


ATTENUATING THE PROBLEM OF MORAL LUCK

How Moral Luck Either Does Not Exist or Does Not Create a Paradox for Our Moral
Systems.

by

Ivan Bock



*Thesis presented in fulfilment of the requirements for the degree of Master of
Philosophy in the Faculty of Philosophy at Stellenbosch University*

Supervisor: Dr Tanya De Villiers-Botha

March 2020

Declaration

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Date: March 2020

Copyright © 2020 Stellenbosch University

All rights reserved

Abstract

In the 1970's Bernard Williams and Thomas Nagel formally introduced the problem of moral luck. Moral luck can be understood as the seeming paradox between the control principle and the moral judgements we confer on others. The control principle states that an agent can only be held morally responsible for an action if, and only if, said agent had control over it. Contrary to this, we often do judge people for many things out of their control. The consequences of our actions, the circumstances we find ourselves in, and our own characters are all things we either wholly or partially lack control over, yet, we hold people responsible for these things. This lack of control and accompanying moral judgements are what is referred to as "moral luck", and we must therefore either conclude that agents cannot be held responsible for their actions, or that we can hold people responsible for things out of their control, both being framed as problems. Here, I will attempt to give a solution to the problem of moral luck. I will do this by discussing some of the most influential writings on the problem, each section of the thesis focusing on a separate type of luck, addressing the mistakes philosophers have made while inferring that moral luck is real. I will argue that each type of moral luck only exists because we have misunderstood important concepts, and once we revise our conception of control, agency, and responsibility the problem of moral luck disappears. In particular, I will argue that 1) Resultant luck is only a problem because we are focusing on the consequences of actions rather than the intentions of the agent, 2) Circumstantial luck is only a problem because we fallaciously transfer the luck of the world onto moral considerations, and 3) Constitutive luck is only a problem because we are misapplying the concept of control onto character. The thesis will also include a section on relevant implication if I am successful in solving the paradox, including theoretical and practical implications. My conclusion will thus be, contrary to the thesis of moral luck, that we can still hold agents morally responsible without having to reject the control principle, however, this is only possible if we accept revisions to important moral concepts.

Abstrak

In die 1970's het Bernard Williams en Thomas Nagel die probleem van Morele Geluk formeel aangevoer. Morele Geluk kan verstaan word as die skynbare paradoks wat ontstaan tussen die beheerbeginsel en die werklike morele oordele waaraan ons ander onderwerp. Die beheerbeginsel bepaal dat 'n agent slegs moreel verantwoordelik bepaalde optrede gehou kan word indien, en slegs indien, daardie agent die nodige beheer oor daardie optrede gehad het. In teenstelling hiermee, beoordeel ons mense dikwels vir dinge wat buite hul beheer is. Die gevolge van ons optrede, die omstandighede waarin ons onself bevind en ons eie geaardhede is alles dinge waaroor ons, in geheel of gedeeltelik, nie beheer het nie, maar tog hou ons mense verantwoordelik vir al hierdie dinge. Hierdie gebrek aan beheer en die gepaardgaande morele oordele wat nogtans gemaak word, word 'Morele Geluk' genoem, en ons moet klaarblyklik dus óf tot die gevolgtrekking kom dat agente nie verantwoordelik gehou kan word vir hul optrede nie, óf dat ons mense verantwoordelik kan hou vir dinge wat buite hul beheer is. Albei opsies word as probleme voorgehou. Ek probeer hier om 'n oplossing vir die Probleem van Morele Geluk te ontwikkel. Ek doen dit deur enkele van die mees invloedryke werke wat die probleem aanvoer te bespreek en dan die foute wat elke filosoof begaan het toe hulle afgelei het dat Morele Geluk 'n werklikheid is, uit te wys. Elke afdeling van die tesis fokus op 'n aparte soort Morele Geluk, soos ontwikkel in die literatuur. Ek argumenteer dat elke soort Morele Geluk slegs oënskynlik bestaan te danke aan die feit dat die betrokke filosowe die ter sake konsepte telkens verkeerd verstaan het. Ek wys dan uit dat sodra ons ons opvatting van “beheer”, “agentskap” en “verantwoordelikheid” uitklaar en hersien, verdwyn die probleem van Morele Geluk. Ek voer spesifiek aan dat 1) “Resulterende geluk” slegs ontstaan indien ons op die gevolge van optrede eerder as op die agent se voornemens, konsentreer; 2) “Omstandigheidsgeluk” ontstaan slegs indien ons alledaagse geluk verkeerdelik aan morele oorwegings toeskryf; en 3) “Konstitusionele geluk” slegs 'n probleem is omdat ons die konsep van beheer op die verkeerde wyse op “geaardheid” toepas. Ek spreek ook die teoretiese en praktiese implikasies van my oplossing van die oënskynlike paradoks aan. Ek kom tot die gevolgtrekking dat ons agente steeds moreel verantwoordelik kan hou vir hulle optrede sonder om die beheerbeginsel op te skort. Dit is egter slegs moontlik indien ons die nodige hersienings aan die ter sake belangrike morele konsepte aanbring.

Table of Contents

| | |
|--|----|
| Chapter 1: Introduction..... | 1 |
| 1.1 Introduction..... | 1 |
| 1.1.1 Thomas Nagel and Moral Luck..... | 2 |
| 1.1.2 Four Types of Moral Luck..... | 4 |
| 1.1.3 Structure of This Thesis..... | 6 |
| 1.2 Groundwork..... | 6 |
| 1.2.1 The Metaphysical Position of Naturalism..... | 7 |
| 1.2.2 Natural Determinism..... | 9 |
| 1.2.3 The Concept of Responsibility..... | 9 |
| 1.2.4 The Control Principle..... | 10 |
| Chapter 2: Moral Luck and its beginnings..... | 11 |
| 2.1 Introduction..... | 11 |
| 2.1.1 Bernard Williams..... | 11 |
| 2.1.2 Thomas Nagel..... | 14 |
| 2.2 Conclusion..... | 24 |
| Chapter 3: Resultant Luck..... | 26 |
| 3.1 Introduction..... | 26 |
| 3.2 The Problems of Including Resultant Luck..... | 27 |
| 3.3 The Virtues and Merits of Resultant Luck..... | 41 |
| 3.4 Conclusion..... | 50 |
| Chapter 4: Circumstantial Luck..... | 52 |
| 4.1 Introduction..... | 52 |
| 4.2 Not Moral, but Epistemic..... | 52 |
| 4.3 Conclusion..... | 62 |
| Chapter 5: Constitutive Luck..... | 65 |
| 5.1 Introduction..... | 65 |

| | | |
|------------|--|-----|
| 5.2 | What We Cannot Control | 66 |
| 5.2.1 | Creating the Intuition | 66 |
| 5.2.2 | Mental States and Voluntary Control..... | 68 |
| 5.3 | Moral Agents and Identity..... | 74 |
| 5.4 | Conclusion..... | 85 |
| Chapter 6: | Putting It All Together | 87 |
| 6.1 | Introduction | 87 |
| 6.2 | Agents, Control, and Responsibility..... | 88 |
| 6.3 | Conclusion..... | 92 |
| Chapter 7: | Implications..... | 93 |
| 7.1 | Introduction | 93 |
| 7.2 | Theoretical Implications..... | 93 |
| 7.2.1 | Responsibility..... | 93 |
| 7.2.2 | Moral Agency | 94 |
| 7.2.3 | Free Will | 95 |
| 7.2.4 | Actions and Consequences..... | 96 |
| 7.2.5 | Negligence and Obligation..... | 98 |
| 7.3 | Practical Implications | 99 |
| 7.3.1 | Criminal Justice..... | 99 |
| 7.3.2 | Egalitarianism | 100 |
| 7.4 | Conclusion..... | 101 |
| Chapter 8: | Conclusion | 102 |
| References | | 108 |

Chapter 1: Introduction

1.1 Introduction

It could be argued that one of the defining features of the human race is our capacity for moral behaviour. We have moral values, think about our actions in moral terms, and base our sense of what people deserve on moral considerations. Of course, we are not the only creatures that display some sense of moral behaviour. Psychologist Frans de Waal (2013) argues that primates have, at the very least, a proto-morality. That is, they display patterns of behaviour indicative of what we might call “moral reasoning” or “moral behaviour”. Furthermore, many other creatures display the type of behaviour and social organisation that are precursors to the type of morality that we have. The part of our morality that makes us unique is our capacity for inspecting, understanding, and revising of our moral principles, beliefs, and practices. We aim to understand the principles that drive behaviour and judgement and hope to be consistent with our judgements. It is this inspection and evaluation of our moral systems that make our morality so unique. Nevertheless, it is this same inspection and evaluation that has led to many of the greatest problems and debates between people regarding our moral practices. I would argue that our morality is one of the greatest achievements of our species, but at the same time, it is also one of the greatest curses¹.

Although the evaluation of our moral systems has arguably led to great growth in human societies, it has also paved the way for criticism and problems. I am not saying that criticism is necessarily bad, but it has, at times, led to problems that threaten to undermine our very moral practice. There have been many of these problems over the ages. From the Euthyphro dilemma, all the way to the modern debates on free will and responsibility, we as a species have been overwhelmed by questions pertaining to morality and its difficulties. If we are to continuously better our world and moral systems, we must address these issues and ensure that they do not undermine the institution of morality itself. One such problem is moral luck, which is what I will be addressing in this thesis. I wish to discuss the problem of moral luck as found in the works of Thomas Nagel. According to Nagel, our moral systems generally hold that we should not punish agents for those actions that are not under the control of the agent in question. Yet, those same systems also allow for the moral

¹ “Curse” might be a strong word. Perhaps challenges is better suited. Nevertheless, there are pros and cons to our moral evaluations.

judgment of such actions. This seems to lead to a contradiction at the heart of our morality—we should not punish people for actions that are beyond their control, but we do (and cannot help but do) so all the time. In this thesis, of the four types of moral luck identified by Nagel, I will argue that two forms do not exist, and a third does not create a paradox. In addition, I will argue that the fourth is already subsumed by the other three. It is through this that I hope to show that our moral systems can remain intact, and that moral luck is not, in fact, the problem for them that Nagel makes out to be.

1.1.1 Thomas Nagel and Moral Luck

In 1979 Thomas Nagel introduced his conception of moral luck to the field of moral philosophy in his paper simply titled, *Moral Luck*. Nagel's argument was that our moral systems presuppose a Control Principle², that is, that they state agents may only be judged for actions within their control. We tend to assume that people cannot be praised or blamed for something that they are not responsible for. A simple example would be if someone knocked a glass of water out of your hand and it then spilt on another person. In such a case we would conclude that you are not morally responsible for this result, as you could not have done otherwise. The fact is simply that the water spilt was not under your control. Nagel then goes on to explain that, contradictorily, our moral systems also *do* seem to judge people for actions outside their control (I will mention specific examples below). These two opposing principles create a paradox within our moral systems. Nagel explains that we as agents frequently lack the kind of control necessary for responsibility, due to factors outside of our control. He refers to such factors in so far as they pertain to morality as "moral luck". "Luck" here refers to any part of an agent's life that seems to influence moral judgments related to her, but which is also outside of the agent's control. Nagel identifies four different types of luck that hinder our responsibility; namely, Resultant luck, Circumstantial luck, Constitutive luck, and Causal luck (Nagel, 1993). His argument has been extremely influential, and many different arguments have been put forward, both for and against his view.

On the one hand, there are arguments that deny the actuality of moral luck (e.g. Rescher, 1993; Richards, 1993; Zimmerman, 1993; Moore, 1997; Wolf, 2013), and arguments that accept moral luck

² This control principle is a result of two things. First, intuitively we do not want to ascribe blame to those actions outside our control, since doing so will result in a morality which is not agent-centred. Secondly, and more relevant to this discussion, the ethical work done by Kant has had a profound effect on morality. He argued that the agent was perfectly rational, and able to control all the moral facets of his/her life (Kant, 2011). Kant's view of what morality is, and his view on agency, has had a lasting effect.

and revise our views on morality accordingly (e.g. Adams, 1985; Browne, 1992; Fischer, 1998). The debate is still ongoing and unresolved, and I believe this is due to theorists using too broad an approach or a too specific an approach to addressing the issue of moral luck. Most of these theorists either only deal with one type of luck or try to solve all the types of luck in the same way. What is needed is a systematic evaluation of the problem of moral luck in which each type of luck is assessed individually in order to determine if it actually is a factor that constitutes a violation of the Control Principle.

The importance of this problem is usually confined to the philosophical world due to most people seeing it as more of an academic issue, or an issue that is only important to those who have the time to muse over abstract theoretical issues. I, however, believe the issue to be highly important for both theoretical and practical reasons. On the theoretical side, this issue lies at the foundation of almost all of our theories about morality and ethics. How are we to build ethical systems if we are not sure to what extent we can rightfully hold one another responsible? It is not possible to talk about ethics, justice, and law if we have not yet established a foundation for human agency and moral responsibility. This issue must first be dealt with before we can move on to more specific problems. The problem of moral luck also has important consequences for normative ethics. The way we see the agent, and the level of control which we deem necessary (or possible) for any given agent, will influence the ways in which we can rightfully judge that agent.

Before we can decide what is right and what is wrong, we must first know which aspect of moral action should be the focus of moral judgment, actions or intentions. If we decide the focus should be on actions and their consequences (such as the actual events that take place, rather than what someone meant to do) we could possibly reject the Control Principle. If intentions are the focus (i.e. what the agent *means* to do is relevant, not how things actually turned out) we need to discuss how the environment is relevant, if at all, to moral judgments. The chapter on Resultant luck will make clear how the discussion on moral luck might determine the answer to these issues. For now, it is enough to understand that the Control Principle might serve to exclude the results of our actions from moral judgements and thus understanding the issue of moral luck will influence our understanding of the Control Principle. On the practical side then, our view of human agency and the meta-ethical claims we hold to be true directly inform how we treat one another on both a personal and a societal level. The way we react to one another when we do something wrong (or right) is based on how we believe we should be held responsible, given our metaphysical commitments. If we want to make sure we treat each other ethically, we must first understand the problem of moral luck. We must understand which factors are relevant when judging an agent. Most notably, our judicial systems are also based on metaphysical and normative commitments such as these. Usually, the way we treat a criminal or

a negligent person is derived from a view of agency that presupposes the possibility of control and a lack of luck. Cases such as a negligent driver who hits a child unintentionally or cases of manslaughter all need to be metaphysically and ethically informed before we make our judgements. We need to determine whether it is right (metaphysically and ethically) to treat people in this way. To make sure the laws of our countries are as just as possible, we must make sure we understand the problem of moral luck. Due to all the different fields and practices that would be influenced by the outcome of this evaluation, I believe moral luck as an issue should be of the utmost importance; research and work within this field will serve an important role in re-evaluating our ethical and political systems.

1.1.2 Four Types of Moral Luck

As noted above, the problem Nagel created for our moral practices was formulated in terms of four different types of moral luck. Moral luck encompasses any factor that hinders the agent's control over his/her actions in such a way so as to take away his/her responsibility for those actions. These factors include, among other things, how you are raised, the environment you find yourself in, and how that environment both influences and responds to your actions. As already mentioned, the four types of moral luck described by Nagel are Resultant luck, Circumstantial luck, Constitutive luck, and Causal luck (Nagel, 1993).

Resultant luck is the type of luck relating to how things turn out, or what follows from your actions. An example would be someone who successfully murders a person as opposed to someone who only attempts to murder a person but fails. Both intended to do the same thing, but we treat them differently due to factors beyond their control which led to the one successfully killing someone and the other not. These factors can include things like different types of victims (strong or weak individuals), the type of care or treatment they receive afterwards (such as hospital care or the lack thereof), and the weapon used and how serious the wound was.

The second type of luck, Circumstantial luck, is the type of luck that has to do with where you find yourself (your circumstances). An example of Circumstantial luck would be a person who happens to live in Nazi Germany who acts per Nazi norms. It may be that had this person not lived in Nazi Germany, he would never have had the intentions or desires that caused him to act as he did. We treat the two versions of this person differently even though the only reason *they* act differently is due to external factors beyond *their* control. The circumstances you happen to find yourself in, to a large extent, determine the challenges and fortunes that you face. Why, then, in these instances, do we judge you when you fail or succeed with regards to your intended actions?

The third type of luck, Constitutive luck, is the type of luck which determines who you are; arguably, our personalities, wills, desires, and capabilities are to a large extent formed by the world around us. We may have no direct control over who we are, and yet we hold one another responsible for actions that came about due to our personality, even though we may not be responsible for it. Some people have extremely hard lives, some people have privileged lives, and this shapes the sort of people they are. Why should we say a thief is a bad person when, for the sake of argument, his behaviour and attitudes were caused by a hard life full of struggle? We praise people who support and assist others, but, again, for the sake of argument, why should we praise them when such behaviour and attitudes were caused by a life of privilege where they had the opportunity to develop such a character? Examples such as these are meant to demonstrate that the people we are, and therefore the moral judgements accrued to us, are influenced by many factors out of our control.

Finally, Causal luck is the type of luck that has to do with how one's actions are caused by the circumstances that precede them. In this regard, if determinism³ is true, then all our actions are caused by events that happened before us. These events are beyond our control, so why hold ourselves responsible for actions that inevitably result from them?

As mentioned above, these four types of luck hinder our classic view of moral responsibility and challenge our basic moral and political systems. The classic view of moral responsibility states that agents can only be held responsible for those things under their control. Each of the types of luck capture a different way in which agents lack the necessary control for moral responsibility.

If Nagel is right, and we cannot account for moral luck in its various forms when it comes to our actions, we as humans seem to lose responsibility. This is a problem for two reasons: first, it challenges our current conception of people and responsibility (i.e. that we are free and in control of our actions) by creating conflict between our moral systems and the natural environment (since the environment seems to diminish the amount of control agents have). Secondly, moral luck eliminates one of the important motives we have for according moral desert (just punishment and reward) and hence undermines the need for ethical systems. The purpose of this thesis is to evaluate each of these forms of luck and argue that they are i) attributable to faulty reasoning and hence do not actually exist; or, ii) exist but they do not actually hinder our responsibility.

More specifically, I will argue that Resultant luck only arises due to our focusing our moral judgments on the wrong aspect of human behaviour when it comes to morality; we should not focus on the *actual*

³ Determinism is the metaphysical view that every event is sufficiently and necessarily caused by events that came before it. For more on this see section 1.2.3.

outcomes of our actions but on the *intended* actions in any given situation. Circumstantial luck will be dealt with by showing that it does not hinder our responsibility as agents can only be judged within the context of given circumstances, the basic idea being we can only punish people for those things they demonstrably intended. Finally, Constitutive luck will be argued against by showing it only arises due to a category mistake. Here we assume the existence of two different categories (the person and their personality) when there is, in fact, only one (the agent). The final form of Luck (Causal luck), will not be individually assessed in this thesis, since I will argue that it is fully addressed by dealing with the other forms of luck. My overall conclusion will be that moral luck does not pose a true threat to moral responsibility, but in order to conclude this we will still require revision of some concepts as the problem arises due to the incorrect understanding of important concepts.

1.1.3 Structure of This Thesis

Before I discuss the four different types of moral luck, I will briefly introduce the important concepts that are needed to address the problem. These include the conception of “responsibility” I will employ in this thesis, as well as the metaphysical views presupposed in the thesis, as well as their justifications. In Chapter Two, the problem of moral luck as put forward by Thomas Nagel will be explained in detail. In each subsequent chapter, I will focus on one of the first three types of moral luck identified. I will use the most influential writings on the topic of moral luck to discuss the issue and attempt to solve it. As mentioned above, the fourth type, Causal luck, will not be discussed as a separate instance of moral luck, since I believe it is fully addressed by addressing Circumstantial luck and Constitutive luck. I will then dedicate a small section of the thesis to synthesising the three different chapters into one coherent solution. After all of this has been done, I will conclude my thesis with a discussion of the implications of what has been said, and how our moral and political systems will need to react if my thesis is taken to be correct.

1.2 Groundwork

In order to better understand the debate surrounding moral luck, I will lay the groundwork with some clarifications regarding some preliminary assumptions and important concepts that are central to understanding moral luck, or that are specific to my thesis. These will include clarifying the metaphysical assumptions that inform my discussion, the type of responsibility discussed here, and the central concepts that I will employ.

1.2.1 The Metaphysical Position of Naturalism

An important presupposition of this paper is that of philosophical naturalism. By “naturalism” I mean the position that there is no first philosophy, and that philosophy should be consistent with our best science (Ritchie, 2014:1 – 27). These two statements on their own do not fully explain what is meant by naturalism, and they still leave the term vague. Since the focus of this paper is moral luck, I do not have the time or space to explain all the nuances of this view; rather, I will briefly explain what I mean in broad terms and will then explain why I accept this view and why it is relevant here.

Firstly, by taking up a naturalist position I reject all supernatural claims. I do not do this because of an unsupported bias, rather I do this for pragmatic reasons. When I say I reject all supernatural claims, I mean all claims that appeal to forces beyond nature, claims that contradict the natural order, and claims that are unjustified or unfalsifiable/unverifiable. The reason for this is twofold:

a) The scientific method “works”:

As explained above, by taking up a naturalistic position I also claim there is no “first philosophy”. This means I accept there is no guaranteed method for gaining absolute knowledge. This, however, is not an issue for the naturalistic philosopher since they are not concerned with absolute knowledge. The important thing for the naturalistic philosopher is that knowledge be useful and pragmatic. One of the best ways to demonstrate this usefulness is by looking at the ability of scientific theories to make predictions. This allows us to test theories made in science by looking at whether predictions come true. The predictions made by science are not only useful in testing the theories, but they also demonstrate that the main purpose of science is to explain and then create theories that can be used. The scientific method has led to theories which have been used to make predictions which were later shown to be true (a deflationary sense of “true”). One example of this is how our theory of the Big Bang and cosmological inflation⁴ allowed us to predict a largely uniform universe which should have background radiation at the edges. These predictions were later verified and thus demonstrate the usefulness of science. A second example of how the scientific method is more useful than others comes from the field of medicine. If we compare the efficacy and influence of medical science to that of alternative medicines, we clearly see much more progress towards better medical understanding and efficacy within medical science than in the alternative (MHNRC, 2015).

It might be true that science can never give us absolute knowledge, but this is not needed. The scientific method has produced knowledge which has indeed helped our understanding of the

⁴ For a thorough explanation of the Big Bang theory and the predictions made by it you can read “Big Bang” by S. Singh (2004).

universe. It has given us knowledge which we can use, test, and challenge. The claim that science cannot give us absolute knowledge is of no concern. It is undeniable that science has given us some form of knowledge which is at least “true” in one sense, and most certainly useful. The issue of a naturalistic approach to moral luck will become pertinent when discussing issues pertaining to agency and determinism.

b) Supernatural claims are “useless”:

When I say supernatural claims are useless, I do not mean they do “nothing”. Clearly, such claims make a large impact on how we see the world and how we shape our lives. The problem is, these claims are not valuable in gaining any useful knowledge. Since these claims lie outside the natural world, and thus beyond our understanding, we are unable to test them. Many different people claim to understand the supernatural, yet what we see is a plethora of different ideas and contradictory theories all having the same minimal amount of support without any way of testing these theories and deciding which ones are true. Simply having an idea of something does not constitute understanding it; what we see within the world of supernatural “experts” is a bunch of people who have decided on a personal view of the supernatural without any objective evidence backing this up. Not only this, there is also no way of knowing how these supernatural claims affect our world. If we approach the topic of free will, for instance, through a naturalistic lens, we can discuss the biological condition of humans and how the natural world interacts with us. By doing this we can create models of free will which are testable and knowable. If we do not approach this problem naturalistically, we can only appeal to unknowable and untestable “properties” (for example: souls and spirits), which we have no way of testing (and thus knowing if we are right or wrong). More than this, we have not really explained anything, since these properties lie beyond our world, and thus our understanding. If philosophy is to produce knowledge which both helps our understanding and makes it possible to test these theories, we need to approach philosophy from a naturalistic worldview.

It is the combination of these two points which inform my naturalistic worldview. Discussing the issue of moral luck through the lens of naturalism should result in a much more fruitful and useful theory of agency and moral responsibility. For these reasons, I will proceed with this paper in a naturalistic way. All claims will either be supported by the scientific approach or will not contradict the natural order⁵.

⁵ The main reason for including such an emphasis on the topic of naturalism is a combination of my own frustration with people giving answers to problems that are not given proper explanation and justification, and papers such as *How to Apply Molinism to the Theological Problem of Moral Luck* (Hartman, 2014) and *Moral Dilemmas and Moral Luck*

1.2.2 Natural Determinism

In this paper, the metaphysical view of determinism will be accepted. This includes all types of determinism such as physical and biological determinism. The distinction between physical and biological determinism (Lewontin, 1982) is not an important one for this paper; both apply to humans in the relevant ways⁶. What is meant by determinism is the view that each event is determined in a fixed way by preceding events, or, stated otherwise: for every event, there is some antecedent state, related in such a way that it would break a law of nature for this antecedent state to exist yet the event not to happen. This applies to both physical events, and to events such as choices⁷.

The reasons for accepting determinism are multiple, but the most important reason comes from our most recent scientific discoveries and models. There is large agreement within the scientific world that our universe is determined, and even in cases such as quantum physics, we do not have any reason to believe our world (as we live in it and experience it) is not determined. It is true that there seems to be indeterminism found on the quantum level, however, indeterminism does not remove the threat of moral luck since indeterminism (which is another form of luck) also does not allow for responsibility.

1.2.3 The Concept of Responsibility

The type of responsibility discussed in this paper is the type which is relevant to morality. Although responsibility can be defined as “being the cause of” or “origin” of a given event, this is not the type of responsibility relevant to this paper. Responsibility here will be defined as moral responsibility, or the type of responsibility which designates some form of “desert”. Desert denotes the idea of how a person ought to be treated depending on their actions, either eliciting praise or blame. Note that this praise or blame is directed towards the agent because he is morally responsible for the action, not just

(Dougherty, 2004) which approach the problem of moral luck in a very different manner than here (from a Theological approach).

⁶ The reason for the distinction not being important here is that both capture the idea that we cannot be anything other than that which we are. Universal determinism states that the way our world works (due to physics and time) means that every event is caused in a fixed way by preceding effects (including humans). Biological determinism is the idea that who we are is based on our biology (strong nature, weak nurture) and thus also limits the possible choices of humans.

⁷ The distinction between physical events and choices is only a matter of clarification. Since I hold a naturalistic view of the world, I believe mental states (such as choices) are in fact physical states.

for deterrence or reinforcement (as it is found in the contagion theory of desert; Pereboom (2006) being a good example of this). It is important to understand responsibility as it is conceived in the former instance (moral) and not the latter (causal), since moral luck has no relevance to the latter. If responsibility is used in any other way than explained here, it will be made explicit.

1.2.4 The Control Principle

A very important concept to understand for this discussion is the Control Principle. This is the principle which states that for someone to be held morally responsible for their actions, they need to have sufficient control over said actions. This means the agent must be free from coercion or external control, and the agent's actions must follow from his/her own desires, wills, and wants. It is this principle that is threatened by moral luck. The most famous formulation of the Control Principle is found in the ethical works of Immanuel Kant when he describes a type of pure agency where the agent is only judged for those actions which comes from his/her "pure" intentions (Kant, 2011).

Most of our ethical systems that aim to give responsibility to agents assume the Control Principle. It is for this reason that moral luck is such a problem for morality; it seems to negate one of the most important foundations of our moral systems. The importance of the Control Principle will be made clear once we start our discussion on the different types of moral luck.

Now that I have both explained and justified the concepts and presuppositions of this thesis I will move on to the formal introduction of moral luck, and then the discussions of the different types.

Chapter 2: Moral Luck and its beginnings

2.1 Introduction

Chapter 1 In this section of the paper, I will discuss moral luck as it is found in the works of Thomas Nagel and Bernard Williams. The issues that they identify will first be discussed in general to create a clear picture of the problem. After this introduction, the four types of moral luck will each be addressed separately. Each will be explained as they are described by Nagel.

As discussed in the introductory chapter, the four types of moral luck are Resultant luck, Circumstantial luck, Constitutive luck, and Causal luck.

2.1.1 Bernard Williams

I wish to start my discussion with Williams's paper. In his paper simply entitled *Moral Luck* (1981), he uses the story of an aspiring artist to demonstrate a problem with our moral systems and practice in general. Williams starts the discussion with a description of how we usually see morality, i.e. as unique. As he explains, we are all aware that the world is full of good and bad luck, but, somehow, we see morality and those actions that we tend to make moral judgements about as being immune to these features (Williams, 1981:36):

The thought that there is a kind of value which is, unlike others, accessible to all rational agents, offers little encouragement if that kind of value is merely a last resort, the doss-house of the spirit. Rather, it must have a claim on one's most fundamental concerns as a rational agent, and in one's recognition of that one is supposed to grasp, not only morality's immunity to luck, but one's own partial immunity to luck through morality.

The intuition that we have about morality is that it pertains to those actions that we can rightly pass judgement on, due to the culpability of the agents performing those actions. This intuition seems to be widespread and is strengthened by various religious teachings, as well as Kant's famous work on morality. We believe that moral responsibility should not be influenced by the unlucky happenings of the moment and that agents should only be punished for those things that are within their control. This "Control Principle" can be formulated as follows:

Control Principle: An agent should only be judged for those actions within the agent's control.

A corollary of this principle is that, usually, any action that is outside of the control of the agent should not be liable to moral judgement.

