Moral Luck, the Failure of Control, and the

Ability and Contribution Theory

Joseph Metz

Undergraduate Honors Thesis

Department of Philosophy

Duke University, Durham, North Carolina, 2012

Introduction:

We do not want arbitrariness in our moral judgments. In the first half of my thesis, I will discuss moral luck and the problem it poses for assigning moral responsibility in a non-arbitrary manner. Once I have discussed this problem, I will present one general model, which I will call the "Control Theory," for how to resolve this issue that I take to be an intuitive solution because it upholds the traditional line of thought that agents should only be held morally accountable for what they can control. In order to fully discuss this model, I will establish a conceptual framework of what it means to "control" a decision or the outcome of an action. Following this, I will describe three versions of the Control Theory in order to provide the idea of assigning moral responsibility based on control the strongest possible chance for success. I will argue against these three frameworks that are based on control and will discuss the implications of the failure of using control to successfully resolve the problem of assigning moral responsibility in cases involving moral luck.

In the second half of my thesis, I will present one response to the problem of assigning moral responsibility in cases involving moral luck, the Responsibility by Ability Theory, and I will then describe the limitations of this theory. This will lead me to modify the Responsibility by Ability Theory to form the Responsibility by Ability and Contribution Theory, which I will then defend as a viable solution to the issue of assigning moral responsibility in cases involving moral luck.

§1. THE FIRST HALF

I. The Types of Moral Luck

      A. Introduction

I will now outline four types of moral luck, as defined by Thomas Nagel in his paper "Moral Luck,"[1] and I will illustrate the problems that they pose for assigning moral responsibility.

B. Resultant Luck

The first type of moral luck is resultant luck, which Nagel characterizes as differences in moral judgments that depend on luck in how events turn out. As an example of resultant luck, he presents two scenarios involving a truck driver who has failed to recently check his brakes. Unluckily for the truck driver and for a child, in one of the scenarios, a small child darts in front of his truck, and he hits the child because he is unable to stop. In the other scenario, the child does not dart in front of his truck. Nagel highlights the arbitrariness of the moral judgments of the truck driver in the two scenarios when he argues that the truck driver's degree of blame for not checking his brakes depends on whether he needs his brakes to stop suddenly, and whether he needs to stop suddenly depends on the actions of the child, which are outside of the truck driver's control. For instance, the truck driver would deserve some blame for failing to check his brakes, but this amount of blame would be nowhere near the amount he would have for hitting the child with his truck. Hence, there is arbitrariness in the moral judgments of the driver's failure to check his brakes because the driver's degree of blame depends on something that is outside of his control.

As Susan Wolf points out in her paper "The Moral of Moral Luck,"[2] every day there are instances of minor negligence that usually do not cause bad outcomes, but

---

[1] Nagel, Thomas. "Moral Luck." In *Mortal Questions*. Cambridge: Cambridge UP, 2003.

[2] Wolf, Susan. "The Moral of Moral Luck." *Philosophic Exchange* 31 (2001): 4-19.

occasionally, they have disastrous consequences. For instance, drivers neglect to check their brakes, building inspectors allow buildings with minor violations to pass inspection, smoke detectors go too long without being checked, and so on. Usually, these small instances of negligence cause no harm, but occasionally, they backfire. How should we assess moral responsibility in these cases? Should we hold people who end up causing harm equally morally responsible as people who do not? These examples illustrate the difficulty of non-arbitrarily assigning moral responsibility for the outcomes of our actions in situations involving resultant luck.

C. Constitutive Luck

The second type of moral luck, constitutive luck, focuses on the luck in our inborn temperamental dispositions, which we do not choose. There is luck in whether we are born with a tendency to be tightly wound or more relaxed, kind or cold, brave or cowardly, and so on.

For example, suppose someone has the quality of being ungenerous, and suppose this agent's neighbor is very sick and asks for a ride to the hospital. The ungenerous agent does not give the neighbor a ride and instead, gives a small amount of money that is insufficient to cover the cost of a cab ride to the hospital. Nagel argues that such qualities like being ungenerous may have partially resulted from previous decisions and could be partially changed by current decisions. However, Nagel says having such a quality is "largely a matter of constitutive bad fortune."[3]

For Nagel, constitutive luck specifically concerns luck in the moral judgments of these temperaments and dispositions. In the example above, there is moral arbitrariness

---

[3] Nagel, "Moral Luck," 454.

4

because the agent will probably be judged as morally blameworthy for his stinginess, even if he suppresses outwardly showing it while he gives money to his neighbor. This agent will probably be blamed because of what his temperament and dispositions are like, which does not account for any effort he makes to suppress and overcome them.

In this paper, I consider luck in our constitutive factors in a way that goes beyond Nagel's usage because I consider how these constitutive temperaments shape the options of our decisions. Specifically, our temperaments and dispositions can affect which options we consider taking and can affect how we weigh the viability of these options. Applying this broader consideration of luck in our constitutive factors to the earlier example in this section, the quality of being ungenerous may have heavily influenced the agent to oppose any options of his decision that required him to give of himself, and the very act of giving some money may actually be commendable because it took a large mental effort to decide counter to his miserly temperament and give anything at all. In addition to Nagel's concern about the arbitrariness in the moral judgments of the agent's temperaments, this case also seems to involve moral arbitrariness in the moral judgments of the agent's decision. This second form of arbitrariness could be present because the agent will probably be judged as morally blameworthy for failing to decide to give enough money cover the full cost of his neighbor's cab ride, which does not account for the temperamental obstacles that the agent may have had to overcome in order to decide to give any money to his neighbor. Thus, there are difficulties in non-arbitrarily assigning moral responsibility for our decisions in cases involving luck in our constitutive factors.

D. Circumstantial Luck

The third type of moral luck, circumstantial luck, pertains to luck in an agent's circumstances, which affect the moral tests that that agent faces. Our circumstances affect what the options of our decisions are. For example, Nagel discusses how a German who became a guard in a concentration camp in Nazi Germany might have led a harmless life if he had left Germany for Argentina because of business reasons before Hitler came into power. Similarly, Nagel also discusses how someone who lived a harmless life might have become a guard in a concentration camp if he had not moved from Germany to Argentina. Since our circumstances affect which options exist for us to choose between in our decisions, I argue that circumstantial luck affects our decisions and raises issues when assigning moral responsibility for them.

E. Luck in Having an Undetermined Will

The fourth type of moral luck that Nagel discusses involves the will. Nagel argues that acts of the will "are the product of antecedent circumstances outside of the will's control."[4] Since he takes this to be the case, he questions how we can be morally responsible for the acts of the will.

Albeit interesting, I will not discuss this type of moral luck in this paper for two main reasons. First, the issue of free will is complex and requires more space in order to be fully discussed than I can give it in this paper. Second, I want to show that even when we set aside possible problems for having a free will, like causal determinism, serious issues still arise when we use control to assign moral responsibility.

F. The Dilemma Raised by Moral Luck When Assigning Moral Responsibility

---

[4] Nagel, "Moral Luck," 454.

A dilemma then arises when assigning moral responsibility for our decisions in cases involving constitutive factors or circumstantial luck and when assigning moral responsibility for the outcomes of our actions in cases involving resultant luck. Should we assign moral responsibility to ourselves based only on what we can control, or should we assign it based at least partially on what we cannot control? One problem with assigning moral responsibility for what we cannot control is that luck enters into our moral judgments. For instance, the moral responsibility of the truck driver from an earlier example then becomes dependent upon the actions of the involved child, which the truck driver has no influence over.

II. The Control Theory[5]

A. Introduction.

I argue that the intuitive response to this dilemma raised by moral luck is to avoid the problem of luck entering into our moral judgments by taking the first horn of the dilemma and only holding ourselves to be morally responsible for what we can control. I dub this response the "Control Theory."

B. The Three Forms of the Control Theory

i. The Complete-Outcome/Complete-Decision Control Theory

The strongest possible form of the Control Theory is the Complete-Outcome/Complete-Decision Control Theory, henceforth the Complete Control Theory, which only assigns moral responsibility to an agent for the outcome of one of her actions

---

[5] For the reader's convenience, please see Appendix A for a table that details the different variations of the Control Theory and their respective bases for assigning moral responsibility to agents for their decisions and for the outcomes of their actions. Appendix A also includes the definitions of the various types of control that agents can have over decisions and over the outcomes of actions.

or for one of her decisions if she has complete control over that outcome or decision. I will fully define what it means to have this "complete control" shortly, but it is sufficient for now to say that complete control requires that the outcome or decision does not depend on factors that are external to the agent.

ii. The Partial-Outcome/Partial-Decision Control Theory

A weaker form of the Control Theory assigns degrees of moral responsibility for an agent's decisions and for the outcomes of an agent's actions based on that agent's respective degrees of partial control over those decisions or outcomes. I name this the Partial-Outcome/Partial-Decision Control Theory, or the Partial Control Theory for short. Roughly, a higher degree of partial control over a given outcome or decision equates to a higher degree of moral responsibility for it.

iii. The Partial-Outcome/Narrow-Decision Control Theory

The Partial-Outcome/Narrow-Decision Control Theory, henceforth the Partial-Narrow Control Theory, assigns degrees of moral responsibility for the outcomes of an agent's actions based on that agent's degrees of partial control over those respective outcomes, just like the Partial Control Theory does. The Narrow Control Theory differs from the Partial Control Theory because the Narrow Control Theory assigns degrees of moral responsibility for an agent's decisions based on that agent's respective degrees of "narrow control" over those decisions. "Narrow control" over a decision differs from "partial control" over a decision, and I will explain the difference in the following section on control.

iv. Arguing Against These Theories

While these three theories initially seem plausible solutions to the dilemma raised by the problem of moral luck, I will argue against each theory. First, we never have the complete control required by the Complete Control Theory. Second, assessing the degree of partial control relies on the flawed concept of "purified control," and I will conclude that it therefore is untenable. Third, the Narrow Control Theory allows arbitrariness in its moral judgments, which defeats the purpose of the Control Theory, and it is therefore an untenable position for a Control Theorist to hold.

In order to analyze these three theories and argue why they fail, I need to first define what it means to control a decision, an action, or the outcome of an action.

III. Control

A. Introduction

An agent potentially has control over the three following factors. First, an agent can control her decision and form an intention. Second, an agent can control the actualization of her intention through her actions. Third, an agent can control an outcome by means of her actions.

B. Luck

Luck or chance, which results from causal chains that do not pass through the agent's decision or intention, can interrupt the realization of an action once an intention has formed. For example, suppose a pitcher forms an intention of throwing a ball to a catcher. The realization of that pitcher's intention through the pitcher's action can be interrupted if he has a muscle spasm while he is attempting to throw the pitch. However, the luck that can interrupt the realization of an action through an intention parallels the luck that can interrupt the realization of the outcome through the action. For example,

the pitcher could control his action of throwing the ball, but the outcome of the ball reaching the catcher could be interrupted by a bird that flies in the ball's path and blocks it. Because the interruption of intention to action parallels the interruption of action to outcome, and for the sake of simplicity, I will assume unproblematic passage from decision to action. There is one case in my discussion of "purified control" in which I will consider the interruption of decision to action. That is the only case where I will consider luck in actions separately from luck in the outcomes of actions. For the rest of the paper, I will not separate them because the issues that arise between decision to action have parallels between action to outcome.

I would also like to note the distinction between luck and moral luck. The luck I just described deals with causal chains that are outside of an agent's control. Moral luck concerns arbitrariness in moral judgments in which an agent's degree of blame for a decision or for the outcome of an action depends on causal chains that are outside of her control.

C. The Levels of Control Over the Outcome of an Action

Next, I will precisely define what it means for an agent to have control over a decision and over the outcome of an action, starting with an outcome. An agent has complete control over the outcome of an action if and only if all of the possible causal chains that are sufficient to change the outcome pass through the agent's relevant intention and all subsequent causal chains that are downstream of the agent's intention and are sufficient to change the outcome accord with the agent's relevant intention. The possibility referred to in these "possible causal chains" is possibility consistent with the physical laws of the actual world. An agent's relevant intention for the outcome of an

10

action is the intention to bring about a specific outcome. An agent has partial control over an outcome if and only if some of the possible causal chains that are sufficient to change the outcome pass through the agent's intention and some of the other causal chains that are sufficient to change the outcome do not. An agent has no control over the outcome of an action if and only if none of the causal chains that are sufficient to change the outcome pass through the agent's intention.

