima facie* intuitions about the case at hand and provides motivation for favoring the objective view over the subjective view.

In addition to the surface appeal of the objective view, there exist two strong criticisms against the subjective view: it leads to agent moral infallibility and it leads to the deflation of moral deliberation.³⁰ I will attempt to show that each of these criticisms results from a larger problem for the subjective view, which I will call the "manipulation worry".

The first criticism considered, which is notably leveled against the subjective view by Michael J. Zimmerman, is that the subjective view results in agent moral infallibility.³¹ A formulation of Zimmerman's argument is as follows:

P1) *The subjective view*: If S has the belief that S's best option is X (call this

³⁰ There are undoubtedly more criticisms of this view. Nevertheless, these two criticisms strike me as the most threatening challenges to the subjective view and most important for the ideal view of the circumstantial element of moral obligation to avoid.

³¹ Zimmerman 2008 p. 5.

N.B. It appears that Zimmerman may be using the term Subjective View to refer to the view that many others call the "putative view". The putative view asserts that an agent S ought do X if and only if S believes that it is right to do X. Contrast this to the subjective view which asserts that an agent S ought to do X if S believes that X is his best option (where best option refers to that act which will maximize moral good as determined by proper moral theory). In other words, in the putative view both the agent's moral facts and the circumstantial element of his moral obligation are both functions of the agent's beliefs. Under the subjective view, on the other hand, the moral facts are fixed, and what he ought to do is a function of his beliefs about his circumstances given these moral facts. Still, since the use of the term "best option" in his definition of the subjective view precludes this view from being the putative view, I will assume that Zimmerman is in fact dealing with the subjective view.

belief B), then S ought to do X

P2) If S has belief B, then S knows that he believes B

P3) S knows the subjective view is true

C) S knows that he ought to do X (from 1,2,3)

From this point, Zimmerman concludes that there is a resulting moral infallibility insofar as the agent will always know what he ought to do.³² While this argument does provide a good argument regarding the deflation of moral inquiry (a point I will revisit later), I do not believe that it establishes a troubling brand of moral infallibility for the subjective view. Nevertheless, by exploring the failures of this argument, we will be able to discover a more serious criticism of the subjective view, through which we can reconstruct a stronger and more successful version of the moral infallibility criticism.

The first criticism of Zimmerman's argument is that it simply does not follow from S's knowing what he ought to do in a given situation that S will satisfy that obligation or that S will satisfy that obligation so long as he attempts to satisfy his obligation. While moral deliberation may become deflated under the subjective view insofar as an agent will always know what he has a moral obligation to do, moral obligation itself does not appear to become deflated because it may still be difficult for the agent in question to live up his moral obligation, despite knowing what that obligation is. Consider, for example, a hypothetical agent, Jim, who becomes very destructive when he drinks cocktails. As such, Jim knows that he ought to refrain from drinking at the

³² Zimmerman 2008 pp. 13-14.

dinner party he will be attending later that night. Nevertheless, despite the knowledge of what he ought to do and despite the fact that he decides to act in accordance with his obligation, he fails to refrain from drinking.³³ Cases such as these are commonplace. Quite often we act contrary to what we know we ought to do (e.g., “I ought to exercise more”, “I ought to procrastinate less”, etc.). It seems that Zimmerman drastically overestimates the connection between an agent’s knowing what he ought to do and that agent’s following through on this knowledge so as to act in accordance with his obligation. For a successful argument that the subjective view leads to moral infallibility, Zimmerman must connect his conclusion that S knows what he ought to do with the desired conclusion that S will do what he ought to do (or S will do what he ought to do so long as he attempts to satisfy his obligation).

Furthermore, Zimmerman’s objection rests upon P3 (i.e. that the agent knows that the subjective view is true). But this seems like an unrealistic and unfair condition to rest the moral infallibility objection upon. After all, nothing about the subjective view entails that if it were true then agents would know that it was true. Thus, an argument for the moral infallibility criticism which rests upon the condition that the agent in question

³³ Some will argue that this is not a counter-example whatsoever to Zimmerman’s argument. After all, if the agent in question is an addict, then perhaps it is not in his power to refrain from drinking (i.e. he cannot refrain from drinking). If this is the case, then it is in fact not his obligation to refrain from drinking, since we have defined an agent’s obligation to preclude those acts which are not metaphysically available to him. Nevertheless, I believe this line of thinking is mistaken. When thinking about addictions we tend to think of them as being the result of some *akratic* behavior. It would be wrong to say that our agent cannot refrain from drinking, for surely if he kicks his addiction the next day, it would be improper to describe that action as not being metaphysically available to him. Thus, we mustn’t take the fact that our addict has a lapse of judgment to mean that he could not have refrained from drinking.

knows the subjective view to be true seems to carry little serious weight.³⁴ As such, the moral infallibility objection, in its current formulation will not do.

Nevertheless, Zimmerman and others who run this line of argument are on to a significant point. Through examining arguments that the subjective view leads to moral infallibility, we realize that the subjective view permits an agent to exert a significant amount of control over his or her moral obligation. As stated in the argument above, under the subjective view, what agent S ought to do is a function of what he or she believes to be his or her best option (i.e. belief B). That is to say, given that S has belief B, S has a moral obligation to do X. The problem with this aspect of the subjective view is not that agents know their beliefs (as suggested in the original formulation of the infallibility argument). Rather, the problem stems from another relationship between agents and their beliefs, namely that agents can exert significant voluntary control to determine their beliefs either directly or indirectly.

The view that agents can voluntarily determine what beliefs they possess and do not possess is known as *doxastic voluntarism*, an admittedly contentious view. Still, if doxastic voluntarism is true, then the subjective view becomes susceptible to a host of problems as a result of its being vulnerable to what I will refer to as a “manipulation worry”. A theory of moral obligation is vulnerable to the *manipulation worry* if and only

³⁴ If we disregard the first criticism of this moral infallibility argument (i.e. that it overestimates the connection between an agent’s knowing what he ought to do and that agent doing what he ought to do) then this argument does establish that under the subjective view, *if* an agent knows that the subjective view then that agent will possess some form of moral infallibility. Though limited in applicable scope, this criticism would be theoretically troubling. Still, this is not the strongest possible formulation of an argument for the subjective view’s leading to moral infallibility. Furthermore, the first criticism seems to successfully undermine this argument.

if that theory allows an agent to, without violating his moral obligation, author (either in part or in whole) his own moral obligation by exerting voluntarily control to intentionally determine his moral obligation, rather than having to live up to a moral obligation independent of his control. According to what I will call *The Principle of Non-Manipulation*, a correct theory of moral obligation will be immune from the manipulation worry.

At this point, however, one might object that the Principle of Non-manipulation begs the question against the subjective view by demanding a universal constraint upon moral obligation that is fundamentally inconsistent with a view that renders moral obligation a function of some individualistic condition, such as an agent's beliefs. But this concern is misguided. The subjective view is only inconsistent with the Principle of Non-manipulation in virtue of certain cases of doxastic voluntarism. If agents could not exert such control over their beliefs, then the subjective view would not be subject to a manipulation worry, while still maintaining a non-universal, non-objective nature. That is to say, the Principle of Non-Manipulation in no way entails any sort of universal, objective moral standards. For example, take the facetious hypothetical view of moral obligation according to which an agent's moral obligation is a function of a randomly generated computer print out. For example, person A's print out stated the he should perform act X, person B's print out stated that she should perform act Y, etc. Despite the fact that moral obligation, under this view, is contingent upon non-universal and individualistic factors, this view in no way violates the Principle of Non-Manipulation.

As a result, the Principle of Non-Manipulation is not question begging against the subjective view nor does it impose a standard of universality.

Since the manipulation worry for the subjective view rests upon the truth of doxastic voluntarism, we must now briefly explore doxastic voluntarism, determine its credibility, and understand how it, in conjunction with the subjective view, results in a manipulation worry. There are two types of doxastic voluntarism: direct doxastic voluntarism and indirect doxastic voluntarism. According to direct doxastic voluntarism, an agent can voluntarily cause himself or herself to hold a particular belief (or prevent himself or herself from holding a particular belief) by willing himself or herself to hold that belief. This type of doxastic voluntarism clearly demonstrates how an agent's having control over his beliefs would lead to a manipulation worry for the subjective view:

P1) *Subjective view*: If S has a belief that some action X is his/her best option (call this belief B), then S has a moral obligation to perform X.

P2) *Direct doxastic voluntarism*: S can directly and voluntarily choose to hold belief B

C) S can directly and voluntarily choose to have a moral obligation to perform X
(1,2)

For example, if moral theory dictates that an agent should always keep her promise and Jane has promised that she will return the book that she has borrowed from her friend by Friday, then under the subjective view Jane will have a moral obligation to

engage in the action that she believes will get the book back to her friend by Friday. But if direct doxastic voluntarism is true, then Jane can simply will herself to believe that carelessly leaving the book on a random park bench will get the book back to her friend by Friday, and thus she will meet her moral obligation by acting in such a way. The problem is clear. Given direct doxastic voluntarism, Jane can voluntarily and intentionally choose her moral obligation for herself.

Nevertheless, the subjective view advocate has a clear response to this problem; it would be outrageous to endorse direct doxastic voluntarism because it seems highly questionable that agents can simply will themselves into holding or not holding a belief. This leads us to indirect doxastic voluntarism, according to which an agent can indirectly choose her beliefs by voluntarily engaging in certain actions that will lead to the agent's acquiring these beliefs. For example, consider an agent, Robin, who for some odd reason strongly desires to believe that *cogito ergo sum* is Latin for "I think, therefore I am." Unfortunately for Robin, given the fact that she only speaks English, she cannot will herself to hold that belief. Nevertheless, she can engage in actions that will bring it about that she will come to hold this belief (for example, studying Latin and learning the Latin phrase for "I think, therefore I am"). Thus, even if we reject the extremely implausible direct doxastic voluntarism, there are still beliefs that an agent can exert doxastic voluntarism indirectly. So long as there exists some act X, such that an agent can voluntarily perform X with the intention of bringing it about that she holds some belief B and her performing X brings it about that she holds B, then the subjective view will be plagued by the manipulation worry because of indirect doxastic voluntarism.