The intuitive appeal of this principle is demonstrated when we consider the case of two people (X and Y) who both spill coffee on someone: X, while carrying some coffee, does not take proper notice of her environment. While focussing on something else, perhaps her phone, X carelessly trips over a desk and spills her coffee on her colleague sitting there. Now, I believe most people would agree that X is at least responsible for being negligent. She is guilty of not taking proper note of her surroundings and should, therefore, be held morally responsible for neglect in this regard. In a second case, Y is also carrying a cup of coffee, but is concentrating on his surroundings and trying his best not to spill any coffee. However, as he passes his colleague, an office chair is pushed into his path and causes him to spill coffee on his colleague. Although the colleague might initially react in the same way in both instances, most would agree that the second case is an accident and that Y is not to be blamed. Both X and Y spill coffee, yet they are treated differently. The only difference is that while it was within X's control to not spill her coffee (by being more vigilant), it was not within Y's control since he could not avoid the chair.

As much as this moral difference seems intuitive, some philosophers have also explicitly argued that it must be a fundamental part of morality. Immanuel Kant (2011:17) famously stated:

Even if by some particular disfavour of fate, or by the scanty endowment of a stepmotherly nature, this will should entirely lack the capacity to carry through its purpose; if despite its greatest striving it should still accomplish nothing, and only the good will were to remain (not, of course, as a mere wish, but as the summoning of all means that are within our control); then, like a jewel, it would still shine by itself, as something that has its full worth in itself. Usefulness or fruitlessness can neither add anything to this worth, or take anything away from it.

Kant thus argues that even if a person is unable to act, or if their actions do not turn out as intended, their *will* remains fully in their control, and we are therefore able to hold them responsible for what they intended. This free "will" is an important aspect of Kant's philosophy, and indeed our own conceptions of what it means to be human and what it means to be moral. How can we judge, blame, reward, and punish people if what they do is not in their control? When someone does something, it is only the fact that we believe that that action was "part of them", that they *meant* to do it rather than doing something else, that we feel we may judge them. When our wills are "our own", we believe that we are *responsible* for them. It is this fully-free *will* of the human that is called into question

when we realise that the world might not afford us the amount of control required by our moral systems. We now have to ask whether we do indeed have such control.

Williams (1981) famously argues that we do not have such control, or, at least, that we do not have enough control over our actions to be fully responsible for them. To illustrate this, Williams tells the story of a fictional artist named Gauguin, who had the responsibilities one might expect a grown man to have, such as holding down a job and caring for a family. However, he also felt that he would not reach his full potential as a painter leading the life that he currently led (Williams, 1981:38). It is important to note that Williams stipulates that Gauguin is not morally bankrupt. He knows he has these other responsibilities, and he believes he ought to fulfil them. Nevertheless, his desire to become a great painter eventually pushes him to neglect these duties. He hopes that by starting a new life he may become a great painter, but he has no way of knowing whether or not he will be successful. It is here that Williams argues Gauguin has no way of justifying his decision (to become a painter) as a good one until he has gained success. But, as has already been noted, Gauguin cannot know if he will be successful. Hence, Gauguin has no way of knowing whether his decision is a good one or not. If Gauguin fails, his decision would have been a bad one, if he succeeds, his decision would be a good one (*ibid.*). It is important to note that Williams does not cast this decision in moral terms; rather, he states that Gauguin has no way of knowing if his decision is rational, and therefore, a good one. Williams then argues that this insight also has bearing on moral decisions and actions.

It is with the story of Gauguin that Williams believes he is able to demonstrate the problem we face. The problem is not that we have to wait until after the fact to make our moral judgement, rather, it lies in how the circumstances which lead to either a good or bad judgment come about, and also in how the effects of our decisions play out. Williams argues that it might seem theoretically possible to find a criterion which allows us to judge the rightness/wrongness of an action before we make it, but that such criteria would still seem to require some information about the future. The problem with the future, at least from our perspective, is that it is uncertain, and the reason for this uncertainty lies in all the different factors that are at play in any decision that we make, as well as its effects. No one person is able to control all aspects of his life; so, whether Gauguin becomes a successful artist or not will be partially up to him, but also partially up to how lucky he is in finding the right place to work, choosing the right moment to act, and working in the particularly popular style of art of that time, to name a few. Williams thus argues that Gauguin has no way of knowing whether or not he will be successful, and that to some degree his rational justification for his actions depends on his success or failure (Williams, 1981: 40 – 41). Something similar applies when it comes to the morality of our actions. One can summarise all of this as follows: Gauguin should only hold himself responsible for those things within his control. He decides to leave his family and pursue a career in art, whether he

did the right or wrong thing (in neglecting his responsibilities) is also partially dependent upon his failure or success as a painter, yet these are not wholly within his control. This is due to the fact that people judge people not just for what was “rational” at a given time, but also for what *actually happened*, and moral judgement is therefore susceptible to post hoc factors. Hence, Gauguin is responsible for something not within his control. It is this story that captures the apparent paradox within our moral systems: we ought only to blame ourselves for what is in our control, but we do blame ourselves for what is not.

Williams’ approach demonstrates both the problem with deciding whether one ought to do something, and the problem of retroactively judging something as good or bad (whether morally or not). This is because whether something is judged as good or bad is (intuitively) at least partially dependent on the consequences of our actions, and we are thus only able to truly know if we are justified in doing something after we have done it. The amount of control we actually have over our actions and their effects does not seem to be enough for us to endorse our intuitive moral principles, which require that we only be held morally accountable for that which we voluntarily do. Williams focus is largely from the perspective of the moral actor and his self-assessment. Nagel approaches the problem of moral luck from more of a spectator’s position (i.e. from the position of others who accord moral praise or censure). He elaborates on the problem of how the constitution of the world (and ourselves) limits our control over our actions, and suggests that this should lead us to conclude that we cannot hold one another morally accountable.

2.1.2 Thomas Nagel

Thomas Nagel was the first author to distinguish between four different “types” of luck. This thesis will be split into the same four types of luck, and most philosophers who write on moral luck still use Nagel’s distinctions. Although he distinguishes between the four types, they are still related to one another and all refer to the same issue: the seeming paradox between the Control Principle and the fact that we tend to include factors out of the agent’s control. Nagel starts his discussion by admitting that the Control Principle seems to be intuitively appealing. But he argues that, although the Control Principle seems to be intuitively true, he does not agree with Kant that intentions are all that matter when making moral judgements. He believes it also seems intuitively true that what actually happens should also have an affect on moral judgements, and therefore, we have a conflict between the Control Principle and the way we actually make moral assessments (Nagel, 1993: 58). He further strengthens this intuition by distinguishing between two different forms of blame. He says, although we may

judge some things as being bad (such as death or harm), these are not necessarily within the scope of moral judgment. As Nagel explains:

“[...] but when we blame someone for his actions we are not merely saying it is bad that they happened, or bad that he exists: we are judging him, saying he is bad, which is different from him being a bad thing.” (ibid)

It is this second form of *bad* that the Control Principle is concerned with. We want to only blame and judge people as bad for doing what is within their control.

Since it is undoubtedly true that much of what we do depends on factors beyond our control, Nagel asks whether we should not simply deny the Control Principle. If we can clearly show examples where it does not hold, we should reject it and find another, more refined principle by which to judge others. The problem, Nagel argues, is that we cannot do this since the Control Principle is not only a foundational part of our moral systems, it is still an intuitively necessary requirement for justified moral judgements. It is not the principle of control itself that is the problem here (i.e. it is not an “ethical” or “logical” mistake), but the paradox that arises when we consider the facts about the way in which the world is constituted and the principles our moral systems imply (Nagel, 1993: 59). The principle of control is not false, and hence is not something we can simply get rid of. Therefore, Nagel believes, we truly do have a paradox at the heart of our moral systems. With this in mind, he goes on to describe four separate instances of luck which each, in their own way, takes away one aspect of an agent’s control over her actions and their results. It is Nagel’s thesis that if we accept these to be true, we lose moral responsibility. What follows will be an introduction and explanation of each of these.

2.1.2.1 *Resultant Luck*

The first form of luck is called Resultant luck. Nagel identifies this as the first of two types of luck that fall within the sphere of external factors that affect the causes and effects of our actions. Resultant luck is concerned with how things turn out, or, put more simply, how your actions play out in the real world and “actually” occur. To demonstrate this, we can consider the case of two different truck drivers:

Both of these truck drivers find themselves driving on the road. Known to both, their brakes have not been checked and serviced in a while. It is clear that both drivers are negligent. If the story were to stop here, most would agree that they should be blamed equally for this negligence. However, as luck would have it, one

of the drivers makes it to his destination safe and sound, while the other is in an unfortunate accident where a child has fallen into the road, and due to the poor brakes of the truck, the driver was not able to come to a stop in time. (adapted from Nagel, 1993)

The story above poses an interesting question: are both drivers equally blameworthy or is the second more blameworthy, given the results of his negligence? It might be necessary to make the same clarification Nagel made: I do not dispute the fact that what happened was a bad thing. The loss of a child will almost certainly never be a good thing. However, the question at hand is, was the one driver “worse” than the other? Both were equally negligent; one did not drive more poorly than the other. The fact that one hit a child rather than not was beyond his control. Had a child fallen in front of the other, the same thing would have happened. Is the one more culpable than the other because his negligence led to someone being hurt? Our first intuitions seem to support this idea, and our legal systems do punish people in such cases more severely (the one who hit the child might receive jail time, whereas the other would probably only receive a fine if he were caught, or at the very least a reduced punishment relative to the other driver).

Something similar happens in instances of drunk driving. An individual who is caught behind the wheel by a police officer is never punished as severely as one who causes an accident. Yet both intend the same thing, and both have the same amount of control in the situation. The fact that one is pulled over before she could cause an accident is merely a matter of luck, and yet we blame them differently¹.

Nagel argues that cases like the above demonstrate that while we hold the Control Principle to be true, we also blame people for the results of their actions, even when these results were not within their control. Moral blame, according to Nagel, is thus subject to external factors (if we blame people for the results of their actions rather than their intentions). Luck (or otherwise stated, a lack of control) often characterises the results our actions and therefore undermines the Control Principle when it comes to moral judgements.

¹ I will add here that I have used punishment as a synonym for moral blame. I recognise we can treat people differently and yet blame them equally, but both of these cases above do elicit different moral judgments from people as well. People do not only treat a drunk driver who caused an accident differently from one who did not, but they blame them differently as well. We do not have the same level of moral outrage against a friend who we find out drove drunk than we do for one who caused an accident while being drunk.

2.1.2.2 *Circumstantial Luck*

A second type of luck Nagel introduces is that of Circumstantial luck. This is the type of luck in the circumstances in which you find yourself. Circumstances (the world around you) play a large role in who you are and what you do. It plays a role in the development of your personality and, of course, different circumstances require different reactions. As Nagel explains:

The things we are called upon to do, the moral tests we face, are importantly determined by factors beyond our control. It may be true of someone that in a dangerous situation he would behave in a cowardly or heroic fashion, but if the situation never arises, he will never have the chance to distinguish or disgrace himself in this way, and his moral record will be different, through morality.
(Nagel, 1993: 65)

Again, to demonstrate how this diminishes your control when it comes to moral matters, Nagel uses the Second World War as an example. He explains that the ordinary citizens of Germany, during the rise of Nazism, had the option to either oppose the Nazi government (a commendable act) or to join them (an action which is now looked upon with disdain). This “test”, as Nagel calls it, was one not shared by the people in other countries, but that is not to say they would not have acted as badly as many Germans did, were they to have been presented with the same circumstances (Nagel, 1993: 65). The same is true of Germans who were lucky enough to have emigrated out of the country before the rise of the Nazi party. It would seem that whether one is morally judged as a good or bad person also depends on how lucky one is in the circumstances one finds oneself in. Or, as Nagel (1993: 66) puts it: “We judge people for what they actually do or fail to do, not just for what they would have done if circumstances had been different.”. After reflection, it seems the Nazis were not necessarily worse than me or you, for we may likely have done the same; yet, have either of us ever been condemned for this?

An interesting example that demonstrates the same principle is that of admired or famous individuals. If one considers famous athletes or cultural stars, they are lauded, not only because of their skills, but also because those skills happen to be valued by society. The reason why Lionel Messi is so famous and admired is not only because he is a skilled football player (although this is, of course, true), but also because the society he finds himself in enjoys football and considers the sport important. Upon reflection, one can thus see that Messi is admired (partly) because he was lucky enough to possess a skill admired by society. By contrast, somewhere out there is a person who can spit the furthest distance of all, but since society tends to deem this behaviour as uncivil, she is not respected. This is

simply a matter of luck (in this case, bad luck) since, if society had believed spitting to be an impressive skill, she would have been famous.

Examples such as the above show how deeply entrenched in the external world our moral judgements are. This is problematic, since the external world is largely outside of our control. Thus, if we continue to hold on to the Control Principle, we cannot hold people accountable for a large part of their actions, because these actions lie within a “causal” net of external circumstances outside of our control. The world pushes and pulls us in different directions, placing us in different circumstances. We are only able to react to those circumstances we are presented with, so why should we be held accountable for those actions when others (through no merit or fault of their own) are not?

2.1.2.3 *Constitutive Luck*

The two forms of luck discussed above focus on what happens to the agent or what the agent happens to do. They emphasise actions and conditions, and, as Nagel notes (1993: 67), they seem to remove moral responsibility from the agent and place it in the world. So, what if we place judgement back on the agent? Do we regain responsibility if we ignore actions and circumstances and just look at the agent’s reasons for his actions? Nagel (1993: 64 – 65) argues we do not, since luck does not only infiltrate the world around us, but the very core of who we are:

A person may be greedy, envious, cowardly, cold, ungenerous, unkind, vain, or conceited, but behave perfectly by a monumental effort of will. To possess these vices is to be unable to help having certain feelings under certain circumstances, and to have strong spontaneous impulses to act badly. Even if one controls the impulses, one still has the vice. [...] To some extent such a quality may be the product of earlier choices: to some extent it may be amenable to change by current actions. But it is largely a matter of constitutive bad fortune. Yet people are morally condemned for such qualities, and esteemed for others equally beyond control of the will: they are assessed for what they are like.

If a murderer murders someone, even if only through luck or misfortune, he is surely still a bad person. Why should we not hold him responsible? The answer, Nagel argues, lies in how we come to be the people we are. Nagel notes that our most inner decisions are not immune to outside influences. He states: “Factors beyond the agent’s control, like a coughing fit, can interfere with his decisions as surely as they can with the path of a bullet from his gun.” (Nagel, 1993: 64). Our decisions and mental states are just as vulnerable to external factors. This itself should cast some doubt on the idea of a wholly free “self”, however, there is an even stronger sense in which the self is not in our control.

Even when we focus only on the *intentions* of people, we must still ask where these intentions come from.

To make this clear, we need to discuss some general assumptions most people will agree with. The first concerns how we make choices. Although the exact nature of choice and what influences it might be under contention, we can agree that choices are informed by the person's wills, wants, and desires. What you want, and the type of person you are, will ultimately inform the choices you make (provided that you are not being coerced or manipulated). If you are the type of person who easily gives in to desires (that is to say, you do not have very good self-control), you are a huge lover of chocolate cake, and you are presented with just such a cake, chances are you will "choose" to eat the cake. So, what we see is that actions are informed by choices, and choices are informed by intentions and our perception of the best means for such intentions to be successful. This explains why it might seem to be a good idea to change the focus of moral judgment from actions to intentions. If our actions are largely affected by external factors, we should take one step back and focus on that which is internal to "us". The second assumption we must look at is that of how intentions come to be formed. This assumption is somewhat more difficult to discuss, since most people do not think of intentions as being formed; intuitively, we feel like *we* simply "have" intentions. We live our lives, always wanting some or other thing, and we then act in such a way so as to gain said thing. However, no matter how much it might feel like intentions are simply there, they are indeed formed by other factors. Again, without being too controversial, these factors can be divided into two main contributions: (1) environmental, and (2) biological. It is not important how these two factors exactly interact within this thesis; all that is important is to realise that neither are much within our control, and both affect who we are. It is, of course, true that one is able in some way to change one's environment to something else, and in this regard, one does have control over it. We manipulate our environments all the time. Nevertheless, we are (by and large) not in control of how the environment affects us. Since conception, all biological factors have interacted with the environment (such as education and your history) and have thus been expressed in certain ways. This complex process eventually produces something akin to a person, with feelings, wants, needs, and desires. All these things collectively inform your intentions, or what you as a creature wishes to do. Your intentions are "determined" by the personality you have (which you never chose but which emerged), and the environment your personality finds itself in. In other words: what you mean to do is a product of factors beyond your control. Even if it seems like intentions are a product of choice, they are a product of situated personalities (embedded within the environment).

From the above one could conclude that the intentions you have are a product of an emotional being within an environment, where these emotions are controlled by chemicals in the brain which are

reactions upon reactions going all the way back to when you were born. So, although you are a unique person with wills and wants, these are not in your control. You can no more choose your intentions than you can choose your favourite ice cream flavour. You like what you like, you want what you want, and you feel the way you feel. There is no point in time where you are able to leave the bounds of who you are and change these factors.

The above can be demonstrated when we consider agents who both did the same thing, but who have very different histories. Let us consider two thieves; the one was born in a poor community, let's call him Peter₁, and the other in a well-off community with all the benefits one would expect from an upper-middle-class upbringing, let's call him Peter₂. Due to his circumstances, Peter₁'s parents had to struggle to raise him. In fact, times were so tough that there was a time when Peter₁'s parents had to steal from others just to survive. Peter₁ was therefore raised in an environment where stealing was a method of survival. From a young age, he experienced theft not as a bad thing, but rather as a way to get out of a tough situation. Furthermore, the community that Peter₁ had grown up in had lived in these circumstances for many generations. The collective environmental pressures had acted in such a way that those who were more aggressive and had less empathy were better equipped to survive in such a place. For this reason, those with these "bad qualities" had more children; and as time went on, these traits were strengthened as generation after generation produced children, and only those with these traits survived². Peter₁, then, was doomed from the start to be the type of person more prone to theft. Peter₂ on the other hand, was raised in a community with all of the modern luxuries. He was safe, was provided with food and family. Peter₂ never identified with any thief. The only thieves he was exposed to were those from outside of his community. They were seen as the other and as a threat to his community life. Theft, therefore, was experienced as something destructive and negative. We can imagine the same sort of environmental pressures from Peter₁'s history playing a role in constructing the type of people who would live in such a modern community. Those with stronger community bonds, more empathy, and who are more prone to cooperation would have more children and thus populate a larger portion of the community. Peter₂, therefore, was born and raised as the type of person who is easily opposed to theft.

These are two very different histories, both out of the control of the agents (particularly the "genetic history" of their forefathers), which affect the persons they are. These histories affect the desires they

² It is important to note that I am exaggerating the evolutionary processes in this thought experiment to help aid the formation of the intuition that we do not have sufficient control over who we *are*. Furthermore, I am not claiming that any such factor needs to be genetic but could also be cultural. I am not claiming that any such process has occurred or that it is possible. It is only for the sake of argument.

have, the possible methods they consider and utilise to meet a desire, and the sort of responses they have to circumstances in life. The question then becomes: do these histories make a difference to how we judge the agents? To continue the story then, let's imagine that at some point in their adult lives both Peters steal from another person. They both desire (without coercion) the items they wish to steal, and both freely undertake the steps to fulfil this desire and succeed. In general, and ignoring the problem of moral luck for a moment, these two individuals did the same thing and caused the same amount of harm. Most would then conclude that both are equally guilty. But this thesis concerns itself with the problem of moral luck, and we must then ask: Is it not easier to understand why Peter₁ would want to steal, and does the history of Peter₂ not make it worse that he would do such a thing without any pressures from his "personal history"? Once we understand the personal histories of both agents, their characters differ. The reasons why they wanted to steal are different; their perspectives on the action differ, and the ease of perpetrating the crime differ. Should these factors not make a difference in our moral judgment? If much of who we are is affected by things not in our control, why should we be judged for who we are?

This is the idea captured by Constitutive luck. Even when we no longer look outside the person, but inside, we still find a lack of control over what is the case and what happens. Hence, changing our focus to intentions rather than actions and their outcomes does not solve the problem; in fact, it almost makes it worse. We now start to even doubt the very nature of the self and how we view what we do and who we are. If we require control to hold someone responsible for their actions, we are still unable to do so.

As you might have noticed, many of these ideas and issues are closely related to the free will problem, and some who have tried to solve this problem have argued that we do have some form of control over our personalities. These will be discussed later and I intend to show why they fail. For now, I believe the above is thorough enough to show that we can at least doubt the amount of control we have over who we are and our personalities and hence the choices we make.

2.1.2.4 Causal Luck

As stated in the introduction of this thesis, I will not focus on Causal luck mainly because I believe the moral problems found in Causal luck are all solved by my discussions on Circumstantial luck and Constitutive luck. This is not to say that the distinction between Causal luck and the others is not a valuable one. There are definitely cases or phenomena not captured by the other types of luck; however, I believe these cases to be somewhat different from those that fall within the moral considerations that are the focus of this thesis. To see why this is the case, let us imagine for a moment

that the problem of moral luck as it is found in the first three types of luck has been solved. We have explained how the lack of control in your environment (as we experience it) no longer poses a threat to responsibility. We have also explained how, even though the agent doesn't construct his/her personality, the agent can be held responsible for those intentions/actions that "originate" from the agent's personality. What then is left to be explained?

It seems the only other causal factors to be accounted for are those not found in our phenomenological experience. By phenomenological experiences, I mean the way in which we as humans construct and experience the world from a subjective stance³. If we choose to view the world from a highly scientific and reductionist position, the concepts of humans, plants, buildings, and roadways all seem to disappear and are replaced by atoms in motion. A sea of energy and particles is all that "truly exists" when we take such a stance⁴. The world in which we find moral questions is the world as viewed from the human perspective. The distinction between me and you, between cars and bikes, only works once we as humans have divided the world into separate understandable concepts informed by how *we* experience the world. This is the distinction between the subjective world, as experienced by us after we have constructed our conceptual images, and the *world as it is*. This subjective world is the world where Resultant luck, Circumstantial luck, and Constitutive luck poses problems. Moral perspectives are human perspectives. There is, of course, another world "behind" our world (the scientific image), but this world does not form part of our moral perspectives. Atoms are not responsible (in the moral sense) for anything, people are (or so one hopes).

To make the distinction clear between those factors that are moral and those that, although "real", are not, I will describe three different cases in which an agent is "caused" to act by external factors:

- (1) John is born in a community where certain acts of violence are encouraged. These acts include war, fighting, and a general lack of compassion for those outside your close community. Had John been born in a different community, he would not have been encouraged to be such a violent person. One day, John comes across a stranger from another community, and when he realises this, he decides the best course of action is to strike the person in the face.

³ Daniel Dennett (drawing on Wilfrid Sellars (1962)) also makes the distinction between what he calls the "manifest image" and the "scientific image" which are directly related to the distinction I am making here. The main idea is that we can distinguish between the "only physical" world and the "world from our perspective". The "only physical" world refers to the scientific image and the "world from our perspective" is the manifest image.

⁴ I use here "energy" and "particles" in as vague a way as possible. I am not making any claim about the fundamental nature of physical reality. Whether the most basic particles are atoms, or quarks, or strings: the point is the world we "live", and construct moral systems in, is not the world of "fundamental nature" but one of human perspective.

- (2) John is born in a community where he is taught to respect others. He is told to be kind to those close to him, and to respect those who are different. John has been instilled with an empathetic moral code. One day, John comes across a stranger from another community, and when he realises this, he decides the best course of action, despite his upbringing, is to strike the person in the face (perhaps due to something the person said).
- (3) John is born in a community where he is taught to respect others. He is told to be kind to those close to him, and to respect those who are different. John has been instilled with an empathetic moral code. One day, John comes across a stranger from another community, and when he realises this, he decides the best course of action is to respect and listen to the person. Unfortunately, as John extends his hand in an attempt to shake the person's hand, a random particle from a faraway place makes its way into John's brain and causes neurons to fire in such a way that his hand swings violently and strikes the stranger in the face.

Usually, we would conclude that hitting a person through the face when they do not pose any threat to you would be a morally bad action, and that a person that does such a thing ought to be punished. Ignoring the problem of moral luck for a moment, most people would conclude that John was morally wrong in both (1) and (2). But with (3) very few would conclude that John is "bad"⁵. In (3) the distinction between agent and non-agent is not clear. This is due to the fact that we are mixing up two different approaches to viewing the world, and so our boundaries become vague and blurry. In case (3) we cannot say that it was John who decided to hit the stranger as it was clearly another object (the electron) that infiltrated John's brain and removed John as an agent from the situation. It is only once we introduce the problem of moral luck to the equation that cases (1) and (2) become difficult.

In (1) the cause which "removes" responsibility is the community or the circumstances (a phenomenological construct) in which the individual finds himself. In (2) the cause is the person's personality (a phenomenological construct), or more specifically his will and desires, which are the cause of his actions that he might not have control over. In (3) the cause is not part of our phenomenological perspective but must rather be viewed from the reductionist perspective of science.

⁵ It could be argued that John (1) and (2) both acted in accordance with their respective moral codes. I do not wish to justify any claim about what is right and wrong here, however, this thesis does not accept the normative claims of moral relativism. "Normative moral relativism" claims that there are no objective or absolute moral norms, rather, there are many conflicting, personally justified, moral norms relative to the culture or individual. For the sake of argument, this thesis will assume that both John (1) and (2) are morally wrong. For more information on Moral Relativism see Gowans (2019).

We are no longer looking at causes which form part of our “experienced world”, but which are only understood, and indeed recognised, once we look at the world behind the manifest world.

Causal luck is such a broad concept that it could include all of the factors found in the three examples, but once we introduce the ideas of Circumstantial luck and Constitutive luck, much of what we would call Causal luck can be accounted for. What is left over are those causes found in the scientific (reductionistic) approach. This is important because morality and, as I hope to argue, responsibility, are only meaningful concepts once we move away from the reductionist scientific approach and towards the subjective⁶ (manifest) perspective of people. Hence, what is left of our concept of Causal luck is not *moral* luck but rather just physical luck, or stated otherwise: pure randomness. The problem of randomness and its effects are much more relevant to the classical discussion on free will than it is here. For this reason, this thesis focuses on moral luck and not on the question of whether we have free will or not.⁷ I will briefly address the issue in the chapter on theoretical implications. Causal luck then, as it has been described here, does not pose the same sort of moral problem as the other three types of luck.

2.2 Conclusion

To summarise, we seem to find the Control Principle intuitively plausible, yet, we also seem to lack control in three distinct yet related ways. First, we seem to lack control over how our actions turn out. Second, we seem to lack control over the opportunities and challenges we face. Finally, we lack control over the people we are and the intentions we have. Taken together, these three types of luck undermine the Control Principle and therefore moral responsibility. The next three chapters will discuss each of the forms of moral luck mentioned above. Hopefully, I have given a detailed enough account of moral luck that the next chapters can focus mostly on discussion and critique, rather than having to do any more basic explaining.

In these three chapters, I will argue that each type of luck does not pose a true threat to moral responsibility. In the chapter on Resultant luck, I will argue that by changing the focus of our moral

⁶ “Subjective” here does not mean the approach itself is subjective (dependent on the individual), but rather that the approach centres around the subject (agent).

⁷ For a brief and rough account of my view on free will the work of Fischer is a great place to start. On his view we should be agnostic about whether we do or do not have Free Will, however, we still conclude we can be held responsible whether we do or not. Free Will then is not the determining factor of responsibility but rather something else is (Fischer & Ravizza, 1998). For my account of what this “something else” is see my chapter on Constitutive luck.

judgements from consequences to intentions we can avoid the paradox of moral luck. In the chapter on Circumstantial luck, I will argue that we mistakenly infer *moral* luck, when in fact there is only *epistemic* luck. Finally, in the chapter on Constitutive luck, I will argue that our intuitions have led us astray and we have created the problem by claiming there are two separate categories (the person and their personality) when in fact there is only one category (the agent).

Many philosophers have attempted to solve the issue of moral luck by using the same argument for all three types of luck, or they have only focused on one and left the others. Robert Hartman (2018) argues that most philosophers fail in their attempt to resolve the paradox of moral luck because of this⁸. I believe in order to solve the problem, each type of luck must be evaluated separately, and the solutions should then be brought together so that we can properly understand all the mistakes we have made in order to agree with the conclusions centered around moral luck. In order to do this, I will use and discuss some of the most influential writings on moral luck. By doing this, I hope to be able to convince you of two things. First, I hope to convince you that we require a revision of some of our moral concepts, including responsibility and moral agency. And second, I hope to convince you that once we have made these revisions, we will see that moral luck does not pose any threat to moral responsibility. After this, I will focus on the implications that follow if my thesis is taken to be correct.

⁸ Hartman (2018) argues that philosophers who discuss only one type of luck believe they are successful when they have unfortunately only discussed one part of moral luck, and to thus conclude that morality is immune to luck is mistaken. When authors do discuss all four types, they mostly agree that there is at least one type of luck which still poses a threat.

Chapter 3: Resultant Luck

3.1 Introduction

Now that we have laid down the foundations necessary to understand and analyse the different forms of moral luck, we can start with the first: Resultant luck. In this chapter, I will use many different sources to help understand but to also critically evaluate the arguments in support of and against Resultant luck. As already mentioned, Resultant luck describes the cases in which an agent is being held morally responsible (is being judged and ascribed desert) for the results of an action; however, certain relevant parts or perhaps the whole of these results are or were not within the agent's control. If we then also agree with the Control Principle, which states a person can only be held responsible for actions within their control, cases of Resultant luck demonstrate a seeming paradox between our moral principles and our moral evaluations. As a reminder, "responsible" here does not refer to any simple understanding such as causally responsible. Rather, we are discussing moral responsibility, and it is the sort of responsibility which entails moral desert. Whether this desert is reward or punishment, or praise or blame, the reason we would ascribe this desert to the agent is because they were judged to be morally right or wrong and therefore deserve such desert because they are morally responsible. It is my hope that I am able to demonstrate that this paradox only arises because we are focusing on the wrong phenomena (namely the actions/results) rather than focusing on the agent's character or intentions. The thought experiments and/or examples that are given to demonstrate the paradox of Resultant luck direct our attention to the wrong parts of the story and thus allow our intuitions to make a bad decision/judgement. I will argue that once we clarify which criteria should be relevant to moral evaluations the paradox disappears. If I am correct in arguing that intentions and/or moral character are the relevant criteria, other implications will also follow; such as understanding that actions play a "demonstrative" role in moral evaluations, and that we require a much more robust theory of moral negligence and moral obligation so as to counter certain objections to the idea that intentions are what matter when judging someone.