I will reuse the example of the pitcher throwing a ball to the catcher to illustrate these three levels of control the pitcher has over the outcome of his action of throwing the ball. The pitcher has complete control over the outcome if there are no possible causal chains that are sufficient to change the outcome that do not pass through his intention. For instance, there cannot be any causal chains involving a bird that flies in front of the pitch. It is not enough if the pitcher throws the ball, and it ends up reaching the catcher. Complete control requires the stronger claim that there are no possible causal chains in which a bird or something else interferes with the pitch.

Furthermore, the second condition of complete control, which is the condition that subsequent causal chains that are downstream of the agent's intention and are sufficient to change the outcome accord with the agent's intention, covers cases like the following. Suppose a third person deliberately waits for the pitcher to throw the ball and then steps in and intercepts the ball before it reaches the catcher. In a case like this, all of the causal chains that are sufficient to change the outcome pass through the agent's intention, but the actions of the third person are downstream of the pitcher's intention and do not accord with the pitcher's intention. The pitcher has complete control over the outcome of

11

the pitch only if this second condition is also met that all causal chains downstream of the pitcher's intention accord with his intention.

The pitcher has partial control over the outcome of the pitch if some of the causal chains that are sufficient to change the outcome pass through the pitcher's intention, like the pitcher's action of throwing the ball. However, there are other possible causal chains that do not pass through the pitcher's intention that are sufficient to change the outcome, like the wind blowing the ball off course or a bird interfering. Even if the wind does blow the ball off course, the pitcher could still have partial control over the outcome if he at least controlled his throwing of the ball, for instance. Technically, the causal chain that brings about the outcome of the pitch does pass through the pitcher's intention because the pitcher's intention to throw the ball causes the ball to be thrown. The interfering bird and wind make this a case of partial control and not complete control because the actions of the bird and the movement of the wind do not pass through the pitcher's intention. In other words, the causal chain involving the pitcher is necessary for the bird to even be able to interfere, but the causal chain of the pitcher is not sufficient to determine the outcome of the throw.

The pitcher would have no control over the outcome if no possible causal chain that was sufficient to change the outcome passed through the agent's intention. It is not enough that what the pitcher does has no effect on the outcome. Rather, nothing he could possibly do could alter the outcome. It seems strange to think of a pitcher having no control over a pitch, so perhaps a better example of having no control would be a fan yelling at a TV during a baseball game. No matter how much or in what way the fan yells at the TV, he will not change the outcome of a pitch.

D. The Levels of Control over a Decision

i. Introduction

I will now define what it means for an agent to have control over a decision. Agents can have complete control, partial control, narrow control, or no control over a decision. I characterize a decision as being composed of two parts that agents can have control over. First, an agent can deliberate over her options, which involves assessing and weighing her options and then deciding which option to take. Second, an agent can have control over what her options are that she is deciding between.

ii. Complete Control

Complete control over a decision factors in what the options of an agent's decision are and how the agent weighs these options. An agent has complete control over a decision if and only if all of the possible causal chains that are sufficient to change what her options are in that decision or change the weight she assigns to those options pass through her relevant intentions.

There are two relevant intentions. The first intention is the intention to deliberate over one's options and make a decision. The second is to choose the array of options that one is deliberating between. This second intention of choosing one's options is extreme. Admittedly, only an omniscient and omnipotent being could satisfy this condition of control, so an objector could ask why I made this condition so stringent.

I made this condition so difficult to satisfy because of circumstantial luck. A German making a decision in the circumstances of being in Nazi Germany does not have as much control over his decision as he would if he could choose whether to make that decision in Nazi Germany. For instance, he might have to choose between the option of

joining the Nazi ranks and the option of resisting the Nazis. His choice is constrained by living in Nazi Germany. The strongest amount of control he could have over this decision would be if he could add more options to his decision, like the option of living peacefully in a different place or during a different time. Since being able to choose his options is the highest amount of control he could have, I set this as the requirement for complete control.

I will now give an example to help clarify this idea of complete control. Suppose that an agent is deciding which book to read. This agent has complete control over this decision only if all possible causal chains that are sufficient to change the agent's options or how the agent weighs those options pass through either of the agent's two relevant intentions. In this example, none of the options that the agent is considering of books to read could result from chance or causal chains that do not pass through the agent's intention to choose that option. Since the options of books to choose from is constrained by the agent's knowledge of which books exist, the agent must have enough knowledge about every book to know that it exists as a potential option and enough knowledge to be able to weigh the option of reading that book. Furthermore, the options of which books to read and the weighing of these options must not depend on factors that involve luck, so the knowledge of each of these books must not depend on causal chains that do not pass through the agent's intention, like causal chains that only pass through the intentions of other agents. For instance, this agent's option of reading *Moby Dick* cannot depend on a friend telling this agent about the existence of the book *Moby Dick*.

iii. Partial Control

An agent has partial control over a decision if and only if some, but not all, of the possible causal chains that are sufficient to change what options she chooses between or change the weight she assigns to those options pass through her intention to choose an option or her intention to deliberate amongst her options.

To illustrate partial control, I will reuse the example of the agent deciding which book to read. The agent has partial control if some of her options of books to read depend on causal chains that pass through her intention to choose those books as options, and some do not. For example, she could still have partial control over her decision even if a friend gives her the option of reading *Moby Dick*, as long as she brings about some other options on her own. Similarly, if her friend makes recommendations and affects how the agent weighs some of her options, the agent still has partial control if some of the other causal chains that are sufficient to change the weighing of her options pass through her intention to deliberate amongst her options.

An objector could ask whether a human agent could ever have the ability to choose any of her options of a decision because all of the options of a decision seem to be constrained by an agent's circumstances.

I offer the following example to show that agents can have some ability to choose the options of their decisions. Suppose that an agent is deciding between donating money to Charity A and donating money to Charity B. After researching both charities, the agent realizes that both are corrupt, and she is dissatisfied with these being her only options. In response to her dissatisfaction, she forms an intention to expand her options. She uses the resources of a local library to find out about charities without issues of corruption. She thereby learns about Charity C and Charity D, which are both reputable,

and she then deliberates between these two new options. Technically, the specific

charities that she learns about are constrained by circumstantial factors, like what

resources her library has to offer. In spite of these circumstantial factors, this agent still

has some ability to choose the options of her decision because the causal chains that pass

through her intention to add more options also pass through her intention to choose her

options. Thus, agents can have some ability to choose their options.

iv. Narrow Control

Narrow control only applies to decisions and differs from partial control over

decisions because narrow control does not involve an agent's amount of choice of what

the options of a decision are. An agent has narrow control if and only if at least some of

the possible causal chains that are sufficient to change what weight an agent assigns to

her options pass through her relevant intention of deliberating amongst her options.

Recall that deliberating amongst her options involves assessing and weighing her options

and deciding which option to take.

Causal chains that are sufficient to bring about, remove, or alter the options that

an agent deliberates between are irrelevant for narrow control. Hence, narrow control

does not factor in whether an agent chose his options, and the requirements of narrow

control are therefore less stringent than those of partial control.

To demonstrate narrow control, reconsider the example of the agent choosing

which book to read. This agent has narrow control if at least some of the causal chains

that are sufficient to affect how she weighs the options of reading the various books pass

through her relevant intention to deliberate amongst her options of books to read. Again,

for the purposes of narrow control, it is irrelevant whether the agent chose any of her options of books.

Complete control and partial control assess whether an agent chose all or some of the options of her decisions, respectively. An objector who accepted the idea of degrees of control over a decision could still argue that choosing the options of one's decisions is not an important part of having control over one's decisions. This objector would therefore not be satisfied with the Partial Control Theory's assessment of decisions. I consider narrow control over decisions because it accounts for the views of this objector since it does not assess an agent's degree of control over her choice of options and only assesses her degree of control deliberation amongst her options. However, as I briefly discussed in the section on complete control over decisions, assigning moral responsibility based on narrow control allows arbitrariness to enter into moral judgments because of circumstantial luck, and later in section **VI,** I will present a more complete argument for the presence of this moral arbitrariness.

v. No Control

An agent only has no control over a decision if she does not satisfy the criteria for partial control. In other words, an agent has no control over a decision if and only if none of the possible causal chains that are sufficient to change what her options are or the weight she assigns to those options pass through her relevant intentions.

vi. Decisions That Are Not Deliberated

We do not deliberate many of the decisions that we make on a daily basis because we unconsciously make them, and they are the results of our habits and routines. For

instance, suppose that I always eat cereal for breakfast, so when I wake up, my decision of what to eat is automatic and not deliberated.

I make a note of decisions that are not deliberated because they, like deliberated decisions, can have moral implications. Suppose that a person has a racial bias. On a day-to-day basis, he might not deliberate whether he ought to have this bias, and he may not even be aware that he has it.

Decisions that have not been deliberated can still be assessed by the same criteria that I used for deliberated decisions. Some habits result from a conscious decision, and these decisions can be assessed. For example, a person could decide to volunteer at a soup kitchen once a week. The decision to volunteer on a given week might be unconsciously decided by the routine of volunteering, but the decision to start volunteering could be assessed as a consciously deliberated decision.

On the other hand, some habits do not result from conscious decisions. For example, the person with the racial bias may not have ever consciously decided to stereotype people with certain ethnicities, and he may not even be aware that his does this. If someone makes him aware of his bias, then he faces a deliberate decision of whether or not to try to correct it. However, if he is totally unaware of his unconscious bias, it is difficult to say that he has any control over it.

Thus, for decisions that are not deliberated, if the decisions result from an earlier conscious decision to form a habit, then the agent's control over that conscious decision can be evaluated. If an unconscious decision does not result from some consciously decided routine and if the agent is unaware of it, then the agent does not have any control over that decision that was not deliberated.

IV. My Argument Against the Complete Control Theory

A. An Example

Now that I have defined the various levels of control an agent can have over decisions and over the outcomes of actions, I will apply these definitions and argue why it is undesirable to hold any of the three Control Theories as a response to the problem of assigning moral responsibility in cases of moral luck, and I will begin with the Complete Control Theory.

Recall that the Complete Control Theory assigns moral responsibility to an agent for a decision or outcome of an action if and only if that agent has complete control over that decision or outcome. For an example of how the Complete Control Theory works, consider the case of Lifeguard, who sees Victim drowning in the ocean. Lifeguard has complete control over his attempt to rescue Victim if and only if all of the chains of causation that are sufficient to change the intended outcome pass through Lifeguard's relevant intention and all subsequent causal chains that are downstream of the agent's intention and are sufficient to change the outcome accord with the agent's relevant intention. For example, if Lifeguard has complete control over his rescue attempt, there cannot be any chains of causation that pass through a nearby, hungry shark's intention that are causally downstream of Lifeguard's intention and result in the shark eating Victim. If Lifeguard has complete control over his rescue attempt, then the Complete Control Theory assigns him moral responsibility for the rescue attempt. If Lifeguard does not have complete control over it, then the Complete Control Theory cannot assign him moral responsibility for it.

B. A Preliminary Criticism of the Complete Control Theory

One initial criticism of the Complete Control Theory is that it apparently cannot distinguish between the moral responsibility of Risky Lifeguard and Safe Lifeguard if neither one can completely control the outcome of his action. Risky Lifeguard is profusely bleeding before he starts swimming, and he does not use a nearby, instant bandage. Hence, when Risky Lifeguard swims out to rescue Victim, the extra blood in the water greatly raises the likelihood that a shark attacks Victim. Safe Lifeguard is healthy and does not increase the small chance of a shark attack. Since there is a chance of a shark attack for both Safe Lifeguard and Risky Lifeguard, neither one has complete control over the outcome of the rescue attempt. Since neither has complete control, the Complete Control Theory does not assign moral responsibility to either one for the outcome of the rescue attempt. The criticism of the Complete Control Theory is that Safe Lifeguard's rescue attempt is morally better than Risky Lifeguard's rescue attempt because Risky Lifeguard may lower Victim's chances of survival, but the Complete Control Theory cannot account for the moral difference.

C. A Response to This Initial Criticism

A Complete Control Theorist could respond that while this version of the Control Theory cannot assign different amounts of moral responsibility to Risky Lifeguard and Safe Lifeguard for the outcomes of their rescue attempts, it theoretically could still assign different overall levels of moral responsibility for Risky Lifeguard and Safe Lifeguard. Deciding to jump into the water when he knew he was bleeding instead of using the nearby, instant bandage gives Risky Lifeguard more moral responsibility than Safe Lifeguard. As long as Risky Lifeguard and Safe Lifeguard have complete control over their decisions, then they have different levels of moral responsibility because their

decisions were morally different, and they therefore have different overall levels of moral responsibility.