Consider the following example which highlights a manipulation worry for the subjective view, as a result of indirect doxastic voluntarism:

John promises to donate a kidney to his friend Brian. According to the proper moral theory, it is morally wrong for agents not to keep their promises. John, knowing this moral fact, believes that his best option is to donate his kidney to Brian. Thus, John has a moral obligation to donate his kidney to Brian, according to the subjective view. But on the eve of the operation, John becomes frightened and contemplates not giving Brian his kidney. Then John sees an ad for an experimental drug on television. The drug allows the drug taker to alter his future beliefs in the following way: the agent thinks of a belief that he desires to hold, the agent takes the pill, the agent then slips into a state of unconsciousness, and upon waking the agent holds the desired belief. John orders this pill and prior to taking it he thinks of the following belief: “Brian has unfortunately passed away prior to the surgery.” John slips into a state of unconsciousness and upon waking he holds the belief that Brian has passed away, and as a result, he no longer holds the belief that donating a kidney to Brian is his best option.

Now that John no longer believes that donating a kidney to Brian is his best option, he no longer has a moral obligation to follow through on his promise and donate that kidney, under the subjective view. Thus, John has intentionally determined his moral

obligation through voluntary action, and, as such, the subjective view has been exposed to the manipulation worry.

One may argue that perhaps the subjective view is not exposed to a manipulation worry, because there will be some factor which renders an agent's altering his beliefs morally wrong. For example, perhaps the agent will believe that it would not be his best option to alter his beliefs so as to change his moral obligation. In such a case, the agent would not be able to manipulate his moral obligation without violating his moral obligation at T1, thereby precluding this from being an instance of the manipulation worry.³⁵ While it is true that such cases will not qualify as instances of the manipulation worry, this defense is contingent upon an agent's holding such beliefs. If an agent does not have the belief that it is his best option to not manipulate his beliefs in this fashion, then the subjective view has no built-in-defense against an agent manipulating his moral obligation³⁶. As such, this response on behalf of the subjective view is at best an argument for the infrequency of the subjective view's violating the Principle of Non-Manipulation and does not defend against the criticism that the subjective view is vulnerable to the manipulation worry.

³⁵N.B. A theory of moral obligation is vulnerable to a "manipulation worry" if and only if the agent in question can, **without violating his moral obligation**, author (either in part or in whole) his own moral obligation by exerting control to voluntarily determine his moral obligation rather than having to live up to a moral obligation independent of his control.

³⁶ I will use the phrase "built-in-defense" throughout this paper in the following way. A theory of moral obligation has a built-in-defense against the manipulation view if and only if, under that theory, an agent can voluntarily engage in acts with the intention of determining his future moral obligations and these acts do in fact determine his future moral obligations, however he cannot engage in such acts without violating a moral obligation of his at the time of engaging in that act.

At this point, subjective view holders might respond that the objective view is also subject to the manipulation worry. In fact, it may seem that any theory of moral obligation will be subject to the manipulation worry, which would motivate a *reductio ad absurdum* argument against the Principle of Non-Manipulation.³⁷ After all, agents have the ability to manipulate the actual state of the world, of which moral obligation is a function under the objective view. Consider the following case, which concretely presents this worry:

Jack sees an injured animal locked in a cage. He also sees a key next to the cage which he believes will open that cage. Given that the proper moral theory indicates that we should treat all beings with care and respect, Jack's best option is to open the cage with the key and free the animal. But, in hopes of rendering this action not his moral obligation, Jack swallows the key. As a result, it is no longer his best option to open the cage.

At first glance it may seem that this case exhibits a manipulation worry for the objective view. After all, at a given point in time (T1), Jack's best option is to open the cage and free the animal. At a later point in time (T2), Jack engages in an action (i.e. swallowing the key) which, in turn, changes his best option and, in effect, changes his moral obligation. Thus, this case demonstrates that the objective view permits Jack to engage in voluntary act with the intention of determining his moral obligation and these

³⁷ I would like to thank Professor Dustin Locke for bringing this worry to my attention.

acts do in fact determine his moral obligation. It seems that the objective view also is susceptible to the manipulation worry.

But this analysis is incomplete. It is true that Jack can determine his moral obligation at T2 under the objective view. That is, by swallowing the key John can make it true that at T2 he will no longer have a moral obligation to free the caged animal. Nevertheless, it is important to note that he cannot do so without violating his moral obligation at T1 (i.e. to free the caged animal). Thus, the objective view is not subject to a manipulation worry because it has a built-in-defense against the manipulation worry.

These two cases highlight an interesting point. The subjective view is subject to a manipulation while the objective view is immune to the manipulation. Though both views permit, in certain instances, an agent to voluntarily and intentionally determine their moral obligation at T2, the objective view permits this only if that agent violates his moral obligation at T1. Thus, the objective view has a built-in-defense against the manipulation worry because it will always be morally wrong to determine one's moral obligation under the objective view (as a result of the necessary condition that the agent violates his moral obligation at T1 by doing so). The subjective view, on the other hand, has no built-in-defense against the manipulation worry because it will be possible for an agent to determine his moral obligation at T2 and it will permit him to do so without violating his moral obligation at T1. The question now arises as to why the objective view has this built-in-defense and the subjective view does not. What deeper characteristic of each of these views is responsible for the one's susceptibility to and the other's immunity from the manipulation worry? Answering this question will be of

immense importance to discovering the ideal view of the circumstantial element of moral obligation. I will return to offer some thoughts about this question at the end of Chapter 2. For now, let us move forward with the knowledge that the objective view is immune from the manipulation worry and the subjective view is susceptible to this worry, and, as a result, any attempt to dismiss the Principle of Non-Manipulation by appealing to the claim that any view on the circumstantial element of moral obligation will violate this principle is misguided.

We have now seen that the subjective view is subject to a manipulation worry and thus violates the Principle of Non-Manipulation. Of course, all this tells us so far is that we must choose between endorsing the Principle of Non-Manipulation and endorsing the subjective view. At this point, we must examine problematic consequences stemming from this violation of the Principle of Non-Manipulation in order to motivate the Principle.

Several important consequences plague the subjective view as a result of its being exposed to a manipulation worry. The first consequence is the revamped moral infallibility worry. Recall that, according to the subjective view, if S has the belief that X is his best option (call this belief B), then S has a moral obligation to perform act X. Further, given S's control over his beliefs and the susceptibility of the subjective view to the manipulation worry, it is possible that S could, without violating his moral obligation, make it such that he believes that C is his best option, where C is some action such that S will be able to do C (virtually) any time that S attempts to do C.³⁸ In turn, S can make it

³⁸ There are obviously cases in which an agent will attempt to do C and some force outside his control prevents him from performing the act he attempts to (e.g. a case of

such that his moral obligation is to perform some act such that he could not fail to perform so long as he attempts to perform it. Thus, it is possible for S to make it such that he cannot fail to perform his moral obligation, so long as he attempts to.³⁹

The second consequence of the manipulation worry for the subjective view is the deflation of moral deliberation. Under the subjective view, an agent need not deliberate as to what his moral obligation is in certain instances because he could simply determine what his moral obligation is himself. Given the manipulation worry, an agent can avoid genuine moral deliberation by simply sidestepping the need for introspection and deliberation and instead determine his moral obligation for himself rather than searching for what it is that he believes is his best option. As such, the deflation of moral deliberation stems not from the fact that an agent's moral obligation is a function of his beliefs and that he knows these beliefs (as Zimmerman claims), but rather from the fact that an agent can manipulate his moral obligation under the subjective view and eliminate the need for serious moral deliberation altogether.⁴⁰

In conclusion, the subjective view appears to have deep flaws and manifests consequences unbefitting a desirable theory of moral obligation. Thus, our analysis of the problems plaguing the subjective view gives us good reason to reject the subjective view

spontaneous paralysis). The idea here is simply that an agent, S could make his moral obligation to do C, where C is an act that will be incredibly easy for S to perform, such that S could never fail to perform C except for outside forces that render C metaphysically unavailable to him.

³⁹ Some may argue that this is not genuine moral infallibility, which would be that an agent could never fail to satisfy his moral obligation (without the conditional "so long as he attempts to"). While I disagree with this analysis, I also think it is unimportant because the subjective view renders an agent in possession of a troubling ability for moral flawlessness (whether we choose to call this ability a brand of moral infallibility or not is a semantic debate).

⁴⁰ Zimmerman 2008 p. 14.

as the correct view of the circumstantial element of moral obligation. It is important to note that these problems result from the subjective view's susceptibility to the manipulation worry and resultant violation of the Principle of Non-manipulation.

Section Two: *A Critique of the Objective View*

Now that we have examined some of the problems facing the subjective view, we must also examine the problems which critics of the objective view claim that the objective view is subject to. Though the objective view appears to be more in-line with our intuitions and the subjective view has been shown to fall to the manipulation worry, the objective view also faces significant problems. In particular, the objective view contains what I will call the "inaccessibility condition" and this condition will have problematic consequences. In this section, I will first describe the inaccessibility condition. I will then discuss a consequence of this condition, which subjective view holders have taken to be an argument against the objective view. Then, I will demonstrate that while this consequence is not, in fact, problematic for the objective view, the inaccessibility condition is not an innocuous entailment of the objective view as it leads to a substantial problem—what I will call the "unreasonableness worry".