Solving the problem of Resultant luck requires us to draw a line. It already assumes a moral agent and then asks us where we need to, or even whether we can, draw the line from attempt to execution and onwards when it comes to judging actions. Note that we are talking about judging *actions* and not an agent's character. At least, this is the way most discussions surrounding Resultant luck speak about moral blame. We are judging the agent, but in the "praise-and-blame" sense rather than in the

“virtue” sense. When judging an agent in the “virtue” sense we are analysing, and then evaluating, the agent’s character in light of some standard of “good character”. The “praise-and-blame” sense of judging someone asks the less abstract question of whether a person ought to be rewarded or blamed for an action. It is an evaluation of the action itself and then a consequent blaming of the agent in some sense. This is the primary type of responsibility that most of the writings on Resultant luck focus on. Our judgment in this regard might have an effect on how we assess the agent’s character, but the focus here is specifically on judging the agent’s action. To clarify: rather than asking what the agent’s motivations were and asking if this was how a “good” agent would have acted, we first assess whether an action was good or not (relative to some moral system such as Utilitarianism or Deontology) and then praise or blame the agent accordingly. I hope to argue that it is the focus placed on the action itself that has led so many to conclude that the paradox is real and a problem. Some philosophers do deal with the problem in a different way. Rather than focusing solely on Resultant luck by assessing action and agent, they also focus on the effect of luck on the “moral character” itself. Thus, they argue that luck should or should not be included in our moral evaluations because of the effects this inclusion/exclusion would have on moral character in general.

The rest of this chapter will be split into two different sub-parts; the first will use the more traditional concept of responsibility and Resultant luck, the second will focus on moral character development and the impact of “luck” thereon. It is my hope to show that allowing for luck in these regards leaves us with either a contradictory and inconsistent form of morality, or we diminish the “moral” part of the moral agent. I thus hope to convince you that Resultant luck should be excluded from our moral evaluations on all grounds.

3.2 The Problems of Including Resultant Luck

As discussed before, the problem of Resultant luck lies in the fact that our actions (the results off our intended decisions in this case) seem out of our control. Using the two different ways of evaluating moral blame from the section before, we can see that we are either saying agents are responsible for their actions and thus responsible for something out of their control, or actions demonstrate something about a moral character and that this demonstration is thus out of our control. Either way, the results of our actions, which are not fully in our control, have an impact on resulting moral evaluations. The question thus becomes: “If we hold one another responsible for our actions, how much influence should the actual results of those actions have on our moral responsibility?”. Take note that this question already places very little emphasis on intentions. Most theorists who write on, and support the conclusions of, Resultant luck accept the premise that people should be held responsible for their

actions and that intentions only add another factor to the equation. For them, intentions are not all that matter, actions matter as well, and in most cases, we are not wholly in control of our actions. In day to day life, we treat actions as direct indicators of moral blame/desert. Most of us are fine with this until we start to notice the consequences of this belief. Bad people are praised for lucky coincidences, while good people are blamed for bad luck. I have already discussed the famous example of the two drunk drivers elsewhere, but two more examples will help to explain this. In the case of a “bad” person being praised for something “good”, we can imagine a thief who, while breaking into a house intending to rob the place, accidentally triggers the alarm, sending emergency personnel to the house where they find the owner on the floor and treat him for a heart attack. We can imagine that had the thief not triggered the alarm, the owner would have died due to the heart attack, but thanks to the luck of someone breaking in they were found and treated in time. The robber definitely intended to do something bad, but thanks to luck, they happened to save a life. If the consequences of your actions matter in the moral sense, the case can be made that they should be praised for the life they saved. Or imagine a person who, while trying to save people from an armed attacker, fires her weapon and accidentally kills an innocent bystander. She intended to do good and wanted to save people, but unfortunately, her actions led to someone dying. If consequences matter, then this person should be negatively evaluated. We do not like these consequences, hence the “problem” of Resultant luck.

The question of Resultant luck, to me, can then be addressed using two different methods based on how one approaches morality in general. These two broad methods are moral realism and moral anti-realism. The first states that there are “truths” to be known about moral statements and propositions, that there are moral “facts” that we can discover and learn. The second claims that there are no moral facts out there and that morality is either a construct or grounded in emotional preferences or so on. These are only very broad ways of dividing the positions on morality. Each of these divisions has its own divisions which make very different claims. This means that there are probably just as many disagreements between moral anti-realists as there are between moral realists and anti-realists. Ethical subjectivism, moral relativism, non-cognitivism, and emotivism are all examples of moral anti-realism which make very different claims. For this thesis though, such specific commitment is not needed. For now, we only need to agree with one of the following:

- (1) Either you believe in a metaphysically realist conception of morality where its laws and principles exist in a mind-independent manner and are discovered; with the question then becoming: “How do we discover the truth about the moral worth of our actions and intentions, and what method do we use to discover them, whether they be empirical facts or perhaps psychological facts?”.

- (2) Or, you believe morality to be a construct, a tool to help organise a unique social feature of a particular creature (us). We might still agree that there are facts that help inform our moral values (such as the fact that cutting me will hurt, or the fact that welfare systems either work or do not), but the moral value itself, the “ought” claim, is something *we* place value in (mind-dependent). If this is the case, the question becomes: “Is there a good reason to include our actions (and particularly their results) in our moral evaluations?”¹

I prefer the second line of reasoning and questioning. I find no reason to believe that morality is an intrinsic “real” feature of our universe. Because of this, I will consider the two different sides of this debate: on the one hand are philosophers who believe it is bad and unnecessary to include facts about how our actions turn out in our moral considerations about moral desert; on the other are those who believe this sense of luck is already included in our moral assumptions (luck is then already part of our generally-shared moral intuitions) and adds to the overall “worth” of morality (this “worth” referring to the goals of morality, whatever they may be, and how well the moral system achieves this goal). Examples of this “worth” may be agents taking ultimate responsibility for their rational decisions (Deontology) or that outcomes maximise the utility in the world (Utilitarianism).

The easiest way to defend the position that Resultant luck, as with all kinds of moral luck, should not be included in our moral assessments, is by demonstrating that once we do so errors and logical problems arise. Perhaps one of the best examples of this would be Steven Sverdlik’s paper *Crime and Moral Luck* (1993). To make the strongest and clearest case, Sverdlik chooses to use the most basic of examples. These would be examples where all the relevant factors to the thought experiment are accounted for and any factor not directly relevant is ignored. Although this might seem impractical and too simple a case to be relevant, Sverdlik argues that such examples are all that are necessary to show the inconsistencies that arise when we include Resultant luck in our assessments. This is sufficient since any inconsistency within a moral view will lend credence to the idea that said moral view should be rejected (Sverdlik, 1993: 182 – 183). For these reasons, we do not need to have a complete theory of harm, desert, intent, and so on, because these things have no bearing on whether luck should be included when we judge someone. As an example (and something which will again be dealt with in more detail in a later chapter) many philosophers differentiate between different forms of responsibility, such as moral responsibility (where responsibility is about just desert with

¹ Other than actions, we could also include a general sense of the agent’s character (such as with Virtue Ethics) or we could focus on intentions (that which the agent “meant” to do, such as with Deontological Ethics). Here, we ask if actions should be considered over and above these other factors, or perhaps actions alone (such as with Hedonistic Utilitarianism). For a discussion of each see Hursthouse & Pettigrove (2019), Alexander & Moore (2019), and Driver (2019) respectively.

regards to moral character) and ordinary responsibility (where you are the cause of something) in order to demonstrate that luck does not have an effect on moral responsibility but might on ordinary responsibility. Sverdlik (ibid.) argues that such clarification isn't necessary since if we can show that any form of responsibility cannot be consistently applied without contradiction when luck is included, we have reason to not include luck. The same holds true for other important concepts in moral philosophy such as harm or desert. Whether we believe psychological harm or physical harm is more important, or whether desert should always be just or should be proportional to the response we want from the agent (such as deterrence), these things do not have bearing on the problem of Resultant luck as it primarily deals with responsibility and control.

Sverdlik uses two unique terms to refer to the two sides of the debate: a) the non-equivalence theory – which states that actions (and therefore luck) are part of our moral evaluations and can, therefore, make a difference, and b) the equivalence theory – which states that actions (and therefore luck) should not be included in our moral evaluations (1993: 182). The word “equivalent” here refers to our moral evaluations of the agent. In order to demonstrate that Resultant luck should not be accepted within our moral judgments (what Sverdlik calls the Equivalence Theory) Sverdlik uses a unique case of “transferred intent” to show that once we accept and allow luck to be part of the equation, the concept of “just desert” becomes problematic and the results are counterintuitive to most who accept the non-equivalence theories. “Transferred intent” refers to cases where one person wishes to harm another, but by accident harms a third. Or as Sverdlik (1993: 183 – 184) puts it:

A intends to shoot B fatally but in fact accidentally kills C instead.

Note the importance of the selected phrasing: A has the full intention of shooting and killing B, but by pure luck (chance, factors beyond A's control) A shoots, and kills, C. The intention of A to shoot B was transferred beyond A's control to C, and C thus becomes the recipient of the intended harm. Barring any further details, the equivalence theorist will say that A is just as blameworthy in killing C as she would have been had she killed B, not because the harm is the same, but because the intention behind the action stays the same. Sverdlik goes on to argue that the non-equivalence theorist has no way of stating the difference in blameworthiness between A killing B and A accidentally killing C, even though there is a clear difference in execution and result since the harm (the results) were the same in both (1993: 184).

In its current form, the thought experiment does not pose a problem to either theorist since most would argue that A is still guilty of killing someone, something they should be blamed for. But what is A guilty *of* according to the two different theorists? For the equivalence theorist, intentions are what matter here, and since the intention to kill (to do harm) is still present, and it does not matter for moral

blameworthiness how this intention plays out in the world, A is guilty of wanting to kill and attempting to implement said desire, whether B or C is the recipient of this intent. A has equivalent blameworthiness since the morally relevant factor (intention) is equivalent. The non-equivalence theorist holds to the idea that outcomes have an impact on the moral judgment (thereby including luck) and in the above cases there is a difference in who receives the harm, but no difference in the type or amount of harm itself. Sverdlik does say that the non-equivalence theorist can also appeal to intentions in order to demonstrate that A is equally to blame either way, but he argues that this can only be one part of the story since the non-equivalence theorist is committed to saying that harm (and what actually happens) has an effect on moral blame. Sverdlik thus argues that when the non-equivalence theorist endorses the view that what matters here is the intention, they are endorsing the equivalence theory, and thus undermining their own position (ibid.). If they chose to argue that the two cases are in fact the same on the basis that in both cases A had the intention to kill B, then they are arguing that the morally relevant factor in both cases would be the intention that A had, but this is the main thesis of the equivalence theory. To stay true to their own theory, the non-equivalence theorist must argue that since there is equal harm whether A kills B or A kills C, there is also equal blame. For the non-equivalence theorist, A is blameworthy of inflicting harm on C.

From the above we can see that both the equivalence theorist and the non-equivalence theorist have ways of accounting for the intuition we have that A is equally to blame no matter if he kills B or C. The non-equivalence theorist is including the results of the actions (as they claim it should be) and intention becomes somewhat irrelevant. As an example, recall the case of the two drunk drivers, both are negligent but only the one causes any real harm. By including the results of their actions, factors beyond their control and intentions are now morally relevant. Once we include the results of our actions, the non-equivalence theorist has a way of accounting for the fact that we believe the cases (the intentional and accidental killers) to be morally equivalent, since in both cases a person (B or C) is being killed (for simplicity sake we assume the harm is the same in both “killings”). This does seem to save the non-equivalence theorist for the moment, but as Sverdlik argues, we now have another problem, and perhaps one worse than the first (ibid.).

The classic name for the argument that Sverdlik makes is called an *argument ad absurdum*. He uses the same logic used by the non-equivalence theorist to show that, when applied to other cases the result is absurd, demonstrating that the logic should be rejected. In the case above, the non-equivalence theory states that equal harm means equal blameworthiness, irrespective of intent. While it allows for our intuitions, in this case, to conform to the rational conclusion, the assumption leads to counterintuitive results when applied to most other cases. Ignoring the thought experiment from above for a moment, would we be happy in saying that an accidental killing is equal to an intended

killing simply because the harm was the same? Most would say no. Intention matters and to say that equal harm means equal blameworthiness can be shown to be highly counterintuitive. Take the following two cases as examples:

- 1) P1 decides to inflict harm on another person A. To do this P1 takes a firearm and shoots A. A dies and P1 is held responsible for this act.
- 2) P2 decides to practice their marksmanship at a shooting range. While there, P2 shoots at a target but by sheer luck the bullet ricochets off the target and hits person A, thereby killing A. A dies but P2 is not held to have deliberately killed A.

In both cases, the harm is exactly the same and even the method by which the killing happens is the same, but most would argue that the two agents are not equally responsible. Many would even argue that P2 is not morally blameworthy while P1 should be both judged and punished. The cases above show that equal harm does not mean equal blameworthiness, and similar cases for “good” acts can be constructed as well.

The cases above demonstrate that the non-equivalence theory (equivalent intent does not equal equivalent blame, as harm must also be taken into account) has difficulty with holding either a consistent view of blameworthiness and the factors which are relevant to it, or it has inconsistent and counterintuitive consequences. Sverdlik (1993: 191) then goes on to argue that the equivalence theory does not have these problems. In the first case, we are able to say that the two agents are equivalent because their intentions were the same. It does not matter that A accidentally killed the wrong person as the moral condemnation is focused on the fact that A wanted to kill someone. Even in the second thought experiment, we are able to say (consistently) why we would blame and punish P1 and not P2, while the non-equivalence theorist has trouble doing so. P1 had intent to do harm while P2 did not. These are not morally equivalent scenarios since the intent/character of the agents were not the same. The harm was the same, but as already discussed, accepting this as the morally relevant criterion leads to counterintuitive results.

A final point needs to be made here. Sverdlik does not argue that the equivalence theory is more intuitive than the non-equivalence theory, rather, it is logically more consistent and has a better way of dealing with different scenarios (*ibid.*). Furthermore, in no way has it been argued that harm and method do not affect moral responsibility, but rather that these are factors that further contribute once agent and intention have been considered, something I focus on in the chapter on Circumstantial luck (when actions become indications of moral character) and the chapter on theoretical implications of this thesis (where it is argued that actions might make a difference to how the agent ought to respond to the harm they caused). These factors might also have more of an affect on punishment and

restitution than they do to moral blameworthiness or character judgement (as mentioned above), and since punishment is not yet part of the discussion here these factors have been left out (Sverdlik, 1993: 191 – 192).

The inconsistency of including luck can also be demonstrated using a more recent discussion by Sara Bernstein (2019). In her paper, Bernstein argues that Resultant luck is not only luck in *how* things turn out, but also the *way* things turn out. That is, luck might affect what the actual consequence of an action might be, but it can also affect the way in which an effect happens. The following is an example that Bernstein gives of this type of Resultant luck (2019: 154):

Careless Driver: Galen plots to murder Gillian with poison. As he is driving to buy the materials for the poison, he fails to see a pedestrian walkway, and strikes and kills a pedestrian. Upon rushing out of the car, he realizes the pedestrian was Gillian.

Important to note before moving on, Bernstein presupposes the reality of Resultant luck in her discussion, and so accepts the conclusion that I argue against here. Her discussion does not serve to argue for or against Resultant luck but rather to demonstrate the difficulty in proposing a coherent way of dealing with cases of Resultant luck. Here, I wish to use her example as another way of demonstrating the inconsistency of including Resultant luck in our moral assessments.

Bernstein argues that the intention to kill was still there in the Careless Driver case. Furthermore, the death of Gillian is still “counterfactually dependent” on Galen. That is, had Galen not had the intention to kill Gillian, she would not have been killed by him. This “counterfactual dependence” is used by Bernstein to demonstrate that it is hard to create a consistent method of judging agents in cases of Resultant luck since we would argue that the consequences of an action can increase the agent’s moral responsibility; however, cases where the *method* (the actual causal chain of events) of execution is affected by luck seem to grant less moral responsibility to the agent had all things gone according to plan (2019: 155). I do not wish to either support or contest this claim here. Rather, I wish to use the concept of “counterfactually dependent” to demonstrate that including the results of an agent’s intentions leads to an inconsistent view of responsibility.

To do this, consider the following case:

Agent A, while walking on the side of the road, waves at a car passing by in an attempt to be friendly. This wave causes the driver of the car to slow down ever-so-slightly. Because of this, the driver encounters more red traffic lights than he would have. This makes the driver 5 minutes late for work. As a result, the driver is asked to stay late after work as his boss wants to have a word with him about his tardiness. Now leaving the office slightly later than he

usually does, and being irritated and frustrated by all the events of the day, the driver does not focus on the road and hits a pedestrian while driving.

The question I would like to ask is: is A morally responsible for the killing of the pedestrian? Just as was the case with the Careless Driver (as well as with all the cases mentioned in this thesis) the result of this thought experiment is “counterfactually dependent” on A’s waving. Had A not waved, the driver would not have been late or frustrated, and the pedestrian would not have been killed. We can even change the initial motivation for the wave as being a malicious one. Perhaps A wanted to distract the driver but failed. Still, the eventual results of A’s intention/action is the killing of a pedestrian. What I am trying to convince you of here are two things: 1) “Counterfactual dependence” is true of every action ever, and 2) because of this, it would be inconsistent to include the results of our actions as the determining factors in moral judgements.

For the first, consider that every result of actions that ever occurs is counterfactually dependent on some agent’s intentions/actions, even when we do not know who this agent is, or the agent themselves do not even know. The point is, our intentions and actions have hugely unknown consequences once we stop considering only the closest and most obvious results, and to say we are responsible for them because they are “counterfactually dependent” on us would mean that we have to be responsible for so much more than we would have ever thought. This consideration leads to my second point, if an agent’s moral responsibility depends on the results of their intentions, where do we draw the line and stop considering the “results”? Why stop at point 2 in the causal chain and not 3 or 4 or even 20? If our only criteria for consideration are results and counterfactual dependence, then we have to consider much more than anyone will be willing to consider. All of this is avoided when we focus on the agent’s intentions. In the case of the Careless Driver, we would conclude that Galen was responsible for being a careless driver and lucky that we never got to find out about his true intentions.² We would not say that he is morally responsible for killing Gillian, since it might also be the case that had he not killed her she would have killed a pedestrian herself, and Galen would, therefore, be responsible for saving a life. The criterion of “counterfactual dependence” (which is presupposed in cases of Resultant luck) does not narrow the scope of moral consideration, but broadens it to such an extent that we lose a consistent way of saying what agents are and are not responsible for. This is another reason to not include the effects of our intentions in our moral considerations.

All of the arguments above attempt to demonstrate that including Resultant luck in our moral assessments leads to inconsistencies. This gives us a reason to believe including Resultant luck is a

² More on this type of “luck” in my chapter on Circumstantial luck (Chapter 4).

bad idea. But this only becomes a question/problem because we are trying to understand the apparent paradox of moral luck. The above responds to the problem of moral luck by arguing we should not include luck, and thus argued the paradox goes away. It argues that the problem lies in the fact that people are trying to include luck when they should not. As mentioned in the introduction to this thesis, there is a second line of argumentation which tries to answer the paradox by denying the argument for the paradox itself.

Philosophers like Michael J. Zimmerman believe the arguments presented to support the paradox of moral luck are unsound. He argues the premises of the moral luck argument are false and that we, therefore, do not have a problem once these problems are illuminated (Zimmerman, 1993). Later I wish to show that this line of argumentation also lends credence to the above arguments of Sverdlik, but for now, I will focus on Zimmerman's argument.

The first part of the process is to present the "moral luck argument" in its syllogistic form. Zimmerman (1993: 219) presents it as follows:

1. A person P is morally responsible for an event e's occurring only if e's occurring was not a matter of luck.
2. No event is such that its occurring is not a matter of luck.

Therefore

3. No event is such that P is morally responsible for its occurring.

As Zimmerman notes, the paradox seems to occur because the first two premises seem true and yet the conclusion does not. His claim is that, although the argument seems sound, the premises are not true (ibid.). Zimmerman argues that "control", or in the above case "a matter of luck", is not well defined. There are two forms of control, and each describes the reach of luck differently. Zimmerman concludes that only one of these forms of "control" leads to the paradox, but this version of control is not the type that is relevant or that we are concerned with when it comes to morality³.

Zimmerman describes two types of control (or ways something can be beyond a person's control), these are "Restricted Control" and "Unrestricted Control". Restricted Control can be seen as a more limited control of one's actions/environment. A person can be said to have Restricted Control of an

³ For a recent argument which supports Zimmerman's revision of Control see Monica W. Link (2013). Link argues that if we revise our concept of control, both Resultant luck and Circumstantial luck disappear, however, she does not believe that by doing so Constitutive luck disappears (2013: 103 – 104). I argue Constitutive luck disappears by revising our concept of control and the agent (see Chapter 5).

action or situation if they had the ability to either bring about said action/situation or to not bring about said action/situation (Zimmerman, 1993:220). A person can be said to have Unrestricted Control if they had both Restricted Control over the action/situation and all other event/actions/situations that gave rise to the current action/situation (that said action/situation is contingent upon) (ibid.). These two types of control imply different scopes of luck for the individual or, stated otherwise, the person has more control of more things in the second than in the first form of control.

An example of this would be a confrontation between two individuals. We can imagine that you are walking down the street, and for some bizarre reason you are confronted by a person with an opposing view to your own on a particular topic. We can add to this that the confrontation is about a personal belief, perhaps this is one of those cases of “street epistemology”⁴ about God or abortion that went wrong. The important factors to consider now are as follows:

- 1) You did not expect, nor were you in any control over whether this person would be there.
- 2) You are not in control of your personal beliefs in a way that allows you to not have a strong emotional reaction to them.
- 3) You are not in control of this person (or how they act).

Putting questions of free will aside, we will assume that you are only in control over how to react at this moment. You do not possess Unrestricted Control as you cannot control who would be on the path where you are walking, what they would say and believe, nor any other factors that brought you here (besides that you walked this way in good faith). You *do* possess Restricted Control over how you react. In particular, you have Restricted Control over whether you will or will not punch the confrontational person. You also have Restricted Control over whether you will walk away or stay. Although you do not possess the necessary control to have not been in this situation in the first place, you do have control (even if it is limited) over how you react. A simple way to think about it is to realise that we do not have control over the world or our surroundings, but we do (presumably) have control over how we react to these things. For Zimmerman, this distinction is important and leads to two different versions of the moral luck argument from above. They go as follows:

Argument 1:

⁴ Street epistemology is the name given to a method of questioning someone’s beliefs so that all the parties involved in the discussion gain a deeper understanding of why they believe what they do. It has become a popular method of engaging in discussions with people on the street, particularly for activists interested in changing people’s minds. For a good resource on Street Epistemology see Streetepistemology.com (2019).

1a) P is morally responsible for e's occurring only if P was in restricted control of e.

1b) No event is such that anyone is ever in restricted control of it.

Therefore

1c) No event is such that P is morally responsible for its occurring.

Argument 2:

2a) P is morally responsible for e's occurring only if P was in unrestricted control of e.

2b) No event is such that anyone is ever in unrestricted control of it.

Therefore

2c) No event is such that P is morally responsible for its occurring.

These two arguments argue for the same thing, but because they use different concepts of control they should be analysed separately. Zimmerman spends very little time on the first, as he believes it is now clear that premise 1b is false. As already explained, Zimmerman believes we do have Restricted Control. When you are thirsty and decide to drink from the glass of water next to you, you are using your Restricted Control since, although you might not have had control over whether you have water or do not have water at your disposal, you nonetheless do have the ability to either drink from the water or not (this is Restricted Control). Later I will discuss what issues the problem of free will poses to this point, but for now, we can tentatively move forward assuming we either have free will or that our lack of free will has no effect on Restricted Control. He then goes on to analyse Argument 2 and argues that it is also false; but that it *seems* true because it does contain a moral intuition that Zimmerman also believes is true: people can bear the same blame even if they did not both cause the same event (Zimmerman, 1993:222 – 226).

Before I get to the second argument, I want to analyse the first. Premise 1a does seem true. Depending on exactly how we define Restricted Control, I believe most would agree to it. But to make it clear, here is what I believe to be the sufficient control required for moral responsibility:

The agent must have an awareness of their surroundings in such a way as to make moral reasoning possible. That is to say, the agent must have the cognitive awareness to assess their current surroundings and circumstances and make a decision which corresponds to their own personal agency. For a decision to correspond to one's personal agency it has to be made without coercion or external restrictions. The choice must be "free" and correspond to the personality of the agent in question.

The paragraph above has two main criteria; 1) the agent must be a moral agent such that they can think about actions with regards to some moral framework, 2) the agent must in some way be part of this moral process or, stated otherwise, the moral decisions made must conform and originate from this “moral centre” (the moral character) of the agent. This definition is still somewhat vague, but it is all that is required here as it makes a distinction between a wholly free agent who is able to control their surroundings and the agent that is “free enough” so as to make the connection between their decision and their personal identity clear. The criteria are specific enough to exclude non-moral agents, but also vague enough to be able to be used by both those who believe we have free will and those who believe we don’t but still accept that we can be morally responsible. This notion of control also fits our intuitions well enough. I will focus much more on this notion of control and agency in the section on Constitutive luck as it is more relevant there. For now, it is important that I say I find this notion of control acceptable and “good enough” for moral agency. We have an agent who can participate in and use a moral code, and actions and decisions that are based on, that are inspired by, the agent’s personality in question.

With the above in mind, I believe it is clear that premise 1a is true (even if you disagree, I do believe there are other ways of defining this restricted control so that the premise does come out true). If this premise is then accepted, I will argue that premise 1b is clearly false. There are many ways in which, although the agent does not control their surroundings, they can control how they react, the decisions they make, and therefore the actions that follow.⁵ I cannot control the fact that I was born, who my parents are, and how other people act. But I can, under normal circumstances, control how I react, and so there are events such that I have Restricted Control over them. Once we realise that premise 1b is false, we also see that argument 1 is unsound. I believe I have therefore demonstrated if you want to argue for the paradox of Resultant luck using argument 1, you are doing so using an unsound argument, and the paradox therefore disappears.

The above is enough to argue away the paradox of moral luck, and Zimmerman (1993: 223) agrees, but it is still important to address Argument 2 as it can give us important insights into why the paradox does seem so intuitive. In the second argument, it is clear that 2b is true: people do not have

⁵ As before, we are for the sake of argument assuming that humans do have some form of free will. Later in the chapter on Constitutive luck I will focus much more on this and will discuss a conception of agency which allows for moral responsibility, thus showing that we do not need free will (at least the classic conception of free will) for us to still have Restricted Control (which will also be discussed in the chapter on theoretical implications). For now, it is easier to leave the problem for later.

unrestricted control over their environments and actions. This is not the problematic premise in the second argument. Rather, it is 2a that is the problem. As Zimmerman notes (1993: 222):

After all, 1b [*referring to 2a here*] is clearly false, if only because no one is in control of his being born – an event on which all of his decisions, actions, omissions, and the consequences thereof are contingent. And we all recognize this. Why should anyone think that our received conception of moral responsibility implies otherwise?

Zimmerman has chosen the ultimate event on which everything else is contingent, one's birth, to demonstrate that we do not need control over every event in order to be held responsible for some. In the same way that it can be said that you do not need control over every facet of a situation (like your birth) to be held responsible for those things you do have control over, you do not have to have control over every part of a circumstance to be held responsible for those parts you do have control over. I did not need to have control over the country I live in, or where my parent chose to live in said country, in order to be held accountable for my action of harming my neighbour, even though I had no control over who my neighbour would be. If the claim is that we need Unrestricted Control in order to be held responsible, then we need control of all events on which our actions and decisions are contingent upon. The example of my lack of control over my birth is an extreme example which demonstrates that this claim cannot be consistently ascribed to. Why does it matter that you are not in control of everything? I would argue it does not. It is because of this false premise that we can once again deem the conclusion of the argument false. But why then does the argument still hold such an intuitive appeal? Zimmerman (*ibid.*) argues it is because we tend to praise or blame someone for a good/bad decision/action more than someone who did not make/take said decision/action, even if they did not do so only due to factors out of their control (due to luck).

I have given many examples of the above, but just to refresh our memory, here are a few again:

- 1) We blame the Nazi collaborator more than the non-collaborator, even though the reason for the difference are factors beyond both their control (such as their parents moving/not moving out of the country).
- 2) We blame the drunk driver who hits a child more than the drunk driver who does not, even though the reason for the difference are factors beyond both their control (such as a child being/not being in the road).
- 3) We praise the person who saves a drowning child more than the person who does not, even though the reason for the difference are factors beyond both their control (such as being/not being first on the scene).

The above are hard to justify and can lead people to accept the problem of moral luck. It is hard to argue that the person who is responsible for something (in the causal sense) is just as innocent (in the moral sense) as someone who is not responsible for the same thing (in the causal sense).⁶ The problem of moral luck, derived from examples such as the above, rely on focusing on the agent that does not deserve/receive moral blame/praise and then concludes that the other therefore also does not. This is what gives rise to the paradox. I, however, and Zimmerman, argue that this is a mistake. Rather, we can conclude that the person who is praise-worthy or blame-worthy stays so, and that therefore the non-collaborator/drunk driver is also liable to blame/praise. In our intuitive reaction to these cases, we, therefore, make a mistake of inference. We focus on the non-responsible (because we focus on actions, and they “did” nothing) and extract from that case and apply our findings to the “traditionally responsible” party, but this inference is backwards and should actually run in the other direction.