I note this criticism here because this was my initial objection to the Complete Control Theory, which seemed to prevent it from being tenable at all. I want to be as fair as possible to this theory, so I want to point out that it theoretically can account for a difference in the overall moral responsibility between Risky Lifeguard and Safe Lifeguard. However, I will revisit this criticism in my discussion of the extremity of the position of the Complete Control Theory. My main objection to this theory follows.

D. A Second Argument Against the Complete Control Theory

i. Introduction

My main argument against the Complete Control Theory is that, barring extreme circumstances or omnipotence, we never have complete control over our decisions or over the outcomes of our actions. This implies that the Complete Control Theory cannot ever assign us moral responsibility. The conclusion of my argument focuses on the extreme position the Complete Control Theory faces since it cannot assign moral responsibility to limited beings such as ourselves.

ii. Our Lack of Complete Control Over Outcomes

The first problem for the Complete Control Theory is that we have limitations and never have complete control over the outcomes of our actions because luck in the outcomes of our actions is more pervasive than it first appears. For the following examples, suppose that the agents have complete control over their decisions because I will only be considering the agents' control over the outcomes of their actions.

Consider a game of billiards. Suppose that I intend to strike the cue ball and knock another ball into a pocket. I may not know how to strike the cue ball to achieve this outcome, but it seems like all of the chains of causation that are sufficient to alter the outcome pass through my relevant intention of knocking the ball into a pocket.

However, this is not the case. An earthquake could affect the path of the cue ball after I strike it, which means that some causal chains that were downstream of my intention and sufficient to change the outcome did not accord with my intention. The earthquake could interrupt the link between my intention and my action, but remember that I have assumed that there is unproblematic passage from decision to action, so I interpret the earthquake so that it interrupts the action-outcome link. Thus, there is some luck in the outcome of my action of striking the cue ball when I play billiards.

An objector could quite reasonably argue that the chance of an earthquake occurring during a game of billiards represents the presence of a trivial amount of luck in this outcome, so I essentially have complete control over the outcome of my actions when I play billiards. This is evinced by the assumption that most billiards players would not consider possible earthquakes when they take their shots.

I would agree with the objector that this case involves a miniscule amount of luck. However, this tiny amount of luck means that there are still a few causal chains that are downstream of my intention and sufficient to bring about an outcome that I did not intend, even if the probability that these chains will occur is small. These causal chains that do not pass accord with intention are sufficient to change the path of the ball, and therefore, they are sufficient to change the outcome. Hence, I do not have complete control over the outcome of this action.

I will now consider a more plausible example of a case in which a person appears to have total control but upon examination, does not. It both is a more plausible example and was used earlier because it actually happened. Consider a baseball game. A pitcher theoretically controls the speed, spin, and direction of his pitch. I am trying to be fair to advocates of the possibility of complete control, so suppose that we ignore both the batter and the possibility of the pitcher's action being interrupted. This means that the ball cannot slip out of the pitcher's hand and that the pitcher cannot have a muscle spasm that alters his throw, which gives the pitcher more control than he would ever actually have. In this simplified case, we can apparently infer that the pitcher has complete control over the outcome of his pitch. This, however, is another case in which an agent does not actually have complete control over the outcome of his action. Here, the pitcher can control the release of his pitch, but the pitcher does not actually have complete control over whether that pitch reaches the catcher. In a real baseball game, a bird flew in front of a pitch, which prevented it from reaching the catcher. The causal chains involved in the bird's motion were downstream of the pitcher's intention, and they did not accord with it. Hence, even the outcome of an action as straightforward as throwing a ball to another person still involves luck.

A conceptually similar case, which also happened, shows that whether or not an assassin's bullet kills his target, even if he aims correctly and is at close range, depends on causal chains that are independent of the assassin's intention and action. In this example, Theodore Roosevelt survived an assassination attempt by someone who shot at him from close range and correctly aimed at his chest. Roosevelt was saved from the

bullet because of a lengthy speech he had folded up in his pocket, which absorbed enough of the bullet's momentum to prevent it from being lethal.

I will describe one final example to emphasize that we never completely control the outcomes of our actions. Suppose that someone argues that I can completely control the outcome of one action when I write with a pen, and suppose that no bird or assassin interferes with me, and further suppose that I am in complete control of my muscles and therefore cannot have involuntary muscle spasms. It is just a pen, a piece of paper, and I all alone in a room that we can suppose is earthquake-proof. Can I completely control the outcome of my action of writing? Do all of the causal chains that can bring about this outcome pass through my intention and action?

I still say no because the outcome of writing with a pen, even in these circumstances, depends on causal chains that do not pass through my relevant intentions. My pen could stop functioning mid-sentence, as pens occasionally do, or the ink pouch could burst, which is another outcome that is uncommon, but it certainly happens. Thus, even in the case of writing down my thoughts with a pen, the outcome of my attempt to write is not entirely immune to luck because there are still some causal chains that are sufficient to bring about the outcome that do not pass through my intention.

Thus, since we are limited beings, we never have complete control over the outcomes of our actions. Since we never have this complete control, the Complete Control Theory cannot assign us any moral responsibility for the outcomes of our actions.

iii. Our Lack of Complete Control Over Our Decisions

A Complete Control Theorist may still appear to be able to assign agents moral responsibility for the decisions that agents make, but I will now outline the second

24

problem that as finite human beings, we never have complete control over our decisions because of luck in our circumstantial and constitutive factors.

The following examples illustrate the effects that luck in our circumstances has our decisions. Suppose that you suffer from seasonal allergies in the spring. In this case, your local environment, the surrounding plants, can be causally sufficient to change what your options are and how they are weighted in some of your decisions. I can personally attest that the local pollen count plays a large role in whether I decide to exercise outside or inside a gym and what route I take to get to my classes. Hence, causal chains stemming from my local environment are causally upstream of my intention and affect how I weigh many of my options. Furthermore, as human beings with limitations, we cannot remake the circumstances of all of our decisions. For instance, I cannot realistically add the option of becoming powerful enough over plants to be able to prevent all plants from releasing pollen whenever I am nearby, no matter how much I wish I could. Thus, my local environment can alter both how I weigh the options of my decisions and what those options are, and I therefore do not have complete control over these decisions.

Note that, I am not claiming that the causal chains that arise from our circumstances are causally sufficient to bring about *all* of the options that we choose between or affect *all* of the weighing of our options. Rather, I claim that the combination of these chains that do not pass through our intentions to choose our options or our intentions to deliberate between them is sufficient to affect many of the options of our decisions and affect at least some of how we weigh our options.

25

The next set of examples shows the effects that the presence of luck in our constitutions has on our decisions. Recall that I consider luck in our constitutive factors in a way that goes beyond Nagel's original usage of constitutive luck because I consider how our constitutive temperaments affect what options of a decision we debate taking and how we weigh the viability of these options.

One example of luck in our constitutions is mental illness. I take mental illness to at least partially consist of dispositions resulting from a combination of genetic factors, environmental factors, and chemical imbalances. Mental illness therefore combines elements of luck in an agent's circumstantial and constitutive factors. For the purposes of illuminating the effects of luck in an agent's constitutive factors, I will consider the parts of mental illness that pertain to luck in an agent's constitution and will not consider the parts of mental illness that involve luck in an agent's circumstances. I characterize mental illness as a burden that weighs on a person's mind and can, at times, heavily influence the decisions that that person makes. Mental illness also evinces my claim that we cannot control our inborn character and to some extent, our circumstances, because I cannot imagine someone willingly choosing to be born with or choosing to develop a mental illness.

Consider the example of someone with the mental illness depression. Her depression alone is not causally sufficient to determine all of the options that she chooses between for all of her decisions, but it is sufficient to change which options are present for some decisions and change how much weight many of her options have. For instance, in situations of danger, her depression could result in options of self-preservation seeming less desirable, and her depression could be causally sufficient to

add options of allowing herself to get hurt. When deciding if she is able to live on her own, this agent's depression could be causally sufficient to give the option "no" much more weight than it would if she did not have depression.

Thankfully, many people do not have mental illnesses, but that does not mean that most people are immune to the effects of luck in their constitutions. Rather, this case of mental illness offers an extreme example of the effects that our dispositions can have on our decisions.

For a more commonplace example that illustrates the pervasiveness of the impact that luck in our constitutions has on our decisions, consider the dispositions of optimism and pessimism. While the extent of people's optimism and pessimism varies, everyone or almost everyone has one of these dispositions, and people do not usually choose whether they are optimistic or pessimistic. Having either of these two dispositions affects what options an agent considers and how that agent weighs those options. For instance, an optimistic agent might think that he can finish his undergraduate thesis and his applications to graduate school in the same semester, so when deciding when to write his thesis, he might give more weight to the option of finishing his thesis in the fall, which is when applications are due. A more pessimistic agent faced with the same decision might think that he cannot finish his thesis and applications to graduate school in the same semester, so he might give more weight to the option of finishing his thesis in the spring, which is after applications are due. Thus, luck in our constitutive factors can affect our decisions.

It seems exceedingly unlikely that we ever have complete control over any of our decisions because it is improbable that all of the causal chains that are sufficient to bring

about, remove, or alter the options of a decision pass through our relevant intentions of choosing our options, and it is unlikely that all of the causal chains that are sufficient to alter how we weigh the options of a decision pass through our relevant intentions to deliberate amongst our options. Odds are, some of these chains pass through things in our circumstances or elements of our constitutions that are independent of our relevant intentions.

Thus, we never have complete control over our decisions. Therefore, the Complete Control Theory cannot assign us moral responsibility for our decisions either.

iv. The Extreme Position of the Complete Control Theory

These arguments demonstrate that the Complete Control Theory cannot ever assign us moral responsibility for our decisions or for the outcomes of our actions. Note that the Complete Control Theory cannot assign moral responsibility to agents based on other factors that are outside of our control because that would directly contradict the central premise of the general Control Theory. Thus, the Complete Control Theory cannot assign any moral responsibility to non-omnipotent and non-omniscient beings such as ourselves since we do not have complete control over anything. The inability to assign us moral responsibility constitutes a major flaw for the Complete Control Theory.

A Complete Control Theorist could take a hardline stance and argue that the Complete Control Theory's inability to assign moral responsibility to us reflects an error in our desire to assign moral responsibility to limited human agents and does not demonstrate an error in the Complete Control Theory.

Although I do not have the time or space to argue in this paper that our desire to assign moral responsibility to human agents is correct, I would like to note the extremity of the position that Complete Control Theorist holds.

Examine again the case of Risky Lifeguard and Safe Lifeguard. My initial objection to the Complete Control Theory was that it could not morally account for the difference between the outcomes of the actions of these two people since these outcomes depended upon contingent factors. The response was that the Complete Control Theory could theoretically assign Risky Lifeguard and Safe Lifeguard different amounts of moral responsibility as long as they had complete control over their decisions. However, I have shown that limited human agents do not ever have complete control over their decisions. Hence, the Complete Control Theory cannot assign Risky Lifeguard and Safe Lifeguard different amounts of moral responsibility for the outcomes of their actions, for their decisions, or for factors that are outside of their control. I conclude that the Complete Control Theory cannot account for the moral difference between these people. Neither can the Complete Control Theory account for the intuitive moral difference between someone who hugs puppies and someone who kicks them for fun.

A Complete Control Theorist could assert that our intuitions are flawed and argue that there should be no difference in the moral responsibility between Safe Lifeguard and Risky Lifeguard or between the puppy-kicker and the puppy-hugger. Hence, the Complete Control Theorist must take up the extreme position that all of our moral judgments involving the assignment of different levels of moral responsibility to limited human agents are flawed.

The central issue of this paper is how to assign moral responsibility in instances involving moral luck, and the various forms of the Control Theory are supposed to be ways of assigning moral responsibility in such instances. As a result, I assume that our desire to assign moral responsibility to human agents is justified, even in cases involving moral luck. Since I assume that we want to be able to assign moral responsibility to human agents in cases involving moral luck, and since the Complete Control Theory cannot ever do this, I conclude that the Complete Control Theory is an undesirable position to hold.