The objective view makes an agent's moral obligation a function of the objective element of his situation. The objective element of an agent's situation is comprised of the actual circumstances of that situation, as opposed to the subjective element of that situation, which is comprised of that agent's beliefs to be the actual circumstances of that situation. An important distinction between the objective element and the subjective element of an agent's situation is that unlike the objective element of an agent's situation, an agent always or has access to what the subjective element of that situation.⁴¹ As such, because an agent has access to the subjective element of his situation, and because that agent's moral obligation is a function of the subjective element of his situation under the subjective view, an agent will always have access to what he is morally obligated to do under the subjective view. However, this is not true under the objective view. It is possible for the objective element of an agent's situation to be inaccessible to him. As such, because there are cases in which the objective element of an agent's situation are inaccessible to that agent, and because that agent's moral obligation is a function of the objective element of his situation under the objective view, what an agent is morally obligated to do under the objective view will sometimes be inaccessible to that agent (i.e.

⁴¹ By have "access to" I mean that even if an agent does not presently aware of each component of the subjective element of his situation, that he could come to be aware of each component of the subjective element of his situation. For example, imagine that a car passes Tom as he is walking in town. He takes notice of the car but is not consciously aware of what color that car is and as such is not aware of what he believes the color of the car to be. Still, if someone asks Tom what color the car that just passed him was, he could bring this aspect of the situation to his attention, at which point he realizes that he believes that the car was blue and thus he now comes to be aware this part of the subjective element of his situation. As such, Tom had access to this part of the subjective element of his situation even when he did not know (or was not aware of it) prior to his being asked.

he could not become aware that X is his best option). This entailment of the objective view is what I will call the “inaccessibility condition”.

There are several reasons that one may take the inaccessibility condition to be a problematic characteristic of the objective view. The most superficial response to the inaccessibility condition is that it strikes us as somehow unfair or a symptom of too stringent a moral standard. After all, if it seems unfair to hold an agent to a standard by which he is morally obligated to perform act X despite the fact that he had no access to the fact that X was his best option. Consider for example the following case:

Dr. Bryan believes that drug A will cure his patient and that drug B will kill his patient. Further, Bryan asks several of his colleagues which drug they believe will cure his patient, and all of them agree with Bryan that drug A will cure his patient. Bryan does extensive research which confirms his and his colleagues' beliefs. As such, Bryan gives his patient drug A and much to his surprise and his patient's misfortune, drug A kills his patient. It is later revealed however that the drug manufacturer had mixed up drug A and drug B, and thus Bryan had no way of knowing that by giving his patient the drug labeled drug A he was in fact giving his patient drug B, which would kill the patient.

Advocates of the subjective view would be inclined to say that this example, which exhibits an inaccessibility condition for the objective view, demonstrates an instance in which it would be unfair to hold Dr. Bryan to a standard by which he was

wrong for not giving the patient drug labeled B (i.e. drug A). After all, Bryan had no way of knowing that he should have given drug B and exhausted all resources to ensure that he was engaging in the act which was his best option. Further Bryan believes that he had a moral obligation not to do an act other than giving the patient the drug he did (after all, if he believes that giving the patient drug A is his best option, then he would believe that doing any other action would not be his best option, and as such he would believe that he had a moral obligation not to engage in any other action than he did). In light of this fact, it would seem unfair to say that Bryan was morally wrong to engage in the act that he did. After all, to say that he had a moral obligation to give the patient drug B would be tantamount to saying that he has a moral obligation to do an action which he believes that he has a moral obligation not to do. As such, we must reject the possibility of an agent S having a moral obligation to do X even if S has no access to the fact that X is his best option.

The objective view holder, at this point, however, can rely upon the distinction between an act being morally wrong and the agent who performed that act being blameworthy for that morally wrong act. In other words, they would say that it would indeed be unfair to blame Dr. Bryan for giving the patient the drug labeled A because he had no way of knowing that his best option was some other act. Nevertheless, this does not change the fact that the morally right act would have been to give the patient the drug labeled B (i.e. the one which cures the patient), and as such Bryan did, in fact, engage in a morally wrong act. After all, it would seem absurd to say that had, for some reason despite his beliefs and what his evidence confirmed, Bryan given the patient the drug

which would have cured him that this would have been the morally wrong act. Thus, the distinction between an act's wrongness and an agent's blameworthiness for that act defends the objective view against the claim that its inaccessibility condition leads to an unfair moral standard.

But the critic of the objective view may respond that this defense does not work because the issue arising from the inaccessibility condition is not simply a matter of assigning the appropriate blame. The issue, rather, is that S's having a moral obligation to do an act X, such that S believes that it is morally wrong to do X is an incoherent notion of S's moral obligation. In other words, they will argue, that if X is S's moral obligation, then it follows that it is necessarily not the case that S believes that he has a moral obligation not to do X. The problem with the objective view then is that it seems that it will result in some cases in which S could never satisfy his moral obligation to perform the act which is his best option without violating another moral obligation—namely, an agent's obligation to refrain from intending to perform a morally wrong act. That is, if an agent has a moral obligation to refrain from intending to perform a morally wrong act, then he cannot have another moral obligation to perform an act X, where his performing X would entail his intending to perform a morally wrong act. But if S believes X to be a morally wrong act and X is his best option, it seems that S would have to intend to perform a morally wrong act to satisfy his moral obligation to do X, under the objective view. Thus, some may argue that the inaccessibility condition is problematic because it leads to the objective view rendering S with an incoherent moral obligation, under which

S could not satisfy his moral obligation to perform his best option without violating his moral obligation to refrain from intending to perform a morally wrong act.

Still, this line of criticism will not work. It does not seem right that an agent has a moral obligation to not do what he believes it is morally wrong to do. Though this claim might seem intuitive at first glance, the apparent intuitiveness of the claim that an agent has such a moral obligation may stem from confusing the proposed principle with the *wrong intentions principle*, according to which “if it is [morally] wrong to do an action under certain conditions, it is wrong to intend to do it should those conditions arise.”⁴² The proposed principle seems convincing when we consider the true claim that “S has a moral obligation to not intend to perform a morally wrong act.” It seems that the objective view leads us to reject this claim. After all, if S believes X to be a morally wrong act and X is his best option, it seems that S would have to intend to perform a morally wrong act to satisfy his moral obligation, under the objective view. But the confusion clears up when we consider the following: Whereas the wrong intentions principle claims that it is wrong for S to intend to perform an act X where X is a morally wrong act (i.e. it is wrong for S to intend to perform [a morally wrong act]), the proposed principle claims that it is wrong for S to Q, where Q is intend to perform a morally wrong act (i.e. it is wrong for S to [intend to perform a morally wrong act]). The intuitiveness of the proposed principle vanishes when we compare it to the wrong intentions principle, which is where our intuition actually rests. Thus, it is not true that S’s having a moral obligation to do an X, such that S believes that he has a moral obligation to refrain from

⁴² Davion p. 1.

doing X is an incoherent notion of S's moral obligation. This cannot be the problem stemming from the inaccessibility condition.

Nevertheless, there is a significant problem that arises from the inaccessibility condition. Under the objective view, it is possible for S to have a moral obligation to perform act X *because* X is his best option and for it to be inaccessible to S that X is S's best option. It follows that it is possible that S has a moral obligation to perform an act X, such that S could never perform X for the reason that makes it his moral obligation—that X is S's best option, in the case of the objective view. It is important to distinguish S's not being able to perform X for the reason that makes it his obligation from S's not having a reason to do X. After all, there is a sense in which S will always have a reason to do what is his best option. Even so, there are cases in which S will not be able to perform act X for this reason. Thus, there will be cases, under the objective view, where S will have a moral obligation to perform an act X, such that S will not be able to perform X for the reason that makes X his moral obligation. This is what I shall call the *unreasonableness worry*. According to what I will call the *Principle of Reasonableness*, a correct moral theory will not be susceptible to the unreasonableness worry. Thus, the objective view is at odds with the Principle of Reasonableness.

Upon first glance, raising this issue against the objective view appears to impose the contentious view that an agent should always be able to perform his moral obligation because it is his moral obligation.⁴³ Subjective view advocates might be tempted to argue that S does not have a moral obligation to do the act which is his best option because if

⁴³ In Ross 1939, p. 150, W.D. Ross seems to take exception to the objective view on these grounds: "Further consequences of the objective view are (2) that we can never do a duty because it is our duty..."

S's moral obligation is to perform the act which is his best act, then S will be unable to satisfy his moral obligation because it is his moral obligation, since he will not know what his best act is. Therefore, S will be unable to satisfy his moral obligation for the right reason (i.e. that it is his moral obligation). Of course, the objective view holder will have a strong defense to this objection by once again appeal to the distinction between an act's being morally right and an agent being praiseworthy for his doing that act. That is, the objective view holder may respond that it is true that if S is to be praised for doing his moral obligation X, then S must be able to do X for the right reason—that it is his moral obligation –or at least, that S does not perform X for the wrong reasons. Still it is perfectly acceptable that S can have a moral obligation to do X despite his not being able to perform X because it is his moral obligation.

But this is not the proper analysis of the problem for the objective view. Rather, the problem is that the unreasonableness worry leads to the objective view's inability to be appropriately risk averse and its inability to appreciate the moral significance of ignorance by insuring that S's moral obligation will not be to do something unnecessarily risky. That is, if it is possible that S has a moral obligation to do an act X, such that S could never do X for the reason that makes it his moral obligation (i.e. that it is his best option), then there are circumstances in which S has a moral obligation to do an act X despite the fact that he could never have a reason to do that act.⁴⁴ Consider the following

⁴⁴ Keep in mind the difference between S's never having a reason to perform an act and there being no reason for S to perform that act. There will always be a reason for S to perform X because it is his best option. But it may still be the case that S will not have access to the fact that X is his best option, and thus it is possible that S would never possess a reason to do X. In such a case, the objective view would entail S having a moral obligation to perform an act for which he possesses no reason to perform. In other

example, brilliantly used by Michael Zimmerman, which illustrates the resultant failures of the objective view in such cases, which I assert are the result of the unreasonableness worry.⁴⁵

Dr. Jill has a patient, John, who is suffering from a serious though non-life threatening disease. She must choose between three drugs (A, B, and C) to treat her patient with. All the evidence at Jill's disposal indicates (in keeping with the facts) that giving John drug B would cure him partially and giving him no drug would render him permanently incurable. The evidence, however, leaves completely open whether it is drug A or drug C that would cure him completely and whether it is giving him drug A or giving him drug C that would kill him, though the evidence does indicate that the one will cure and the other will kill the patient (the fact of the matter is that drug A will fully cure and drug C will kill John).

