

Luck and Responsibility

by

Jennifer L. Woodrow

**Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy**

at

**Dalhousie University
Halifax, Nova Scotia
August 2007**

© Copyright by Jennifer L. Woodrow, 2007



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence

ISBN: 978-0-494-31507-1

Our file Notre référence

ISBN: 978-0-494-31507-1

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

DALHOUSIE UNIVERSITY

To comply with the Canadian Privacy Act the National Library of Canada has requested that the following pages be removed from this copy of the thesis:

Preliminary Pages

Examiners Signature Page (pii)

Dalhousie Library Copyright Agreement (piii)

Appendices

Copyright Releases (if applicable)

For my Family

Table of Contents

Abstract	x
Acknowledgments	xi
Chapter 1. Introduction	1
1.1. Moral and Epistemic Luck	1
1.2. Moral Luck: The Phenomena	3
1.3. What is Moral Luck?	5
1.4. Moral Luck: The Varieties	6
1.5. Epistemic Luck	10
1.6. Epistemic Luck: The Phenomena	11
1.7. What is Epistemic Luck?	21
1.8. Epistemic Luck: The Varieties	21
1.9. Moral and Epistemic Luck	27
Chapter 2. Standard Approaches to Moral Luck	29
2.1. The Skeptical Approach	30
2.2. The Internalist Approach	32
2.2.1. Internalism and the Epistemic Apology	34
2.3. The Externalist Approach	42
2.4. Summary of Problems with the Standard Approaches to Moral Luck	56
Chapter 3. Standard Approaches to Epistemic Luck	58

3.1. The Skeptical Approach	59
3.2. The Internalist Approach	62
3.3. The Externalist Approach	66
3.4. Summary of the Problems with the Standard Approaches to Epistemic Luck	81
 Chapter 4. Proposed Resolution to the Problems of Moral and Epistemic Luck	 83
4.1. The Hybrid Solution and Evolution	87
4.2. The Evolution of Normative Control	88
4.3. From Normative Control to Responsibility	90
4.4. Normative Control and the Objective Condition on Responsibility	100
4.5. Normative Control and the Subjective Condition on Responsibility	101
4.6. Responsibility and Luck	103
4.7. Hybrid Epistemology	104
4.8. Hybrid Morality	107
4.9. A Skeptical Solution or a Straight Solution?	109
 Chapter 5. The Two Necessary Facets of Responsibility	 113
5.1. Responsibility and Successful Intentional Acts	116
5.2. Responsibility and Unfulfilled Intentions	124
5.3. Responsibility and Accidents	127
5.3.1. What Justifies Holding Persons Responsible for their Unintentional Acts?	127
5.3.2. How do we Attribute Unintentional Acts to Agents?	134
5.3.3. Responsibility and Bare Bodily Movement	138
5.3.4. Human Causation and Inanimate Causation	144

5.4. Responsibility, Justification, and Truth-Value	148
5.4.1. Doxastic Responsibilism and Control	148
5.4.2. Indirect Doxastic Voluntarism and Epistemic Justification	148
5.5. Full-fledged Epistemic Responsibility	152
5.6. Epistemic Responsibility and Justification	157
5.7. Epistemic Responsibility and Truth-Value	158
 Chapter 6. Applying Hybrid Responsibility to Moral Luck	 160
6.1. Hybrid Responsibility and Resultant Moral Luck	161
6.1.1. Example 1: Good Resultant Moral Luck	162
6.1.2. Example 2: Bad Resultant Moral Luck	166
6.1.3. Analysis of Hybrid Resultant Moral Luck	166
6.2. Hybrid Responsibility and Circumstantial Moral Luck	172
6.2.1. Example 1: Bad Circumstantial Moral Luck	176
6.2.2. Example 2: Good Circumstantial Moral Luck	179
6.2.3. Analysis of Hybrid Circumstantial Moral Luck	180
6.3. Hybrid Responsibility and Constitutive Moral Luck	182
6.3.1. Is Constitutive Moral Luck a Real Problem?	183
6.3.2. Resultant and Circumstantial Moral Luck as Constitutive Moral Luck	187
6.3.3. Voluntariness and Origination	189
6.3.4. Example 1: Bad Constitutive Moral Luck	191
6.3.5. Example 2: Good Constitutive Moral Luck	194
6.3.6. Analysis of Hybrid Constitutive Moral Luck	195
6.3.7. Rejecting Voluntariness and Origination	195
6.3.8. Constitutive Responsibility and Agent Responsibility	197