V. My Argument Against the Partial Control Theory

A. The Partial Control Theory Revisited

Now that I have defined partial control over decisions and over outcomes of actions, I will utilize partial control to fully explicate the Partial Control Theory and to then argue against it. The Partial Control Theory assigns degrees of moral responsibility to an agent for her decisions and for the outcomes of her actions based on her respective degrees of partial control over these decisions and actions.

B. Purified Control[6]

i. Introduction and Description

The agent's degree of partial control is assessed by examining the possible causal chains that are sufficient to change a decision or outcome that pass through the agent's intention and comparing them to the total set of possible causal chains that are sufficient to change that decision or outcome. This extracts the causal chains that are sufficient to

---

[6] I use Bernard William's metaphor of the purification of control from
Williams, Bernard. "Moral Luck," in *Moral Luck: Philosophical Papers* (Cambridge: Cambridge University Press, 1981), 29-30.

change the decisions or outcomes that do not pass through the agent's intention, and it removes them as a basis for the agent's moral responsibility for this decision or outcome. This "purifies" the agent's control over this decision or outcome in that it removes luck and other causal chains that do not pass through the agent's intention. Hence, I call the remaining causal chains the agent's "purified control" over that decision or outcome because all of the remaining causal chains that are sufficient to change the decision or outcome pass through the agent's intention.

If the total set of causal chains that are sufficient to change the decision or outcome pass through the agent's intention, then that agent has complete control. Hence, the comparison between the causal chains that pass through the agent's intention that are sufficient to change the decision or outcome and between the total set of causal chains that are sufficient to change the decision or outcome amounts to a comparison between the agent's actual amount of purified control and the agent's hypothetical complete control. The closer the agent's purified control is to this hypothetical complete control, the higher the agent's degree of partial control is.

In describing this theory, the following two factors seem relevant when comparing an agent's purified control to his hypothetical complete control. First, the probabilities of occurrence of the causal chains that pass through an agent's relevant intentions are significant. For instance, if the only causal chains that do not pass through the agent's relevant intentions and are sufficient to change the agent's decision or outcome also have very low probabilities of occurrence, then the agent's purified control is close to the agent's hypothetical complete control. The second factor is the number of causal chains that pass through the agent's relevant intentions. For example, if all but

one of the causal chains that are sufficient to change the agent's decision or outcome also pass through the agent's relevant intentions, and if the probability of occurrence of that one causal chain is low, then that agent's purified control is close to the agent's hypothetical complete control.

I will not discuss the exact methodology of how purified control is compared to complete control because my criticism of purified control is more fundamental than a critique of the methodology of this comparison. Thus, in order to be as fair as possible to the notion of purified control, for the purposes of this paper, I suppose that there is some reasonable way to compare purified control to complete control using the two factors that I have outlined. I will argue against the fundamental idea of purified control, which is that we can extract from a decision or action the causal chains that are sufficient to change that decision or outcome that do not pass through the agent's intention.

ii. Examples of Purified Control

For an example of purified control over the outcome of an action, consider the case of me writing with a pen. I stipulate for this example that I cannot have muscle spasms or other similar issues, so my movement of the pen is part of my purified control because all of the causal chains that are sufficient to change the movement of my pen pass through my relevant intention of bringing about the outcome of me writing a specific passage. The outcome of the movement of the pen in its entirety is not part of my purified control because there are causal chains that are sufficient to change the outcome of the writing that do not pass through my intention. As an example of a causal chain that does not pass through my intention, the pen's ink pouch could burst. Another possibility is that an earthquake shakes the table that I am writing on and affects where

my hand ends up touching the paper. Both are sufficient to change the outcome of the movement of the pen. Parts of the outcome of the movement of the pen are parts of my control if all of the causal chains that are sufficient to change those parts pass through my relevant intentions, which is the case for the part constituted by the movement of my hand.

For another example, suppose that a pitcher throws a ball to a catcher, and a bird flies in front of the pitch. If I stipulate that the ball does not slip out of the pitcher's hand and make other similar stipulations, then the pitcher's purified control includes the motion of the ball before it leaves his hand. The entire outcome of the pitch is not part of the pitcher's purified control since the causal chains involving the motion of the bird are independent of the pitcher's intention to bring about the outcome of the ball reaching the catcher.

An agent's purified control over a decision includes all of the causal chains that are sufficient to change what the options are or how the options are weighed that pass through the agent's intention. A causal chain is sufficient to change what the options of a decision are if and only if it alters an option, removes an option, or adds an option. An agent's purified control over a decision also includes all of the causal chains that are sufficient to change the weighing of an option and pass through the agent's intention.

For an example of purified control over a decision, consider the embellished case of me choosing my thesis topic. Suppose that I am considering writing about epistemology and logic, and suppose that the causal chains that are sufficient to change or bring about both of these options pass through my intention. That means that these two options are part of my purified control over my decision. Suppose that I am also

contemplating writing about moral luck, and suppose that I am only considering this option because a friend told me about the problem of moral luck. Since the causal chain that resulted in me considering moral luck as another option passed through my friend's intention and not my own, the option of writing about moral luck is not part of my purified control.

The weight I give to the options of writing about epistemology and logic are part of my purified control if all of the causal chains that are sufficient to change the weight I give them pass through my relevant intentions. The weight I give to the option of writing about moral luck is not part of my purified control if some causal chains that are sufficient to change the weight I give to this option pass through my friend's intention and not mine. For instance, my entire weighing of the option of writing about moral luck would not be part of my purified control if my friend biased me by telling me how fascinating moral luck is. Part of my weighing of the option of writing about moral luck is part of my purified control if some of the causal chains that are sufficient to change the weighing of that option pass through my intention.

iii. Summary of Purified Control

In this section, I have outlined what purified control is and how it connects to the Partial Control Theory. An agent's purified control includes all of the causal chains that are sufficient to change a decision or outcome that pass through that agent's intention. The Partial Control Theory assesses an agent's degree of partial control over a decision or outcome by comparing the agent's actual amount of purified control to the agent's hypothetical complete control over that decision or outcome.

In the next section, I will argue against the coherence of the fundamental notion of purified control, which is that we can abstract away causal chains that are sufficient to change a decision or outcome that do not pass through an agent's intention.

C. A Practical Problem with Assessing Purified Control over Decisions

i. Introduction

I object to the Partial Control Theory because purified control sounds plausible in theory, but upon closer examination, it falls apart because of both practical and conceptual problems.

ii. A Practical Problem with Assessing Purified Control over Decisions

One practical issue with assessing purified control is that it is difficult to know of all of the causal chains that are sufficient to affect the weight I assign to each of my options in a decision because I am not consciously aware of many of them.

For instance, my circumstances could put me in an especially good mood, which might make me give more weight to certain options that I otherwise would barely consider. Similarly, my mental state can make certain options seem much more appealing than others. Practically speaking, it is impossible to be aware of all of these potential influences. Since I am not consciously aware of many of these causal chains, I cannot assess whether or not the many unconscious causal chains pass through my relevant intentions, so I cannot assess whether these chains are part of my purified control. Hence, practically speaking, I cannot know how much of the weighing of an option resulted from causal chains that are part of my purified control compared to how much of the weighing results from causal chains that are not part of my purified control.

Thus, it is difficult for me to assess which pieces of the weighing of my options are part of my purified control.

Again, for the purposes of this paper, I assume there is a reasonable way to make the comparison between actual purified control and hypothetical complete control over a given decision. My point with my objection here is that there are major practical issues in assessing what all of the pieces of a given decision are that I could have control over. Thus, even if I theoretically could make a comparison between actual purified control and hypothetical complete control, I still cannot practically assess what my purified control is or what all of the causal chains are that constitute my hypothetical complete control.

iii. A Practical Problem with Assessing Purified Control over the Outcomes of Actions

In addition to the practical problems with evaluating an agent's amount of purified control over a decision, there is a major practical issue with evaluating an agent's amount of purified control over the outcome of an action. This issue arises from the infinite number of possible causal chains that are sufficient to change the outcome of an action. Note that this problem still holds even though we are only considering possible causal chains that are consistent with the physical laws of the actual world.

For example, suppose that during a baseball game, a pitcher tries to assess his purified control over the outcome of throwing a ball to his team's catcher. Once the pitcher releases the ball from his hand, an infinite number of causal chains could interfere with the ball and prevent it from reaching the catcher. A bird could fly in the way, the wind could blow the ball off course, a fan who has jumped down from the stands could

intercept the ball, a nuclear missile could wipe out the stadium, a sniper could shoot the ball out of the air, and so on *ad infinitum*.

Similarly, there are infinitely many possible causal chains that are sufficient to change the outcome of an action that pass through the pitcher's relevant intentions. For instance, there are limitless possible minor variations of the ball's trajectory that still result in the ball reaching the catcher. Also, there are infinitely many possible causal chains that are sufficient to change the outcome of an action that are downstream of an agent's relevant intentions and accord with these intentions. A bird could just miss interfering with the ball, the wind could change the path of the ball without preventing it from reaching the catcher, a nuclear missile could remain in its silo and not wipe out the stadium, and so on.

When assessing his amount of purified control, the pitcher would have to evaluate the number of possible causal chains that are sufficient to change the outcome of his pitch that either pass through his intention or are downstream of his intention and accord with it. He must also assess the probabilities that these causal chains have of occurring. He then has to compare this number of causal chains and their probabilities of occurrence to both the overall number of causal chains that are sufficient to change the outcome and the respective probabilities that these causal chains will occur.

Since the total number of causal chains that are sufficient to change the outcome is infinite, since there are infinitely many causal chains that are sufficient to change the outcome that either pass through his relevant intentions or are downstream of his intentions and accord with them, the pitcher's comparison would be between two infinite factors. The pitcher would lose any sense of perspective of the ratio between the total set

of causal chains that were sufficient to change the outcome and the set of causal chains that were sufficient to change the outcome that either passed through his relevant intentions or were downstream of his intentions and accorded with them. Therefore, the pitcher would be unable to make the comparison between these two infinitely large sets.

Therefore, agents encounter the practical problem of trying to make a comparison between these two infinite terms. Since this comparison is necessary in order for agents to assess their amounts of purified control over the outcomes of their actions, there is a major practical issue with assessing one's amount of purified control over the outcome of an action.

D. The Conceptual Problems with Assessing Purified Control over Decisions

In addition to these problems of practicality, there are further conceptual problems with assessing purified control. Conceptually, when we assess purified control, we are abstracting out causal chains that do not pass through an agent's intention as a way of removing parts of decisions and outcomes of actions that we lack control over. Since both my circumstances and dispositions at least partially depend on luck, conceptually, I must abstract them away when I assess my amount of purified control. However, it is unclear what is left of me after my circumstances and dispositions have been abstracted away. My circumstances and inborn temperamental dispositions affect how I develop. Even though they are not the only factors involved in my development, it is unclear which parts of my development we would abstract away if we tried to remove any traces of luck that stem from my circumstances or dispositions.

Recall that I consider how the luck involved in our constitutions shapes the options of our decisions, which extends beyond Nagel's original usage of constitutive luck.

For an example of the difficulty of removing all of the influence of our dispositions and temperaments, consider the case of a woman who develops depression. Suppose that we wanted to assess her purified control over a decision. We must conceptually remove her depression and the influence it had on this decision in order to assess her purified control because her depression is not part of her purified control, and her depression affects her mental state, which in turn affects which options of her decision she considers and affects how she weighs those options.

Abstracting away her depression could start with abstracting away an unpleasant daily mental state. However, depending on the severity of her depression, this mental state could have been causally sufficient to affect her development and conceivably could have changed her very approach to problem solving. For instance, suppose that she must now examine whether a problem is external to herself or a result of negative thinking caused by her depression, and suppose that she has learned that some situations are not as hopeless as they appear to be. To fully abstract out her depression, we could abstract out this development and then assess how she would have developed without it. However, this seems like a dangerous path for the logic of our reasoning because we quickly arrive at a being who consists of purely potential causal chains. Do we all have the exact same potential, or are we still different even as potential beings? This seems like a fundamentally unanswerable question since we have no way of knowing which is correct.

Furthermore, can we even say that these hypothetical beings bear any meaningful connection to the actual agents in question? For example, consider my disposition of stubbornness. My stubbornness affected my decisions to refuse to give up some beneficial projects and to not know quit some harmful ones, and these projects affected my later decisions and eventually had some overall affect on my development, so this temperamental disposition of stubbornness affected my development. This trait is a fundamental part of my self-identity, and abstracting it away would require removing a basic piece of myself and much of my development to the extent that I would doubt that what was left could be meaningfully applied to myself. Stubbornness is by no means my only inborn temperamental disposition, and without great difficulty, I could envision myself being less stubborn. However, to ask what I would have been like if I had never been stubborn or if it had never affected my development or decisions is asking about a being who no longer bears much resemblance to me.