Our intuition in this case is that Jill should give her patient drug B. After all, the risk of killing the patient is too great given her ignorance and resultant lack of reasons to give her patient drug A. While there obviously is a reason for Jill to give drug A (as it is her best option and would cure her patient) she could never perform the act of giving her patient drug A for this reason and thus it seems that Jill does not possess any reason to

words, S would have a moral obligation to perform an act that it would be unreasonable for him to perform.

⁴⁵ Zimmerman 2008 pp. 17-18.

give drug A.⁴⁶ After all, Jill has neither evidence nor belief indicating which of A or C will cure or kill her patient. But she is certain that drug B will have a positive effect. Yet notice also, that she knows that drug B is not her best option. Thus, the objective view holder cannot condone her choosing drug B, the intuitively appealing choice, because Jill would knowingly be doing what is wrong in this case. As such, the distinction between giving drug A being her moral obligation but Jill not being blameworthy for not giving drug A will not work here.⁴⁷ The objective view would render Jill's obligation in this case to perform an act which it would be unreasonable to do and this view has no way of condoning the intuitively right act, the cautious act insuring against unnecessary risk due to an agent's ignorance.

The problem for the objective view is now clear. The fact that moral decisions are made in a state of ignorance cannot be ignored or else we will be left with a theory of moral obligation which prescribes dangerous and reckless courses of action. Recall, the Alejandro case from the introduction. The objective view would claim that Alejandro had no moral obligation not to riskily speed through the intersection despite his lack of knowledge that he would not be putting anyone in harm by doing so. The objective view's inability to cohere with our intuitions in cases where S's best option is an

⁴⁶ It appears that the only reason in this case for Jill to give her patient drug A is that giving drug A is her best option. But this reason is inaccessible to her and thus it seems that she does not possess any reason to give drug A. Some may argue that she does possess a reason to give her patient drug A over drug B, namely that she is trying to cure her patient fully rather than partially. Notice, however, that this is not a reason for A anymore than it is a reason for her to give her patient drug C. Thus, it is not properly considered a reason for giving drug A. Rather we can say that she has a reason not to give B. Of course, she will have a stronger reason to give her patient drug B.

⁴⁷ That the objective view holder cannot rely upon the blameworthy/morally wrong distinction in this case is made in Zimmerman 2008 p. 18.

unreasonable reckless act, given S's ignorance, results from the objective view's possessing the inaccessibility condition. The resultant unreasonableness worry causes problems for the objective view in cases such as these, and as such we should abandon the objective view.

Section Three: *Motivating the Epistemic View*

As we have now seen, there exist serious criticisms undermining each the subjective view and the objective view. The subjective view violates the Principle of Non-Manipulation, according to which the proper view of moral obligation will not permit an agent to, without violating his moral obligation, author (either in part or in whole) his own moral obligation by exerting voluntary control to determine his moral obligation rather than having to live up to a moral obligation independent of his control. This violation of the Principle of Non-Manipulation leads to the subjective view's resulting in agent moral infallibility and undermining the need for genuine moral deliberation. The objective view, on the other hand, violates the Principle of Reasonableness, according to which the proper view of moral obligation will render it

impossible for an agent to have a moral obligation to perform an act X if S is not able to perform X for the reason that makes X his moral obligation (i.e. that this reason is inaccessible to him). This violation of the Principle of Reasonableness leads the objective view to yield moral obligations that are unjustifiably risky given what the agent possesses reason to do. Thus, we have good reason to dismiss both the subjective view and the objective view.

Further, we have gained insight into characteristics which the correct view of the circumstantial element of moral obligation will avoid. The ideal view of the circumstantial element of moral obligation will violate neither the Principle of Reasonableness nor the Principle of Non-Manipulation. Further, because we have seen that it is the objective view's inaccessibility condition which leads to its being susceptible to the unreasonableness worry, we can rest assured that the correct view on the circumstantial element of moral obligation will not contain the inaccessibility condition. I now aim to motivate the epistemic view by appealing to the claim that the epistemic view captures these desirable qualities while also avoiding the undesirable characteristics to which the subjective view and objective view succumb (i.e. the manipulation worry and the unreasonableness worry).

We must explore what characteristics of the subjective view and objective view lead them to encounter the problems that they do. After all, if the epistemic view truly does avoid these problems, then it must not possess the same problematic qualities as the other two views. In order to clearly depict each of these three views' fundamental characteristics and their respective implications, consider the following chart:

	Actuality Affects Moral Obligation	Belief Affects Moral Obligation	Evidence Affects Moral Obligation
Objective View	YES	NO	NO
Subjective View	NO	YES	NO
Epistemic View	NO	NO	YES

Since the epistemic view is taken to be a cure to the problems faced by the objective view and the subjective view, epistemic view holders will need to demonstrate not only that the subjective and objective views possess serious problems but also that the problems faced by these two views arise from either the objective view's condition that actuality affects moral obligation, the subjective view's condition that belief affects moral obligation, or their shared condition that evidence does not affect moral obligation. For if their problems stemmed, respectively, from the objective view's not permitting belief to affect moral obligation or the subjective view's not permitting actuality to affect moral obligation, it becomes abundantly clear that the epistemic view would also fall vulnerable to these criticisms, on account of the fact that it too shares these qualities. Thus, in considering the epistemic view's ability to avoid the criticisms faced by the other two views, we must be mindful that the epistemic view will only be capable of avoiding those problems that do not result from permitting neither actuality nor belief to affect moral obligation.

Keeping this in mind it appears that we do in fact have good reason to endorse the epistemic view and believe that it will succeed at avoiding the manipulation worry and the unreasonable worry. After all, the subjective view's manipulation worry appears to stem from the fact that the subjective view ties an agent's moral obligation to something which the agent has voluntary control over, namely his beliefs. Thus, it appears that this worry stems from a theory of moral obligation being affected by beliefs—a characteristic which the epistemic view does not share with the subjective view. Further, because the epistemic view renders moral obligation a function of evidence, something which is beyond the agent's voluntary control, it appears that the epistemic view will be immune from the manipulation worry. And, the epistemic view seems promising in regards to avoiding the unreasonableness worry. As saw in our analysis of the objective view, the objective view faced the unreasonableness worry as a result of the objective view's failure to take into account the reasons that an agent possessed for performing certain acts. The epistemic view appears to dodge this worry by making epistemic considerations a determinant of an agent's moral obligation. Further, the inaccessibility condition that led the objective view to face the unreasonableness worry, does not appear to plague the epistemic view because an agent's evidence is accessible to him in a way in which the objective element of an agent's situation simply is not.

Thus, the epistemic view appears to dodge both of these substantial issues that are faced by the subjective and the objective view, and thus provides hope of a satisfactory middle ground in the question of what is the circumstantial element of moral obligation. Further, by attempting to bridge the gap between the subjective elements of the agent's

consciousness and the objective elements of an agent's situation, it appears to capture the desirable qualities of each view without succumbing to their respective weaknesses. While avoiding the manipulation worry and the unreasonableness worry are not sufficient conditions for a view of the circumstantial element of moral obligation being the correct view, we have established avoiding these two worries to be a necessary condition for the correct view. Thus, it seems that we have good reason to endorse the epistemic view in light of the fact that it satisfies these necessary conditions and its competitors, the objective view and the subjective view cannot.

CHAPTER 2

Section One: A Dilemma for the Epistemic View

At first glance, the epistemic view appears to avoid both the manipulation worry, which plagues the subjective view, and the unreasonableness worry, which plagues the objective view; but, when we force ourselves to answer a simple question, we see that the epistemic view faces a crucial dilemma. In this section, I will argue that if we respond to this dilemma in one way, we will see that the epistemic view faces problems paralleling those faced by the objective view, insofar as the epistemic view will also be susceptible to the unreasonableness worry. If, however, we choose to avoid this worry and respond to this dilemma in the other way, we will see that the epistemic view faces problems

paralleling those faced by the subjective view, insofar as the epistemic view will also be susceptible to the manipulation worry.

Recall that according to the epistemic view an agent's moral obligation is defined in the following way: an agent, S, is morally obligated to perform act X if and only if S's evidence indicates that act X is his best option. The imperative question that we must ask is, "what does it mean for a piece of evidence to be properly considered 'S's evidence'?"⁴⁸ There appear to be two possible answers to this question in the context of the epistemic view.