The way we can show this to be the case is by again focussing on the arguments presented before these, those of Sverdlik. Sverdlik argued that it is logically more consistent to not include Resultant luck in our moral systems and used cases of transferred intent to demonstrate this. What was further demonstrated there, and with the rebuttal of Argument 1 above, is that the morally relevant criteria in these cases is intention/Restricted Control. Note that in this case these two roughly refer to the same phenomenon, at least from our perspective. The control that you have is at the very least Restricted Control (this is what we would call sufficient control) or, stated otherwise, you can only control what you intend to do, not everything about the situation. Briefly stated: what matters are your intentions and luckily you have control over them (in the Restricted Control sense). As used here, intentions and Restricted Control refer to the same phenomenon. I had mentioned at the start of the section on Zimmerman that his arguments can be used together with Sverdlik’s to make clear what is going on here. When we use Zimmerman’s conception of Restricted Control along with the arguments presented by Sverdlik, what we find is a clear and coherent account of the morally relevant aspect of responsibility. Once we agree that intentions are what matter when making moral evaluations, we can ask if we at least have control over them. In a later chapter on Constitutive luck, I will take a closer look at this question. But for now, we can say that we do not have complete control over how our actions turn out, but we do have control over our intentions. Linking the two arguments from Sverdlik and Zimmerman, we find that the control we need (Zimmerman’s Restricted Control) are our intentions (what Sverdlik argues are the morally relevant criteria). From this, we can conclude that people can still be held accountable even though they do not have full/Un-restricted Control over

⁶ For a discussion on this distinction between these two senses of responsible, see my chapter on the implications of this thesis (specifically section 7.2.1).

their environments/actions and that this accountability can be said to apply to those with similar intentions, even though the results were different. Imagine another example of two thieves: T1 and T2. T1 corners someone on the street and threatens them with a weapon in order to get the person to give them their money. T1 is not stopped and is successful. T2 does all the same things, however, while in the process, a police vehicle drives by, scaring him and causing him to flee before the act is completed. While T1 was successful and T2 not, we can agree that both attempted to rob someone while threatening them with violence, which is an act that is morally reprehensible. While we might agree that T1 is also responsible for returning the stolen goods, both T1 and T2 are equally responsible for the morally atrocious act of robbing someone, and we can conclude this because they both wanted, and attempted the same thing. The reason these thought experiments are so successful in creating the paradox of moral luck is because they confuse our intuitions and lead us astray in our inferences.

What then should be taken away from the above arguments and analyses? Sverdlik demonstrates that it is logically incoherent to allow luck to infiltrate our moral assessments. Using the cases of transferred intent, Sverdlik demonstrates that it is a bad idea to include Resultant luck in our assessments because including them leads to an inconsistent and problematic moral system. Following this, the arguments from Zimmerman were used to demonstrate that the argument used for the paradox of moral luck was too broad and vague to draw conclusions from. Once we made the relevant factors clear (control) we were able to show that the conclusions of both arguments do not follow since at least one of each of their premises are demonstrably false. Resultant luck is therefore no longer a problem. We can hold people responsible for those things in their Restricted Control, and this remains true in all the cases above.

The authors discussed above all support what Susan Wolf (2013) calls the “rationalist position”. They argue that the paradox relies on unsound arguments, bad inference, and bad logical structure. If we think rationally about moral luck (most focus on Resultant luck) we will conclude that the paradox disappears. Some more recent authors who agree with this position are Darren Domskey (2004) and Edward Royzman and Rahul Kumar (2004). David Concepcion (2002) has summarised most of these arguments as saying that once we have made the distinction between “causal responsibility” and “moral responsibility” clear, we can see that the paradox of Resultant luck disappears. What follows next are those who support what Wolf (ibid) calls the “irrationalist position”.

3.3 The Virtues and Merits of Resultant Luck

There are those who argue that we cannot get rid of luck in ways such as the above because luck is an asset. They argue if we remove luck from our moral frameworks, we lose something important to

morality. Luck, then, has something to add to morality or is an important part of it.⁷ This is accomplished in a few different ways. Some argue we cannot get rid of luck because we require action for morality, and action is always in danger of luck. Others argue that luck adds something good to morality, a certain beauty to the entire practice. There are also those who argue that without the inclusion of luck there would be bad consequences and a general diminishment of moral agents. In this section, I will discuss some of the most influential points.

A good place to start is with the arguments of Martha Nussbaum as she begins with the original thinker who believed luck was simply a part of morality: Aristotle. The arguments put forward by Nussbaum can be split into two, the first being more relevant to Resultant luck and the second to Constitutive luck, but both are closely related. In this section of my thesis, I will focus on the first part and leave the second for my later chapter on Constitutive luck.

Up until now, most of what has been said was framed within the context of analytic philosophy and Deontological Ethics⁸. Kant's great influence on ethics has led many to take such a narrow approach to ethics, but there are those who look at the issue from a very different lens, that of Virtue Ethics⁹. This is a moral framework rather different from Kantian ethics, that focuses much more on the agent in the world and the impacts of the world on the agent. It is from this perspective that Nussbaum approaches the issue of moral luck.

Nussbaum (1993: 73) starts her argument with a poem about a vine tree: "But human excellence grows like a vine tree, fed by the green dew, raised up, among wise men and just, to the liquid sky.". The poem is meant to capture something about human life, and in this case, moral life. Nussbaum argues, just like in the case of the vine tree, we too are shaped by the world outside. The tree requires good soil to grow, just like we require a good home and community to be moral agents. The tree requires sunlight to gain the necessary energy, just like we require good food. And just as the tree is

⁷ Darren Domsy (2005) gives a good summary of all the arguments that have been put forward in support of the idea that moral luck might be somewhat of a problem, but that it is better to have it than not. He argues that they all fail since we do have a solution to the problem of moral luck.

⁸ Deontology is a theory of normative ethics which says the right- or wrongness of an action should be determined using a set of rules rather than the consequences of said action. Immanuel Kant's "Categorical Imperative" is one of the most well-known deontological rules (Kant, 2011).

⁹ Virtue Ethics is a theory of normative ethics which focuses on the virtues of the character and his/her mind. Rather than being about rules and consequences, Virtue Ethics holds that to be a morally good person said person should have a "good" or virtuous character. Although many of the principles are found in the works of Socrates and Plato, Aristotle is most well-known for his development of Virtue Ethics (Hursthouse & Pettigrove, 2019).

at the mercy of the weather, we too can be broken or beaten by the world around us. We are worldly creatures and subject to the world. More importantly to this section of the discussion, it is not simply enough to have good potential as a seed, but that seed must prove itself by growing. Nussbaum argues that we are the same, as morality requires action and not just a good character (1993: 80). From this premise follows a conclusion: that if morality requires action, and action is not immune to luck, then morality is not immune to luck. In other words, because you must act in order to be a moral agent, luck will always intrude on the moral sphere since luck is intrinsically bound to action. No action is isolated from the world around you, and so luck makes its way into our moral evaluations. This argument contains two important claims: 1) morality requires action, and 2) the luck that infiltrates action therefore also infiltrates morality. I will discuss these separately.

Nussbaum starts her discussion on action by first stating that the “good-condition theorist” (this is what Nussbaum calls those who believe morality only requires “good states” or “good intentions”, Sverdlik’s “equivalence theorist”) believes that since these conditions (or mental states) are immune to luck, and under the agent’s control, we do not have any problems with luck since we have removed it from our moral assessment. However, Nussbaum argues, this view is lacking something as it does not fully capture what is going on when we talk about a “moral agent” (1993: 81). To demonstrate this Nussbaum asks if we can call an agent with “good states”, but no action, a “good person”? Imagine if an agent is born and raised, and throughout this, they are taught to be moral agents, but one day they fall into a deep coma. This coma is not one of brain death, but rather, the agent still has all their internal states, their will to be and do good, but these are all internal and the agent is basically in a deep sleep. Nussbaum then asks, “Can such a person be said to be living a good life?”. Nussbaum argues that the “good-condition theorist” must accept that we can praise and blame such a person since all the necessary conditions are present. But this seems counterintuitive and wrong. We most certainly do not do this, and most would argue we should not. As Nussbaum (1993: 82 – 83) puts it: “just as we do not think a fetus, who lives a purely vegetative existence, without awareness, lives a full human life, so we are not going to be willing to praise and congratulate the life of this hopelessly inactive adult.”.

Another analogy used by Nussbaum is one of a runner that does not compete. She argues we only praise those who actually compete, not simply those we believe to have the best abilities. Even though it might be said that these people possess talent, until they compete and demonstrate this talent, there is not yet any room for praise and blame, as these require action. Again, as Nussbaum (1993: 83) puts it:

Just as our assessment of people as runners depends upon there being some actual running (though of course, they depend, too, on our belief that this good running

was caused by their good condition, not by some external force), so too our ethical assessments are based on actual effort and activity, as well as upon the presence of a stable character that is the cause of the activity.

All of the above is meant to serve to convince us that moral character judgement requires action. Nussbaum is not claiming that we cannot focus on intention when making moral evaluations, rather, whether we do or do not focus on intention, we must focus on actions as well, and this is where luck infiltrates our judgements. In the next chapter, I will discuss some examples of where this is not the case, and we do indeed judge people for their intentions alone. But for now, this is not necessary as we can accept the above argument and its claims, without having to accept that luck is an intricate and intrinsic part of morality, and/or that this isn't really a problem. For the moment then, I will agree that moral agents require action, but not because you cannot be a moral agent without action (as Nussbaum is arguing), rather, because *moral judgement* requires action. It is, after all, what morality is for, to make the world better and to organise agents in it in terms of which actions they should and should not perform. *Action*, then, is at the core of morality, otherwise, it is nothing but a theoretical game. For the virtue theorist action is the final enactment of the agent's virtue and is therefore equally important since virtue means nothing without action (Nussbaum, 1993: 84). The problem I have with the above argument is the second claim that is made, that because morality requires action, and action cannot be immune to luck, morality can therefore also not be immune to luck. There is, in a way, a transfer¹⁰ of luck from the one state to the next, and this transfer is what I disagree with. To make this clear I will put the former argument in syllogistic form:

- 1) Being good, being moral, requires action and doing.
 - 2) Actions are susceptible to luck or causes out of the agent's control.
- Therefore
- 3) Moral judgements are susceptible to luck because they require action.

It is the conclusion of this argument I disagree with, as it makes a claim about transference that is either not justified or simply not required. There is thus a mistake of inference.

¹⁰ A helpful concept to understand what Nussbaum is saying is a "transitive property". In formal mathematics, the Transitive Property of Equality states that if something relates to one element and also relates to another, and this second element relates to a third, it can be inferred that the first element also relates to the third. Stated in formal terms: if $a = b$, and $b = c$, then $a = c$. In the above case luck is being treated as a transitive property, being transferred to moral judgement from the world.

Let us take the analogy of a racing car and its engine to make my point clear. For this analogy, we can imagine the car's engine as its "moral core", or that which does the work that the character of the agent would do in day-to-day life. Just as we would judge the moral character of an agent, so we can judge the "performance" of the engine. Just as above, I will agree that this requires action (a real-world performance) and then subsequent assessment. It makes no sense to build an engine, and then, based on its theoretical specifications, make a judgement on its performance. The engine must be put into a car, and this car must then be tested against others. Exactly as is the case with people, the engine is subject to what I will call here "performance luck". How it performs also depends on how the body of the car itself is designed, which tyres are used, the track that is being driven on, and any unlucky accidents like pebbles and loose rubber on the track. The engine must be tested in the real world and here its performance is susceptible to luck. But does this luck of the world entirely determine the performance of the engine? Might it not be said after the testing, "All things considered, the engine performed well/poorly." If the car were to hit a pebble and lose its grip to the road, would we judge the engine? Could it not be said that had the car been designed with a better body, more adept at cutting through wind, the engine would also have performed better? Although the engine requires a car; and external action to be judged, we can still judge the engine fairly by taking into consideration these factors and taking their impact into account in our final verdict. Yes, morality requires action, and yes, action is subject to luck, but this luck can be taken into account, and our final "moral consideration" can be wholly directed at the agent without having to transfer this "worldly luck" onto the morally relevant factors that we should take into consideration.

In the previous section, I believe the above approach was shown to be possible and preferable. The case of the two drunk drivers demonstrates that we require an action (which demonstrates their intention/negligence, in some sense) to judge the driver, but the factors which we would deem related to "luck" are accounted for. We say that both drivers are equally to blame, even though the results of their actions were different, and we do this because we use their actions to identify the relevant moral considerations, i.e. their intentions. In the drunk driver example, it is not the case that either of the drivers had malicious intent¹¹, rather the moral failing in this example was their negligence. They were meant to care about others on the road and therefore not drive while intoxicated, but chose to do so anyway. Both are guilty of this, while only the one is actually responsible for killing someone. For Nussbaum, these two characters are not equally to blame in the moral sense, the one did

¹¹ Intention not as the directedness (the way philosopher of mind use it) but rather as what the character wanted (the everyday notion of intent, to do it on purpose), otherwise stated: what an agent wants/wanted to do, is what the agent intends/intended.

something different from the other. But once we take into account all the arguments above, we can conclude that both agents are equal in their moral blame since their moral fault was the same. I will add again that this does not mean the results have no effect on our evaluations and punishment. They help us detect the underlying intentions of the situation and might have a bearing on what the punishment itself is. I will explain and justify these ideas in a later chapter. These extra factors aside then, morally speaking, both drivers are equally to blame. We can say this because we have assessed and identified those factors which are morally irrelevant (i.e. the results), even though we judge the agent based on these actions.

To summarise then: morality requires actions, and these actions are not immune to luck. But these factors of luck can be accounted for once moral judgement occurs, and thus the transfer of luck can be mitigated through rational moral deliberation. Because of this, luck does not have to determine how our moral judgements turn out. Although Nussbaum is trying to argue that we cannot get rid of luck while performing moral judgements since luck is part of our actions, it seems clear that we can, in fact, do this through careful analysis of events and agents. We have, though, hit a crossroads again since it might be said that this mitigation of luck is harmful in some way. That is to say, we again find ourselves in a situation where it is better to ask whether luck should or should not be included, rather than asking whether it is or is not. I have argued in the previous section that it is better to not include luck than to include it, and this was done by appealing to the notion that moral systems/frameworks must be internally consistent. There are those, however, who, rather than appeal to consistency, appeal to the greater effect of moral systems and argue some are better at doing “moral things” than others. That is to say, they believe that including luck somehow strengthens the effect of morality and makes us better moral agents, whatever “moral” means here. One such author is Margaret Walker. She argues if we eliminate luck from our moral assessments, we risk losing something important, and in the long run, we only diminish moral agents. Her arguments rely on two key analyses, which Walker uses to demonstrate that there will be a diminishing effect on our moral characters once we eliminate luck.

Walker uses the terms “pure agency” and “impure agency” to distinguish between moral systems that exclude luck and those that include luck (1993: 241). Moral systems that try to eliminate luck from the equation, and treat the agent as having a faculty wholly under their control, could be said to accept moral agents as “pure agents”. The opposite, where moral systems allow luck within their judgements, is then said to be instances where agents are treated as “impure agents”. Walker argues that “pure agents” will be far less superior moral agents than “impure agents”, and this is because of two different results that follow from the practice of including luck. The first relates to how agents react

and judge themselves in light of actions, and the second relates to how these agents will accept and view responsibility in light of a lack of control over the results of said actions.

Let us suppose we are talking to someone who has done something, in this case, something “bad”, but again we are dealing with a scenario where the bad that has happened was due to luck (in this case specifically Resultant luck). Walker argues that there are many things we can argue about, but one thing she believes would strike us all as odd would be if the agent said something like the following¹²:

It’s really too bad about what happened and the damage that’s been done, but my involvement was just a happenstance that it was my bad luck to suffer. I admit my negligence (dishonesty, cowardice, opportunism, etc.) and accept such blame as is due these common faults. But it would be totally unfair of you to judge, let alone blame me for unlucky results and situations I didn’t totally control and stupid or masochistic of me to let you. (Walker, 1993: 240)

For Walker, this sort of response is possible as the agent has accepted the “pure agency thesis” and is now able to give such responses, responses which would be deemed by most as inappropriate and unbecoming a moral agent. An easy reply to this point would be that Walker has deliberately chosen a tone that seems wrong and inappropriate for a moral agent to take. I could create my own hypothetical response, which goes something as follows:

I recognise that I have done wrong. It was my decision to be negligent (dishonest, cowardly, opportunist, etc.), and I, therefore, deserve all judgement and punishment that should occur due to this. I recognise that my actions have had bad consequences, but these consequences in themselves should not be the basis of my desert or punishment as they could have been otherwise and I had no control over them. My fault was a moral one, a case of bad agency and neglect, and thus my moral condemnation should be based on this. Of course, the result of my action should be rectified, but this is a practical consideration and not a moral one.

This to me does seem a morally appropriate response given the situation. I will also add that what is missing here is the emotional effect that you can only truly access when in a real conversation with

¹² The argument made here by Walker makes an appeal to intuitions. Intuitions play a large role in how we see things and the conclusions we accept. I do not make any appeals to intuitions since they are often varied, but it is very useful to understand intuitions while discussing cases of moral luck. Neil Levy (2015) wrote an extremely interesting paper on why we have the intuitions we have about Resultant luck. He argues once we understand why we believe what we believe about Resultant luck, we will see that we are mistaken in accepting our conclusions of the paradox.

someone. This is something that is also lacking in Walker's example and adds to the intuition that she is trying to elicit. The important point though is this: the "pure agent" does have a response which is appropriate to a moral agent, despite Walker's claim.

Walker does also argue that taking the view of the "pure agent" will have other effects such as the agent becoming "removed" from the real world and complacent towards the complex chain of events they find themselves in (1993: 241). She argues that the "impure agent" is better equipped and more proficient at navigating the moral terrain as they accept that not everything is in their control and yet still try to mitigate this. The claim is thus that the "pure agent" hypothesis has a disadvantage when compared with the "impure agent" hypothesis as the latter allows agents to be more aware of the world around them and thus allows them to be more willing to accept responsibility. Walker argues these results are positive ones for the "impure agent" view and leads to better moral agents.

I would argue that this is in fact not true. I accept that, since moral systems are not only meant to judge moral character but also to build and foster better moral character, they should include or exclude luck based on which approach facilitates the latter better. Walker argues that since agents who reject "impure agency" have less to be responsible for and do not have to accept judgement based on luck, they become less aware of their moral power and influence (1993: 245). I argue that the opposite is true. If the judgment is based on factors beyond the agent's control, the morally relevant criterion is no longer the agent but are now the result of their actions and contingencies in the world. Since these are not part of the agent, and indeed are things the agent cannot control, it is hard for an agent to relate to the judgement (desert or punishment). If, however, the morally relevant criterion is the agent's character/intention, the agent is much more likely to understand what exactly he or she did wrong. If whether you are good or bad depends on things out of your control, why would you want to control whether you will be good or bad? If what you are judged for depends on who you are, you will want to control and change who you are.¹³ Walker's point then is that "impure agency", on average, leads to virtuous characters, but I would argue that the opposite is true. If what I am being judged for is "myself", I am much more likely to be willing to internalise said judgement, thus forming a better moral character (a virtuous character). Of course, none of this is necessarily factual but rather speculation. My claim is that it seems at least equally plausible for both positions to create moral agents who care about who and what they are. I personally find the "pure agency" view much more consistent and fair with regards to moral agency, but even if both views about agency are true,

¹³ A friend of mine noted that this is only true if the agent in question has a drive for moral responsibility. Although it does not matter for the purposes of this thesis, my own perspective is that moral motivation comes down to a subject decision to be moral. For the purposes of this thesis I will assume agents are appropriately motivated to be moral.

it could then be argued that the “pure agency” view is “better” because it is also compatible with all the arguments in this thesis that were made before this one.

The second line of reasoning Walker gives for the inclusion of luck in moral judgements is that by not allowing for luck we will create agents that “care” less. If we are only responsible for those things in our control, then we will not be responsible for many of the things that we usually are responsible for. As Walker (1993: 245) puts it:

Relationships, situations, and encounters in which emerge uncontrolled and uninvited needs, demands, and opportunities to enable or harm will not be thought to ground morally legitimate claims upon us or in our behalf, in ways we might have thought or hoped.

To use an example from Walker, if I decide to have a child, then the having of this child is my moral responsibility, but should the child become sick, something I did not choose (or control), then I will not be responsible for helping said child. Walker argues that this is problematic and counterintuitive. Another example is that of becoming friends with someone and reaping all the rewards, but should that person be going through a hard time, something not in my control, then I am not responsible for helping my friend. Walker argues that the above cases show that we do need to, and in fact just do, accept responsibility for things out of our control (1993: 245 – 246).

The question is, do we really do these things even though they are out of our control, or is there another reason for doing these things? In the first case, that of the child, I would argue that the things which motivate care and treatment of the child is the implicit agreement to care for this being since they are dependent on you. Before the child is even born, the role of “parent” is accepted and internalised. With this comes the complete care of the child, whether they are sick or not. The responsibility for the child is not based on control or lack of control, but rather on the accepted role of “parent”. Do you want to be a parent? Then you should care for a child when they become sick, whether or not you caused the illness. In the example of the friend going through a rough time, it seems strange to me to say you are responsible for caring for this person in the same way that one is responsible for a child. Rather, the decision of care is based on the fact that you wish to have friends. This also comes with expectations, if not responsibilities. If a friend of yours is going through a tough time, and you do not help, then you are probably no longer their friend. It seems to me to be a “hypothetical imperative”; IF you want to be a good friend, THEN you should care for your friends. Being a parent holds benefits for you (although these might be nuanced) and being a friend holds benefits for you, and, therefore, you assume the roles associated with these “identities”. Certain roles come with certain responsibilities. Being a parent means caring for the child, being a friend means

caring for your friend, and being a driver means you should be conscious of your own abilities and of other people who use the road. The cases used by Walker are examples of responsibility in terms of what is expected of us rather than being causally responsible *for* them. These are two different senses of “responsibility”. When dealing with Resultant luck we are dealing with the latter rather than the former. Walker seems to be equivocating¹⁴ the two, and this is why her analysis is ultimately unsuccessful. In her examples, we are dealing with the responsibility that we owe to others due to our relationship with them, rather than the type of responsibility due to a moral failing on our part. We have to make sure to distinguish between the responsibility associated with what is expected *of us* and what we are morally responsible *for doing* because we are somehow sufficiently causally responsible for it. In conclusion, Walker’s thesis that the “impure agent” is less equipped to be a moral agent is false.

As has been mentioned, authors such as Walker support what can be called the “irrationalist position”. They argue that there is something right about including luck in our moral evaluations even when such a thing cannot be said to be rational. Susan Wolf (2013:9 – 11) argues that although we can agree with the rationalist about moral luck, the irrationalist position does capture something intuitively true, namely: we would expect people to have different emotional reactions about themselves depending on the consequences of their actions. Wolf believes this is why moral luck seems so intuitively true. In the section on theoretical implications (specifically section 7.2.4) I will consider this intuition and argue that we can still hold it to be true without accepting the paradox of Resultant luck.

3.4 Conclusion

In this chapter on Resultant luck, I have looked at many different ways to deal with the problem, but these can be split into two main categories: 1) those who believe that luck diminishes moral responsibility, and 2) those who believe luck increases moral responsibility. Throughout this chapter, I have attempted to support the former and argue that the latter is not true. I hope to have shown in the first section that there are good reasons to believe that factors due to luck should not be included in our moral judgements. I did this by first showing that it is internally inconsistent to include luck in our moral assessments, as this leads to strange cases where people are blamed for things we would

¹⁴ An equivocation fallacy is when two ambiguous concepts, distinct from one another, are used interchangeably as if they mean the same thing, thereby drawing incorrect conclusions based on this misunderstanding. For an interesting read on fallacies see Hansen (2019).

usually not blame them for, and cases where we wouldn't blame people even when intuitively we would have done so otherwise. I also showed that the very argument used to create the intuition of the paradox is based on too broad a conception of "control". Once we define our concept of control, we see that there are in fact factors which we have "control" over, and which can plausibly be the focus of moral judgement. Finally, I looked at arguments that claim that luck can either not be eliminated or that it should not be eliminated. In response to the first claim, I countered by showing that there is a mistake of inference and an unjustified, and unnecessary, transference of the property of luck from the action to the intention and therefore the responsibility. I also tried to answer all concerns raised by Margaret Walker and argued that the very opposite of what she is arguing is true; that we create better agents by focusing on the agent themselves, rather than focusing on the complex chain of causality we live in. All of the above serve to demonstrate that it is better to eliminate luck where we can when judging an agent. This section on Resultant luck is only a small part of the argument, and only once we have considered the other forms of luck can we draw more general conclusions with regard to moral luck. For now, I believe I have given some good reasons not to allow Resultant luck to infiltrate our moral systems.

Chapter 4: Circumstantial Luck

4.1 Introduction

Let us imagine we are all part of a contest, one in which the winners are those who score above a certain grade. Now imagine that this grade is based on a set of challenges each individual has to complete. At the end, all the points scored by each individual in the contest are tallied. For the sake of argument, we can say that those who score higher than 80/100 are rewarded, and those who score lower are punished. So far so good, as all are treated the same and all follow the same rules. But let us now imagine that there are many different challenges, some more difficult than others, but they all count the same amount at the end. Furthermore, no individual can choose his or her challenges; rather, different challenges are assigned to different people. Moreover, the distribution of these challenges is not based on merit or status, they are simply given to people at random. At the end of the day then, each individual will have to complete each challenge they have been assigned, and in total score above 80/100 if they are to receive reward rather than punishment.

Let us now consider two different individuals in this contest: the first receives a series of challenges that, relative to the others, are easy to complete and get a high score on. The second person receives a series of difficult challenges. When all is said and done, the first person does very well and scores above 80/100; the second does not do so well and does not meet the minimum requirements for reward. The question then becomes, was this situation fair? Take note that if the first person had received the same set of challenges as the second, he too would have failed. Had the second received the easier set of challenges, she too would have passed. Each, within the rules of the contest itself, deserved their punishment as each competed without hindrance of their own skills. But, could it be said that the contest itself was fair? Could it be said that although both competed in the same contest and followed the same rules, that the contest itself was not fair since people faced different challenges? This is the question posed by Circumstantial luck, and it really does pose a large threat once we consider its implications.

4.2 Not Moral, but Epistemic

The easiest way to understand Circumstantial luck is by shifting our focus to the environment rather than the agent. It is not the case, as it was with Resultant luck, that our actions and their results

themselves are now no longer under our control, but rather that the world where we find moral judgement and the challenges we face are not in our control. Different people face different situations and challenges where a decision must be made, and even though we might say the rules by which they are judged stay the same, the things they are judged for are part of a larger web of situations that they do not have control over.

Using the example given in the earlier chapter on Resultant luck by Nussbaum (1993: 83) of a runner competing in a race, we can see that it is no longer about the different runners in the race, facing different forms of luck on their way to the race and during the race. Rather, it is now the fact that these runners are being praised for their skill, while someone else, one who cannot afford to compete in such a competition, does not even receive the chance to compete and earn praise. It is important to note that, had this person been allowed to compete, he would have performed just as well, but due to his circumstances, he was never given this opportunity. If, as I have argued above, responsibility (punishment and reward in the moral sense) lies with the agent and her intentions, then we have a problem, as, in this instance, both individuals (the competitor and the non-competitor) have the same ability and intentions, but we reward the one and ignore the other.

The example of the runner above is used by Nussbaum to argue that morality requires action, but in this case, the issue does not relate to how well the runners perform, but rather to whether they even get the chance to perform and hence be subject to the judgements that accompany this. Nevertheless, in this instance, the example is again used to demonstrate that morality and moral character require action – and, in this case, circumstances – and then serves to support the argument that the luck present in the circumstances we face transfers to the resulting moral assessments. As I have made clear in the previous chapter, I agree with the claim that morality requires action. However, I objected to the claim that since such action is not immune to luck, morality is not immune to luck. With regard to circumstances, the example of the non-runner can be used to demonstrate that the agent lacks control over her circumstances and environment. She does not determine whether she gets to run or not. Nevertheless, the same objection that I made regarding Resultant luck can be made here again: the luck of one's circumstances need not be transferred to moral responsibility. To make my case, a thorough examination of how circumstances affect our judgement is necessary here. Once this is done, I will attempt to argue that the luck present in our circumstances once again does not need to be *moral* luck.

Nussbaum (1993: 84) summarises the threat of Circumstantial luck in the following passage:

The good condition of a virtuous character, like good athletic conditioning, is a kind of preparation for the activity; it finds its natural fulfilment and flourishing in

activity. To deprive the person of that natural expression of the condition is to make a difference in the quality of the person's life. It is to make the condition fruitless or pointless, cut off. Just as a runner who gets into good condition and is then prevented from running would be pitied more than praised, so we pity the virtuous person in situations of impediment. Activity, *energeia*, is the coming-forth of that good condition from its state of concealment or mere potentiality; it is its flourishing or blooming. [...] Like an actor who is always waiting in the wings and never gets a chance to appear on the stage, it is not doing its job, and, in consequence, is only in a shadowy way itself.