E. The Conceptual Difficulties with Assessing Purified Control over the Outcomes of Actions

My argument for the conceptual issues with assessing our amounts of purified control over the outcomes of our actions draws upon the conceptual problems with assessing our amounts of purified control over our decisions.

The Partial Control Theory assigns degrees of moral responsibility to an agent for the outcomes of her actions based on her respective degrees of control over those outcomes. Some of the causal chains that are sufficient to alter those outcomes involve that agent's decisions. To illustrate this, consider the earlier example of the truck driver who fails to check his brakes and hits the child. The truck driver's decision to not check

his brakes affects the outcome of his attempt to stop suddenly to avoid hitting the child because if he had had his brakes checked, he would have been more likely to be able to stop sooner. If the truck driver had a small degree of partial control over that decision, then he would have a smaller degree of partial control over the outcome of his attempt to stop than he would have had if he had had a large degree of partial control over his decision to not check his brakes. Hence, an agent's degree of partial control over the outcome of an action depends on that agent's degree of partial control over the relevant decisions that were causally involved in that outcome.

Since an agent's degree of partial control over the outcome of an action depends on her degree of partial control over her decisions that were causally involved in that outcome, and since the Partial Control Theory cannot assess an agent's respective degrees of partial control over her decisions because of the practical and conceptual issues inherent in purified control, the Partial Control Theory also faces issues with assessing an agent's degrees of partial control over the respective outcomes of her actions. Thus, the Partial Control Theory encounters difficulties assessing an agent's respective degrees of partial control over her decisions and over the outcomes of her actions.

F. Conclusion of the Argument Against the Partial Control Theory Based on Purified Control

A Partial Control Theorist who defends the notion of purified control has to offer solutions to these practical and conceptual difficulties. It is not clear to me how this theorist would do so, to the extent that I struggle to come up with his or her possible responses. Thus, I conclude that purified control fails because of practical and

conceptual problems that arise if we abstract away causal chains that are both sufficient to change a decision or outcome and do not pass through an agent's relevant intentions. Thus, the Partial Control Theory assigns moral responsibility based on the flawed concept of purified control, and therefore, it is an undesirable position to defend.

VI. My Argument Against the Narrow Control Theory

    A. Introduction

The Narrow Control Theory assigns agents degrees of moral responsibility for the outcomes of their actions based on those agents' degrees of partial control over the outcomes of their respective actions, and it assigns agents degrees of moral responsibility to agents for their decisions based on their degrees of narrow control over their respective decisions. For those who think that choosing the options of one's decisions is irrelevant to one's respective degrees of control over these decisions, the appeal of the Narrow Control Theory is that it bases its assessment of an agent's degrees of moral responsibility for her decisions based on her degrees of narrow control, not partial control, over those respective decisions. Recall that narrow control only applies to decisions and not to outcomes, and narrow control does not account for whether an agent chose the options of her decisions.

In order to provide a more complete description of how the Narrow Control Theory assigns moral responsibility for decisions, I will now outline a rough metric for assessing an agent's degree of narrow control over a decision. Her degree of narrow control over a given decision only depends on two factors. First, it depends on how many causal chains that are sufficient to how she deliberates between her options also pass through her relevant intention to deliberate between her options. The second factor is

what the probabilities of occurring are of these causal chains that are sufficient to change her deliberation and also pass through her intention.

B. Practical Problems with the Narrow Control Theory

As I discussed in my argument against the Partial Control Theory, there are practical problems with assessing an agent's degree of partial control over the outcome of an action because of the difficulties comparing the set of causal chains that are sufficient to change the outcome of an action and either pass through the agent's relevant intentions or are downstream of these intentions and accord with them to the total set of causal chains that are sufficient to change the outcome because both sets are infinitely large. These practical difficulties with assessing an agent's degree of partial control over the outcome of an action also apply to the Narrow Control Theory since the Narrow Control Theory, like the Partial Control Theory, assigns degrees of moral responsibility to agents for the outcomes of their actions based on their respective degrees of partial control over those outcomes.

In addition to these practical issues, the Narrow Control Theory faces the more fundamental problem that there is arbitrariness in its moral judgments of agents for their decisions and for the outcomes of their actions. I will first demonstrate that the Narrow Control Theory allows arbitrariness into its moral judgments when assigning degrees of moral responsibility to agents for their degrees of narrow control over their respective decisions. I will apply this arbitrariness in assessing an agent's decisions to demonstrate the arbitrariness in assessing the outcomes of an agent's actions. This will allow me to argue that the Narrow Control Theory is an undesirable position to hold.

C. How Arbitrariness Entering into Moral Judgments of Agents' Decisions

To illustrate how the Narrow Control Theory allows arbitrariness to enter into its moral judgments of agents' decisions, consider an embellished version of Nagel's example of circumstantial luck. Suppose that two people were born in Germany in the early 1900's. In 1930, one person, Lucky, and his family move to America, while the other, Unlucky, and his family do not. Suppose that Lucky and Unlucky have nothing to do with the moving or lack of moving. Lucky experiences the Great Depression in the United States while Unlucky experiences the economic and political upheaval in Germany as Hitler gains power.

The circumstances that Lucky and Unlucky face are different, and that difference has nothing to do with what Lucky and Unlucky did. For instance, Unlucky had no causal role in the rise of the SS, but he faces a choice between joining the SS or the other ranks of Nazi soldiers, which entails at least indirectly if not directly participating in the Holocaust; the option of fleeing Germany; the option of trying to save potential victims of the Holocaust; and so on. I think it is quite clear that the options that Unlucky must choose between are drastically different from those options that Lucky must choose between, and I take it to be clear that neither Lucky nor Unlucky chose most of the options of their decisions.

If Unlucky's decisions are evaluated solely on the basis of what he decides, and if these moral judgments do not factor in whether he chose the options of his decisions, then luck sneaks into these moral judgments that are based on narrow control. Luck sneaks into this case because it is more likely that Unlucky will be judged as worse than Lucky since Unlucky faces worse options than Lucky, and he therefore has more opportunities to decide to act in a worse manner than Lucky. Hence, Unlucky is likely to face worse

moral judgments than Lucky because of the luck in the circumstances in which their decisions are made.

I argue that the luck in an agent's circumstances that enters into the moral responsibility assigned by the Narrow Control Theory cannot be avoided by examining what Lucky would have done if he were in Unlucky's circumstances and vice versa. This is because it is impossible to definitively say what a person would be like if he or she had grown up in a different set of circumstances. We could try to remove each person's circumstances from consideration and assess the moral responsibility of Lucky and Unlucky. However, we currently do not have the ability to do this, and it is unlikely that we would be able to in the future because removing a person's circumstances from consideration quickly leads to the earlier problems discussed in section **V D** of trying to assess a hypothetical being who we can say little about, especially when we consider that it is unclear what connection this being has to Unlucky or Lucky.

Note that I am not excusing Unlucky's actions if he ends up becoming a Nazi soldier. Rather, I am saying that there is some degree of luck whether he faces the choice of becoming a Nazi, so only assigning moral responsibility for decisions based on an agent's degree of narrow control over that decision becomes a case of circumstantial luck because the agent's moral responsibility for his decision depends on circumstantial factors that he did not choose to have.

The luck in our circumstantial factors is not always as grandiose as it is in the previous example, but it is still pervasive on a daily basis since our circumstances constrain many of the options of our decisions. Suppose Sleepy is deciding whether or not to drive to work after being kept awake all night by a neighbor's barking dog.

Sleepy's office is too far away for her to walk, and Sleepy would be taking a small risk of causing harm if she drives to her office because her reaction times will be slower due to her fatigue. On the other hand, she risks being fired if she skips work, and she needs her job to support her family. Suppose that her job requires her to be on-site, so she cannot work from home. Hence, Sleepy's decision is a choice between two risky options, and she did not choose what either of these options is. If we assess her degree of moral responsibility for what she decides to do by assessing her degree of narrow control over this decision, then we do not account for her lack of choice over the options of this decision. This is a case of circumstantial luck because Sleepy's moral responsibility for her decision depends on circumstantial factors that she did not choose to have.

For both examples, if we assign moral responsibility to Lucky, Unlucky, or Sleepy for their decisions based on their respective degrees of narrow control without considering whether they chose any of their options, then we allow arbitrariness to enter into these moral judgments because of the constraints that the circumstantial factors of the agents impose on the options of the agents' respective decisions. More generally, even if we have some narrow control over our decisions, there is still luck in what options we are choosing between. Thus, the Narrow Control Theory's assignment of moral responsibility for decisions based on an agent's degrees of narrow control over those respective decisions involves luck in an agent's circumstances. Because of this luck, the Narrow Control Theory's moral judgments of agents' decisions involve moral arbitrariness.

D. Applying This Arbitrariness in Moral Judgments of Decisions to Demonstrate the Arbitrariness in Moral Judgments of the Outcomes of Actions

An agent's degrees of partial control over the outcomes of her respective actions depend to some extent on her degrees of control over her respective decisions that are involved in those outcomes. For the Narrow Control Theory, an agent's degrees of control over her respective decisions involve her degrees of *narrow* control over her respective decisions. Hence, an agent's degree of partial control over the outcome of an action partially depends on her degree of narrow control over her relevant decisions.

Under the Narrow Control Theory, an agent is assigned a degree of moral responsibility for the outcome of an action based on her degree of partial control over that action, and she is assigned a degree of moral responsibility for a decision based on her degree of narrow control over that decision. Therefore, an agent's moral responsibility for the outcome of an action partially depends on her degree of moral responsibility for her relevant decisions that were involved in that outcome. Since there is arbitrariness in the agent's degrees of moral responsibility for her decisions, and since her degrees of moral responsibility for the outcomes of her actions depend on her degrees of moral responsibility for her decisions, there is therefore arbitrariness in an agent's degrees of moral responsibility for the outcomes of her respective actions under the Narrow Control Theory.

E. Conclusion of My Argument Against the Narrow Control Theory

Thus, there is arbitrariness in the Narrow Control Theory's moral judgments of an agent's decisions and actions. Since the central question the Control Theory tries to answer is how to assign moral responsibility in a non-arbitrary manner, and since the Narrow Control Theory has arbitrariness in its assignment of moral responsibility, I conclude that the Narrow Control Theory is an undesirable position to defend for a

Control Theorist or for anyone else who does not want arbitrariness in their moral judgments.

VII. Conclusion of The First Half

A. The Implications of the Failure of the Variations of the Control Theory

In the first half of this paper, I have described the dilemma of how to assign moral responsibility in instances involving moral luck, and I then described the three variations of the Control Theory, which I take to represent the intuitive responses to this issue. The Control Theory preserves the notion that we should only be held morally accountable for what we can control, and these theories strive to accomplish this by assessing our control over our decisions and over the outcomes of our actions. However, the Complete Control Theory fails to ever assign moral responsibility to limited beings like ourselves, the Partial Control Theory encounters serious practical and conceptual issues when assessing an agent's degree of partial control, and the Narrow Control Theory contains arbitrariness in its moral judgments. Therefore, each of these forms of the Control Theory fails to satisfactorily assign moral responsibility in a non-arbitrary manner in instances involving moral luck.

§2. THE SECOND HALF

I. Moving Beyond Control

All three forms of the Control Theory fail to assign moral responsibility in a satisfactory manner in cases involving moral luck. This shows that control is not a good basis for assigning moral responsibility in instances involving moral luck. Our circumstances and constitutive factors at least partially constrain the options of all of our decisions, and the outcomes of all of our actions at least partially depend on causal

factors that do not pass through or accord with our intentions. Thus, the inability to assign moral responsibility in cases involving moral luck is a serious issue since all of our decisions and the outcomes of all of our actions are to varying degrees instances of moral luck. Given the pervasiveness of luck, and given that control is not a satisfactory basis for assigning moral responsibility, how should we non-arbitrarily assign moral responsibility in instances involving moral luck?