The first way of answering the question is to say that an agent's evidence consists of the evidence "of which the agent is aware." The evidence of which the agent is aware, is simply that evidence which S is already in possession of without any further investigation. Contrast this with "an agent's available evidence." An agent's available evidence consists of the evidence such that either S is aware of that evidence or S could become aware of that evidence through further investigation.⁴⁹ As is seen from the definitions above, if an agent is aware of a piece evidence then that piece of evidence is also a piece of S's available evidence. However, there are pieces of S's available evidence of which S is not aware. This distinction will be imperative to the discussion that follows. Recall the earlier distinction between the objective element of an agent's situation and the subjective element of an agent's situation. The objective element of an agent's situation consists of an agent's actual circumstances. On the other hand, the

⁴⁸ For an insightful discussion of this question refer to Feldman 1988

⁴⁹ Let pieces of evidence such as past experiences or facts that an agent could recall if provoked into doing so but is currently unaware of as being part of that agent's available evidence. Once an agent brings such things to his own attention thought, they properly become pieces of the evidence of which that agent is aware.

subjective element of an agent's situation consists of what that agent believes his circumstances to be. This distinction elucidates an important distinction between S's available evidence and the evidence of which S is aware: the evidence of which S is aware will be a part of the subjective element of S's situation, while S's available evidence will only be a part of the subjective element of S's situation if S is aware of that evidence. Those pieces of evidence which are part of S's available evidence but of which S is not aware, on the other hand, will only be a part of S's objective situation.⁵⁰

A concrete example will help illustrate the concept of evidence of which an agent is aware and that of an agent's available evidence. Jane has a science exam tomorrow covering chapter three in her textbook. Due to her procrastination, however, Jane is only able to read the first half of chapter three and neglects to read the other half. Consider the entirety of chapter three to be Jane's available evidence for the correct answers to the question on this exam. Conversely, only the first half of this chapter (the portion of the chapter which she has read) counts as evidence of which Jane is aware for the correct answers to the questions on her exam. The second half of this chapter (the portion of the chapter which she has not read) does not count as evidence of which Jane is aware though it does count as her available evidence for the correct answers to the questions on her exam.⁵¹

⁵⁰ All pieces of evidence which are available to S will be a part of the objective element of his situation. Those pieces of S's available evidence such that S is aware of them will be both part of the objective element of his situation and the subjective element of his situation. Those pieces of S's available evidence such that S is not aware of them will only be a part of the objective element of S's situation. Finally, all pieces of evidence of which S is aware will be a part of the subjective element of his situation.

⁵¹ If Jane is aware of the fact that there is a second half of this chapter, then the evidence of which Jane is aware does indicate that there is a second half to the

The dilemma for the epistemic view is now clear. Moral obligation under the epistemic view (i.e. S has a moral obligation to perform act X if and only if S's evidence indicates that X is S's best option) can either be a function of S's available evidence or a function of the evidence of which S is aware. In some instances these two types of evidence will not indicate the same act as being S's best option. Consider an analogous case to the Zimmerman case that was examined earlier in Chapter 1, section 2 (refer to pages 37-38). Zimmerman's case compares the objective view against the epistemic view by separating the facts from what the agent's evidence indicates. This case will analogously compare these two competing formulations of the epistemic view by separating an agent S's *available evidence* (E) from the *evidence of which S is aware* (e) in the following manner:

<u>According to E:</u>	<u>According to e:</u>
Drug A cures fully	Drug A either cures fully or kills
Drug B cures partially	Drug B cures partially
Drug C kills	Drug C either cures fully or kills

In this case, if S's moral obligation is a function of E, then S ought to give the patient drug A. If, on the other hand, S's moral obligation is a function of e, then S ought to give the patient drug B. This case demonstrates that the two, E and e, are in fact

chapter. Nevertheless, we would not say that this evidence (the second half of the chapter) is evidence of which she is aware. Rather, we would say that according to the evidence of which she is aware, there exists evidence available to her of which she is not aware. In other words, according to the evidence of which she is aware, the evidence of which she is aware is not equivalent to her available evidence.

distinct from one another and will yield different obligation in some cases. The problem for the epistemic view is that if we interpret the epistemic view as rendering an agent's moral obligation a function of his available evidence, then the epistemic view will be plagued by a parallel unreasonableness worry, similar to the objective view. If we interpret the epistemic view as rendering an agent's moral obligation a function of the evidence of which she is aware, then the epistemic view will be plagued by a parallel manipulation worry, similar to the subjective view.

Section Two: *The Unreasonableness Worry for the Epistemic View*

Let us now explore the problematic consequences of defining the epistemic view in terms of an agent's available evidence (E). If we claim that it is S's available evidence that determines S's moral obligation, we realize that S's moral obligation is a function of something which S need not be aware of. Recall that one of the distinctions between an agent's available evidence and the evidence of which that agent is aware is that an agent can be unaware of his available evidence. That is to say, that there are pieces of evidence at an agent's disposal which an agent has not yet brought their attention to (e.g. the second half of chapter 3 in the Jane science exam case). Even if the agent is aware that such evidence exists, that agent is not aware of those pieces of evidence themselves,

insofar as, the agent either is unaware of whether those pieces of evidence are or what they indicate. While an agent always has access to the subjective element of his circumstances, there are objective elements of his circumstances which he does not have access to. Similarly, while an agent always has access to the evidence of which he is aware, the agent does not always have access to pieces of his available evidence. As such, because an agent does not always have access to pieces of his available evidence, if we were to make an agent's moral obligation a function of his available evidence, that agent's moral obligation will sometimes be inaccessible to him (i.e. in a given situation it would be inaccessible to him that his available evidence indicates that X is his best option). Thus, we now see that the epistemic view formulated in terms of an agent's available evidence will entail the inaccessibility condition, similarly to the objective view (previously discussed in Chapter 1, section 2).

As we saw earlier in our analysis of the objective view, the inaccessibility condition will lead to problematic consequence for a view of moral obligation. Under views which satisfy the inaccessibility condition, it is possible for the reason that S has a moral obligation to perform act X (e.g. that S's available evidence indicates that X is S's best option, under the current formulation of the epistemic view) is inaccessible to S. As a result of such a possibility, these views result in cases in which an agent S would be morally obligated to perform an act which it would be unreasonable for him to do given the subjective element of his situation, because he could not possess a reason why he should perform that act. Though there of course would exist a reason for S to perform the

act which his available evidence indicates would be his best option, S will be unaware of any such reason in these particular cases.

The inaccessibility condition (that S can be morally obligated to do X even though the reason why he is morally obligated to do X is inaccessible to him) is a particularly troubling consequence for the epistemic view, even more so than it is for the objective view. As a result of the inaccessibility condition, S need not be aware of his available evidence in order to satisfy the epistemic view's moral obligation (if formulated in terms of S's available evidence). But to say that an agent has satisfied her moral obligation, under this formulation of the epistemic view, even if she does not act *because* of her available evidence (or even if she is unaware of her available evidence entirely!) seems to cheapen the significance of evidence and to undermine the motivation for the epistemic view in the first place. A fundamental characteristic of evidence is that it informs agents. Evidence is important to philosophers and every day people because of this function of informing people. That is to say that evidence's value is in its function. But this formulation of the epistemic view undermines the need for this function by rendering cases in which agents act because of their evidence and cases in which agents simply act in accordance with what their evidence indicates morally equivalent. To highlight this version of the epistemic view's disservice to the significance of evidence's function, consider the following example:

A criminal defendant is sentenced to life in prison despite the fact that a criminal investigation could turn up no evidence pointing to his guilt. Nevertheless, thirty

years after the sentencing, it is discovered that the initial criminal investigation was deeply flawed and evidence had been available heavily indicating that the defendant did commit the crime for which he was sentenced.

While it is true that the jury's verdict is in accordance with what the available evidence indicated, the jury did not in any way act because of the evidence. Thus, it is pure coincidence that their decision was in accord with what the available evidence indicated. Further, by undermining the function of evidence in this particular case, the standard of making a decision in accordance with one's evidence is completely undermined as a valuable standard. The desirable standard seems, rather, to be that of making a decision *informed* by one's evidence. After all, what is so good about a jury making a decision that happened to be in accordance with undiscovered evidence? At that point, we might as well say that the standard should be the jury making a decision in accordance with the truth. After all, ought we say that a jury should make a decision in accordance with evidence of which the jury is completely unaware and which is defective or should we that a jury should make a decision in accordance with truth? By being completely indifferent to whether or not an agent is aware of the evidence which renders an act X his moral obligation, this formulation of the epistemic view cheapens the significance of evidence, which is the functional role of evidence. Further, this indifference undermines the appeal of the epistemic view over the objective view, which was to provide an agent with an accessible moral obligation and to build reasons into the actions which agents ought to do.

Consider the following example demonstrating the effects of using evidence of which S need not be aware of as the basis of S's moral obligation. This case demonstrates the distinction between an agent's acting because of evidence as opposed to her simply acting in accordance with what the evidence indicates, and the troubling consequences of using the latter as the determinant of what an agent morally ought to do:

Dr. Jane is trying to cure a patient. The fact of the matter is that drug A will cure the patient and drug B will kill the patient. Further, Jane's available evidence indicates that drug A will cure and that drug B will kill the patient. Nevertheless, Dr. Jane has a very superstitious nature and this nature overrides her desire to act based on her evidence. She decides that she will do whatever the magic-8 ball in her office tells her to do, rather than make an informed decision based upon what her evidence indicates would be her best option.⁵² Fortunately for the patient, the magic 8 ball affirms that Jane should choose drug A in accordance with her available evidence.

Intuitively it seems odd to say that Dr. Jane has satisfied her moral obligation in this instance. Her act appears to be incredibly reckless and unreasonable, and yet because of the entailed inaccessibility condition, the epistemic view renders Dr. Jane's action in this case to be morally right.

⁵² Let us stipulate that Dr. Jane's magic-8 ball is in fact not magic. That is, the magic-8 ball has neither a demonstrated ability nor potential ability to predict one's best option. Let us assume that Dr. Jane has never used the magic-8 ball before, so her decision to use this method is neither encouraged nor deterred by past results.

Still, the defender of this formulation of the epistemic view has a familiar response to this criticism of the epistemic view. While it is true that Dr. Jane's acting in accordance with her magic 8-ball is not praiseworthy, she nevertheless was morally obligated to give her patient the drug which her available evidence indicated would cure her patient and she still satisfied this moral obligation by giving her patient drug A, regardless of her reasons for doing so. In other words, whether an agent acts *because* of what her evidence indicates will determine whether the agent is blameworthy or praiseworthy for the acts that she engages in. But what that agent has an obligation to do in a given situation is purely a function of what that agent's available evidence indicates. Thus we can call Dr. Jane's actions reckless and careless and we can say that she only satisfied her moral obligation due to incredible luck, but we can maintain (as this formulation of the epistemic view insists) that her moral obligation in this case was to give the drug which her available evidence indicated was her best option (i.e. drug A), an obligation which she has satisfied.