Nussbaum thus argues that because we need the world for morality, because we need action, and because the world is different for each one of us, our moral evaluation/standing will be different, even through no fault of our own. It is important to again remind ourselves that most would agree that morality ought to be fair, that it ought to apply to everyone equally. We want to be held, and hold people, morally responsible for those things that are in our control. If what we are responsible for is determined by circumstances which we do not have control over, there is an opportunity for people to be treated in different ways through no fault of their own. This leads to a morality which is not fair, and which treats people differently. Returning to the example of the would-be Nazi collaborator and the true Nazi collaborator, mentioned in the introduction, we must ask why the one should be judged (the true collaborator) with moral indignation while the other (the non-collaborator, which theoretically could be any person) is not. We are again faced with a moral judgement which is based on factors not in control of the agents; this seems unfair. When we say the circumstances that are presented to different agents are morally relevant, we allow luck to infiltrate our moral evaluations.

The above claim about people being differently judged only because they were given different opportunities is easy to support using examples. Norvin Richards (1993: 73) uses the example of a drowning victim and the hero that saves him. Richards tells us a story of a person (H) who finds out that someone was saved from drowning by a bystander the day before, and that this bystander is now being praised. This person argues that, had they been on the scene, they too would have jumped in and saved the drowning victim. It was only because he had the misfortune of not knowing of the crisis that was happening that he did not respond and step in. As Richards (*ibid.*) puts it: "So you missed your chance to be a hero, and, thus, your chance to deserve the praise bestowed on the woman who was at the right place at the right time." The same sort of problem can be demonstrated for negative evaluations as well. Consider person V who finds himself standing in front of a rude, obnoxious, and racist man. This man is screaming and professing the inherent inferiorities of certain racial groups. This makes V very angry, and as time ticks on, his anger is converted into true resentment, which

causes V to take a swing and knock this man down. Now we can grant that this might be an action that will bring about some relief to V, but none-the-less, it was an unnecessary use of violence and V is negatively judged for this. As news of this incident spreads, we notice that a few people profess that had they been in that situation, they too would have hit the man, yet, these people are not morally blamed in the same way as V. To use the words of Richards, they “missed their chance” to be quick to anger and violence, and, thus, their chance to deserve the blame bestowed on V, who was in the wrong place at the wrong time.

For the sake of argument, we will use the first example given, but the argument applies to both kinds of situations. In this case, we find it strange that H complains about not receiving any praise for his potential heroism. We think it wrong for H to want praise since, as H should also understand, he did not do anything good. He did not do anything at all. But do we find it strange because we do intuitively allow luck to make a moral difference, or do we do it for other reasons? Richards argues that if we support this intuition with the claim that we do not generally morally judge people for what they *would have done* we would be mistaken. There are many cases in which someone’s counterfactual action or motive *is* the object of moral evaluation. To make his case he asks you to suppose that someone you thought was your friend abandons you in a time of need. Now imagine that, while discussing a friend’s conduct with another person, you realise that, had this second person been with you during this time of need, *she* would have acted as a true friend and helped. Richards suggests that this realisation would cause you to praise this second friend as a good person? Or, conversely, he suggests, “Suppose you were convinced that a repairman would have stolen your wallet had he not realized he was under scrutiny. Might not that be sufficient basis for treating him rather harshly?” (Richards, 1993: 173). Richards argues that these examples demonstrate that we do indeed judge people for what they would have done counterfactually.

The cases above serve to demonstrate that the problem of Circumstantial luck does not capture all of the morally relevant criteria such as counterfactuals and moral character. They do not serve to counter the argument for Circumstantial luck, rather they are meant to show that we do blame people for hypotheticals. At the same time, we would still say that there is a moral difference between the hero and the potential hero, between the villain and the potential villain. We now have to ask whether this difference is due to luck or perhaps something else. Richards argues that this difference is one of demonstration (1993: 174). He argues if you have the sort of character that would make you a hero, then at some point this would be demonstrated. Or as Richards (*ibid.*) puts it: “My central contention will be that if the potential agent is as much like the actual one as we are imagining, then there will be something else in his behaviour which will call for the same response.”. In other words, his claim is that *if* you are a hero, you will *act* like one, and *if* you are a villain, you will *act* like one. Richards

thus argues that the luck does not lie in the moral evaluation itself, but rather in when this evaluation occurs and to what extent we can be certain of it.

In order to demonstrate this Richards (1993: 174 – 175) goes back to the classic example of the Nazi collaborator and the non-collaborator. As before, the collaborator is such because his parents stayed in Germany and the non-collaborator is such because he moved away before the rise of the Nazi party. We then describe the counterfactual as being that, if the non-collaborator had stayed in Germany, he too would have been a Nazi and would have committed the same atrocities as the collaborator. But in order to say this, we would also have to acknowledge something about both the collaborator and non-collaborator's characters and personalities. The reason why they both would have collaborated if presented with the same circumstances is because both share traits that make them susceptible to such political movements. We could thus imagine that both share a similar desire to please those in authority, a general lack of compassion, and little aversion to the suffering of other people they deem "lesser" or "inferior".

If we keep all of this in mind while creating our hypothetical thought experiments, we can start to see where the relevant criteria lies when judging people. Richards states that the reason we treat the two cases differently is because 1) the protagonists' behaviour does not demonstrate their moral desert equally, and 2) we should judge people only according to our understanding of their behaviour and character (1993: 175). Richards thus argues that there are cases where we do treat people differently even though this is only due to different circumstances, but at the same time, there are cases where we treat people similarly even when the one is only hypothetically accountable. The difference between these two sorts of cases lies in the *knowledge* that we possess. As with Resultant luck, actions do serve to demonstrate a moral character/intention, and in such cases, the opportunity/circumstance allows for a demonstration of character. We might not always require the same level of demonstration, but we do need a legitimate performance of character (whether that is an action or a sufficient extrapolation of intent) to make a fair moral judgement. Richards (*ibid.*) puts it as follows:

Now, I have already suggested that it is not the harm a person does that determines his desert, but the features of character which he enacts. Thus, it is someone's indifference to my suffering that calls for attempts to change him, or to deny him future opportunity to harm me, or to show him that such behaviour will be costly in the future. The harm he actually does me in his indifference is only a sign of how indifferent he is and how dangerous his indifference is: the same reactive measure may be appropriate even though he has done me very little harm.

If we take this line of argument seriously and apply it to the previously mentioned examples, we can see why we would sometimes treat the actual and counterfactual perpetrators differently, while other times we treat them equivalently. In the case of the Nazi collaborator and non-collaborator, the reason we would treat them differently is because there is no sufficient demonstration given that the non-collaborator would truly have been a collaborator had things been different. The same is true for the hero and would-be hero of the drowning victim, as someone saying that they would have done the same does not sufficiently demonstrate that they would have. Of course, we can imagine that an individual could be very convincing in their attempt to make such a counterfactual claim, and if we do believe them, then some praise might be in order. If, for instance, a child was to profess her reverence for the hero and went on to say that she would have done the same, then some form of moral praise might be warranted. We can also see why you would praise your second friend in light of the first friend's failure, as the level of demonstration for friendship might not be so high as for the other moral virtues. Importantly, I do not mean to say that we treat counterfactuals exactly as we do actual cases. It might be said that counterfactuals might allude to certain character traits, while actual actions give a much more certain demonstration. Regardless, the point is that counterfactuals might in some manner inform us of moral character and thus warrant some sort of moral judgement.

The same argument is made by Judith Jarvis Thomson (1993: 206) when she says, while discussing the moral difference between a German who is a Nazi and a Canadian who never received the opportunity, "For my own part, I find this Canadian rather hard to get a grip on. What on earth is supposed to fix it about him that he would have acted as badly as the German did had he been forced to take the same test?". Thomson argues that the paradox of Circumstantial luck does not hold once we take a closer look, as we find that the difference between the cases is a practical one and not a moral one. To demonstrate this, Thomson creates a thought experiment where we have two different judges. She explains that there is "Judge Actual" and "Judge Counterfactual", who are both corrupt (1993: 206 – 207). She argues that if either of the judges were to be offered a bribe, both would accept it, but, so far, no one has offered either of them a bribe. We can then imagine a defendant does offer "Judge Actual" a bribe and that the judge accepts it. Had the defendant been tried before "Judge Counterfactual", this judge would have also accepted the bribe, but by the luck of the draw, it is not so. Thomson goes on to argue that it is the case that the moral *record* of "Judge Actual" is different to that of "Judge Counterfactual", but this moral record does not necessarily reflect the true moral standing of the judges. Thomson then asks whether we would regard "Judge Actual" with a moral indignation that was not *also* due to "Judge Counterfactual" had we known about this counterfactual acceptance of the bribe. We might treat them differently for practical reasons regarding our actual knowledge, but as Thomson (1993: 207) puts it: "Would you have God throw [Judge] Actual into a

deeper circle of hell than [Judge] Counterfactual?”. I believe this question gets to the heart of the moral intuition at play when it comes to Circumstantial luck. Yes, their moral *records* might differ, but all things considered, their moral *standing* does not, and might only seem so to us because we do not know all the facts.

I would add as well that if we take this thought experiment one step further, we can show that Circumstantial luck really does not hold at all when we have the relevant information at our disposal. Suppose you were to over-hear “Judge Counterfactual” say that if someone were to offer him a large enough bribe, he would, in fact, accept it and act accordingly. This fact alone might serve to warrant a judgment of moral indignation. He is morally on par with “Judge Actual”. We can also imagine that this information causes us to decide to set a trap and catch “Judge Counterfactual” in the act of bribe acceptance. Why do we want to first catch him in the act? Because we need some sort of objective demonstration of moral character, but this is again a *practical* consideration and not a moral one. Catching him in the act causes his moral record to reflect his moral standing, i.e. it allows for his moral character to be clearly and objectively demonstrated. His moral record provides proof for his moral standing.

Now it might still be argued that luck is still present in all these cases. Something really does seem to be out of the control of these agents (their circumstances), and we usually do judge them differently. However, Nicholas Rescher makes the argument that the reason we tend to judge the protagonists differently in all these thought experiments is due to the experiments themselves, rather than due to the existence of actual moral luck. He claims that philosophers are under-describing the situations, leaving out important factors which do contribute to moral worth, which affects our moral judgements (1993: 154).

Rescher starts his analysis by again stating the claim of those who accept the paradox of Circumstantial luck; some are given the opportunity to be a hero or villain, while others are not, and this makes a difference to their moral record (ibid.). The standard intuition is that we should let a would-be thief “off the hook” as he is different from the actual thief in some important ways. The claim is then that this is unfair, as they are morally equivalent. But Rescher rejects this claim. As he puts it:

Plausible though this may sound, it gets the matter wrong. The difference between the would-be thief who lacks opportunity and his cousin who gets and seizes it is not one of moral condition (which, by hypothesis, is the same on both sides); their moral *record* may differ, but their moral *standing* does not. (ibid.)

“Moral record” then, as used here, refers to the moral judgment we actually accord them, whereas “moral standing” denotes the moral judgement that is truly deserved. Rescher argues that from the vantage point of someone who can see all and know all, the moral records and standings of these two people would be the same. When Rescher (1993 :155) then says, “The difference at issue is not moral but merely epistemic”, he means to say the type of luck at play here is not moral luck, but Epistemic Luck.¹ We happen to know less about the would-be thief’s moral character than about the actual thief’s.

Why then do we have these intuitions of unfairness when presented with these thought experiments? As stated before, Rescher argues it is because philosophers under-describe the cases and leave out important factors that contribute to moral judgment. To demonstrate this Rescher uses the example of another philosopher who describes a case of two rescuers, one who succeeds and one who fails. In this thought experiment, the philosopher argues that there is a difference in the moral judgement we make for the rescuer who enters the burning building and saves someone, and the rescuer who, while attempting to save the person, drops them from a great height and they unfortunately die. Rescher (1993: 158) then states: “The difference our author speaks of is indeed there. But only because of a lack of specificity in describing the case.” We need to know why the second rescuer dropped the victim and not the first. If we extend the thought experiment, we could add, for example, that this is due to carelessness or incompetence. Perhaps the rescuer suddenly became malicious. Or perhaps, even though the rescuer took due care, a piece of burned timber fell on him and caused his grip to slip. Each of these additions to the thought experiment makes a difference in how the rescuer is judged. In the first two additions, we would have good moral grounds for judging the two rescuers differently, but in the last, we can see that the difference is due to luck, and so the moral judgment should not (and does not) differ.

The above demonstrates Richards’ claim that the cases that tend to show a moral difference due to Circumstantial luck are under-described and should not be taken to demonstrate a rational response to such cases. It is possible to remove the pull of luck on our moral intuitions by taking a step back and rationally evaluating the situation, trying our best to separate morally relevant from morally irrelevant differences. When we do blame people differently, it is due to a lack of demonstration of

¹ David Concepcion (2002: 460) argues that there is still luck involved which is not just epistemic. He argues there is also “ethical luck” because we do *judge* an agent morally differently based on what is and is not demonstrated (which is not in their control). This *judging* is ultimately an ethical act, Concepcion argues. He says this might only be a semantic distinction (between ethical and moral), but if accepted we can conclude that there is no moral luck, only epistemic and ethical luck.

character and hence of epistemic certainty. Once this certainty is obtained, we are able to say we would judge them in the same way. Importantly, these considerations are still based on the moral character of the protagonists and not simply on the results of their actions (but still, the results of their actions may give important insight into their intentions and motivations).

It is also a good idea to evaluate the language that we use when constructing these hypothetical moral tests and placing agents in these hypothetical situations. The way we use our language helps form certain intuitions, but this also means that we can be led astray by language. When we create hypothetical scenarios for individuals and then make claims about them, it is not only important to not under-describe the important factors at play, but also that we make sure we know *who* we are talking about. To explain what I mean, let us use a classic example of Circumstantial luck:

If *you* had been in Germany, that is, if *you* were given the chance, *you* would have chosen to be a Nazi.

I have placed emphasis on “you” in this counterfactual because I am not sure who this refers to. It is most certainly not the character/agent that you are now, as you are not a Nazi, and even more importantly, the person you are has negative feelings towards the Nazi party. Just as was the case for Thomson, when she could not grasp how we know the Canadian would be a Nazi, I do not know what part of “you” it is that we are making a claim about here. *You* are not a Nazi, nor would *you* be one, should I give you the choice now. When the above statement is made in support of Circumstantial luck, the *you* in question here is not the relevant *you* when we make a moral evaluation about *you*. Nothing about the thought experiment captures anything relevant about *your* moral character.

It may be said that you would be a different “you” had circumstances been different, and I am fine with this claim. But then we should not make any moral judgement about you based on this counterfactual “you”, as in theory you could be just about anything if enough had been different. What if I were to construct the most bizarre and causally different history for “you”, one where “you” are not even human anymore. Suppose the matter (particles) that currently form you had had a very different history. Rather than being human, “you” turn out to be a virus and a deadly one at that. Would it be fair to say that *you* are guilty (for the sake of argument we will imagine a virus can be guilty) for the deaths that this hypothetical “you”-virus causes simply because *you* could have turned out this way had things been different? I do understand that this thought experiment seems far-fetched and silly, but it does the same thing the other thought experiments do. It tries to make claims about *you* on the basis of something that is clearly *not you*. In the cases where the hypothetical “you” does sufficiently relate to *you*, I am fine with making such judgement (like in the case of the good and bad friend from before), but this is because the hypothetical “you” *is* in some way the same moral

character. If our thought experiments become so far removed from the character in question, they only serve to lead us astray in our intuitions. We should be cautious not to confuse the hypothetical “you” for the real *you* when making such counterfactual claims. If we are to create a counterfactual, the character in question should relate to you in *all* the morally relevant ways.² As far as I can see, all the problematic thought experiments that seem to support the idea of Circumstantial luck make this mistake. A philosopher who agrees with this analysis and who argues that our counterfactuals have to be relevant in order to draw conclusions from them is Rik Peels (2015; 2017).

Before moving on to the next type of luck, I would like to look at a unique counter-argument used by Anders Schinkel (2009). In his discussion, Schinkel argues that authors he calls “epistemic reductionists” are ignoring cases where circumstances do not only play a demonstrative role but actually make a difference to moral judgements. Schinkel (2009: 267) explains that epistemic reductionists are “critics of the notion of ‘moral luck’ that maintain that all supposed cases of moral luck are illusory; they are in fact cases of what I describe as a special form of epistemic luck [...]”. Richards and I would, therefore, be epistemic reductionists.

Schinkel (2009: 273) argues that although there are cases in which a moral character was demonstrated through circumstances (opportunities), there are also cases where circumstances do not only demonstrate a certain character, but also have an effect on the character the agent has. As Schinkel (*ibid.*) puts it: “Luck sometimes plays a role, not just in what we get to know about someone, but also or especially in *who someone is or becomes*.” To support this, Schinkel explains that an agent, while behaving in a secretive manner in order to avoid being perceived as having something to hide, actually does end up acting as if they have something to hide. The circumstances (other people perceiving you as someone with something to hide) has an effect on the agent's character and “causes” their character to change. Schinkel (2009: 274) summarises his argument by stating: “Situational luck can play a role in two ways: chance circumstances may expose what is already there [...], but they can also partly shape a person and thereby make a real – more than epistemic – difference.” The first way is the “way” that has been discussed throughout this chapter on Circumstantial luck. It is the second way which Schinkel argues is true moral luck.

The above, however, gets things very wrong. As was made clear at the beginning of this thesis, as well as at the beginning of this chapter, Circumstantial luck (what Schinkel calls “Situational luck”)

² Much of the work done here relates to the philosophy of modal logic and counterfactuals. I am ultimately arguing that for a counterfactual to be morally relevant it has to be a “closest possible world”, or a world where the relevant moral factors are the same for both the real and the counterfactual world. For an interesting discussion on counterfactuals see Starr (2019).

denotes all those cases where the moral status of an agent is affected by the opportunities and challenges the agent in question faces (which the agent does not have control over). It assumes an (at least temporarily) stable character which must act within a set of circumstances. The moment we change the focus of causation from character-to-circumstances to circumstances-to-character, we are talking about a different type of luck. The examples Schinkel gives do not, therefore, describe a type of Circumstantial luck, but rather, they demonstrate a type of Constitutive luck (as should be clear using the definitions given at the beginning of this thesis). If the claim is that agents are judged for their characters, and these characters are influenced by things outside the agent's control, the claim, therefore, refers to Constitutive luck and does not undermine the position of "epistemic reductionism" with reference to Circumstantial luck. To make it clear, I am not claiming that Schinkel is wrong when he argues that circumstances have effects on moral character, rather, his examples of such cases do not undermine the arguments against Circumstantial luck, but only describe a different type of luck. The next chapter will focus on Constitutive luck and attempt to solve the problems caused by it. Here, I will conclude that Schinkel has not given good reasons to doubt the "epistemic reductionist" position.

4.3 Conclusion

A summary of all the arguments against Circumstantial luck can thus be made quite economically, as they all boil down to the same thing. Before I make such a summary, I would like to quote Rescher's (1993: 162) summary of the arguments against Circumstantial luck, as I do not believe many have been able to put it so eloquently as he does:

For the moral point of view, the crucial thing is to earn an E for effort. Whether our circumstances are straightened or easy, whether our childhood is protected or brutalized, whether our chances in life are many or few – all this is not a matter of our own choosing but lies 'in the lap of the gods.' But what matters from the angle of morality is what we make of the opportunities at our disposal, such as they are, however meagre. Of these to whom little is given, little can be expected, and of those to whom much is given, much. Our horizons for moral action may be narrow or wide – that depends on the visissitudes of facts. For the moral point of view, however, it is – to reemphasize – effort that counts. Those who confront a steep slope cannot be expected to make heading comparable to that of those who find an easy path before them. Situations and circumstances are realities of luck, but the

rational evaluation of moral blame or credit is of course designed to take this into account.

Rescher's point is thus that circumstances only make an *epistemic* difference to our moral judgements and not a truly moral difference. We make moral judgements based on what we know about people's moral character, but what we know does not necessarily accurately reflect their true moral character. Importantly he also touches upon an important point when he says (ibid.): "Of these to whom little is given, little can be expected, and of those to whom much is given, much.". We do in many cases also account for this epistemic difference when judging people. It is morally relevant how old you are (because this might make a difference to your moral capacity) or whether you had loving or abusive parents (as this might make a difference to your character).³ The moral tests you face are part of a larger circumstantial web and the harshness of our judgments are also contingent on this. In an effort to make morality fair, our moral considerations attempt to account for all the extra factors, including circumstances. But as Rescher argues, what truly matters is how we attempt to act given our particular surroundings.

Just as before, in the chapter on Resultant luck, I am happy with accepting the claim that morality requires action. For Resultant luck this was because one of the functions of morality is to make a difference in the world, we must act and do in order to be considered moral agents. For Circumstantial luck it is because we require circumstances/events/tests to allow us to demonstrate and extrapolate moral character. The important part, though, is to realise that these are all practical considerations and that the morally relevant criteria in both cases are still intentions/moral character. Martha Nussbaum does a very good job of creating the moral intuition that we do judge/evaluate people only for those moral characters that are demonstrated through circumstances, and since circumstances are out of our control, our moral evaluations are affected by luck. However, I hope it was made clear through the use of the arguments given by Norvin Richards, that this is not true. We do indeed sometimes judge people for what they *would* have done, but, crucially, this is only the case when we have sufficient reason to believe the counterfactual does in some way capture the true moral character in question. Moreover, Thomson, with the use of examples, further supports the claim that it is not *moral* luck that affects our judgements. Using the case of the two Judges, we see that they are in fact

³ As framed here, it might be argued that a moral character might be "defective". Defective is a strong word and could cause some to feel like this is an incorrect way of phrasing the statement. However, I would argue, if the mind and its accompanying character are seen as a sort of machine able to make moral "calculations", and such a character is consistently demonstrated to make "bad" decisions, then this machine can be phrased as defective. This idea will be discussed in detail in the chapter on the theoretical implications of my thesis.

not morally different, only practically different. And, as I hope I have shown, we only need more information in order to make an accurate moral evaluation, not subjected to luck. Nicholas Rescher makes this point even stronger with his assessment of the type of thought experiments used in this debate. The times we fall for the paradox of Circumstantial luck are the times we are presented with under-described situations. It is only because we do not take into account all the morally relevant criteria that we are able to create the intuition of luck's intrusion into the moral sphere. Hence, in the case of Circumstantial luck, we are not, in fact, talking about moral luck but rather Epistemic Luck. It is our position as moral judges that makes a difference, not moral luck itself. Finally, I also made the argument that our language also leads us astray in these cases. Due to the vague nature of words and the strange way we like to refer to people in counterfactuals, we create moral intuitions about people who are not sufficiently similar to the morally relevant character in question. All of this serves to show that we have not taken a rational enough approach to Circumstantial luck. We have created problematic thought experiments that have led us astray, causing us to infer the existence of moral luck when in reality it is only Epistemic Luck. Yes, life is unfair. No, morality need not be. Indeed, we allow for this fact by taking a closer look at particular situations and clarifying the issues at play. Circumstantial luck thus does not pose any true threat to morality and does not violate the Control Principle, and the paradox is only there when we misdescribe cases.

Chapter 5: Constitutive Luck

5.1 Introduction

In the two previous chapters, most of the arguments and conceptual analysis focused on how we conceive of actions and the different moral tests we all face. We assumed a moral agent in some or other form and tried to understand how this agent navigates a terrain of unknowns and different situations. In many ways, we simply ignored our conception of the moral agent herself and placed our focus on the world around her. This is, I argue, a much easier task than the one that will be attempted in this chapter. The old response, “Don’t hate the player, hate the game”, captures a common theme in the foregoing chapters on Resultant luck and Circumstantial luck. We have our conception of the “player” – the moral agent – navigating the terrain of the “game” – our conception of moral responsibility – and we then try to understand and set the rules of the game in such a way so that the player is treated with fairness in relation to other players. In this chapter, however, we turn our focus to the “player” herself. Now, even if the game were fair and everyone was allowed to play, we must ask what it means to be a “player” when every move and action one makes have been determined by events and circumstances that one had no control over. We have to ask what it means to be a “player”, an agent that has responsibility for her actions during the game, rather than just another “pawn” in the game.

I do not mean to say that the arguments presented in this chapter will be somehow more difficult to understand. Rather, this chapter will devote a large amount of attention to the conceptual analysis of the *concept* of the moral agent. The difficulty herein lies in the fact that we all have an extremely robust and intuitive sense of ourselves, our ability to make decisions, our emotions and experiences, and our “moral nature”. Because all of these intuitions are so closely related to our identity, we are often bound to this conception of ourselves, and we find it very hard to put these intuitions aside when thinking about the problems that arise due to Constitutive luck. It is a case of the agent inspecting the agent; the brain trying to understand the brain. In this chapter, I will try to argue that our intuitive and classical way of understanding the moral agent is flawed, and, since I am talking about each one of us as moral agent, I will ultimately be challenging the very way we see ourselves. Hopefully, as the chapter progresses, I will succeed in convincing you that the way you see yourself and how we talk about ourselves and our mental states, in general, are at least partially wrong and require revision. If I am successful, I hope to be able to demonstrate that the problems arising from our general

conception of moral agency disappears once we change the way we think about ourselves. I also hope to show that moral responsibility is still a valuable concept.

5.2 What We Cannot Control

Before we can discuss the revision of our concept of the moral agent, we must first understand the traditional conception of the moral agent. Luckily, many of the popular writings on Constitutive luck discuss the issue using the traditional conception of the agent, arguing that the agent does not have the necessary control required for moral responsibility. For this reason, the current chapter will start with the works of Martha Nussbaum (1993) and Robert Merrihew Adams (1985), discussing the issue of Constitutive luck and how it relates to our conception of the moral agent. After this, I will focus on a revision of our concept of the agent and the consequences that this has for our conception of the agent and responsibility.

As discussed before, we are no longer arguing that the agent lacks the necessary control over his environment or circumstances in order to be held morally responsible. Rather, in this instance, the lack of necessary/sufficient control is now found in the agent's own mental states. We are shifting our focus one level deeper. The question can be stated as follows: how can we hold the agent responsible for her willings, wants, desires, emotion, or intentions when those things are not in the agent's control? No one chooses their emotional states or their desires. We want what we want and we feel the way we feel. It is important to note that we are not discussing the agent being held responsible for actions and consequences when they do not have control over them (as was the case in the previous chapters), rather, we are discussing how it is that we can hold someone responsible for a desire or emotion (mental states) when these things arguably are not in the agent's control, again seemingly contradicting the Control Principle.

5.2.1 Creating the Intuition

Perhaps the best example and discussion of the agent in relation to Constitutive luck is that of Nussbaum's *Luck and Ethics* (1993). Nussbaum spends a relatively small amount of time on the issue of Constitutive luck. This is understandable, since, by the time she reaches this part within her own discussion, she has already made it clear that she and Aristotle are convinced that luck does play a role in moral responsibility. Because of this, Nussbaum only discusses how luck infiltrates the character of agents, in doing so, once again diminishing eudaimonia. It is this discussion that is so useful in understanding Constitutive luck.

Nussbaum (1993: 97) supports her discussion on Constitutive luck with a lengthy quote from Aristotle, discussing the character of the elderly and how it differs from the young. I have left out some parts and included only those necessary for our discussion (ibid.):

Because they [the elderly] have lived many years and have been deceived many times and made many mistakes, and because their experience is that most things go badly, they do not insist upon anything with confidence, but always less forcefully than is appropriate. [...] Furthermore, they are excessively suspicious because of their lack of trust, and lacking in trust because of their experience. [...] And they are small of soul because they have been humbled by life: for they desire nothing great or excellent, but only what is commensurate with life. And they are ungenerous. For property is one of the necessary things; and in, and through, their experience they know how hard it is to get it and how easy to lose it.

This quote from Aristotle demonstrates the sort of character that a long life can give you. The biggest reason adults and children are so different is because they have learnt to be different. The passage briefly mentions what can be deemed a “good character” and then moves on to explain how this character is hard to keep when life happens. For Nussbaum (and Aristotle) then, no character escapes the transformations that life imparts on it. We live in a world where every experience, every loss, every success, and every event has an impact on our character. Different people are raised with different values and different perspectives, and these have an effect on the choices people make and the sort of people they turn out to be. Ultimately, what Nussbaum and Aristotle are arguing is that we have about as much control over who we are and the intentions that we have, than we do over the results of our actions and the circumstances we face.

It is easy to understand how Constitutive luck is different from Resultant luck, but it is important to understand the difference between Constitutive luck and Circumstantial luck. The two ideas sound very similar and they do both deal with something we lack control over. It is also true that many of the examples and discussions surrounding Constitutive luck include details about the environment and influences from things external to the agent. The best way to understand the difference is to focus on what exactly the agent does not have control over. With Circumstantial luck, it is the tests we face and the situations we find ourselves in. Circumstantial luck is all about opportunity: not everyone faces the same challenges or have the same opportunities to be good. With Constitutive luck, it is the personality and character we have that we lack control over. Circumstances are only relevant here in so far as they influence the character that you have, but circumstances also encompass much more than with Circumstantial luck. Circumstances then also includes the genes you inherit from your

parents, the culture you were raised in, the social class you belong to, and every other factor that played a role in developing the character you have. The problem of Constitutive luck can then be understood as the lack of control (the lack of influence) an agent has over all the factors that contribute to her character and the subsequent moral blame that can be accorded said character, even when the agent had no control over it. As Nussbaum (1993: 96) puts it: “For misfortunes can ‘pollute’ good activity in two ways: by disrupting the expression of good dispositions in action, or by affecting the internal springs of action themselves.”. The first way misfortunes “pollute” good activity here refers to Resultant luck and Circumstantial luck. The second way refers to Constitutive luck; the “internal springs” are those things which drive us to action: our desires, preferences, and intentions – what can be generally referred to as our character. Misfortunes then “pollutes” our character by shaping and influencing it in ways we cannot control or mitigate. We are shaped by our makeup and our past, and every subsequent event we are judged for was inspired and derived from these factors. How can we be held responsible for our actions and character when none of them were in our control from the start?