In this half of my thesis, I will begin by presenting one answer to the problem of assigning moral responsibility in instances involving moral luck, which I will call the Responsibility by Ability Theory. While there are benefits to this solution, its requirements are too great for limited human agents such as ourselves. I will discuss our perception of ourselves as agents with a causal influence and will apply it in order to modify the Responsibility by Ability Theory and form the Ability and Contribution Theory. The Ability and Contribution Theory is a virtue theory, and I will discuss the aspects of this virtue theory that are relevant to moral luck. Finally, I will conclude by arguing how the Ability and Contribution Theory addresses the issue of assigning moral responsibility in cases of moral luck while simultaneously ensuring that people who have been harmed do not have to endure the blows of fortune alone.

II. The Responsibility By Ability Theory

A. How We View Ourselves As Part of the World's "Causal Web"

Part of understanding the prevalence of luck in the outcomes of our actions and in our decisions is recognizing that we, as limited humans, are part of the world's vast causal web. I use the metaphor of a web to illustrate how interconnected the causal chains are that are sufficient to change our decisions or change the outcomes of our

49

actions. The outcomes of my actions can affect other people's decisions or the outcomes of their actions, and the actions and decisions of others can change my decisions and the outcomes of my actions.

Because of the interconnectedness of these causal chains, I cannot point to a decision or the outcome of an action and say that I am the only thing that is causally responsible for bringing it about. Since we are not the sole causes of the outcomes of our actions or our decisions, if we assign moral responsibility for our decisions or for the outcomes of our actions, then we are assigning moral responsibility to agents for causal chains that do not pass through our intentions.

B. A Better Method

A better method for assigning moral responsibility is to assign agents a universal responsibility to help other agents who have been harmed based on each agent's ability to help others. I call this the Responsibility by Ability Theory. The idea for this theory is that blaming agents for unintentionally causing harm through the outcomes of their actions is unfair because we cannot point to a decision or outcome and say that that agent caused it in the absence of other causal influences. Instead, that agent's decision or outcome was partially due to causal chains that did not pass through the agent's relevant intentions. Rather than trying to assign blame for an agent's decisions or outcomes of actions, and rather than trying to utilize control or a partial degree of control, which both fail, the Responsibility by Ability Theory gives agents a "global" obligation to help other agents who have been harmed by their own decisions and actions and by those of others, which inevitably involve varying degrees of luck. This duty is "global" in the sense that

it does not depend on whether the agent was causally involved in the creation of the harm. All agents who are equally able to help have an equal moral responsibility to help.

i. An Example

I will now use an example to further explain the Responsibility by Ability Theory.[7] Suppose that Swimmer, an able-bodied person who knows how to swim, is standing on a pier next to a lake in which a child is drowning. Suppose that the child fell into the lake because a nearby person, Clumsy, accidentally tripped into the child, and the collision knocked the child into the lake. According to the Responsibility by Ability Theory, Swimmer's reaction should be to focus on rescuing the drowning child since Swimmer knows how to swim and is therefore able to help the child. Swimmer should not focus on discerning whose fault it is that the child is drowning in the first place.

According to this theory, focusing on whose fault it is that the child fell into the water would be a mistake because if we assume that Clumsy did not intend to trip into the child, then it is clear that some of the causal chains that were sufficient to lead to Clumsy's tripping into the child did not pass through Clumsy's relevant intentions, so assigning moral responsibility to Clumsy would entail assigning him moral responsibility for causal chains that did not pass through his relevant intentions. Even if Clumsy did intend to knock the child into the water, the outcome of Clumsy's action still partially depended on luck, and furthermore, punishing Clumsy would not help save the drowning child. That is not to say that Clumsy would be excused of blame if he intended to knock the child into the lake, but rather, this theory places a primary emphasis on helping those

---

[7] I base this example on Peter Singer's famous example from his paper "Famine Affluence, and Morality."

who have been harmed before considering the blameworthiness of the agents who are causally involved in the creation of the harm.

Suppose that Nonswimmer, who does not know how to swim, is also standing on the pier. Since Nonswimmer is unable to swim and rescue the child, he does not have a moral obligation to jump into the water and try to rescue the child. However, if Nonswimmer has some other way of helping, then he has an obligation to try to carry out that form of help. For instance, Nonswimmer might be able to signal to nearby people who do know how to swim that the child is drowning and needs their assistance.

Before I go any further, I have been using the term "moral responsibility," and I need to precisely define how the Responsibility by Ability Theory uses this term in order to fully explicate the nature of this theory's global responsibility.

C. Redefining Moral Responsibility

The term "moral responsibility" as it is typically used is an assessment of an agent's fault or blameworthiness for a decision or for the outcome of an action based on that agent's degree of control over that decision or outcome. In the earlier half of this thesis, I argued that assessing an agent's degree of control fails, so I do not include this baggage of blameworthiness based on control in my definition of moral responsibility. Instead of being based on control and fault, "moral responsibility" as the Responsibility by Ability Theory uses it is based on acknowledgement of a moral obligation. The moral responsibility of an agent to help someone is an acknowledgement that one is able to help and that one has an obligation to help if possible. Because the Responsibility by Ability Theory recognizes the contingency in our decisions and actions, there is no blame if an

agent tries to help someone but fails because of causal chains that do not pass through his relevant intentions.

D. Summary of the Responsibility by Ability Theory

According to the Responsibility by Ability Theory, all agents have a moral responsibility to help others, and this responsibility is based on each agent's ability to help.

This theory recognizes that human agents are part of an interconnected set of causal chains, which means that we cannot claim sole causal ownership of our decisions or of the outcomes of our actions because there are causal factors that are external to our intentions that also affect these decisions and outcomes. Rather than trying to assign moral responsibility for decisions and outcomes that depend on contingencies, the Responsibility by Ability Theory resolves the problem of assigning moral responsibility in instances involving moral luck by assigning agents a universal duty to help those whom they can.

Since human agents are physically incapable of helping everyone, the Responsibility by Ability Theory needs to establish a method for agents to prioritize whom to help. According to the Responsibility by Ability Theory, agents should prioritize helping people based on the severity of the harm suffered. Those who have suffered greatest degree of harm should be helped first.

The amount of responsibility that an agent has to help others depends on that agent's ability to help others. While that ability to help others is contingent on luck, since there is no blame associated with moral responsibility, agents are not blamed or praised for the contingent extent of their ability to help others.

III. My Modification of the Responsibility by Ability Theory

A. Introduction

I will now present my objections to the Responsibility by Ability Theory, but I will not argue that this theory is entirely flawed beyond redemption. Instead, I will argue that the global responsibility to help others when one is able can be good, but this universal responsibility needs to be tempered because it is too strong for limited human agents as it currently stands. Thus, I will modify this theory to form a final theory of how to assign moral responsibility in instances involving moral luck, and I will conclude by defending this new theory.

B. Why the Responsibility by Ability Theory Needs to Be Modified

A crucial issue with the Responsibility by Ability Theory is that the global responsibility it places on human agents to always help others when they are able is too strong of an obligation for limited beings. Because of the large extent of harm and suffering in the world, if our primary obligation were to help others, we could easily be obligated to devote our entire lives to assisting others. Initially, the idea of devoting ourselves to helping others may not sound bad, but I will elaborate on what this devotion to others would entail to show why it is too strict of an obligation.

The following includes a list of potential obligations that an agent might have if that agent were always doing everything in her power to help others. Any wealth this agent accumulates that goes beyond what she needs to sustain herself would have to be donated to charities that provide food, shelter, and other such resources for others because this agent would be able to help these people who needed this food and shelter. That might not sound bad, but consider the following case.

Suppose that the agent is choosing between two job offers. She would love the first job and would hate the second, and although both provide enough money for her to be able to live comfortably, the second pays significantly more than the first. If the work done for these two jobs helps people equally but the second pays much more than the first, then the agent would have to take the second job because the extra money she earned could be used to provide more help. This would still be true even if the agent would derive much less personal satisfaction from her work if she took the second job than if she took the first. She might not even be able to take either job offer if she could spend all of her time helping others by working in various charity organizations.

Even if an agent worked flat out to help others, the amount of help required would always exceed the amount of help that the agent could offer. Hence, that agent could never take time to develop personal talents or focus on achieving personal goals since the time spent pursuing those personal interests and goals could always be spent working to help others, unless those talents and goals turn out to be the most effective way for her to help others. If that agent chose to focus on cultivating her personal development of talents and chose to spend time pursuing personal goals, this agent would be ignoring her moral obligation to help others.

While spending time and energy helping others is usually positive, I argue that we should not be morally obligated to help others to the extent that we cannot focus on any personal development. Thus, the global responsibility to help others assigned by the Responsibility by Ability Theory is too strong because it does not account for the meaningfulness of leading interesting lives and achieving our personal goals.

On the other hand, I do not wish to advocate an entirely selfish position and hold that we have no global duty to help others. For instance, I do not wish to argue that the development of our personal interests and goals should always take priority over any duty to help others. Consider the earlier case of Swimmer, who is standing on a dock next to a lake in which a child is drowning. I do not wish to advocate a position in which Swimmer ignores any global duty to help the drowning child simply because Swimmer has a personal goal of not getting his clothes wet.

The question therefore arises of how to limit the global responsibility to help that is mandated by the Responsibility by Ability Theory so as to preserve an agent's ability to live an interesting life and ability to try to achieve his or her personal goals, while simultaneously not morally allowing an agent to entirely devote himself or herself to personal pursuits at the expense of any global obligation to help others.

IV. The Responsibility by Ability and Contribution Theory

A. Introduction

In this section, I will first propose how to limit an agent's global responsibility. In doing so, I will present and argue for the Responsibility by Ability and Contribution Theory, or the Ability and Contribution Theory for short. Then, I will present the Ability and Contribution Theory as a virtue theory and will describe the elements of it that are relevant to assigning moral responsibility in instances involving moral luck.

B. Limiting Our Global Responsibility to Help Others

i. Introduction

My proposal for how to limit out global responsibility consists of two distinct parts. In the first part, I will argue for accepting personal moral responsibility for our

decisions, actions, and outcomes as a means for us to focus on where to direct our efforts to help others. Although this personal responsibility offers us a way to channel our efforts to help others, it alone does not limit our responsibility to help others, which leads to the second part of my proposal. In the second part, I will describe how our personal and global responsibility to help others is balanced with the weight of personal fulfillment.

ii. Part One of Limiting Our Overall Responsibility to Help Others

a. A Second View of Ourselves as Agents

The first part of limiting our global responsibility lies in more fully exploring the earlier idea of how we, as human agents, view ourselves, given the extent that luck influences our decisions and the outcomes of our actions. Although the Responsibility by Ability Theory is correct in saying that we do view ourselves as inextricable pieces of a causal "web" in which we cannot point to a given decision or outcome and say that we alone were entirely responsible for it, we also do not view ourselves as purely "passive" beings. Despite the contingency in the outcomes of our actions and in our decisions, we do not consider ourselves to be "passive" in the sense that we have no effect on the world because in addition to this first view of ourselves as pieces of an interconnected set of causal chains, we also have a second view ourselves as agents with a causal influence. We are limited agents in that we alone are not sufficient to bring about our decisions or the outcomes of our actions. Nonetheless, we still view ourselves as making a causal impact on the world.

I argue that at least part of the reason why we see ourselves as agents, albeit limited agents, comes from the way that we experience the world as discrete individuals.

We are "discrete" in that we experience the world via our sensory perceptions in a separate manner from each other, and we do not have something like a hive-mind in which we share thoughts and experiences. One important part of this discreteness is that we separately experience the causing of harm to others. Even if we are helping or hurting others as part of a group, we still have the actual experiences as individuals, and we cannot literally experience this as a group, even if we identify with that group.

For instance, members of a firing squad might collectively kill a prisoner, but each member of the firing squad experiences contributing to this harm as a discrete agent.

A second important part of our discreteness is that we have individual goals and desires. In the case of groups, a group of individuals might share a collective set of goals. However, those goals could have differing levels of importance and mean different things to different individuals within the same group. Also, the manifestation and realization of those goals could take different forms for different individuals. My point is that even if I am a member of a group and that group has goals, I have my personal take on what exactly those goals mean, what it would take to realize them, and what the achievement of them would signify.

In the context of the firing squad, each member has a different set of goals that motivates him or her to participate in the execution. For example, one of them might participate in order to make money while another does so because he wants to ensure that prisoners are punished for their crimes. Even if several members of the firing squad had the same general goal that motivated them to participate in the execution, like wanting to ensure justice, this general goal of ensuring justice means somewhat different things for each agent. Thus, we have personal goals and desires, even in the context of groups.