But as we saw in the our criticisms of the objective view, this response of distinguishing what an agent has a moral obligation to do and what an agent is properly praised for doing, will not get to the heart of the problem for the epistemic view either. The problem resulting from a view of the circumstantial element of moral obligation's being plagued by the unreasonable worry is that it leads that view to yield the wrong answer in cases in which an agent ought to take a more cautious course of action than will be prescribed by the view plagued by the unreasonableness worry. This is the motivation behind accepting the Principle of Reasonableness. Consider the following

example illustrating that the epistemic view's violation of the Principle of Reasonableness will result in an agent's moral obligation being overly risky and unreasonable under this view:

Dr. Anderson arrives at the hospital at which she works. An orderly informs her that she has left in her office the report regarding which of her patients have requested "Do Not Resuscitate" (DNR) status. As she is walking towards her office to read the report, she is suddenly informed that two patients are experiencing cardiac arrest. Dr. Anderson, knowing that if she goes to her office to read the report she will only have time to save one of the patients, rushes to these two patients, Allan and Brianne, and resuscitates both. Following the incident, Dr. Anderson returns to her office and reads the report, according to which Allan has requested DNR status.

Assuming that given the proper moral theory, it is morally good for Dr. Anderson to respect the wishes of each patient's wishes regarding their DNR requests, her available evidence would indicate that her best option is to save only Brianne. Thus, according to the epistemic view, Dr. Anderson ought to not have saved Allan and thus has done something morally wrong in the case at hand. However, it would have been unreasonable for Dr. Anderson to only have saved Brianne given what she knew at the time of her action and given the evidence of which she was aware of at the time. This case demonstrates that due to the inaccessibility condition, the epistemic view, formulated in terms of S's available evidence, violates the Principle of Reasonableness and as such, it

leads to agent's having a moral obligation that is too risky in certain cases. Thus, it becomes clear that S's moral obligation cannot be to do what S's available evidence indicates is best.

Section Three: *The Manipulation Worry for the Epistemic View*

Let us now consider the other option for the epistemic view in this dilemma—that S's moral obligation is a function of the evidence of which S is aware. I suggested earlier that interpreting the epistemic view in this way leads the epistemic view to be susceptible to the manipulation worry, similarly to the subjective view. I will now attempt to prove this point.

Recall that a theory of moral obligation is vulnerable to the manipulation worry if and only if that theory allows an agent to, without violating his moral obligation, author (either in part or in whole) his moral obligation by exerting voluntary control to intentionally determine his moral obligation rather than having to live up to a moral obligation independent of his control. Since this formulation of the epistemic view

renders an agent's moral obligation a function of the evidence of which that agent is aware, and the evidence of which an agent is aware is, at least in part, capable of being voluntarily determined by the agent in question, it will be possible for that agent to intentionally exert control over his actions so as to determine his moral obligation. In order to see if the epistemic view is susceptible to a manipulation as a result of this, we must ask the following question: Are there cases, under the epistemic view, in which an agent can exert this intentional and voluntary control to determine his moral obligation without violating his moral obligation by exerting this control? An answer yes to this question will result in the epistemic view facing a manipulation worry. An answer no to this question will demonstrate that the epistemic view has a built-in-defense against the manipulation worry and will therefore not violate the Principle of Non-Manipulation.

Let us use the following example to explore the issue as to whether or not the epistemic view, like the objective view, has a built-in-defense against the manipulation worry, which also plagues the subjective view:

Bob's friend, Zach, is in dire need of a kidney. It turns out that Bob is the only available match, and, as such, if Bob does not give Zach his kidney, then Zach will die. In addition, the evidence of which Bob is aware of indicates as much. Further, let us stipulate that according the proper moral theory and what the evidence of which Bob is aware indicates his best option is for him to give Zach his kidney.⁵³ But on the eve of the operation, Bob becomes frightened and

⁵³ Let us take Peter Singer's moral theory to be the proper moral theory in this case: "if it is in our power to prevent something bad from happening, without thereby sacrificing

contemplates not giving Zach his kidney. Then Bob sees an ad for a drug on the television. The drug allows the drug taker to alter the evidence of which he is aware in the following way: the agent thinks of an experience, the agent then slips into a state of unconsciousness, and upon waking the agent has memories of the experience such that the agent cannot distinguish these memories from memories of a genuine experience that he has had. Bob orders the pill and takes the pill, ensuring that he will have the memory of getting a phone call telling him that Zach has passed away, attending his funeral, etc when he awakes from his unconsciousness (a few hours prior to the surgery).⁵⁴ Thus, the evidence of which Bob is aware, when he awakes from his unconsciousness, indicates that it would no longer be Bob's best option to give Zach his kidney, as this evidence indicates that Zach has passed away.

This case appears to be very problematic for this formulation of the epistemic view by demonstrating a case in which an agent can manipulate his moral obligation. After all, Bob voluntarily engages in actions that will alter the evidence of which he is aware of at a future point in time in order to intentionally determine that a particular act

anything of comparable moral importance, we ought, morally, to do it." This theory is discussed in Singer 1972.

⁵⁴ This is a parallel to a Cartesian dream case in which an agent cannot distinguish a convincingly real dream from an actual experience. I am making the assumption that the agent's memory of the events of this dream, since they are indistinguishable to the agent from a memory of a real experience, count as an agent's evidence in the same way that a memory of an actual experience would count as evidence. If, however, one takes exception to this internal approach to evidence the case can be adapted to meet an external approach to evidence. To do so, assume that prior to taking the pill the agent can hire a team of actors to ensure that when he awakes from the pill they will create convincingly real situations that lead him to falsely believe that Zach has died.

X (i.e. that of giving his kidney to Zach) will not be his moral obligation at that future point in time.

Still, the epistemic view proponent has a plausible defense available to him. They will argue that while all of the above is true, Bob cannot alter his moral obligation at T3, by engaging in the act of taking this drug at T1 without violating what the epistemic view determines is his moral obligation at T1. Thus, the epistemic view has a built-in-defense against a manipulation worry in this case, similar to the objective view. To see this line of defense more easily, consider the following three conditions which are necessary and sufficient for a manipulation worry for the epistemic view formulated in terms of the evidence of which an agent is aware:

Condition 1 (the “intentionality condition”): S can voluntarily engage in an act Q at T1 with the knowledge (or at least informed belief) that Q will ensure that at T2 the evidence of which S is aware will not indicate that X is his best option at T2. This condition ensures that S can intend, at T1, to determine his moral obligation, at T2.

Condition 2 (“the determinacy condition”): An act, Q, occurring at T1 will ensure that, at a T2, the evidence of which S is aware will not indicate that X is his best option at T2. This condition ensures that S’s act Q, at T1, will determine his moral obligation, at T2

Condition3 (“morally permissible condition”): The evidence of which S is aware at T1 does not indicate that Q is not S’s best option at T1. This condition ensures that S’s performing Q at T1 will not cause him to violate his moral obligation under the epistemic view. It is the violation of this condition that results in a theory’s having a built-in-defense against the manipulation worry.

At this point, the epistemic view holder will argue that condition 3 (“the morally permissible condition”) is not satisfied in the case above. That is to say, while Bob can take the pill with the intention to determine his moral obligation at T2, and his taking the pill will ensure that he will not have a moral obligation to give his kidney to Zach at T2, the epistemic view is immune to the manipulation worry because Bob’s taking the pill will result in the violation of condition 3. Thus, the epistemic view has a built-in-defense against the manipulation worry. Further, the epistemic view proponent can give a very principled reason for why condition 3 will not be satisfied (an exploration of this reasoning will be very insightful into finding the underlying condition which causes the manipulation worry. Just as we know that the unreasonableness worry is caused by the inaccessibility condition, we are on our way to discovering the analogous condition which causes in the manipulation worry).

To see why the epistemic view might have a built-in-defense let us examine Bob’s actions more closely in the case above. Condition 1 demands that Bob can take the pill with the informed belief that his taking the pill will have the effect of altering his

evidence at a future point in time, T2, such that he will have defective evidence at T2. In other words, in order to satisfy condition 1, Bob, at T1, must be aware that he is insuring that he will have defective evidence at T2. Bob does seem to satisfy condition 1 because he has seen the ad and is aware of the consequences of taking this drug. Further, Bob intends to have the evidence, when he wakes up, indicating that Zach has died. Thus, at T1 Bob intentionally chooses to have evidence which is defective, and which will prevent him from having a moral obligation to perform a specific act X at a future point in time. But what bearing does this intentionality have on his ability to satisfy condition 3? Condition 3 demands that Bob's taking the pill at T1 does not violate his moral obligation by performing an act that is not his best option at T1. Keeping in mind that under the epistemic view Bob has a moral obligation at T1 to do what his evidence indicates is his best option at T1, it appears that Bob is unable to satisfy condition 3 in light of his satisfying condition 1. After all, satisfying condition 3 would mean that at T1, according to the evidence of which he is aware at T1 Bob's intentionally choosing to have defective evidence would not be such that it is not his best option. But it surely is the case that Bob's evidence would indicate that it is not his best option to bring it about that he has defective evidence. Thus, satisfaction of condition 1 will preclude the satisfaction of condition 3 under the epistemic view formulated in terms of the evidence of which S is aware. As such, it appears that the epistemic view has a built-in-defense against the manipulation worry, similar to the objective view.