5.2.2 Mental States and Voluntary Control

Hopefully, I have made clear what is meant by Constitutive luck, at least in an intuitive sense. Even when we have accounted for all factors that the agent does not have control over when we are making moral evaluations, we still need to account for the moral character of the agent herself. Now that we have formed the intuition of what Constitutive luck entails, it is time to take a closer look and discuss exactly what we are not in control of and why this is an issue when it comes to moral judgement. The work of Robert Merrihew Adams, in particular his article titled *Involuntary Sins* (1985), will be used to do this. In it, Adams makes clear exactly which parts of the agent are not under her control, and he discusses some initial responses which help illustrate the problem.

Adams (1985: 3) starts his article by giving an example of a mental state for which someone is held morally responsible and then reminds us of another moral principle which seems to contradict this responsibility. He argues that anger without just cause is a sin. It is important to note that, although Adams does write from a religious perspective, he does not require you to hold such views. “Sin”, for Adams, can be understood as any action, omission, or state that is inherently blameworthy (ibid.). With this in mind, he argues that anger which is not necessary or “just” is morally blameworthy and deserves moral judgement. Adams also argues that this moral intuition is at odds with another intuition we have, namely, that we can only hold people accountable for *voluntary* actions and

omissions. “Voluntary” here refers to control, and it will be helpful to first understand how these concepts are used and relate to one another before moving on.

Adams (1985: 8), when discussing what he means by voluntary control, writes: “To say something is (directly) within my voluntary control is to say that I would do it (right away) if and only if I (fully) tried or chose or meant to do so, and hence that if I did it I would do it *because* I tried or chose or meant to do it, and in that sense voluntarily.”. Something is thus within my voluntary control if I wanted to do it and thus “chose” to do it. It is a claim about the agent having to choose to do something, meaning the disposition or action which came about originated from or was caused into being by a deliberate decision from an agent, as opposed to being caused by factors not related to “choosing” (such as other physical causes). If we thus understand voluntary control as being opposed to a matter of luck or being caused by things not within the agent's control we can see that voluntary control captures the same idea as the Control Principle. Understood this way, and focusing on the issue of Constitutive luck, we can say the following about our intuitive conception of responsibility: a person can be held responsible for a mental state or action if, and only if, said state or action came about through the agent’s voluntary control. Take note: actions are also included here since actions “flow” from the mental states of the agent, but for the sake of argument, the rest of this section on Adams will focus on the mental states alone.

Adams (1985: 3) states the following about anger: “Exaggerated or senseless anger, an anger that is not justified by a proportionate provocation, is morally offensive; and one who is guilty of it is liable to blame.”. The justification for such a statement is not required here. Provisionally, this can be accepted as a moral wrong without argument. Most would agree that an appropriately disproportionate emotional outburst of some kind or another, without appropriate justification, is a moral wrong and people should be judged for it (whatever that judgment is). Should someone accidentally bump into you, and you have an aggressive reaction, most would agree that you did something morally wrong (they might say you should have first investigated why they bumped into you). What we have, then, is a mental state which most would agree is morally blameworthy, but which Adams argues is not voluntary (*ibid.*). As he puts it (*ibid.*): “For anger is not in general voluntary. It is apt to be manifested in voluntary behaviour, and we may make voluntary effort to control it; but we cannot choose to be angry or not as we normally can choose to sit or stand.”. Once we understand voluntary behaviour, and how it relates to the Control Principle and responsibility, we see that anger as described above is morally wrong, but we can also not hold the agent morally responsible for this wrong. We are, barring further analysis, once again forced to conclude that the paradox of Constitutive luck is a problem. Moral responsibility is once again lost. Having concluded

the above, Adams continues to discuss three different responses to the problem as stated and argues that all of them fail. These three responses will be discussed next.

The first response, or as Adams (1985: 4) puts it, “alternative”, to what has been said above, is to claim that these involuntary sins (the mental states) are not the morally relevant factors for responsibility; rather, the actions which came about because of these motives (the mental state) are the relevant criteria. That is to say, we are not responsible for the mental state of unjustified anger but rather for the actions that came about because of them (which are arguably voluntary). We blame the agent for screaming or throwing something rather than the anger “encapsulated” in said action. This response aims to convince us that we are not being held responsible for something not in our control (the anger) but rather for something that is (our actions). Adams responds to this by stating: “On this view we are not accountable for a motive unless we have voluntarily consented to it by acting on it. But the ethics of the heart cannot be reduced to an ethics of conduct in this way.” (ibid.). He argues that it would be wrong to reduce all responsibility for mental states to their resultant actions since we can conceive of cases where there is no action from the agent and yet we would still say they are morally responsible.

The above claim is supported using two examples. In the first, Adams argues that even if I do not act on the anger I have for you (by not hitting or insulting you for example) I am still responsible for the anger I have in the first place. It might be true I could receive some praise for my self-restraint, but the anger remains a moral fault. The unjust anger I feel causes harm by perhaps harming our relationship or by making it impossible for me to think fairly about you when necessary (of course, it is not the thinking, which is an action, that is the problem, but rather the *way* I think about you). The agent is still responsible (morally) for these consequences even if they are not consequences of any voluntary action (ibid.). The second example used by Adams (1985: 4 – 5) is an agent guilty of self-righteousness. In cases where one agent has sinned (Adams’ terminology) against another through voluntary action, and the other does not react in any way because of self-righteousness, it could be argued that both agents are still morally responsible, even if the one acted while the other did not. Adams (ibid.) explains the sin of the second agent as follows: “He wants to be above the other person, to be in the right against the other person. The passion with which he clings to this superiority is what energizes all his wit and will-power to do his duty as he sees it – and to see his duty as he has done it.”. The point is that the agent does not act in any way here. The agent is not guilty of any voluntary

action, he is only guilty of believing¹ he is better than the other. Adams argues that this is a mental state that the agent is morally responsible for even if there is no action. As he puts it (Adams, 1985: 5): “What we think is rather that he ought not to have had the bad motive in the first place.” This “bad motive” is again something that simply seems to happen, rather than anything the agent voluntarily wills into being. For Adams, what these two examples demonstrate is that we cannot appeal to the action as being the morally relevant criterion in these instances, since in both cases the only fault is a disposition. We are still forced to conclude that we are holding agents accountable for involuntary states. The thesis that it is not the mental state but the action inspired by it which is the morally relevant criteria is shown to be false using these two examples.

I believe there is another argument which demonstrates the failure of the above claim/response to the problem of Constitutive luck. Recall the definition of voluntary control Adams (1985: 8) gave: “To say something is (directly) within my voluntary control is to say that I would do it (right away) if and only if I (fully) tried or chose or meant to do so, and hence that if I did it I would do it *because* I tried or chose or meant to do it, and in that sense voluntarily.” It seems strange to say “I (fully) tried or chose or meant” to do an action when the action itself was *also* inspired by a mental state, something I do not have voluntary control over. As an example, imagine two individuals in the same situation, each being confronted about an issue, but one finds himself becoming more and more angry while the other has a calm demeanour. We can now imagine, should each act according to their wills, the first is much more likely to punch his interlocutor in anger, while the second is much more likely to engage in some form of civil, if robust, discourse. Using the above definition of voluntary control, we should conclude that neither of these actions were voluntary, since at least part of why the agents acted as they did was *because* (the same emphasis as Adams placed here) they were angry or calm, both of which are mental states themselves and not under voluntary control. It would be inconsistent, therefore, to claim that although we cannot hold an agent responsible for an involuntary mental state, we can hold them accountable for an action which is the result of the involuntary mental state.

The second reaction to our apparent lack of control over our morally relevant responses that Adams discusses is the view that mental states are in fact voluntary because they are “operations of the will” (1985: 6 – 7). Since the “will” is intuitively understood as that part of us which wants and decides, some argue by including mental states as operation of the will we are able to say they are voluntary since the will is understood as voluntary. Adams responds to this by arguing there is a fallacy of

¹ Belief is not an action as an agent cannot be said to *try* to believe a thing, only that they do or do not. Try to believe that you are 100 meters tall or that the sun is blue, and you will realise that you cannot will a belief. For an interesting discussion about belief see Schwitzgebel (2019).

equivocation (my phrasing) when speaking of the will as “voluntary” in relation to how “voluntary” has been used within discussions of moral responsibility. As mentioned before, Adams (1985: 8) argues that in a voluntary action: “[...] the action itself is an object aimed at by the agent in a way that is aptly described as ‘trying, choosing, or meaning’.”. Using this conception of “voluntary”, we can see that mental states do not qualify, even if we are able to, in some sense, say they are “voluntary” because they are part of the will which itself is voluntary. Adams (1985: 9) gives two reasons why desires and emotions (mental states) are not within our voluntary control as defined above: a) “they are not cases of trying, choosing, or meaning, as is illustrated by the fact that if I simply desire to do something, it remains a question whether I will try or choose or mean to do it”, and b) “it is rarely true that one would have a particular desire or emotion right away if and only if one (fully) tried or chose or meant to have it.”. For Adams then, mental states cannot be voluntary since we cannot “will” them into being by actively trying. They are still things that happen to us or things that happen by us trying other things (such as remembering something which causes anger or controlling our breath to calm down). Although it might be said that mental states are operations of the will, we do not have the necessary control over these operations to claim they are voluntary. Adams (1985: 11) concludes his argument by again making the claim that although mental states are not voluntary, we do still need to hold the agent responsible for them:

Our desires and emotions, though not voluntary, are responses of ours, and effect the moral significance of our lives, not only by influencing our voluntary actions, but also just by being what they are, and by manifesting themselves involuntarily. Who we are morally depends on a complex and incompletely integrated fabric that includes desires and feelings as well as deliberations and choices.

The argument that mental states are voluntary fails once we understand the concept as it relates to control and responsibility. We are still left with the problem of Constitutive luck if we agree with Adams that we should be held responsible for our mental states and that these states are not voluntary.

The third and final response considered by Adams is the claim that we can only be held morally responsible for those mental states that are indirectly voluntary by being results of voluntary actions. On this view, you are responsible for a state of mind if, and only if, some voluntary action let to or caused the state, thereby making it an indirectly voluntary action. If you have thus practised meditation which allows the fostering of a calm persona, and then find yourself in a situation of great stress and yet retain the virtue of rational deliberation, you would be morally responsible for this mental state and hence be praised. The same is true for negative mental states. If an agent had deliberately omitted actions which help with controlling his temper, and then lashed out at an innocent

person, he would be morally responsible for this anger, even if he did not at that instance voluntarily become angry. Adams argues that although this argument does capture something true – namely that we ought to try and be “better” through voluntary actions – it still only captures part of the moral equation (1985: 12). To make his point, Adams (1985: 12 – 13) asks the following question:

“Suppose you have just realized that you are ungrateful to someone who has done a lot for you – perhaps at great cost to herself. Far from responding to her sacrifices with love and gratitude, you have made light of them in your own mind; and if the truth be told, you actually resent them, because you hate to be dependent on others or indebted to them. Surely this attitude is blameworthy. Must we assume that you have caused it, or let it arise, in yourself by actions that you have voluntarily performed or omitted?”.

Why is it required that such a state be caused by voluntary actions when the state itself, and the consequences that follow from it, are enough to warrant moral blame? We could not respond by saying that you should have tried beforehand to foster a character that is grateful since to try and be better implies knowledge of the imperfection, and we would not say you are only blameworthy if this is the second time you were ungrateful. Again, even if we agree that to try is an important part of moral responsibility, we would also agree that it is not the only part. It is just as important to have the right mental states, not just trying to have them. Adams (1985: 14) concludes this section by stating: “These considerations confirm the intuitively plausible judgement that what we chiefly blame in the present immoral state of mind is not the imprudence of the previous voluntary omissions.” There might be cases where we blame an agent for voluntary actions, but examples such as those used throughout this discussion have demonstrated there are instances where we would blame an agent for a state of mind even when no voluntary action was present.²

The three different responses to Constitutive luck that have been considered and discussed in this section have all been demonstrated to fail. In each case, it was demonstrated that there are always mental states, which are involuntary, that are the focus of moral evaluation. With regards to the first response: we do not only blame the agent for the actions that result from a state; even if there is no action, the state can still be wrong. The second fails because to speak of mental states as part of the will also does not solve the problem, since this conception of “voluntary” does not fit the definition required for moral responsibility. Finally, to consider only those states which follow from voluntary

² For an example of a philosopher who agrees with this analysis (namely that we cannot make sense of an agent controlling their character) see Gregg Caruso (2019).

actions ignores a large and important part of moral responsibility, that is, we should not only *try* to be but also *be* better. All of this serves to support the problem of Constitutive luck. Mental states have been, and should continue to be, objects of moral evaluation. Yet, we cannot escape the strong pull of the Control Principle, in this discussion framed as “voluntary control”. It would seem then that we have to conclude one of two things: 1) that we cannot hold people morally responsible for their mental states, or 2) that we can hold people accountable for involuntary actions. The first seems problematic because so much of what we do and the reasons for how we act are the consequences of our mental states. To say we cannot be held morally accountable for them is equivalent to saying we cannot be held morally accountable for anything, whether they be our mental states or the actions that flow from them. The second seems problematic because it opens the door to all kinds of problems as discussed throughout this thesis. We want to only hold people accountable for things they can control, for things that are voluntary. Anything more and we seem to lose the essence and practical reason for moral responsibility. Next, we will consider a third option, but it will be very different to those proposed here. For this option, we require a revision of our concept of the “moral agent” and what it means to “have a mental state” and what it means for an agent to “lack control of their character”.

5.3 Moral Agents and Identity

Up until now, this chapter has approached the issue of Constitutive luck in much the same way as Resultant and Circumstantial luck. Looking at the agent, we ask if there are things we tend to hold her responsible for even when she does not have sufficient control over these things. The question/problem in this chapter can be stated as follows: do agents have enough control over their characters such that we may hold them morally responsible for their character faults/strengths? At face value, this might seem like a good question to ask. But what if we are being led astray here by our language use, and also by our, all too often, mistaken intuitions? Is it not a strange question to ask whether an agent has control over “her” character when, arguably, said character is what we call the “agent” in the first place? I would think that, were I to ask you whether a computer has control of its code, you would think it a strange question, since it is the code that makes it a computer and does all the controlling in the first place (as well as its internal configuration, but these are simply different parts of the same thing, just like character and the brain). The question itself is problematic. This is the issue addressed in this section on Constitutive luck. In it, I hope to demonstrate that the problem of Constitutive luck is not a problem of a lack of control on the part of the agent, but a fundamental misunderstanding of what an agent is and what it is that does the “controlling” in the first place.

Nicholas Rescher (1993: 155) states the above objection as follows: “After all, it makes no sense to say things like, ‘Wasn’t it just a matter of luck for X to have been born an honest (trustworthy, etc.) person, and for Y to have been born mendacious (avaricious, etc.)?’ For it is just exactly those dispositions, character traits, and inclinations that constitute these individuals as the people they are.”. For Rescher then, it makes no sense to make a claim about X’s constitution – X is trustworthy/untrustworthy or kind/cruel – and then question why X should be held accountable for being X. If I am talking about an agent, what exactly am I discussing? It cannot be something which then “has” (as in “owns”) certain character traits, since I cannot conceive of what this something might be other than those character traits (what *is* an agent without character traits?). The problem of Constitutive luck asks something strange of us: it asks that we not only accept that there is something like “you”, but also that there is something like a “proto-you” which might cause or influence this “you”. The idea seems suspiciously like Descartes’³ ghost in the machine⁴, but now the ghost itself requires an internal ghost as well. The question either makes no sense or it leads us to an infinite regress, postulating an ever-deeper agent responsible for all those things that make up an agent one level up. Discussing the concept of a person, Rescher (1993: 155 – 156) claims that to hold “people responsible for their moral character (rather than seeing this as something added fortuitously *ab extra*) is an inherent part of the fundamental moral presumption involving treating a person as a person.”. He argues that in order to treat people as people, we must accept that character traits are part of their constitution, otherwise, we are treating them as something other than a person. Fundamentally then, people are their characters rather than characters being something that people have in addition to themselves. Imagine we have a perfectly intact human body, but this body is not alive. Would we call this a person? I would think not. Now what if this body was alive, but rather than being “controlled” by any mind, it is controlled using electrical signals from a computer. Would we call this a person? Again, I think we would not. Even though all the physical requirements are there to call it a “human” body, nothing about it can be said to be a person. It is only once this body has an accompanying personality, a mind, a phenomenological experience, that it can be said to be a person (or moral agent, if it is capable of moral reasoning).

³ René Descartes (2019) popularised the idea that the mind and body are two separate entities. The ghost in the machine referred to here is another way of stating that the mind is the thing doing the controlling, and the body is the thing being controlled. One of the major criticisms of Descartes is that his idea of the mind controlling the body leads to an infinite regress, where the “actions” of each entity in question must be explained by another before it. The same criticism applies here. For a thorough discussion on Descartes’ work see Hatfield (2019).

⁴ The phrase “ghost in the machine” to describe Descartes’ conception of the mind was introduced by Gilbert Ryle in 1949 (Ryle & Tanney, 2009).

Shifting focus from the “person” to the issue of control again, we see that proponents of Constitutive luck gets things all wrong. Rescher (1993: 156) frames the question that a proponent of Constitutive luck might ask about control over “our” personalities as follows: “But what then of one’s inclinations, dispositions, and character? Are not these too issues outside our control?” Again, the question misunderstands the nature of our personalities and the concepts of control as it relates to moral responsibility. Our character and personality, those things which make up who we are, are exactly those things which the concept of control does not apply to. We are not able to give an adequate account of what does the controlling if not the character. We are also not able to give an adequate account of what an agent might be other than something with a character.

Perhaps a reason why we make these mistakes is because of language and how we talk about one another. Sentences such as: “He is an angry person” and “She has a jealous disposition”, make a point of distinguishing between the agent and their personality. We often first refer to someone as “he” or “she” before attributing a character trait to them, often describing such traits as things they *have* rather than things they *do*. Language such as this has an effect on our intuitions, but much more importantly, our discussions and analysis all happen within the bounds of the language available to us and influence the ways we can talk about things. One way of alleviating the effect of our language is to form a more accurate intuitional understanding of what exactly we *are*. One way of doing this is by discussing *what* we are in light of the history that led to us becoming “agents”.

In his most recent book *From Bacteria to Bach and Back* (2017), Daniel Dennett tries to tell the story of how things like us, things with “minds”, could have come about through natural processes. For the sake of argument, we will treat his argument as correct; however, no such commitment needs to be held for my resulting claims regarding agency to hold. I am using his work in this section because it helps us to understand what exactly an agent is and what the relationship between an “agent” and “her” character is, at least conceptually, even if he should be incorrect about *how* agency came about.

Dennett (2017: 254) postulates the following about us:

Perhaps we are just apes with brains being manipulated by memes in much the way we are manipulated by the cold virus. Instead of looking only at the prerequisite competences our ancestors needed to have in order for language to get under way, perhaps we should also consider unusual vulnerabilities that might make our ancestors the ideal hosts for infectious but nonvirulent habits (memes) that allow us to live and stay mobile long enough for them to replicate through our populations.

The word “meme” might be an unfamiliar one. For our purposes we will use the term as used by Dennett (2017: 206): “There is no term readily available in the technical language of the scientific image that aptly encapsulates what kind of a thing a meme is. Leaning on the ordinary language of the manifest image, we might say that *memes are ways*: ways of doing something, or making something [...]”. Memes are very similar to genes and can be understood as “genes of the mind” (hence the name). They are culturally acquired items – ideas – which are expressed in certain ways through agents, and they, similarly to genes, are transferable from one person to another. They are also able to “mutate” (change) and thus relate to cultural evolution. Exactly as we can imagine genes mutating and being passed on to other people, thus causing genetic evolution, memes can mutate and spread causing cultural evolution (although they do this through different mechanisms). Memes encapsulate a large number of things, but for our discussion, they will include all those ways of doing things which we might call the “mind”. This is not to say that agents are memes, rather, agents are constituted by memes in the same way that our bodies are constituted by genes. Minds can be understood as a collection of memes, and an agent as a way of referring to this mind as *doing/acting*. Using this conception of meme, we can argue that the agent is all those things which drive or influence a biological machine (that is, our brains and body) to do certain things (such as actions). Ways of feeling, thinking, and intending are all different memes, non-genetic replicators, and they are another part of what we call a “human”. The agent, then, cannot be understood as something that has character, but rather the character is another part of the agent. Again, there is no way of conceiving of an agent if we do not presuppose that the concept already encompasses a character. Dennett (2017) thus argues that agents like us evolved from mindless brains, which eventually developed a means of copying/imitating other brains. Memes then hijacked this copying mechanism and have used our brains as hosts ever since (metaphorically). Our minds are then a collection of these memes, organised into a relatively coherent experience we call “the mind”. Understood this way, the locus of control is the mind (the collection of memes) itself, rather than something else somehow influencing the mind. An agent who finds herself in a situation where she intends to act will base her decision on the dispositions she has, her emotional state at the time, and the intentions she has. The “cause” of the action is, therefore, the collection of a mind (memes) and a body (genes and physiological processes) we call the agent (of course, these two are not entirely distinct entities). The agent thus causes the event which follows from her actions, and what she intends to do and how she intends to do it, is “controlled” by her. The moral responsibility for these actions are still ascribed to the mind since this is the part which is influenced by moral deliberation and judgement.

The above is meant to help develop the plausibility of the claim that there is no “you” without all the accompanying character traits that form part of who you are. It is also meant to help develop an

understanding of where the locus of control lies, namely with the mind itself (which includes one's character traits) and not someplace "prior" to such traits. If we now ask the question: "Does the agent have control over her character?", we can understand that this question makes a category mistake. By category mistake, I mean to say that the question implies or assumes two separate categories, namely the "person" and her "personality", when in fact they are one and the same thing: the agent. As Rescher (1993: 157) puts it while discussing character traits: "[...] these factors are not things that lie outside oneself but, on the contrary, are a crucial part of what constitutes one's self as such.", and therefore, to ask if "we" have control of "them" is a mistake. Rescher (ibid.) supports this by claiming it is incoherent to conceive of an agent without these factors. I have decided to use Dennett in an attempt to demonstrate that we can coherently understand an agent as *being* these factors. Rescher (1993: 155) strengthens his argument by saying:

"It makes no sense to envision a prior featureless precursor who then has the good (or bad) luck to be fitted out with one particular group of character traits rather than another. In person theory, as in substance theory, there is no appropriate place for bare particulars that, having a priori a nondescript (propertyless) prospect identity, can then be filled out with properties a posteriori."

He argues that a "person" without any personality does not make sense. Such a thing is conceptually problematic.

In order for us to then solve the problem posed by Constitutive luck, we must rid our discussions of these problematic conceptions of agents. If we have a proper understanding of "moral agents" we can see that the problem of Constitutive luck does not pose a real problem. "Identity must precede luck" (ibid.) can be the slogan of those who agree with this solution to the problem. There must first be an agent who lacks control before we can say luck is present, and such an agent must already contain all those character traits which are supposedly out of "her" control. If we return to the language used in the section on Adams, we can understand voluntary control as including the faculties of character in discussion here. Character traits are thus understood as part of "trying, choosing, or meaning" (Adams, 1985: 8), they *are* different ways of meaning or trying, and under Adams' definition, they are a part of voluntary action by being prior to them rather than being something that needs to be brought about through voluntary action. The very thing that is meant to be doing the controlling, the locus of control, are all those parts that make up the agent (all the memes).

To further understand this conception of the agent, I would like to focus on a metaphysics of agency proposed by philosophers who write on free will and luck. Specifically, I will focus on Randolph Clarke's (2005) discussion on causation and agency. Clarke's discussion is framed in such a way that

it discusses and addresses the type of luck that has been called “Causal luck” in this thesis. As mentioned, Causal luck is not addressed in this thesis as I have argued that it is not true “moral” luck; however, I believe, once we have clarified some other concepts, his discussion is still useful in discussing Constitutive luck.

Clarke starts his discussion with a distinction between what he calls event-causal libertarianism and agent-causal libertarianism (2005: 408 – 409). Libertarianism is the view that endorses incompatibilism, and because of this, it claims that whatever an agent does (causes) must be undetermined by any prior occurrences or events for them to possess free will. Event-causal libertarianism holds that this “causing” consists of certain events, such as the agent having certain reasons for acting or having certain intentions to act. In contrast to this, agent-causal libertarianism holds that causation does not consist in causation by events, but by an enduring substance called the “agent” which does the causing. On this account, an agent can be understood as the ultimate source of his free actions (ibid). Understood this way, we can see that Rescher and myself are arguing for an agent-causal metaphysical account of moral agency. There is one difference between what Clarke argues for and what I argue for. Clarke’s conception of this agent-causal account is libertarian with regards to free will, while I do not hold such a commitment to the necessity of indeterminism. The reason for this lies in the difficulty agent-causal libertarians have in giving a rational account of how free actions are made (how exactly the agent causes them).

The above problem is demonstrated when we consider the following example given by Clarke (2005: 411):

Suppose that on some occasion a certain individual, Leo, is deliberating about whether to tell the truth or to lie. He has reasons favoring each alternative and an intention to make up his mind now. Suppose that there is a nonzero probability that Leo’s having the reasons favoring telling the truth (together with his having the indicated intention) will nondeterministically cause his deciding to tell the truth; and suppose that there is, as well, a nonzero probability that his having the reasons favoring lying (together with his having the intention) will instead nondeterministically cause his deciding to lie. Then, given all prior conditions, it is open to Leo to make the former decision and open to him to make the latter instead. Finally, suppose that, in fact, Leo decides to tell the truth. His decision is caused by him, and it is nondeterministically caused by his having reasons favoring the action decided upon (and his having the intention to make up his mind).

There are two different ways philosophers attempt to demonstrate the problems with the case above: one focuses more on determinism and indeterminism within the same timeline, while the second is a modal argument analysing possible worlds. I will focus here on the latter. Suppose there are two equally possible worlds that contain the same histories and laws of nature, let us call them W1 and W2, where W1 is the world in which Leo decides to tell the truth and W2 the world where he decides to lie. Assuming indeterminism is true, we would conclude that the two worlds only diverge in their histories once Leo makes his decision. Before the decision, everything about the two worlds is the same. Following this, there is no difference in the two worlds to account for the difference in the decision that Leo makes, therefore, which decision Leo eventually makes is a matter of luck. Clarke (2005: 412 – 413) argues that if the difference is a matter of luck, then the agent (Leo) does not freely make his decision. Note that it is still true that Leo causes the decision to tell the truth or to lie as him telling the truth is caused by him having reasons to tell the truth. The decision itself is not due to chance, it is due to Leo, but which decision he makes is a matter of luck. Clarke (2005: 415) argues that the best way to frame the problem is by stating that “there is no explanation of *why* Leo decided to tell the truth rather than deciding to lie and no explanation of *why* Leo caused a decision to tell the truth rather than causing a decision to lie.”. Nothing we have said has explained, even though we claim that Leo caused the action, how or why he decided to cause the one and not the other. It would be a mistake to say the difference is due to a difference in his intentions or in the outcome he perceives is the best, because we already stated that the causal histories of W1 and W2 are the same up till the actual decision. Clarke (2005: 416) argues that the difference can be understood as a difference in will:

“What would be true is that the difference between the actual world, where Leo decides to tell the truth, and world *W* [what I refer to as W2], where he decides to lie, is a matter of Leo’s exercising his free will one way in the actual world [W1] and another way in *W* [W2], and likewise for the other difference.”

The problem I have with Clarke’s answer is that he does not justify it by explaining the mechanism of how this “will” is exercised, rather, he gives an account of a discussion between him and Leo in which Leo explains that the difference is in how he exercised his free will. This does not explain the difference, it simply states, once again, that there is one. The following is a part of this discussion by Clarke (2005: 418):

Me: But when the worlds diverge, isn’t it just a matter of luck that you actually exercise your control one way rather than the other? I mean, in *W*, you determine that you make the other decision, you make that decision happen, you originate that

decision. Isn't the difference between your originating the actual decision and your originating the alternative one just a matter of luck?

Leo: Not if my determining, making happen, and originating the actual decision is my exercising my free will. Then the difference is a matter of my exercising my free will differently. And since my decision was brought about the way you agree it was, I'm convinced that my determining it, making it happen, and originating it was my exercising free will in acting.

Clarke thus argues that the difference in the decision is explained by the difference in the execution of Leo's free will. But what explains the difference in the execution of Leo's free will? It seems Clarke has only pushed the problem back one step, rather than solving it. As I understand it, the indeterminism implied by Clarke still possess a problem for the agent's *will* to both still be able to *do* something and be *free*. If nothing determines Leo to act in one way rather than another (such as his character and intentions), his decision is a matter of luck. I cannot conceive of how something can cause an event in a completely undetermined way. If something causes an event, this something has to in some way determine what happens. As mentioned, I do not accept the reality of indeterminism. It might be argued that no absolute determinism exists, but I do, at the very least, hold that probabilistic determinism is true. That is, given some event, some effects are more likely to occur than others because the event determines them in some appropriate manner. Perhaps your childhood, circumstances, and emotional states do not absolutely determine what you will choose in any given situation, but they do make some decisions more likely than others, and in this regard, they do determine the decision is some way. Following this, the events and circumstances that occur before an agent makes a decision have an influence on which decision the agent makes; they are determining factors. The important part is to make sure we distinguish between event-causation and agent-causation. Something can be understood as being caused by events if no agent played a role in their occurrence (event-causation). Something can be understood as being caused by an agent if said agent played a role in causing the event (agent-causation). Notice that in both cases what happens is still influenced by what happened before, the difference is that the cause for the action in the first case is not a moral agent, whereas in the second the cause is. The *type* of thing doing the causing is what determines whether it can be held responsible or not, not whether it is undetermined or determined.⁵

⁵ Agents could be understood as events, but a unique sort of events. Nevertheless, what can be called an "event" is a passive cause while what can be called an "agent" is an active cause. Agents have deliberative power and they therefore play an active role in the events they cause by "choosing". It is therefore useful to distinguish between "events" and "agents".