One final important part of our discreteness as agents is that we feel morally responsible for some decisions, actions, and outcomes of actions because we view ourselves as causally contributing to them, and we experience these feelings of moral responsibility in a way that goes beyond any general moral responsibility to help others whom we have not harmed. This view holds even if we recognize that these decisions, actions, and outcomes are influenced by other causal factors beside ourselves. Consider again the case of the truck driver and the child. The truck driver will still probably feel at least partially responsible for killing the child if the child darts in front of his truck and the truck driver is unable to stop in time, even if the truck driver has just had his brakes checked and is not acting negligently.

b. Applying This Second View to Choose Whom to Help

Since there are too many people for us to help at once, we need some way of choosing whom to help. Given that we need some way of choosing whom to help, given that using control to assign moral responsibility fails, and if we want to maintain this view that we make a causal difference in the world, then I argue that the most viable option is to apply this second conception of ourselves as discrete agents to decide where to direct our efforts when deciding whom to help.

c. Adding a Second Moral Obligation to Help

In addition to our global moral obligation to help others when we are able, I will argue that we also have a second moral obligation, which is to accept personal moral responsibility for our decisions and for the outcomes of our actions, even though these decisions and outcomes are dependent on contingent causal factors that do not pass through our relevant intentions. I name this modified version of the Responsibility by

Ability Theory that adds the moral obligation to accept this personal moral responsibility the Responsibility by Ability and Contribution Theory, or the Ability and Contribution Theory for short. The Ability and Contribution Theory uses the same notion of moral responsibility that the Responsibility by Ability Theory for the Ability and Contribution Theory's general moral obligation to help others, and this notion of global moral responsibility does not carry the concept of blame. However, the Ability and Contribution Theory's notion of moral responsibility for our personal moral responsibility for our decisions, actions, and outcomes of actions does retain some of the notion of blame. In section **IV D**, I will describe how blame is partially retained by this personal moral responsibility.

d. The Justification of Personal Moral Responsibility

I begin my justification of taking personal moral responsibility for our decisions, actions, and outcomes by noting our belief that we can exert a causal impact on a decision, action, or outcome. We also experience ourselves as making an impact on certain decisions, actions, and outcomes. This experience may be mistaken because we are often unaware of all of the other causal factors involved in these decisions, actions, and outcomes, which can lead to a mistaken belief that we had a causal impact when we actually did not.

For an example of this mistaken belief, consider the following example of a belief that amounts to a superstition. Suppose that a person keeps track of the number of people who die each day, and suppose that she notices that when she wears a certain outfit, the worldwide deaths that day decreases. If she accepts personal moral responsibility for her actions, then she might mistake correlation for causation and feel morally responsible for

causing the number of deaths to increase when she decides not to wear that particular outfit, even though what she wears has no actual causal effect on the number of deaths on a given day. This example shows the fallibility of our assessments of which decisions, actions, and outcomes we have a causal impact on.

Nonetheless, even though our first-person perception of making a causal impact can be mistaken, we cannot simply set aside this personal perception and somehow adopt an "objective" third-person perspective in which we can clearly see all of the factors involved in each causal chain. As humans, we are constrained to view the world from a first-person perspective, and we have no way of identifying which exact pieces of a decision, action, or outcome we are causally responsible for. Because of these factors, if we assume that we make some causal impact on the world, then I argue that we do not have any viable way of specifying which decisions, actions, and outcomes are "ours" and which we should take personal responsibility for other than to use our intuitions to assess when we make a causal impact. Once we have established that have made a causal impact, then we bear a degree of personal moral responsibility for that decision, action, or outcome based on the extent to which our intentions reflect favorably or unfavorably on our character. Our intentions reflect on our character because they arise from our character.

For an example of when we would not accept personal moral responsibility, suppose that I open a door to a classroom, and to my surprise, I walk in on Villain taking people hostage. Suppose that Villain, who has a gun pointed at a group of hostages, becomes distracted upon seeing me. This distraction leads some hostages to try to escape, and suppose that several hostages are severely wounded in the escape attempt.

My causal contribution to the harming of these hostages may have been quite high if the hostages would not have tried to escape if I had not opened the door. However, I would not bear moral responsibility for this bad outcome because my intentions in opening the door to the classroom do not negatively reflect on my character.

e. Applying Personal Moral Responsibility to Focus Our Efforts

We can use our personal moral responsibility for our decisions, actions, and outcomes in order to prioritize whom to help, and it is necessary to make some prioritization because we are unable to help everyone. All other concerns being equal, if we are choosing whether to help someone we harmed or someone we did not, and if both victims received a roughly equal amount of harm, we should focus on helping the person whom we harmed. Later in this paper, I will outline a more complete metric of how to decide whom we should help.

Note that this personal responsibility to help does not replace our global responsibility to help. Instead, this personal responsibility offers us a way to choose whom to help when we cannot help everyone. We still have a global obligation to help others, and our personal moral responsibility for our decisions, actions, and outcomes offers us a way to decide whom to help when we have caused harm and all other deciding factors are approximately equal.

Although this personal moral responsibility for our decisions, actions, and outcomes offers us a way to decide whom to help, the Ability and Contribution Theory as it currently stands does not yet limit our overall amount of moral responsibility to help others. This need to limit our overall personal and global responsibility to help others

leads me to the second part of my argument for how the Ability and Contribution Theory tempers our responsibility to help others.

iii. Part Two of Limiting Our Overall Responsibility to Help Others

As a virtue theory, the Ability and Contribution Theory argues for the ethical value of personal fulfillment and uses this to counterbalance our overall responsibility to help others. The Ability and Contribution Theory does this within the framework of the broader philosophical tradition first established by Aristotle and continued by philosophers like Bernard Williams that an agent's personal well-being and flourishing rightfully have a special importance to that agent. This flourishing includes duties to help others, but these duties do not require the agent to neglect his own well-being. For an example of this tradition, consider the follow passage from Bernard Williams: "There can come a point at which it is quite unreasonable for a man to give up, in the name of the impartial good ordering of the world of moral agents, something which is a condition of his having any interest in being around in that world at all."[8]

Consider again the way that we view ourselves as discrete agents. From the perspective of the Ability and Contribution Theory and its broader philosophical framework, our goals and desires are important to us as discrete individuals. "Desires" here goes beyond our basic needs of food, water, and shelter. You and I might share goals or desires, and we may even have similar thoughts about what those goals and desires mean, but my goals matter to me in a way that is distinct from how they matter to you. I can respect your goals and even make your goals my own, but I still view your goals through my personal perspective, and you still view your goals through your

---

[8] Williams, Bernard. "Persons, character and morality," in *Moral Luck: Philosophical Papers* (Cambridge: Cambridge University Press, 1981), 14.

personal perspective. Hence, each of our sets of goals and desires are discrete and personal.

The fulfillment of our personal goals can be a source of meaning in our lives.[9] It is important to examine whether these are desirable or worthy goals to have, and if they are worthy, then personal development by means of satisfying these goals can be a valid reason not to always help others. Note that the Ability and Contribution Theory considers living an interesting life and the personal development gained by fulfilling one's goals to have ethical significance. This means that the balancing of our global and personal responsibilities to help others with our personal fulfillment of our goals as part of our personal development is a balance between two sets of ethical concerns. This distinguishes the balance made by the Ability and Contribution Theory from the balance described by Susan Wolf in her paper "Moral Saints," in which Wolf describes weighing ethical considerations of duty with non-ethical considerations of personal development.[10]

The reason why the Ability and Contribution Theory limits our personal and global responsibilities to help others is that it is impossible for limited beings such as ourselves to help everyone, and if we only had these responsibilities to help others, we would be required to devote our lives solely to helping others. It is undesirable for an ethical theory to have an insurmountable gap between what is morally required and what is possible for the relevant agents to attain. Morally requiring more than what an agent can accomplish risks alienating that agent from that morality. Here, I draw upon Peter Railton when he argues, "[p]art of constructing an answer [to the question of why we

---

[9] For example, see Bernard Williams' arguments in *Utilitarianism: For and Against*.
[10] Wolf, Susan. "Moral Saints," *The Journal of Philosophy* 79-8 (1982): 419-439.

should be moral] is a matter of showing that abiding by morality need not alienate us from the particular commitments that make life worthwhile…"[11]

Since the Ability and Contribution Theory balances our concern for others with our concern for personal development and fulfillment, this theory avoids the issue that the Responsibility by Ability Theory faces of alienating agents from morality by having too large of a gap between its moral requirements and what it is possible for agents to accomplish.

iv. Elaborating on This Argument

a. Introduction

Here, I need to further explain my reasoning behind one step in my argument for accepting personal moral responsibility for our decisions, actions, and outcomes of actions. The step in question is the progression from viewing ourselves as causally responsible for certain decisions, actions, and outcomes to accepting moral responsibility for them. To explain my reasoning for making this step, I will first consider two objections to making this progression.

b. The First Objection

First, an objector could argue that simply because we do view ourselves as having a causal influence on the world does not mean that we should view ourselves as having this causal influence since we might actually just be deluded and not have any actual influence.

c. My Response

---

[11] Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality," *Philosophy and Public Affairs* 13-2 (1984): 164.

I would reply that I assume for the purposes of this paper that we do make some causal difference in the world. Much of what we do is dependent upon contingent factors, but that does not preclude us from having any influence. For instance, the options of our decisions may be greatly constrained by our circumstances, but we can still exert a causal influence by trying to expand the options that we are choosing between. Similarly, the outcomes of our actions also largely depend on contingencies, but we still seem to have some effect on these outcomes, and for the purposes of this paper, I assume that this sentiment is correct. I cannot point to a decision, action, or outcome and say which pieces of it I controlled since the notion of control fails. However, I can point to certain decisions and outcomes and say that I made some difference to them, even if I cannot say with exact certainty what that difference was.

For example, a pitcher surely has some influence over the outcome of a pitch, even if he cannot say assess exactly what his impact was on the pitch. If we did not think that we had any effect at all over a decision, action, or outcome of an action, then we would probably not call it "our" decision, action, or outcome. This leads the objector to his or her second objection.

d. The Second Objection

The objector could then argue that just because we view a decision, action, or outcome as our own does not mean that we are justified in believing that we made a significant enough causal impact on it to be morally responsible for it, especially if we consider that "our" decisions, actions, and outcomes largely depend on contingencies.

My response to this objection hinges on the Ability and Contribution Theory being a virtue theory, so I will now present this virtue theory in order to both respond to this objection and further define the Ability and Contribution Theory.

C. The Ability and Contribution Theory as a Virtue Theory

I would like to first note that I will not develop the Ability and Contribution Theory as a virtue theory in its entirety. Developing a full-blown virtue theory goes far beyond the scope of this paper. This means that I will not discuss all of the virtues and vices that are considered by this theory. Instead, I will focus on how the virtuous person, according to the Ability and Contribution Theory, makes decisions, acts, and responds to the outcomes of her actions. In doing so, I will describe how the virtuous person approaches the global and personal responsibilities and the value of personal fulfillment of the Ability and Contribution Theory. That is not to say that these are the only moral considerations of the virtuous person, but to reiterate, I am only focusing on the aspects that are relevant to assigning moral responsibility in situations involving moral luck.

Earlier, I argued that we each have our own set of personal goals and desires. Our goals and desires are realized by means of our decisions and our actions. We view some decisions, actions, and outcomes of actions as our own. Our decisions about what our goals should be, our decisions of how we should try to achieve our goals and satisfy our desires, what we actually do to achieve them, which ones we actually achieve, and what the consequences are of doing so all affect the development of our moral character. Since the Ability and Contribution Theory is a virtue theory, the development of moral character takes on an ethical significance.

The Ability and Contribution Theory holds that it is morally admirable to care about one's causal contribution to the harming of others, despite the contingencies involved in this harm. There may be a biological basis for the tendency to take responsibility for one's causal contribution to harm. Nico Frijda suggests that this tendency is widespread in humans[12] and suggests that it "renders the individual careful in her or his dealings with fellow individuals, and diminishes chances of reprisals. Guilt can be seen as a mechanism for the regulation of interpersonal behaviour…"[13] If feeling guilty for causally contributing to harm makes us more cautious and aware during our interactions with others, which in turn facilitates our cooperation with other people, then it makes sense that we would view it as morally admirable to care about our causal contribution to the harming of others.