Though this analysis appears to be right upon first glance, I have now come to believe that it is not true that the satisfaction of condition 1 precludes the satisfaction of

condition 3 under the epistemic view. For this point to become clearer, consider the following question: Why would Bob's ensuring at T1 that he will have defective evidence at T2 violate his moral obligation to do what the evidence of which he is aware at T1 indicates is his best option at T1? It is important to keep in mind what Bob's best option at T1 is. Bob's best option at T1 is that act which at T1 maximizes the moral good as determined by the proper moral theory and the agent's actual circumstances. In other words, the question is: Why does Bob's ensuring at T1 that he will have defective evidence at T2 violate his moral obligation at T1 to do the act X such that the evidence of which Bob is aware indicates that X will maximize the moral good at T1? There are several possible answers which I can think of, but upon careful analysis none of these answers are satisfactory.⁵⁵

The first answer would be to claim that there is a moral fact according to which an agent should ensure that he does not have defective evidence (i.e. that it is a moral good to have non-defective evidence)? If this were true then at T1 it follows that it is S's best option to engage in acts which ensure that he will not have defective evidence. Since the ad informing Bob that taking the pill will result in his having defective evidence, Bob has evidence of which he is aware, at the time of his taking the pill (T1), that taking the pill is not his best option (because this evidence indicates that his taking this pill will give him

⁵⁵ Of course, if someone is capable of satisfactorily answering this question, then they will demonstrate that the satisfaction of condition 1 does preclude the satisfaction of condition 3, under the epistemic view. In turn, they will demonstrate that the epistemic view is immune from the manipulation worry. If no satisfactory answer to this question exists then the epistemic view is plagued by the manipulation worry because it does not possess a built-in-defense. After careful consideration, I cannot think of a satisfactory answer and as such see no reason to believe that any such built in defense exists for the epistemic view.

defective evidence and it is his best option to ensure that he does not have defective evidence.).

While this argument is sound, it rests on the entirely unconvincing claim that it is morally good to have non-defective evidence and that, as a result, an agent's best option in a given situation is to ensure that he does not have defective evidence. After all, the claim that S should not have defective evidence is an epistemic norm. To claim that S has a moral obligation to ensure that he does not to have defective evidence is to build an epistemic norm into his moral obligation. But there seem to be instances in which it is completely amoral for an agent to bring it about that he has defective evidence. Consider for example if I could make it such that I had defective evidence indicating that George Washington was born in Maine rather than Virginia. I see no reason to say that I have violated any moral obligation in this case (such as a moral obligation not to make it such that I have defective evidence) by engaging in the act that brings about my having such defective evidence. I certainly have violated an epistemic norm in my action, but it seems clear that my action is neither moral nor immoral. Thus, for the epistemic view to rely upon a moral principle according which it is morally bad to intentionally bring it about that we have defective evidence is misguided. We simply have an epistemic obligation to not intentionally bringing it about that we have defective evidence. Neither having defective evidence nor S's intentionally bringing about his having defective evidence is morally bad, and as such S's evidence cannot indicate at T1 that either his having defective evidence at T2 or his bringing about his having defective evidence at T2 is not his best option at T1.

Another possible answer is to say that there exist some cases in which S has a moral obligation not to bring about his having defective evidence at a latter point in time, not because of some epistemic obligation, but rather because doing so would violate some other moral obligation. For instance, S has a moral obligation to ensure that he will not engage in any act such that this act will prevent him from doing his moral obligation at a latter point in time. And if S's ensuring that he will have defective evidence at T2 prevents him from satisfying his moral obligation at T2, then he has a moral obligation at T1 not to ensure that he will have defective evidence at T2. That is to say that Bob's engaging in an act at T1 that results in giving himself defective evidence at T2 will result in his inability to satisfy his moral obligation at T2 and as such, he will violate his moral obligation to not make it such that he could not perform his moral obligation at T2. This principle, while very plausible, is in no way related to the question at hand, though it at first might seem to be. Recall that S's moral obligation at T2 is purely a function of what his evidence indicates at T2. As such, at the point in time when Bob engages in the act of determining that X will not be his moral obligation at T2, he does not have a moral obligation at T2. His only moral obligation at the time of his act of determinacy is to do what his evidence at T1 indicates his best option.⁵⁶ Thus, so long as his best option at T1 does not preclude intentionally bringing it about that he has defective evidence at T2, then Bob's act of determinacy (i.e. taking the pill) will not violate any moral obligation (and this brings us right back to our original question: why would Bob's best option at T1 preclude his intentionally bringing it about that he has defective evidence at T2?). As

⁵⁶ Let "act of intentionality" mean that act which an agent engages in that satisfies the intentionality condition.

such, while it may be true that S has a moral obligation at T1 not to engage in an act which will make it impossible for him to satisfy his moral obligation at T2, this principle is not relevant to the situation at hand, since Bob's moral obligation at T2 is a result of his evidence at T2. Bob's act of determinacy occurs at T1 and in no way prevents him from being able to satisfy his moral obligation at T2—it simply determines what that moral obligation at T2 will be.

Our options for answering our question—why does Bob's act of intentionality violate his ability to satisfy the morally permissible condition?—are clear. Either we can say that Bob has a moral obligation not to ensure that he will have defective evidence in the future. But this has been shown to be an unsatisfactory response because Bob's obligation to refrain from ensuring that he will have defective evidence in the future has been shown to be a purely epistemic obligation and not a moral obligation.⁵⁷ The second option is to say that Bob's failing to satisfy his epistemic obligation (to refrain from ensuring that he will have defective evidence in the future) will entail his violating some moral obligation. But since Bob does not have a moral obligation not to ensure defective evidence in and of itself, Bob's doing so could not violate any moral obligation at T1, unless his doing so would violate his ability to satisfy a moral obligation at some other point in time. But Bob's moral obligation at the other point in time under the epistemic view will be the result of what his evidence will indicate at that point in time. That is, S's

⁵⁷ This precludes Bob's best option at T1 from being to satisfy this epistemic norm and it precludes his evidence from indicating that it is his best option at T1 is to satisfy this epistemic norm. After all, S's best option is a moral obligation. If S does in fact have a moral obligation to satisfy this epistemic norm it will be because his failure to do so will result in his violation of some moral obligation. Thus it will be for moral and not epistemic reasons.

moral obligation at T2 would need to be violated by his determining, at T1, his T2 evidence will be defective. But this cannot be the case since his moral obligation at T2 is a function of that T2 evidence, defective or not. And we have shown that the act of intentionality cannot violate his moral obligation at T1 because there is nothing immoral, in and of itself, about his intentionally bringing it about that he will have defective evidence at a future point in time. Thus, there exist no grounds upon which to say that Bob has violated his moral obligation at T1 by satisfying the intentionality condition. As such, the satisfaction of condition 1 and condition 3 under the epistemic view are not mutually exclusive and as a result the epistemic view formulated in terms of the evidence of which an agent is aware does violate the Principle of Non-Manipulation.

As we now clearly see, the epistemic view in either formulation cannot be successful. Should an agent's moral obligation be a function of his available evidence, then the epistemic view will be plagued by the unreasonableness worry, stemming from the epistemic views possessing the inaccessibility condition. If, on the other hand, an agent's moral obligation is a function of the evidence of which she is aware, then the epistemic view will be plagued by the manipulation worry just as the subjective view is. In truth, the epistemic view is not immune from both the problem facing the objective view and the problem facing the subjective view, but rather can avoid either the problem facing the objective view or the problem facing the subjective view but in doing so must succumb to the other. At this point we have reached a stalemate.

Section Four: *Concluding Thoughts*

We have now discovered conclusive reasons for dismissing the objective view, the subjective view, and both formulations of the epistemic view. Thus, we have reached a stalemate. But two important questions exist in light of this stalemate: What is the condition causing the manipulation worry (call this condition Z)? We have identified that it is the inaccessibility condition, contained by both the objective view and the epistemic view formulated in terms of an agent's available evidence (and not possessed by the subjective view), that causes the unreasonableness worry. What is the condition possessed by both the subjective view and the epistemic view formulated in terms of the evidence of which an agent is aware (and not possessed by the objective view), that causes the manipulation worry? If we can discover what condition Z is, then we have identified two conditions that the ideal view of the circumstantial element of moral obligation will not possess—namely, condition Z and the inaccessibility condition.

The second important question is what could this ideal view of the circumstantial element of moral obligation possibly be? After all, while these four views (i.e. the two formulations of the epistemic view, the subjective view, and the objective view) are not the only possible views, it appears that any view not considered within this paper will have to be some sort of fundamentally-similar variation on one of these four views. In other words, we have considered S's moral obligation as a function of each of the following: his available evidence, the evidence of which he is aware, his actual circumstances, and his beliefs. It appears that we have exhausted the possibilities of things that an agent's moral obligation can be a function of, and therefore it appears that it will be difficult for any view to avoid both of the two worries established in the critiques of the four aforementioned views (i.e. manipulation worry and the unreasonableness worry).

Let us, however, investigate into the first matter at this point, putting aside this second question for the time being. What is condition Z? Our analysis in Chapter 1 section 1 (refer to pages 33-35) as to why the objective view had a built-in-defense against the manipulation worry whereas the subjective view was vulnerable to this worry will shed light on the answer. The objective view was immune from the manipulation worry because for any act Q, such that S's performing Q at T1 entailed that S would not have a moral obligation to perform act X at T2, S would violate his moral obligation at T1 by performing Q at T1. This was the objective view's built-in-defense against the manipulation worry. Bearing this in mind let us revisit the three necessary and sufficient conditions for the manipulation worry that were discussed in Chapter 2 Section 2 (refer to

page 66). But let us revisit these three conditions in a way that takes them from the specific context of the epistemic view formulated in terms of the evidence of which an agent is aware, and let us put them in a generalized context, neutral to any view:

Condition 1 (“the intentionality condition”): S can voluntarily engage in an act Q at T1 with the knowledge (or at least informed belief) that his doing Q at T1 will ensure that at T2, $f(x) \rightarrow \sim X$. This condition ensures that S can intend, at T1, to determine his moral obligation, at T2.⁵⁸

Condition 2 (“the determinacy condition”): S’s performing act Q at T1, will ensure that $f(x) \rightarrow \sim X$ at T2. This condition ensures that S’s performing Q at T1 will determine that his moral obligation is not to do X at T2.