Again, it is the fact that the agent is the locus of control that she can be said to be responsible. To explain this I turn to what has been called a Frankfurt case.

Above, the work of Dennett was used to help facilitate an understanding of *what* an agent is. Specifically, his work was briefly discussed so that our concept of the agent, and what it means to “have” a character, is better understood. The argument then shifted to discussing this agent as the “locus of control”, arguing that it does not make sense to speak of an agent as “controlling” their character since the character in question should be understood as the locus of control, that is, the thing doing the controlling. Following this, the work of Clarke was used to explain the metaphysics of agent-causation and how it differs to event-causation. Next, I wish to argue that it is not important to focus on what has caused the agent to act in a certain way (which was the issue discussed above); rather, it is important that the event in question was sufficiently caused by the agent before “control” can be inferred. In order to support this claim, I wish to use something called a Frankfurt case. Frankfurt cases are thought experiments first developed and introduced by Harry Frankfurt in his essay titled *Alternative Possibilities and Moral Responsibility* (1969). Originally, the thought experiment was meant to demonstrate that agents can be held morally responsible for their actions even when said agent does not have alternative possibilities. It is argued by some that if an agent lacks alternative possibilities, that is, the metaphysical possibility for an agent to have multiple alternatives available to them for actions, as opposed to being determined to act in a single way, they cannot be held morally responsible for their actions (this is the view on free will that Clarke holds). Here, however, I wish to use the thought experiment as a way of arguing that the locus of control is the agent's character (although, the argument does still challenge Clarke's agent-causal libertarianism, arguing for what might be called agent-causal compatibilism).

The argument for alternative possibilities is usually phrased as follows:

If an agent is only able to act in one way, that is, if the agent does not have the option to have multiple different ways of acting, they do not really have control over their actions. If I am faced with the choice of choosing between a carrot and a piece of cake, but because I am determined (through whatever means), I will always choose the one over the other, do I have true control over my choices (actions)? Intuitively, for my choice to really be mine, and for me to have control over it, I should be able to choose either the carrot or the cake. I should have alternative possibilities with regard to actions. If for whatever reason, I can in fact only

choose the one, my choice does not seem a free one. Stated otherwise, I do not seem to possess control over my choices.⁶

Intuitively, this argument seems true. If I program a robot to have a choice in its actions (leaving aside the issue of whether this is possible), the robot should in theory not be determined by its programming (or by the laws of physics) to make a *particular* decision, but should rather be able to consider and ultimately choose out of an array of possibilities. If the robot makes a decision, and we conclude that it could not have made any decision but the one that it did, we would conclude that it did not have ultimate control over its actions. It is this conclusion which is challenged using Frankfurt cases. What follows is my own example of a Frankfurt case:

Imagine a mad scientist who is completely obsessed with forcing people to make healthy diet decisions. Focused on this obsession, the scientist develops a mind-control chip which, when implanted, is able to override an agent's choices and force him to make a different, in this case healthy, choice. Now imagine two different individuals, both implanted with this chip. P1, while making lunch, decides to make some deep-fried potato chips. As P1 makes this decision, the mind-control chip is activated and overrides P1's original decision. P1 now decides to make himself a salad. P2, while making lunch, decides to make a salad. Because P2 never decided to make an unhealthy meal, the mind-control chip is never activated.

Assuming for the sake of argument that agents can be praised for good dietary decisions, would we praise both agents? In the case of P1, we could argue that ultimately his decision was not in his control. Most would agree that we should not praise him as his decision to eat healthily was determined by factors outside of his control. The interesting case is that of P2. The fact that P2 could not have decided to eat an unhealthy meal does not seem to take away any of her responsibility for choosing to eat a healthy meal. Even though P2 did not truly possess alternative possibilities (the chip would have kicked in had she decided on an unhealthy meal), she none-the-less *decided* to eat a healthy meal. The difference between these two cases, I would argue, is that in the first the locus of control is the chip in P1's brain, while in the second, the locus of control is P2's character (her wills, wants, needs, and desires). As long as the actions were appropriately determined by the agent's character (all the ways of trying, meaning, and choosing; memes), we can conclude that the agent can be held responsible for said actions, even if they were determined to do so.

⁶ For an example of this argument as it is used in the discussion on free will, see Robert Kane's *A Contemporary Introduction to Free Will* (2005).

The above is meant to demonstrate that responsibility can be ascribed to that which controls whatever actions are being judged. In the case of moral actions, we require something that can have *moral* control (“control” meaning to choose based on relevant considerations, “moral control” meaning to have the ability to choose based on moral consideration), and for people, this is their character. All things considered then, it does not make any sense to ask if an agent has control over her character since her character is part of the agent, and this, along with all the other things that make up the mind, does the controlling. The agent is the locus of control, and because of this, we understand the metaphysics of agency to be of the agent-causation type. Constitutive luck, therefore, does not pose a threat to moral responsibility as responsibility can still be ascribed even when an agent is determined by preceding event as long as it could be said that the agent was the one who decided to do it. Our metaphysical view is that the agent is an enduring, mostly coherent, substance understood as the locus of control for actions/events which they can be held responsible for because the decision to do the act relates in some appropriate way to the agent’s character.

Before concluding this chapter, I would like to discuss a further distinction made by David Concepcion (2002) which may decrease the amount of things the agent can be held responsible for. Concepcion (2002: 458 – 459) agrees with the argument that the concept of control cannot rationally be ascribed to the agent’s volitions (the power of using your will). He does, however, argue that the agent’s character (his dispositional states) are in fact out of the agent’s control. Concepcion believes that only acts of *volition* do not have to be said to be controlled (they are by definition doing the controlling) by something for agents to have responsibility, but the agent’s emotions and dispositional states do still need to be controlled in order for the agent to be morally responsible for them. This is because Concepcion argues that character is something different to volitions. I am not certain if I agree with this. I do not understand how one can conceive of “volitions” without their accompanying character traits which influence (drive) the volition. I believe it is consistent and rational to conceive of the entirety of the agent’s character to be involved in volitions (or at the very least to be the origin of volitions). “Volition” is ultimately the power or faculty of using one’s will, the ability to do what you intend (at least theoretically), and something must possess this power (this something being the agent’s character). The “thing” which ultimately drives your will and actions is the collection of desires and emotion (in short, one’s cognitive faculties) which come together in a roughly coherent state of experiencing and doing. Perhaps the distinction between volitions and will is relevant, and this leads to the agent being responsible for even less things. I do not believe this is a large concern for my argument as the agent can still be said to be responsible for what he intends. Nonetheless, I believe that conceiving of the agent as consisting of all their cognitive faculties (all those things that

form part of the character) is rational and justified, and they can, therefore, be held responsible for volitions *and* accompanying character traits.

5.4 Conclusion

At the start of this chapter, we discussed the way people have classically understood the problem of Constitutive luck. Using the work of Martha Nussbaum, we were able to see that so much of what and who we are seems to be influenced by factors outside of us. Children become adults and, in the process, develop a character influenced by the world. At the same time, we hold people responsible for the sort of person they are and the ways in which they act. Together, these form the intuition that we hold agents morally responsible for character traits not in their control. This is Constitutive luck.

I used the work of Robert Adams to discuss three different responses to the problem of Constitutive luck and why they fail. Each tries to solve the problem by focusing on control and how agents might be said to regain it with regards to their personalities. The first response considered was that agents are not held responsible for their personalities (which they do not have control over), rather, they are held responsible for the *actions* that flow from them (which they do control). Adams argues that this response fails because it neglects the fact that we often do, and should, hold people accountable for character faults even when no action is present. I also added that it is strange to claim that an agent has control of an action which follows from a mental state when that state itself is not in the control of the agent. The second response discussed was the idea that character traits are voluntary because they are operations of the will. Adams argued this response also fails because it uses a different definition of the word “voluntary”. It cannot be said that a character trait is voluntary, even if it could be said that it is an operation of the will, since the agent still does not choose her character traits (which is the issue at hand). Finally, we considered the response that character traits can be understood as indirectly voluntary if past voluntary actions led to their development, thereby making them indirectly voluntary. This response fails because it excludes a large amount of what we would usually hold the agent responsible for, namely character traits that already exist, even if not as the result of voluntary action. No matter how you choose to frame character traits, they are not something which the agent has control (or voluntary control) over.

Having reached the conclusion above, we next tried to solve the problem of Constitutive luck in a different manner. Rather than trying to regain control, it was argued that to speak of “control over one’s character” is to make a fundamental mistake. Using the work of Nicholas Rescher, we discussed the nature of the agent and what it means to have mental states. It was argued that it makes no sense to talk of an agent without character traits, and thus it is a mistake to argue that an agent should have

control of these character traits. For us to speak of an agent “having control” we must, prior to the discussion, already assume a character (that is, an agent with mental states). Using work by Daniel Dennett, I tried to help create the intuition that to speak of an agent we must include all those ways of experiencing and doing (namely memes), and by doing so, we cannot conceive of a mind (a thing that “does stuff”) without including all the accompanying character traits. The work of Clarke was also used to justify the metaphysical view of agent-causation in order to distinguish between mere *events* and *moral agents*. I also used Frankfurt cases to demonstrate what I mean by the notion of locus of control and how this relates to agents and their character. All of this, taken together, was meant to demonstrate that the problem of Constitutive luck only becomes a problem if we misunderstand the very nature of agents and control. We should not speak of agents lacking control over “their” personalities since this makes a category mistake. There are not two separate categories: person and personality; rather, there is only one: the agent, which should be understood as a single entity. The entire argument is eloquently summarised by Rescher (1993:155) when he says: “Identity must precede luck.” Therefore, Constitutive luck does not pose a true threat to moral agents.

Chapter 6: Putting It All Together

6.1 Introduction

The three chapters that came before this one each discussed a single instance of moral luck, at least as they were described and differentiated by Nagel.¹ This was done for the sake of argument as it is much easier to discuss each separately. However, this separation does not indicate an intrinsic difference between the different types of moral luck, nor is it helpful to keep the discussions separate. Although easier to discuss separately, in order to gain a full understanding of moral luck, why we believe there is a paradox, and where we have made mistakes, we need to discuss all types of luck together. Hence, this chapter will attempt to bring all that has been said together and sketch a coherent view of the moral agent and those things that she can legitimately be held responsible for.

In this chapter, I will discuss the three types of luck in a different order than I have up till now. First, the conclusions about Constitutive luck will be discussed, allowing us to develop a coherent understanding of the agent, since this helps us understand how both control and responsibility relate to the agent. Next, my conclusions regarding Circumstantial luck will be discussed, allowing us an understanding of the difference between worldly/epistemic luck and moral luck, as well as the relationship between moral judgement and character. This allows us to understand that luck does exist, but it is not moral and therefore does not undermine responsibility. Finally, the conclusions about Resultant luck will be discussed so that we can understand why it is more important to focus on intentions rather than the actual consequences of our actions when it comes to making moral judgements. In this way, I can develop an argument, each part building on the previous, so that we can assess the general implications of my arguments thus far and draw conclusions from it regarding the role that moral luck plays (or does not play) in morality.

¹ Recall that Causal luck did not receive its own chapter, as I argued in the introduction of this thesis that Causal luck either only refers to physical luck and any “moral” luck is subsumed by Circumstantial and Constitutive luck.

6.2 Agents, Control, and Responsibility

Before we addressed what it is that the agent can control and what she can be held responsible for in Chapter 5, two things were established: 1) what we mean by an agent (in this case a moral one) and 2) what it means for an agent to control something. Intuitively, we understand people to have a part that controls their actions and another part which is controlled. Using the work of Martha Nussbaum (1993) in Chapter 5, we saw that character traits or mental states are often seen as being beyond the control of the agent. If we accept the Control Principle, we hold that agents can only be held morally responsible for their character traits in so far as these traits are somehow under their control. As demonstrated by Robert Merrihew Adams (1985), this claim leads to a problem. Adams argued that there is no conception of control or of what constitutes character traits that allow us to conclude that agents have a sufficient level of control over their characters to satisfy the Control Principle. Characters are not *indirectly* voluntary by being operations of the will, since there is no way to make sense of the idea that agents can “try” or “will” their characters into being. In addition, we cannot conclude that characters are the result of *voluntary* actions, since this is not true in all cases, and it only pushes the problem back another step, as those voluntary actions would have had to be inspired by other involuntary character traits. We thus concluded that we cannot say agents have control over their characters.

All is not lost, however, as we can still get rid of the problem of Constitutive luck by revising our view of the agent. We have seen that Nicholas Rescher (1993) argues that the problem of Constitutive luck makes a fundamental category mistake by claiming that there is both the person and their personality. Rescher says there is only one category, namely the agent. When we speak of agents, it makes no sense to try and speak of something as being an agent without the accompanying character. Our personality is part of what constitutes us as the people we are. Using the work of Daniel Dennett (2017) and Randolph Clarke (2005), we tried to gain a better understanding of this. Using Dennett’s “just so” story of how we came to have the minds that we do, we can understand that the things we are are specifically machines with minds, minds which, as per their identity, include all those things we would call character. If we talk about something without character, we are not talking about moral agents. Clarke was used to help understand agent-causation and why we do not have to worry about determinism as a factor which diminishes responsibility. Building on this, I used a Frankfurt case to demonstrate that an agent’s character should be understood as the thing controlling (at least in part) the actions of the agents, rather than it being controlled by something else. When we make decisions and choose to act, our characters are part of this process and are an important part of our agency. All of this taken together means that to ask if agents have control over their characters is to ask the wrong question. It is only once an agent *has* a character that we can start to ask what they have control over.

In conclusion, the problem of Constitutive luck gets the issue backwards, and agents can only be expected to control things once they have a character. There is no such thing as an agent without – or prior to – character. Morality is our means of organising social creatures and the part of us that reacts to and interprets moral “rules” are our characters.

Now that we have a better understanding of the agent and what exactly it is that we expect to be in control, we can move on to understanding the relationship between the agent and their circumstances. The problem of Circumstantial luck states that different people face different challenges and tests, as well as different opportunities for success, and that this is not fair since the praise or blame accrued to them would have also been accrued to others, had they found themselves in the same circumstances. Again, we lack control over our circumstances and this violates the Control Principle, which undermines our moral responsibility. This claim was challenged by Norvin Richards (1993) who argued that people are often praised or blamed for counterfactual situations. Although in general we are judged for the things we actually do, if a counterfactual is convincingly similar in all the relevant ways, an agent can indeed be legitimately judged for said counterfactual. The example given by Judith Jarvis Thomson (1993), and subsequently expanded by me, of the two Judges open to bribes, was used to demonstrate this. What all these arguments are meant to show is that it is not *morality* which is susceptible to luck in these instances (examples of Circumstantial luck), rather, it is the certainty with which we can judge particular individuals that is susceptible to luck. The individual’s moral standing (the true moral judgement they deserve) is not itself subject to luck, but rather his moral record (what he has been judged for, given what we know about him). This is not moral luck but Epistemic Luck. The amount of certainty we can have about an agent is determined by their particular circumstances, not their true moral standing (Rescher, 1993). The conclusion we must draw then, is not that morality is unfair. Rather, we must draw the conclusion that life in general is unfair, and it is our responsibility to make sure moral judgements treat all people equally, as far as this is possible. Circumstantial luck thus gets the issue wrong by claiming that agents are judged for the *circumstances* they find themselves in, and because they cannot control these circumstances, they cannot be held responsible for them. Agents are, in fact, judged for their *actions* and *moral characters*, and although we often require circumstances in order to demonstrate these things, we are still being judged for what *is* in our control (character² and actions).

² Again, to remind the reader, I am not in fact saying that we do have complete “control” over our character (as demonstrated in the chapter on Constitutive luck), but we none the less still do not *lack* control over our characters such that we lose responsibility. The problem of Circumstantial luck is wrong with regard to what it is that we lack control of.

A brief summary of what has been said up until now is in order before moving on to the final part of the argument: (1) we have moral agents who, conceptually, possess all the relevant psychological states and faculties required for making moral decisions and taking action. Without these states/traits, we are not talking about moral agents. (2) It is this collection of states and faculties that do the “controlling”, and thus, it is this collection of states and faculties that are held morally responsible (agent-causation). (3) Moral judgement requires actions to judge and hence circumstances that allow one to act so as to demonstrate moral standing. (4) These circumstances and the accompanying judgements themselves do not determine the agent’s moral standing, but only his moral record. Next, I discuss exactly what it is that agents are judged for and also how relevant the actual consequences of an agent’s actions are for legitimate moral judgements.

When we considered the problem of Resultant luck, we analysed two different ways of arguing about it. The first was to ask if it is philosophically consistent to include actions in our moral assessments and whether the argument for moral luck in this regard even makes sense. The second was to ask whether we have good reason to include luck in our moral assessments; that is, would this lead to a better result. When we looked at the first question, we concluded that it is inconsistent to include the results of our actions (those things not in our control) when making moral judgements. Using Sverdlik’s (1993) examples of transferred intent, we were able to argue that including results in such judgements led to cases where we judged agents that intuitively are morally good as “bad” and agents that intuitively are morally bad as “good”.³ On the other hand, when we focus on intentions when making moral judgements, we have a coherent and intuitively plausible way of determining who should be blamed for what and to what extent.⁴ Hence, intentions, rather than results, are much better indicators of an agent’s moral character and therefore their moral responsibility. Furthermore, once we made clear what exactly we mean by “control”, using the arguments of Zimmerman (1993), we found that the apparent paradox that gives rise to moral luck disappears; since it is predicated on an invalid conception of “control”, one far too broad to be usefully applied to moral responsibility. Once we concluded that agents only need “restricted control” (control to either do or not do something, rather than having control of all the factors that led to the current situation), we saw that they do

³ The example we used was two men who each shot someone, but the first did so out of malicious intent while the second had their bullet bounce off a sign and hit someone else. If we include the result of their actions (that is, if we include luck) then these two agents are equally blameworthy even when we intuitively do not think so.

⁴ Again, using the example of the two men who each shot someone; when we do not include the results of their actions but rather focus on their intentions, we are able to say why the one is morally culpable while the other is not, since the first had malicious intent while the second intended to hit a target and was only unlucky.

indeed possess it, and can, therefore, rightly be held to be morally responsible. Taken together, these two arguments led us to conclude that agents do have the necessary control required for moral responsibility as they have “restricted control”, that is, the ability to choose how to react given a particular situation, and that this control is encapsulated by the agent’s character/intentions, since these are the parts which can be said to be doing the “controlling”.

Considering the next line of arguments in favour of including the result of actions in our moral judgements, we concluded that they too fail at supporting the idea that the results of actions are legitimate moral indicators. It was argued that Nussbaum’s (1993) view, that the luck present in the results of our actions is necessarily transferred to our moral judgements, makes a mistake of inference. Although it might be true that moral judgements require a demonstration, and therefore require action, it is not true that the luck present during this demonstration needs to be transferred to the moral judgement. This is because we are able to take account of all these factors and through rational moral deliberation we are able to account for the effects of luck. Furthermore, Margaret Walker’s (1993) argument that “impure agents” (those who accept the effects of luck on their responsibility) are better moral agents was also shown not to be true. Her first point, that “pure agents” have intuitively better responses to moral situations and judgements because the rationalist response seems insincere, was shown to be false and based on a strawman depiction of the “pure agent”. “Pure agents”, agents who are only responsible for what they intend and *will*, are just as able to accept the moral blame due to them, even if they would be blamed for less. Her second argument, that “pure agents” will not be responsible for those things we would traditionally hold agents responsible for, was also demonstrated to fail on the basis that she committed a fallacy of equivocation. Walker confused cases where we owe people certain responsibilities due to the relationships we have with them for cases where an agent is responsible for something because they “caused” it in some way. Once this is understood, we are able to see that “pure agents” would still be responsible for things such as their friendships and children, and the only things they would be less responsible for would be those things not in their control.

To statements 1 to 4 above, we can also add the following from our discussion on Resultant luck: (5) agents should not be held accountable only for the results of their actions; (6) agents do have sufficient control for moral responsibility, but only with regards to their characters/intentions; (7) there are no good reasons to conclude that agents should indeed be held accountable for things beyond their control, although, we may still conclude that the results of their actions play, at the very least, a demonstrative role in moral considerations.

6.3 Conclusion

It is now possible to sketch a preliminary conclusion on the basis of all that has been said about moral luck thus far. We are able to conclude that moral agents contain all the psychologically relevant aspects of control required by the concept of “agent” and that such agents can rightly be held morally responsible for their characters and intentions. The results of their actions and the circumstances in which they find themselves all serve as indicators of the agent’s moral character, but *they* do not determine an agent’s moral standing. Actions might also make a difference when it comes to the punishment or reward of the agent (as will be discussed next), but they should not have an effect on the moral judgement itself. All the claims that this thesis thus make are:

- 1) There are moral agents who, conceptually, contain all the relevant psychological states and faculties required for making moral decisions and taking actions. In the absence of these states/traits, we are not talking about moral agents.
- 2) It is this collection of states and faculties that are in “control”, and thus, it is this collection of states and faculties that should be held morally responsible (agent causation).
- 3) Moral judgement requires actions and therefore circumstances under which to act, so that moral standing can be demonstrated.
- 4) These circumstances and the accompanying moral judgements do not in themselves determine moral standing, but rather, only an agent’s moral record.
- 5) Agents should not be held accountable for the results of their actions only.
- 6) Agents do have sufficient control for moral responsibility, but only with regards to their characters/intentions.
- 7) There are no good reasons to conclude that agents should indeed be held accountable for things beyond their control, although, we may still conclude that the results of their actions play, at the very least, a demonstrative role in moral considerations.
- 8) Therefore, the apparent paradox of moral luck does not pose a true threat to moral responsibility.

The next chapter will focus on all the important implications of these conclusions. I will consider both the theoretical and practical implications. After this, my main conclusion will follow.

Chapter 7: Implications

7.1 Introduction

In this chapter, I will discuss a few implications of the arguments that this thesis makes. These implications all work from the assumption that these arguments and the conclusions that I draw from them are correct, and will therefore not attempt to justify any of them any further. I will rather focus on justifying the implications themselves. The chapter will be split into two sections. Firstly, I will consider the theoretical implications of my arguments against moral luck. With “theoretical implications” I mean those that have an effect on the concepts that are used within moral philosophy, as well as the metaphysical and ontological assumptions and normative commitments that we employ in moral philosophy. These will include our conception of the moral agent, the types of responsibility an agent might be said to have, and views on the role that actions should play in our moral judgements. Secondly, I will discuss the practical implications of my arguments. “Practical implications” are all those that have an effect on how we treat people in practice and how we structure societal institutions. These will include the role of prisons/correctional institutes and whether certain egalitarian practices are warranted. All of this is done in order to demonstrate the relevance of this thesis, as well as to demonstrate the effects our philosophical ideas have, or perhaps should have, on the world.

7.2 Theoretical Implications

7.2.1 Responsibility

As was made clear at the start of this thesis, the primary type of responsibility under discussion here is moral responsibility. This is the type of responsibility that warrants moral praise or blame (moral desert). It is distinct from being causally responsible for something. One of the main implications of this thesis, specifically in terms of the discussion on Resultant luck, is that these two are distinct senses of the word “responsible”, and that most people confuse the two when making moral

judgements. A useful example to illustrate the difference is given by Judith Andre.¹ Imagine someone accidentally breaks a vase in your home. In what sense is this person responsible? Andre (1993: 127) argues it is usually only in the “causal” sense. He is responsible for breaking the vase, but he is not morally responsible, and is, therefore, accorded no moral judgement. One is not said to be a bad person because one does a “bad” thing (e.g. breaking a vase). Rather, one is only responsible in the sense that one caused the vase to break. Now it might be said that one has a moral responsibility to rectify one’s mistake, but, again, here one is morally responsible *to do* something and not *for doing* something. We must take care to always distinguish between these types of responsibility.² Further implications of this point will be discussed in section 7.2.4.

7.2.2 Moral Agency

Although I have already made clear the required revision to the concept of a moral agent in the chapter on Constitutive luck, I think it is important to highlight the most salient points here. I made the argument that we cannot conceive of a moral agent without all the relevant character traits of the agent already included in the concept. To speak of a moral agent is to speak of something that already has dispositional states, as it makes no sense to envision a moral agent without these. In the metaphysics of agency, this view is called “volitionism”, and it claims that an action can only be understood as being caused by the agent if it was an act of volition (an act of character and will). For examples of this, see Ginet (1990) and McCann (1998).

Another implication of this view of agency is the recognition that agents have personal histories. As has been discussed, no one is exempt from the effects of time. This means we have to take into account the history of an agent when judging her character. Although a bad character remains a bad character, how this character is judged should be relative to their history, as this history is relevant to the intentions that they have and the reasons they have for doing something. This is not to say that an agent will deserve greater moral blame than another *because* of their histories. Rather, their respective histories are important considerations when evaluating a moral character. If two agents both steal from someone, but the one has had to steal her entire life just to survive, whereas the other only steals for pleasure, we would be justified in judging the former less harshly than the latter. This is not because of a difference in history, but because of a difference in moral character which can only be

¹ David W. Concepcion (2002: 455 – 457) also places emphasis on this distinction and argues that once we properly differentiate between “causal responsibility” and “moral responsibility” the paradox of Resultant luck disappears.

² For a much more in-depth look at the different types of moral responsibility, see Talbert (2019).

understood once histories have been considered. An agent who steals because they have concluded it is good despite being taught otherwise, has a “worse” character than an agent who steals because they have been taught to do so. Returning to the example above, perhaps the moral wrong committed by the first is the willingness to consider theft a legitimate option and the moral wrong committed by the second is the malicious intent to take from another. These are two different motivations and require different judgements. The first agent has a character that considers theft a viable option for survival, ignoring cases where we would agree theft is acceptable, we can argue that it is morally wrong to view theft in such a neutral light, and therefore, that this agent is morally responsible for this. The second agent has a character that finds pleasure in theft and commits it out of malice. Arguably, this is a much worse moral character (perhaps we can frame it as being less virtuous), and therefore, the second agent is more morally culpable. These attitudes and character traits are only understood because we have considered their personal histories.

The main reason this distinction is included here is because of the strong intuitions people have about agents. If we are going to understand and agree with all the other implications found in this thesis, we must view the agent as an irreducible system of wills, wants, and character traits with a personal history.³ Further implications of this are discussed in sections 7.2.3 and 7.3.3.

7.2.3 Free Will

Recall that at the beginning of this thesis, I discussed Causal luck and stated that it is closely related to the issue of free will. Also, I argued that once we eliminate all those factors already captured by the other three types of moral luck, we are left only with those phenomena not present in our subjective experience of the world. Understood this way, Causal luck describes the idea that we are not in control of how preceding events cause antecedent events, or stated otherwise, we lack free will. This thesis does not concern itself with the problem of free will as it argues that responsibility is related to character and intentions, not with alternative possibilities or free will (see Chapter 5), and the main focus of this thesis is moral responsibility.

Throughout this thesis, I have argued that agents do have the necessary control in order to be held morally responsible. Importantly, if my view on Constitutive luck (as well as the section on Causal luck) is right, we might conclude that moral responsibility follows from the fact that agents can act in accordance with their will, and moral blame and praise can ultimately affect how we act in the

³ For a full discussion on all the philosophical work on agency, see Schlosser (2019).

future. On this view, we do not require absolute control over causes and effects, rather, we only require the appropriate amount of control over our actions as a product of the will. The *will*, therefore, does not need to be “free” in the traditional sense, only free from coercion so as to make it that whatever the agent does will ultimately be in accordance with their character and intentions. This view was justified in Chapter 5 where I argued responsibility can be ascribed to moral characters without the requirement of alternative possibilities.

All of this taken together means that this thesis is in agreement with those who state that we do not require free will in order to have moral responsibility, or those who argue that free will is in some way compatible with determinism. I take no position of whether we do or do not have free will,⁴ and I am silent on the issue since, for the purpose of this thesis, this is not important for according moral responsibility.

7.2.4 Actions and Consequences

The chapter on Resultant luck made the argument that an agent cannot be held morally responsible (in the praise and blame sense) for the consequences of her actions, since these consequences are not under her control. I did, however, touch upon the idea that actions and their consequences might serve as demonstrations of will/moral character. In this sense, actions and their consequences are important, but epistemically so, not morally so. This all still pertains to moral responsibility, but there is also a more practical type of responsibility, which is affected by the consequences of actions.

Recall the example of the agent breaking a vase by accident. We concluded that the agent is not morally responsible for breaking the vase, but he is still causally responsible for breaking it. Now, it might be argued that, given our relationship with others and a certain expected level of courtesy between individuals, the agent is still responsible for replacing the vase. That is, we might be justified in saying that people are responsible for rectifying certain consequences based solely on the fact that our actions do have consequences. Although these consequences are not in our control (and therefore do not inherently demonstrate moral responsibility *for doing something*) they are still part of our responsibility, as we are causal influences (agent-causation). This is a different sense of responsibility because it is not responsible for doing something but responsible to do something. To be mindful of others around you, to respect their wishes, to inform yourself before making decisions, and to fix what you broke are all things we are morally responsible to do. This same line of reasoning can be

⁴ For a thorough review of all the major works on free will, see Franklin & O’Connor (2019).

applied to those cases where agents are equally responsible in the moral sense, but we would argue that they deserve different punishments.