Therefore, under the framework of the Ability and Contribution Theory, we take personal moral responsibility for our actions and decisions both because these affect the development of our personal moral character via our goals and desires and because a concern for our causal contribution to the harm of other reflects favorably on our moral character.

Thus, I refute the second objection that we are not justified in feeling morally responsible for our decisions, actions, and outcomes because these decisions, actions, and outcomes reflect on our moral character and because it is morally admirable to care about our causal contribution to the harming of others.

---

[12] Frijda, Nico H. "The Place of Appraisal in Emotion." *Cognition & Emotion* 7-3 (1993): 357-387.

[13] Frijda, "The Place of Appraisal in Emotion," 368.

To provide a more complete picture of the Ability and Contribution Theory, I will discuss further aspects of it as a virtue theory. Since decisions, actions, and outcomes of actions are partially matters of chance, the virtuous person as defined by the Ability and Contribution Theory is not necessarily the person who has the best decisions, actions, or outcomes. This would mean that the virtuous person is virtuous only because she is lucky enough to have her decisions, actions, and outcomes turn out well.

Rather, the virtuous person uses every ounce of influence she can muster to try to make a positive impact on the world. This "positive impact" includes the helping of others when she can. When her decisions, actions, or outcomes backfire, the virtuous person recognizes the contingency involved. Nonetheless, she also recognizes that she has causally contributed to the harming of another, and she accepts personal moral responsibility for her decisions, actions, and outcomes. Part of accepting this moral responsibility is acknowledging that she has causally contributed to the harming of another being and acknowledging that she has an obligation to amends with the injured party.

These amends could include helping the victim recover, assisting the victim's closest relatives, and showing remorse for the harm caused. Even if both parties recognize the luck involved in the harm, the virtuous person's remorse shows her recognition that in spite of her constraints as a limited agent, she still has enough of a causal impact to harm others, even if that harm is unintentional.

When another agent accidentally harms the virtuous person, the virtuous person recognizes the luck involved in the creation of that harm, and because of this, the virtuous person is more willing to forgive that other agent.

The virtuous person does not necessarily devote her entire life solely to helping other people, for she also lives an interesting and personally satisfying life. She cultivates her talents and strives to achieve her goals. Her moral character reflects what her goals are, how she decides to pursue those goals, whether or not she takes personal moral responsibility for her actions and the outcomes of her actions, what she decides to do in response to what the outcomes of her actions are, and whether she helps others when she is able to do so.

D. An Objection about Blame

An objector to the Ability and Contribution Theory could ask if this theory can assign moral blame to agents who deliberately cause harm since the successful causation of that harm depends on causal chains that do not pass through the agents' relevant intentions.

I would respond that the Ability and Contribution Theory still retains the concept of blame in its assignment of personal moral responsibility, although this concept is modified to account for the luck inherent in our constitutions and our circumstances. Recall that the Ability and Contribution Theory assigns us degrees of personal moral responsibility for our decisions, actions, and outcomes based on the extent to which our intentions reflect favorably or unfavorably on our character. Since an agent's intentions arise from her character, we can gauge an agent's character by examining her intentions.

There is luck involved in what our character is because of the luck involved in our temperamental dispositions. These temperamental dispositions are not the only factors that compose our character, but they do affect our character. For instance, suppose an agent has the temperamental disposition of being ungenerous. This lack of generosity is

70

not the only factor that is relevant to that agent's character, but this trait does make up a part of his character.

Similarly, there is luck in how our character develops because the luck in our circumstances constrains the options of our decisions. For example, the ungenerous agent might not have the option of using his personal wealth and abilities to lessen his ungenerous temperament by helping others if he barely has the means to support himself.

"Blame" usually carries the connotation of control. Hence, since the Ability and Contribution Theory assigns degrees of moral responsibility based on the extent to which our intentions reflect positively or negatively on our character, and because of the luck involved in what our character is and how it develops, the personal moral responsibility assigned by the Ability and Contribution Theory loses much of its concept of blame.

However, the personal moral responsibility assigned by the Ability and Contribution Theory does not lose its entire concept of blame because it is mistaken to say that we have no influence over our character and its development. Some of the factors involved in our character and its development are matters of luck, but other factors depend on our causal impact on certain decisions, actions, and outcomes of actions.

As I argued earlier, we cannot point to any one decision, action, or outcome and specify which exact pieces of it we control. However, we can recognize when we have some impact on a decision, action, or outcome, even if we cannot precisely say what the exact nature of that impact is. Similarly, we may not be able to specify exactly which pieces of our character and its development we caused, but when we examine the

aggregate of all of our decisions and actions that affected our character, we can still recognize that we have some impact on our character and development as a whole.

Since we have some impact on our character and its development as a whole, the Ability and Contribution Theory retains some connotation of blame in its concept of personal moral responsibility. "Blame" for the Ability and Contribution Theory is a way of notifying an agent that his character is flawed and signifies to the agent that he ought to use whatever causal impact he has to try to improve his character. Since the Ability and Contribution Theory is a virtue theory, assigning blame is a way of motivating an agent to try to make his character like the ideal character of the virtuous agent. Because of the luck involved in improving one's character, the Ability and Contribution Theory cannot hold malicious agents to be at fault if their respective characters doe not end up improving in spite of the efforts they make, but this theory can still encourage them to try to improve their characters. Thus, blame under the Ability and Contribution Theory is not an assessment of an agent's fault, but this theory can use this forward-looking sense of blame to motivate malicious agents to try to make what impact they can to improve their characters.

V. Juggling All of Our Responsibilities

This model of the virtuous person suggests how we can limit our global moral responsibility to a more manageable level. According to the Ability and Contribution Theory, we need to balance three competing claims for our efforts. We have a global responsibility to help those whom we are able to and a personal moral responsibility for our decisions, actions, and outcomes of actions. Thus, we must balance helping others out of our global duty to help with our personal moral responsibility for the outcomes of

72

our decisions, actions, and outcomes. Since we cannot help everyone, we must prioritize.

When deciding whom to help, we must consider what role we played in the relevant

victim's harm, and we usually should focus our efforts on righting the wrongs that we

contributed to. When the harm and other ethical factors in two cases are roughly equal,

we have more responsibility to address the case in which we causally contributed to the

harm.

That said, there are exceptions in which it is clear that we must prioritize our

global responsibility over our personal responsibility. For instance, if I knock someone

over while rushing to jump into a lake during an attempt to rescue a drowning child, it

seems intuitively clear that I should focus on saving the child before I turn my efforts to

helping the person I knock over. However, my broader point is that in cases where the

stakes are roughly equal, I should prioritize helping those whom I have harmed.

In situations in which we have not caused harm, our primary responsibility to help

those who have been harmed based on our global responsibility is based on the following

factors and potentially others: the severity of the harm suffered by the victim, our ability

to help, what personal sacrifice we would make if we helped, and our relationship to the

victim. To reuse an earlier example, this metric of assessing responsibility means that

even though Swimmer did not causally contribute to the harm suffered by the drowning

child, Swimmer still has a moral duty to help the child because the harm is severe,

Swimmer is able to help, and Swimmer would not be making a great personal sacrifice by

jumping into the water.

While both of these responsibilities to help others are important, it is an

undesirable position to hold for the Ability and Contribution Theory if it requires us to

devote our entire lives to helping others. Hence, we balance our global responsibility to help others and our personal responsibility for our decisions, actions, and outcomes by recognizing the ethical importance of living interesting and fulfilling lives and striving to achieve our goals.

VI. Conclusion

There is no easy and exact formula for juggling our personal and global responsibilities to help others and the value of striving to achieve our goals, but if we can find this balance, then the Ability and Contribution Theory provides a way of resolving the issues that arise for assigning moral responsibility in cases of moral luck. The Ability and Contribution Theory assigns moral responsibility to agents for their decisions, actions, and outcomes, which are all often matters of luck, and it assigns a moral responsibility to help others when the agents are able to, and the amount that these agents are able to help is also largely a matter of luck. However, there is no blame attached to this global moral responsibility to help others.

One advantage of the Ability and Contribution Theory being a virtue theory is that agents are only blamed or praised for how their intentions reflect on their moral character. A second advantage of the Ability and Contribution Theory being a virtue theory is that it offers a way to limit the obligations of agents to help others because it accounts for the ethical value of living a fulfilling and interesting life. There is no hard and fast set of rules that systematically assess exactly when it is better to focus on one's personal goals and when to make the extra effort, even when inconvenient, to help another.

I conclude this paper by explaining why I have been focusing so much of my discussion of the Responsibility by Ability Theory and the Ability and Contribution Theory on helping others. I make this emphasis on helping others because of the problem of what I call "victim luck." "Victim luck" arises when we assign little to no moral responsibility to agents who unintentionally cause harm. Assigning these agents little to no responsibility protects them from being assigned an arbitrary amount of moral blame based on an unlucky outcome. However, if we purge the moral responsibility from the agent for this unintended and unlucky outcome, then the victim of this unintentional harm is now stranded in that no one is morally obligated to help that victim.

The Responsibility by Ability Theory protects victims from being stranded like this by having the global responsibility to help those who have been harmed. The Ability and Contribution Theory goes beyond the Responsibility by Ability Theory and gives agents a personal moral obligation to help those whom these agents have harmed, even if that harm was unintentional. This ensures that the victims of bad luck are not left to deal with the ill effects of chance. At the same time, these agents with this personal moral responsibility are only blamed for how their intentions behind their decisions, actions, and outcomes reflect unfavorably on their character, and these agents are not blamed for the actual unlucky decisions, actions, and outcomes. Thus, the Ability and Contribution Theory ensures that no one has to endure the blows of fortune alone.

Works Cited

Frijda, Nico H. "The Place of Appraisal in Emotion." *Cognition & Emotion* 7-3 (1993): 357-387.

Nagel, Thomas. "Moral Luck." In *Mortal Questions*. Cambridge: Cambridge UP, 1979.

Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13-2 (1984): 134-171.

Singer, Peter. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1-3 (1972): 229-243.

Smart, J. J. C. and Bernard Williams. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 1973.

Williams, Bernard. "Moral Luck," in *Moral Luck: Philosophical Papers,* 20-39. Cambridge: Cambridge UP, 1981.

Williams, Bernard. "Persons, character and morality," in *Moral Luck: Philosophical Papers,* 1-19. Cambridge: Cambridge University Press, 1981.

Wolf, Susan. "Moral Saints." *The Journal of Philosophy* 79-8 (1982): 419-439.

Wolf, Susan. "The Moral of Moral Luck." *Philosophic Exchange* 31 (2001): 4-19.

Appendix A

Variation of Control Theory	Type of control that is the basis of assigning moral responsibility for the outcomes of actions	Type of control that is the basis of assigning moral responsibility for decisions
Complete Control Theory	Complete Control	Complete Control
Partial Control Theory	Partial Control	Partial Control
Narrow Control Theory	Partial Control	Narrow Control

Definitions:

Complete Control over the Outcomes of Actions: An agent has complete control over the outcome of his action if and only if all of the possible causal chains that are sufficient to change the outcome pass through the agent's relevant intention and all subsequent causal chains that are downstream of the agent's intention and are sufficient to change the outcome accord with the agent's relevant intention. The possibility referred to in these "possible causal chains" is possibility consistent with the physical laws of the actual world. An agent's relevant intention for the outcome of an action is the intention to bring about a specific outcome.

Partial Control over the Outcomes of Actions: An agent has partial control over an outcome if and only if some of the possible causal chains that are sufficient to change the outcome pass through the agent's intention and some of the other causal chains that are sufficient to change the outcome do not.

Complete Control over Decisions: An agent has complete control over a decision if and only if all of the possible causal chains that are sufficient to change what her options are in that decision or change the weight she assigns to those options pass through her relevant intentions. There are two relevant intentions. The first intention is the intention to deliberate over one's options and make a decision. The second is to choose the array of options that one is deliberating between.

Partial Control over Decisions: An agent has partial control over a decision if and only if some, but not all, of the possible causal chains that are sufficient to change what options she chooses between or change the weight she assigns to those options pass through her intention to choose an option or her intention to deliberate amongst her options.

Narrow Control over Decisions: Narrow control only applies to decisions and differs from partial control over decisions because narrow control does not involve an agent's amount of choice of what the options of a decision are. An agent has narrow control if and only if at least some of the possible causal chains that are sufficient to change what weight an agent assigns to her options pass through her relevant intention of deliberating amongst her options. Deliberating amongst her options involves assessing and weighing her options and deciding which option to take.