Condition 3 (“the morally permissible condition”): $\sim [f(x) \rightarrow \sim Q, \text{ at T1}]$. This condition ensures that S’s performing Q at T1 will not cause him to violate his moral obligation, under the given view.

⁵⁸ Let us clarify what is meant by $f(x) \rightarrow \sim X$. Let $f(x)$ denote whatever it is that determines an agent’s moral obligation in the given view on the circumstantial element of moral obligation. For example, in the objective view $f(x)$ will be an agent’s circumstance, in the subjective view $f(x)$ will be what an agent believes to be his circumstance, in the epistemic view in terms of an agent’s available evidence $f(x)$ will be what an agent’s available evidence indicates his circumstances are. Now, let the phrase $f(x) \rightarrow \sim X$ denote the fact that $f(x)$ indicates that X is not S’s best option. In other words, $f(x) \rightarrow \sim X$, indicates that X is not S’s moral obligation.

Now, it seems clear that the reason why the objective view is not susceptible to the manipulation worry is because under the objective view, S's satisfying condition 2 entails his failure to satisfy condition 3. In other words, we know that under the objective view:

$$[[S\text{'s doing } Q \text{ at } T1] \rightarrow [f(x) \rightarrow \sim X, \text{ at } T2]] \rightarrow [f(x) \rightarrow \sim Q, \text{ at } T1]$$

In fact, for any view under which the antecedent entails the consequent above, the result will be that this view is immune from the manipulation worry. If, however, under a given view the consequent can be false while the antecedent is true, then that view will be plagued by the manipulation worry. But, while this is a good test to see whether a particular view is susceptible to the manipulation worry or not, it is not where the explanation of *why* some views are susceptible to the manipulation worry whereas others are not rests. After all, it seems that every view will be capable of satisfying the antecedent (i.e. satisfying condition 2). So the explanatory difference between views susceptible to the manipulation worry and those that are immune to the manipulation worry, it seems, must rest somewhere else besides condition 2. The next place to look for an explanation is condition 1.

But the only difference between condition 1 and condition 2 is that condition 2 deals with a specific type of act (i.e. determinant acts) while condition 1 deals with an agent's ability to *intend* determinant acts.⁵⁹ That is, the difference between views that are

⁵⁹ Let us define a determinate act as follows, Q is a determinant act if and only if, (S's doing Q at T1) \rightarrow (f(x) \rightarrow \sim X, at T2)

immune from the manipulation worry and those that are plagued by it, is that those view's which are immune from the manipulation worry make it such that S cannot *intend* to perform a determinant act. The issue rests in the agent's inability to intend such acts, not in the lack of such acts existing.

It is now time to consider something which may already seem apparent. The objective view and the epistemic view formulated in terms of an agent's available evidence are plagued by the unreasonableness worry and yet they avoid the manipulation worry. On the other hand, the subjective view and the epistemic view formulated in terms of the evidence of which an agent is aware are plagued by the manipulation worry and yet they avoid the unreasonableness worry. Though this could be coincidental, this observation gives us reason to suspect that perhaps the two worries have some relationship which has up to this point been overlooked.

Let us remind ourselves what causes the unreasonableness worry. It is the inaccessibility condition, according to which S can be morally obligated to do X even though the reason why he is morally obligated to do X is inaccessible to him. In other words, the inaccessibility condition can be expressed as follows: $f(x) \rightarrow X$ & S is unaware that $f(x) \rightarrow X$. Up to this point, our discussion of the inaccessibility condition has been strictly as viewing it in a negative light due to its entailing the unreasonableness worry. But, it now seems clear that it the inaccessibility condition will prevent a view from becoming susceptible to the manipulation worry. The reason that agents in the objective view cannot *intend* determinant acts (and thereby this view is immune from the manipulation worry) is because the inaccessibility condition makes the fact that $f(x) \rightarrow X$

inaccessible to agents under the objective view. For S to *intend* a determinant act, S must have an awareness of the consequences of his action. In other words, Q is a determinant act if and only if S's performing Q at T1 makes it such that $[f(x) \rightarrow \sim X, \text{ at } T2]$. But to *intentionally* perform a determinant act, S must be aware that his doing Q will result in $f(x) \rightarrow \sim X$ at T2. Here in lies the defense. Given the inaccessibility condition, S is unaware that $f(x) \rightarrow X$ and therefore could not intentionally perform a determinant act.

We now see that the same condition (i.e. the inaccessibility condition) that causes the unreasonableness worry leads to the avoidance of the manipulation worry. Therefore if a view is to be susceptible to the manipulation worry, a necessary condition is that this view does not contain the inaccessibility condition. Could it be that condition Z, the condition which causes the subjective view is the denial of the inaccessibility condition?

At this point, I am not sure what to conclude. Upon first glance, it seemed that a view's containing the inaccessibility condition was a bad thing. After all, the inaccessibility condition leads to the unreasonableness worry which in turn results in the unjustifiably risky moral obligations that we have seen through our analysis. But it also appears that the inaccessibility condition renders a view immune to the manipulation worry. While it is clear that the inaccessibility condition is sufficient for a view's being immune to the manipulation worry, could it also be that it is necessary condition for a view's being immune to the manipulation worry (i.e. that the denial of the inaccessibility condition is condition Z)? If this is the case, then we would be forced to choose between the Principle of Non-Manipulation and the Principle of Reasonableness. That is, if the inaccessibility condition is a necessary condition for a view avoiding the manipulation

worry, then it becomes evident that *any* view on the circumstantial element of moral obligation will be susceptible to either the manipulation worry or the unreasonableness worry.

I think there is some reason, albeit a non-conclusive reason, to suspect that the inaccessibility condition is a necessary condition for immunity from the manipulation worry. If this were not the case, then there would be instances in which a view was not subject to the manipulation worry and yet, agents under that view would still have access to the fact that $f(x) \rightarrow x$. Thus, agents in this case can satisfy the intentionality condition for the manipulation worry—that is, under such a view it would be possible for S to intend to perform a determinant act. The question now becomes can there be something inherent to a determinant act which will lead to that determinant act violating S's moral obligation at the time of performing that act. Now that we have cleared the issue of intentionality it seems hard to imagine what that could be.

As I said before, I am not entirely sure what to make of all this. But, it appears that we are at the following fork in the road in our examination of the circumstantial element of moral obligation. Either: 1) The inaccessibility condition is a necessary condition for immunity from the manipulation worry, at which point we cannot accept both the Principle of Non-Manipulation and the Principle of Reasonableness; or 2) We must find some reason why a view would render it possible for an agent to satisfy the determinacy condition only if his performing a determinant act would violate his moral obligation at the time of his action. But, the only cases in which we have seen this occur (i.e. the objective view and the epistemic view formulated in terms of an agent's available

evidence) it has been an epiphenomenal effect of the view having an inaccessibility condition. Thus, we must now search for a reason why a determinant act would be inherently immoral under a view.

BIBLIOGRAPHY

- Davion, Victoria. "Action-Guides and Wrongful Intentions." *Public Affairs Quarterly* 6, no. 4 (1992): 365-374.
- Feldman, Richard. "Having Evidence." In D.F. Austin (ed.), *Philosophical Analysis: a defense by example*, 83-105. Dordrecht: Kluwer Academic Publishers, 1988
- Fischer, John Martin. "Introduction: A Framework for Moral Responsibility." In *My way: essays on moral responsibility*, 1-37. New York: Oxford University Press, 2006.
- Fitelson, Branden. *Studies in Bayesian Confirmation Theory*. Madison: University of Wisconsin-Madison, 2001.
- Hawthorne, John, and Jason Stanley. "Knowledge and Action." *The Journal of Philosophy* cv, no. 10 (2008): 571-590.
- Hieronymie, Pamela. "Responsibility for Believing." *Synthese* 161, no. 3 (2008): 357-373.
- Jackson, Frank. "Decision-theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* 101, no. 3 (1991): 461-482.

- McConnell, Terrance. "Ross on Duty and Ignorance." *History of Philosophy Quarterly* 5, no. 1 (1988): 79-95.
- Montmarquet, James. "Epistemic Virtue and Doxastic Responsibility." *American Philosophical Quarterly* 29, no. 4 (1992): 331-341.
- Moore, G.E. . *Ethics*. Oxford: Oxford University Press, 1912.
- Pink, Thomas. "Normativity and Reason." *Journal of Moral Philosophy* 4, no. 3 (2007): 406-431.
- Prichard, H. A.. *Moral writings*. Oxford: Clarendon Press, 2002.
- Pritchard, Duncan. "Moral and Epistemic Luck." *Metaphilosophy* 37, no. 1 (2006): 1-25.
- Ross, W. D.. *Foundations of ethics; the Gifford lectures delivered in the University of Aberdeen, 1935-6*. Oxford: Clarendon Press, 1939.
- Singer, Peter. "Famine, Affluence and Morality." *Philosophy and Public Affairs* 1, no. 1 (1972): 229-243.
- Wood, Allen. "The Duty to Believe According to the Evidence." *International Journal for Philosophy of Religion* 63, no. 1 (2008): 7-24.
- Zimmerman, Michael J.. "Moral Luck: A Partial Map." *Canadian Journal of Philosophy* 36, no. 4 (2006): 585-608.
- Zimmerman, Michael J.. *Living with uncertainty: the moral significance of ignorance*. Cambridge, UK: Cambridge University Press, 2008.