Take the example of the two drunk drivers again. Both are equally responsible (in the moral sense) because both had the same level of negligence; however, as a matter of luck, the one kills a child while the other does not. As per the argument of this thesis, we would agree that they are equal in their moral standing. They are both equally morally blameworthy. Nevertheless, we do not have to conclude that they should be punished in the same manner. It could be argued that since A actually *caused* something more serious to happen than B, A has a responsibility to rectify these consequences while B does not. This is what Susan Wolf calls the nameless virtue (2013:13). The nameless virtue is the recognition of an agent that their actions have consequences and that they are obligated to rectify them simply because they are *causally* culpable. Imagine two agents, P1 and P2, who both knock over their little brother's dominoes set. In the case of P1, his brother set up 20 dominoes in a row. In the case of P2, her brother setup 100 dominoes in a row. In both cases, P1 and P2 knocked over the dominoes because they were not paying attention. We would thus say both are equally morally responsible for they both were equally negligent. The difference, however, is that P1 is responsible to rectify less than P2. He only needs to re-set 20 dominoes to rectify the wrong while P2 needs to re-set 100. Both are equally morally culpable but not equally causally culpable.

The case of the drunk driver is a difficult example to use since it is hard to say what someone can do to rectify the loss of a life (such a thing seems impossible to some). Perhaps the agent must complete some form of community service or spend their time as part of an awareness program. Both must be punished for their negligence, but only one can (for instance) be used as an example of the consequences (for others) of drunk driving. Either way, the point is that consequences are not entirely irrelevant. Although they may not serve to demonstrate moral responsibility, they should be considered when deciding how an agent ought to be treated and how the agent ought to feel after the fact. Many of the writers on moral luck agree with this distinction and argue that it also helps explain why we intuitively treat the two drivers differently and why the two drivers should *feel* different (Thomson (1993); Rosebury (1995); Wolf (2013); MacKenzie (2017)). It is argued that we believe the one should feel worse than the other because he *did* something worse, and we expect good moral agents to have emotional reactions to their actions. We mistakenly believe that this justifies a greater moral indignation, but it might only justify greater punishment.

7.2.5 Negligence and Obligation

The final theoretical implication I wish to discuss is in some way a response to those who argue that serious problems arise from only focusing on intentions when it comes to moral responsibility. For many philosophers, it is not only the Control Principle which we must focus on when deliberating about moral responsibility, but also the epistemic principle. This principle states that for an agent to be held morally responsible for an action, that agent must have been aware of the potential consequences of a given action.⁵ The worry is that agents might appeal to ignorance in order to argue for their moral innocence if we say agents can only be held responsible for their intentions (wilful ignorance as excuse). This has led some to worry that agents who do harm but who did not intend it will appeal to ignorance, and might have excuses for their negligence and hence escape rightful moral blame (Rosen (2004); Levy (2011)). An example would be an agent who fires a gun into the air. Later it is discovered that the bullet from the gun had hit someone and fatally injured them. The worry is that this agent might appeal to the fact that he did not know the bullet would hit someone and that they should therefore not be held responsible for it since they did not intend it.

Rescher (1993) argues that this “excuse” can be countered with a principle of moral obligation to avoid negligence. In an attempt to argue that agents cannot appeal to ignorance *simpliciter* to escape the moral judgement due to negligence, Rescher argues that agents have some obligations to contemplate the potential consequences of their actions; ignorance cannot be used as an excuse in all cases (1993: 149). Rescher (*ibid*) frames it as agents having the ability to mitigate the effects of luck, and he suggests three ways:

1. Risk avoidance: people who do not court danger (who do not try to cross the busy roadway with closed eyes) need not count on luck to pull them through,
2. Insurance: people who take care to make proper provisions against unforeseeable difficulties by way of insurance, hedging, or the like need not rely on luck alone as a safeguard against disaster.
3. Probabilistic calculation: people who try to keep the odds on their side, who manage their risks with reference to determinable probabilities, can thereby diminish the extent to which they become hostages to fortune.

We can thus use these three principles (I would call them moral obligations) to see if an agent had taken all due care, and determine whether they did indeed fall victim to a type of luck (misfortune),

⁵ The work on this condition is vast and convoluted and cannot be covered here. For a good review and discussion on all the works regarding this, see Rudy-Hiller (2019).

which would negate the moral blame due to them. We are thus left with an account of intentions which includes the standards by which to judge whether these intentions were justified, given the circumstances. This serves as a counter to the claim that agents can use ignorance (because we are focusing on intentions) as an excuse. Ignorance is only an excuse when it is legitimate and not wilful.

7.3 Practical Implications

7.3.1 Criminal Justice

The view of moral agency that this thesis argues for has implications, both theoretical and practical, for our view of criminals and how they ought to be punished. All of these implications will be discussed in this section, as the theoretical implications lead directly to the practical ones.

Many people hold, intuitively, that justice ought to be retributive. Retributive justice is the principle that people should be punished for their wrongdoing, not for any practical reason (such as deterrence), but solely on the bases that they deserve punishment because they did wrong. If an agent hurts someone else, then they deserve to be hurt back. This principle, I argue, is not compatible with what has been concluded in this thesis. The principle of inflicting harm for harm's sake does not take into account the connection between the harm caused by an agent and their moral character. If we only punish someone because they "deserve" punishment, no positive effect on their moral character is guaranteed. Moral judgement, in part, is meant to demonstrate the relationship between an agent's intentions and actions and their moral character. Punishment should serve as a way of fixing any "moral fault" while reward should reinforce the "moral virtue". Simple restitutive justice does not serve any function other than making victims "feel" better because the perpetrator now "feels" bad. This thesis argued that moral responsibility is an indication of the relationship between an agent's character and their actions, and therefore, moral desert should serve to highlight this relationship and to correct/reinforce the character that led to the action in question. No other sense of responsibility and desert allows for this emphasis on moral character that strengthens the agent's sense of moral responsibility and has greater practical effect in creating a more "moral" society. My thesis thus argues that justice ought to be "restorative" or "rehabilitative". My personal preference is rehabilitative justice since this justice encapsulates all that has been said up until now: it treats an agent as having a character with "faults" or "virtues" and aims to deter us from the first and reinforce the second. Rehabilitative justice states that agents should be punished in such a way that the fault which led to the moral wrong is corrected in some way.

All of this serves to inform the practical issue of criminal punishment. Our legal systems and prisons should not simply lock people away. Legal systems should understand and view criminals as moral characters requiring rehabilitation and eventual re-entry into society. Our systems of punishment should thus stop viewing punishment as “pay back”, but rather as a means to hold agents accountable for their actions and eventually to help foster better moral agents.⁶ For two thorough and useful discussions on the implications of moral luck on the criminal justice system, see Bittner (2008) and Moore (2014).

7.3.2 Egalitarianism

In the chapter on Circumstantial luck, it was argued that although individuals do not have all the same opportunities, they are still held morally responsible only for those things pertaining to their character in some sense, and that circumstances, therefore, serve to demonstrate this character. It was thus concluded that Circumstantial luck is not really moral luck but rather Epistemic luck. One practical implication of this may be the argument that, seeing as life is unfair and has an influence on your moral record, we have an obligation to give everyone equal opportunity within society to try and mitigate this “unfairness” from a moral perspective.

The conclusions of this thesis support political egalitarianism.⁷ If we want our moral records to be as close to our moral standings as possible, society should distribute resources in such a way that everyone, as far as possible, receives the same opportunities and equal access. This is because so much of what we are able to do, and what we do actually end up doing, is based on the resources we have available to us. Equal access to food, water, shelter, schools, jobs, and so on will allow for moral characters to face many of the same challenges and access many of the same opportunities. This will allow for our moral records to be much more in line with our true moral standings.

Another economic principle which is supported by this approach is what is called a Universal Income Grant (UIG). Briefly, a UIG would give all citizens of a country the minimum amount of funds every month required to not live in poverty, and in this way allow everyone to contribute as much as they can without having to worry about “surviving” every month. If we want our economic systems to be “moral”, we might have reason to agree that a completely free market capitalist system is not the best

⁶ This discussion only touches the surface of the debate. Here, I only wish to highlight the shift in understanding that is required from a legal perspective. For a more thorough discussion of punishment see Kelly & Bedau (2019), and for theories on criminal law see, Edwards (2019).

⁷ For a summary of the major works on egalitarianism and a discussion of them, see Arneson (2019).

option as such an option leads to great inequality in both resources and opportunities. A redistribution of wealth and funds is supported in this manner. It might be the case that one of the best ways to give everyone equal opportunity is with a UIG.

Of course, all of this should also be considered practically. If the idea of a UIG will lead to a society less well off, it should be abandoned. If such a UIG, in the long run, results in a weak economy and extremely poor government, the desired equality might be undermined. The right to equal moral opportunities/circumstances is only a right because it has positive effects (for both individuals and society). If these effects are not achieved we have reason to endorse a different system with better results. Here, the UIG is used as an example to demonstrate the type of political/economic impact this thesis might have.

7.4 Conclusion

This chapter was meant to discuss some of the implications of this thesis if it is taken to be correct. It also serves to demonstrate the relevance of this thesis as well as what impact our philosophical ideas might have. It is a common occurrence to meet someone not well-versed in philosophy and to find that they believe it to be impractical. Hopefully, the brief discussion in this chapter has demonstrated this to be false. Philosophical ideas have an enormous impact on many other aspects of life. Although we might not always notice it, every idea we have is made up of concepts that are often hard to understand and relate to one another in ways not always expected. Moreover, many of our institutions are predicated on values and perspectives that are philosophical in nature. If we are going to be consistent with our ideas, inferences, and practices, we must update such societal systems. The conclusions we reach about moral luck serve as a foundation for many other philosophical issues to be resolved. In this chapter, I have discussed those I find most interesting and important. The final chapter of this thesis will summarise everything that has been said throughout.

Chapter 8: Conclusion

Throughout this thesis, I have discussed some of the most influential and relevant papers that have been written on moral luck. As explained, the problem of moral luck is the seeming paradox that occurs due to the fact that we believe two contradictory things, both seeming to be intuitively true. The first is the Control Principle which states that an agent can only be held morally responsible for those things within their control. A corollary to this principle is that an agent cannot be held responsible for that which is not in their control. The second is the fact that we do, in general, judge agents for things not in their control, and intuitively we feel this is right.

The justification for the Control Principle is both intuitive and moral. As was discussed, we intuitively believe the Control Principle to be true because it seems right. More than this, the Control Principle also seems like a fundamental part of our moral systems. We can see this in the writings of Kant, Williams, and Nagel who all agree, and argue, that our moral frameworks must presuppose the Control Principle in order to make sense and be coherent. Furthermore, many thought experiments demonstrate our acceptance, and the general validity, of the Control Principle. Recall our discussion of the two people who spill coffee on a colleague. While it was within the one's control to not do so, the second could not but spill coffee. When assessing the moral culpability of these two agents, most would agree that the first is morally responsible for their negligence while the second is not morally blameworthy. We treat them differently because we accept the Control Principle.

Unfortunately, other thought experiments seem to demonstrate that we do also intuitively judge agents for what is not in their control. The discussion of this issue was introduced using the work of Williams who wrote about an aspiring artist in order to demonstrate the conflict between the Control Principle and our moral evaluations. For Williams, the story of Gauguin demonstrates the difficulty of assessing whether an action is a rational one before we have done it. The actual consequences that come about because of one's actions can make a difference to whether one did the right thing or not. For his negligence and decision to follow his dreams to be justified, they must lead to desirable and successful outcomes. The problem is, we cannot know before we have acted whether we will be successful. The justifiability of his decision is ultimately out of his control since the ultimate success or failure of his actions depend on many other factors not within his control. Williams argued that this fact might also have bearings on moral responsibility. In order for us to have a better understanding of the moral nature of the problem, we turned to the work of Nagel. Nagel identified

four different, but related, ways that an agent can lack control in such a way so as to make a moral difference, seemingly contradicting the Control Principle. This thesis attempted to discuss and solve each type as presented, except in the case of Causal luck, as I argued it does not fall within the scope of this thesis and that any “moral” relevance it holds is subsumed by Circumstantial and Constitutive luck.

The first way an agent is said to lack control is in terms of Resultant luck. Resultant luck denotes any time an agent is morally assessed for the consequences of her actions, even when those consequences were not in her control. The famous example used in this thesis and by most writers on the issue is the case of the two drunk drivers. Both are negligent and decide to drive while intoxicated, but only one ends up killing a pedestrian. The only reason the one driver killed someone, and the other driver did not, is due to factors beyond both of their control, namely that someone was in the wrong place at the wrong time and in the way. This difference seems to change how much we blame the two drivers even when it is due to factors beyond their control.

Next, we discussed what Nagel termed Circumstantial luck. This is when an agent lacks control of the circumstances (the opportunities and challenges) he finds himself in and therefore over what it is about them that comes in for moral assessment. The most famous example used in these discussions is that of the Nazi collaborator and non-collaborator. We judge those Germans who were Nazi collaborators as being morally wrong, yet most people (arguably) would have been Nazi collaborators had they found themselves in Nazi Germany. The moral blame accrued to the collaborator and non-collaborator is seemingly only different because of circumstances that were not in either of their control.

Finally, we discussed what has historically been termed Constitutive luck. Constitutive luck is the type of luck (lack of control) one has with regards to one’s constitution, one’s character. Who we are, the emotions we have, and the motivations we possess are influenced by many factors that are not in our control. We are largely a product of a long history of causal processes which have shaped who we are. At the same time, we judge agents for the emotions and intentions they have, even when we cannot reasonably say they had control over them.

Once these three types of luck were discussed and understood, we moved on to discussing and evaluating the most relevant and influential papers that have been written on the topic. In each chapter, I argued that the paradox of moral luck only exists because we have accepted mistaken intuitions, and once we understand all the relevant factors and criteria properly, the paradox of moral luck disappears. What follows will be a brief summary of the arguments discussed and the conclusions reached.

The chapter on Resultant luck started with a brief clarification of the assumption of moral constructivism to help our understanding of the discussion as it would follow. It was argued that most of the writings on Resultant luck discuss the issue as being one of deciding if we should include the consequences of actions, and in so doing, allow luck to infiltrate our moral evaluations, or if we should exclude consequences. The chapter was then split into two sections, the first discussing why we ought to reject Resultant luck, the second discussing those who argue we should accept it. Using the work of Sverdlik, I argued any moral system which includes Resultant luck will be internally inconsistent and will deliver moral evaluations which are unfair and counterintuitive, while moral systems which exclude Resultant luck in this way have a consistent way of arguing why agents are or are not responsible. I also argued, using the work of Bernstein, that including the consequences of our actions in our moral evaluations leads to a problem of scope and thus more inconsistencies while making moral judgements. After this, the work of Zimmerman was used to demonstrate that the problem of Resultant luck only occurs because we have an inadequate conception of the “control” that moral agents require. It was argued that once we refine and limit our conception of control to what Zimmerman calls “restricted control” the paradox of Resultant luck disappears. Following this, the works of Nussbaum and Walker were used to discuss those who argue that by removing luck from our moral evaluations we lose something important to morality. Nussbaum argued that we require action for morality, and since luck is intrinsically bound to action, we will remove this important aspect of morality if we argue that luck should be excluded. I argued that although morality might require action, the luck present can be accounted for and its effects mitigated using rational inquiry. In other words, no such transfer of luck from actions to moral judgements is required. Walker argued that by removing luck from the equation we will remove so much of what agents are responsible for, we will diminish the scope of morality. I countered this argument by arguing that Walker’s argument was based on an equivocation between two different senses of “responsible” (namely the responsibility *for* something and the responsibility *to do* something) and that we might even gain a better way of relating moral responsibility to an agent’s character if we accept the thesis of “pure agency” and not include luck in our moral evaluations. I summarised this by stating: if whether you are good or bad depends on things out of your control, why would you want to control whether you will be good or bad? If what you are judged for depends on who you are, you will want to control and change who you are. All of the above was meant to serve as an argument that Resultant luck should not be included in our moral evaluations and only becomes a problem as a result of focussing on the wrong moral criteria.

The chapter on Circumstantial luck started with a discussion on the work of Nussbaum to help clarify the issue at hand, namely that agents are judged for actions even when these actions were relevantly

determined by events outside of their control (namely the environment). After this, the work of Richards was used to demonstrate that we can and do judge agents for counterfactual intentions and actions. This was done to clarify what the true moral difference between a thief and a would-be-thief (for example) is. Richards argued that the only real difference is the opportunity for demonstration of moral character and, therefore, the amount of certainty we can have about agents' moral characters when making moral evaluations. Just as before, we require action to demonstrate moral character, and because of this, it is our moral knowledge relating to an agent which is subject to luck and not the agent's true moral status. Richards, Rescher, and I thus argue that Circumstantial luck arises from Epistemic Luck, not moral luck. It might be true that we cannot always control the way others see our moral record, but our moral standing is still under our control. The rest of the chapter discussed why so many thought experiments seem to support the idea that Circumstantial luck constitutes moral luck. Using the work of Rescher and Thomson, it was argued that the reason for this is because we are under-describing the cases in the thought experiments and hence not making clear *who* exactly we are talking about what the reasons for their actions are. Once we do make these considerations clear, we can see that cases of Circumstantial luck (when sufficiently described) only demonstrate Epistemic luck and not moral luck.

The chapter on Constitutive luck started with a discussion on the work of Nussbaum in order to help us understand what we mean by an agent lacking control over their character. History and life have a massive impact on the person you are, which implies that who we are (i.e. our character) is not significantly within our control. After this, the work of Adams was used to discuss the problem of control in much more detail. Adams considers three different ways we can conceive of control in order to address the apparent paradox of moral luck, and he concludes they all fail. I agreed with Adams and argued there is no way of conceiving of the agent as having control over her character. Following this, I used the work of Rescher to argue that the problem of Constitutive luck makes a category mistake since, by discussing the issue in terms of the agent lacking control over her character, it presupposes a false distinction between the agent and her character. Once moral agency is properly understood, we can see that there are not two different categories (person and character), but rather only one roughly coherent agent. The work of Dennett was used to help flesh out this conception of the agent. Frankfurt-cases were then used to help describe this agent and their accompanying character as the locus of control where moral responsibility lies. All of this taken together was meant to demonstrate that the paradox of Constitutive luck does not hold, since it presupposes a false conception of the moral agent (a non-agent-causal metaphysics) and what it is that is doing the controlling.

Each of these chapters attempted to resolve the problem of moral luck by focussing on the individual instantiations of it. However, as was mentioned in my introduction, this separation is only practical and for us to understand the solution to the problem we must draw together all the discussions and sketch a coherent view of control, agency, and responsibility before we can conclude that we have gotten rid of the paradox of moral luck. I attempted this in the Chapter 6. The structure of this chapter ran in the opposite direction as the thesis up until that point. This was because it was more useful to start with our concept of the agent and control, and then move on to how they interact with the world, and finally to what the focus of moral evaluations should be. Doing it this way allowed for a better flow and understanding of the discussion. The chapter focused on the conclusions of all those that came before, summarising each in an attempt to sketch a coherent view of moral responsibility. What follows is a further brief summary of these conclusions.

Because of the work done in this thesis on Constitutive luck, we can conclude that the concept of the agent must include all those psychological faculties that are part of the deliberation process, as well as eventual intentions, when making decisions. Furthermore, this concept of the agent must be understood as the locus of control over decisions and must, therefore, serve as our focus for moral evaluations. Next, the chapter on Circumstantial luck concluded that circumstances do not have an effect on one's moral standing, but on one's moral record, thus placing even more focus on the character of the agent in question rather than the results of their actions. Finally, the chapter on Resultant luck concluded that there are no good reasons to consider the consequences of an agent's actions when assessing moral responsibility. Instead, these consequences might only serve as practical considerations for determining punishment or reward. Taken together then, we can conclude that we have a coherent view of agency, control, and responsibility which does not allow luck to infiltrate our moral evaluations. We can, therefore, conclude that the paradox of moral luck does not pose any true threat to moral responsibility.

Continuing with this conclusion, the next chapter discussed some of the most relevant implications of my discussion and analysis. In an attempt to demonstrate the importance of this issue, as well as to raise some potential future research topics, I discussed both the theoretical implications as well as the practical implications of my findings regarding moral luck. In terms of theoretical implications, I argued that we must have a much more robust understanding of responsibility and take care to differentiate between the different types. I also summarised some of the implications for our conception of moral agency and for the free will debate, arguing that this discussion can shed light on both. Finally, I also discussed what role actions should play in our moral evaluations, as well as making the suggestion that agents have an obligation to be epistemically informed before acting to counter some issues when focusing on intentions. Two practical implications were also considered.

The first was that this discussion has implications for justice and criminal law, mainly that justice ought to be rehabilitative rather than retributive. The second implication was a political/economic implication. I argued that, given the results of the chapter on Circumstantial luck, we might be justified in arguing it would be better to support egalitarian practices in society insofar as this would allow for equal moral opportunities.

I have discussed a large number of the most influential writings on the topic of moral luck, and I have discussed a variety of solutions to the problem. I also hope to have given a good overview of what has been said. Significantly, I have argued that each instance of moral luck relies on faulty arguments and intuitions. If my thesis is taken to be correct, we must conclude that moral luck is not a real threat to morality and that we can continue to hold agents morally responsible, provided that we have made the necessary revisions. As mentioned in my introduction, morality is such an important part of what makes us human. Furthermore, our ability to investigate and argue about our moral practices has both positive and negative effects. Hopefully, my discussion and argument as found in this thesis has had a positive effect. I believe it is our rational approach that can save us from this apparent paradox, and more than this, I believe it is extremely important to do so. My final conclusion to the problem of moral luck can be stated as follows: the problem of moral luck states that we can only hold agents responsible for those things in their control; yet, we also do hold people responsible for many things not in their control. However, this paradox between our principles and intuitions is a false one, and once we revise our intuitive understanding of agency, control, and responsibility the paradox disappears, and we are able to conclude that agents can rightly be held morally responsible for their intentions/character.

References

- Adams, R.M. 1985. Involuntary Sins, *The Philosophical Review*, 94: 3 – 31.
- Alexander, L. and Moore, M. 2019. *Deontological Ethics* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/ethics-deontological/>.
- Andre, J. 1993. Nagel, Williams, and Moral Luck, in Statman, D. *Moral luck*. Albany: State University of New York Press. 123 – 130.
- Arneson, R. 2019. *Egalitarianism* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/egalitarianism/>.
- Bernstein, S. 2019. Moral Luck and Deviant Causation. *Midwest Studies In Philosophy*, 43(1): 151 – 161.
- Bittner, T. 2008. Punishment for Criminal Attempts: A Legal Perspective on the Problem of Moral Luck. *Canadian Journal of Philosophy*, 38(1): 51 – 83.
- Browne, B. 1992. A Solution to the Problem of Moral Luck, *The Philosophical Quarterly*, 42: 343 – 356.
- Caruso, G.D. 2019. Why Luck (Still) Undermines Moral Responsibility, *Journal of Information Ethics*, 28(1): 51 – 72.
- Clarke, R. 2005. Agent Causation and the Problem of Luck, *Pacific Philosophical Quarterly*, 86: 408 – 421.
- Concepcion, D.W. 2002. Moral Luck, Control, and the Bases of Desert, *The Journal of Value Inquiry*, 36: 455 – 461.
- Dennett, D. 2017. *From Bacteria to Bach and Back*. Penguin Books.
- Descartes, R. (2019). *Meditations on First Philosophy*. Lanham: Dancing Unicorn Books.
- Domskey, D. 2004. There is No Door: Finally Solving the Problem of Moral Luck. *The Journal of Philosophy*, 101(9): 445 – 464.
- Domskey, D. 2005. Tossing the rotten thing out: Eliminating bad reasons not to solve the problem of moral luck, *Philosophy*, 80(4): 531 – 541.

- Dougherty, M.V. 2004. Moral Dilemmas and Moral Luck, *Proceedings of the American Catholic Philosophical Association*, 78: 233 – 246.
- Driver, J. 2019. *The History of Utilitarianism (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/utilitarianism-history/>.
- Edwards, J. 2019. *Theories of Criminal Law (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/criminal-law/>.
- Fischer, J.M., and Ravizza, M. 1998. *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge: Cambridge University Press.
- Frankfurt, H. 1969. Alternate Possibilities and Moral Responsibility. *The Journal of Philosophy*, 66(23): 829.
- Franklin, C. and O'Connor, T. 2019. *Free Will (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/freewill/>.
- Ginet, C. 1990. *On Action*, Cambridge: Cambridge University Press.
- Gowans, C. 2019. *Moral Relativism (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/moral-relativism/>.
- Hansen, H. 2019. *Fallacies (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/fallacies/>.
- Hartman, R.J. 2014. How to Apply Molinism to the Theological Problem of Moral Luck, *Faith and Philosophy*, 31(1): 68 – 90.
- Hartman, R.J. 2019. Moral luck and the unfairness of morality. *Philos Stud*, 176: 3179 – 3197.
- Hatfield, G. 2019. *René Descartes (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/descartes/>.
- Hursthouse, R. and Pettigrove, G. 2019. *Virtue Ethics (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/ethics-virtue/>.
- Kane, R. 2005. *A Contemporary Introduction to Free Will*. 1st ed. Oxford University Press.
- Kant, I., Gregor, M. and Timmermann, J. 2011. *Immanuel Kant: Groundwork of the Metaphysics of Morals*. Cambridge: Cambridge University Press.
- Kelly, E., and Bedau, H. 2019. *Punishment (Stanford Encyclopedia of Philosophy)*. [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/punishment/>.

- Levy, N. 2011. *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*, Oxford: Oxford University Press.
- Levy, N. 2015. Dissolving the Puzzle of Resultant Moral Luck. *Review of Philosophy and Psychology*, 7(1): 127 – 139.
- Lewontin, R. C. 1982. *Biological Determinism*, University of Utah, The Tanner Lectures on Human Values, March 31 and April 1.
- Link, M.W. 2013. Moral Luck and the Condition of Control, *Southwest Philosophy Review*, 29(1): 99 – 106.
- MacKenzie, J. 2017. Agent-Regret and the Social Practice of Moral Luck, *Res Philosophica*, 94(1): 95 – 117.
- McCann, H.J. 1998. *The Works of Agency: On Human Action, Will, and Freedom*, Ithaca: Cornell University Press.
- Moore, M. 1997. *Placing Blame: A theory of the Criminal Law*, Oxford: Clarendon Press.
- Moore, M. 2014. *Causation and Responsibility*. Oxford: OUP Oxford.
- Nagel, T. 1993. Mortal Question, in Statman, D. *Moral luck*. Albany: State University of New York Press. 57 – 72.
- National Health and Medical Research Council. 2015. *NHMRC Information Paper: Evidence on the effectiveness of homeopathy for treating health conditions*. Canberra: National Health and Medical Research Council; 2015.
- Nussbaum, M. 1993. The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy, in Statman, D. *Moral luck*. Albany: State University of New York Press. 73 – 108.
- Pells, R. 2015. A Modal Solution to the Problem of Moral Luck, *American Philosophical Quarterly*, 52(1): 73 – 88.
- Pereboom, D. 2006. *Living without free will*. 1st ed. Cambridge: Cambridge University Press.
- Rescher, N. 1993. Moral Luck, in Statman, D. *Moral luck*. Albany: State University of New York Press. 141 – 166.
- Richards, N. 1993. Luck and Desert, in Statman, D. *Moral luck*. Albany: State University of New York Press. 167 – 180.
- Ritchie, J. 2014. *Understanding Naturalism*. Hoboken: Taylor and Francis.
- Rosebury, B. 1995. Moral Responsibility and Moral Luck, *Philosophical Review*, 104: 499–524.

- Rosen, G. 2004. Skepticism about Moral Responsibility, *Philosophical Perspectives*, 18: 295–313.
- Royzman, E., and Kumar, R. 2004. Is Consequential Luck Morally In-consequential? Empirical Psychology and the Reassessment of Moral Luck. *Ratio*, 17(3): 329 – 344.
- Rudy-Hiller, F. 2019. *The Epistemic Condition for Moral Responsibility* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/moral-responsibility-epistemic/>.
- Ryle, G. and Tanney, J. (2009). *The concept of mind*. London: Routledge.
- Schinkel, A. 2009. The Problem of Moral Luck: An Argument Against its Epistemic Reduction. *Ethical Theory and Moral Practice*, 12(3): 267 – 277.
- Schlosser, M. 2019. *Agency* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/agency/>.
- Schwitzgebel, E. (2019). *Belief* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/belief/>.
- Sellars, W. 1962. *Science, Perception, and Reality*. London: Routledge and Paul.
- Singh, S. 2004. *Big Bang*. New York: Fourth Estate.
- Starr, W. 2019. *Counterfactuals* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/counterfactuals/>.
- Statman, D. 1993. *Moral luck*. Albany: State University of New York Press.
- Streetepistemology.com. (2019). *Street Epistemology | StreetEpistemology.com*. [online] Available at: <https://streetepistemology.com/>.
- Sverdlik, S. 1993. Crime and Moral Luck, in Statman, D. *Moral luck*. Albany: State University of New York Press. 181 – 194.
- Talbert, M. 2019. *Moral Responsibility* (*Stanford Encyclopedia of Philosophy*). [online] Plato.stanford.edu. Available at: <https://plato.stanford.edu/entries/moral-responsibility/>.
- Thomson, J. J. 1993. Morality and Bad Luck, in Statman, D. *Moral luck*. Albany: State University of New York Press. 195 – 216.
- Waal, F. 2013. *The bonobo and the atheist*. New York [etc.]: W.W. Norton & Company.
- Walker, M. U. 1993. Moral Luck and the Virtues of Impure Agency, in Statman, D. *Moral luck*. Albany: State University of New York Press. 235 – 250.
- Williams, B. 1981. *Moral Luck*. Cambridge [u.a.]: Cambridge Univ. Press.

Wolf, S. 2013. The Moral of Moral Luck, *Philosophic Exchange*, 31(1): 5 – 19.

Zimmerman, M. 1993. Luck and Moral Responsibility, in Statman, D. *Moral luck*. Albany: State University of New York Press. 217 – 234.