6.3.9. Constitutive Responsibility and Societal Responsibility	200
6.4. Hybrid Responsibility and Causal Moral Luck	208
6.4.1. Causal Moral Luck and the Problem of Freewill and Determinism	209
6.4.2. Compatibilism and Causal Moral Luck	210
6.4.3. Incompatibilism and Causal Moral Luck	211
6.4.4. Freewill and Determinism is a Different Problem from Moral Luck	212
 Chapter 7. Conclusion	 216
7.1. Hybrid Responsibility and Resultant Epistemic Luck	216
7.1.1. Example 1: Bad Resultant Epistemic Luck	218
7.1.2. Example 2: Good Resultant Epistemic Luck	220
7.1.3. Analysis of Hybrid Resultant Epistemic Luck	223
7.2. Hybrid Responsibility and Circumstantial Epistemic Luck	223
7.2.1. Gettier Cases and Epistemic Luck	225
7.2.2. Example 1: Good Circumstantial Epistemic Luck	226
7.2.3. Example 2: Bad Circumstantial Epistemic Luck	230
7.2.4. Analysis of Hybrid Circumstantial Epistemic Luck	235
7.2.5. Reliabilist and Responsibilist Virtue Epistemologies	236
7.3. Hybrid Responsibility and Constitutive Epistemic Luck	237
7.3.1. Is Constitutive Epistemic Luck a Real Problem?	237
7.3.2. Example 1: Bad Constitutive Epistemic Luck	242
7.3.3. Example 2: Good Constitutive Epistemic Luck	244
7.3.4. Analysis of Hybrid Constitutive Epistemic Luck	248
7.3.5. Constitutive Epistemic Luck and Agent Responsibility	249
7.3.6. Taking Responsibility for Constitutive Epistemic Luck	253

7.3.7. Constitutive Epistemic Luck and Societal Responsibility	257
7.4. Hybrid Responsibility and Causal Epistemic Luck	263
7.4.1. Causal Epistemic Luck and the Problem of Free Belief and Determinism	264
7.4.2. Compatibilism and Causal Epistemic Luck	264
7.4.3. Incompatibilism and Causal Epistemic Luck	265
7.4.4. Free Belief and Determinism is a Different Problem from Epistemic Luck	267
End Notes	269
Reference List	277

Abstract

Moral and epistemic luck occur when an agent's responsibility is determined, at least in part, by luck. The problem of luck and responsibility is most perspicuously cast as a clash of two intuitions, both of which seem central to responsibility. On the one hand, persons' standings as responsible moral and epistemic agents are almost univocally held to be immune from determination by luck. On the other hand, we invariably count luck in assessing agents' moral and epistemic responsibility.

Luck both must not and does influence responsibility. Does this mean norms of responsibility are incoherent? Arguing from the perspective of participants in practices of attributing and undertaking responsibility, I maintain that neither intuition can be eliminated from our concept of responsibility.

Drawing from current evolutionary theory, I articulate a hybrid conception of responsibility that accommodates both intuitions. This hybrid view incorporates the objective demand that persons do and believe the right things and the subjective demand that persons do and believe the right things for the right reasons. Moral and epistemic luck are then diagnosed as a mismatch between the objective and the subjective dimensions of responsibility.

Acknowledgements

Richmond Campbell deserves more thanks than I can give. His enthusiasm for my project and his confidence in my abilities were there when mine failed. He provided intellectual and emotional support, and both were in high demand. Richmond deciphered my torturous drafts and heard my inarticulate ideas with patience and charity. Afterwards, he responded to both with incisive and insightful commentary. His generosity with his intellect, his time, and his friendship humbles me.

I am also deeply indebted to Sue Campbell and Mike Hymers. Their thoughtful advice, astute criticisms, and helpful suggestions, about this work and others, have been invaluable to me. Sue and Mike exemplify philosophical virtuosity, and I am ever so grateful for all their help. My external examiner, Bruce Hunter, was also tremendously helpful in pointing out areas that required improvement, and for identifying the substratum of numerous infelicities.

I owe thanks to every member of the incredibly supportive philosophical community at Dalhousie. Professors, students, and fellow graduate students each lent sympathetic ears, critical voices and keen minds to the refinement of my ideas. Though my gratitude is to all, I want to particularly thank Meredith Schwartz and James Anderson for their interest, input, and support.

I also thank my parents, Keith and Louise, and my siblings, Karen, Naius and Dylan, for their inveterate trust in my abilities and their assistance throughout the writing of this thesis. Finally, I thank my *aipak*, Nathaniel, upon whom this project and all others depend, at least in part.

Chapter 1. Introduction

“The best laid schemes o’ mice an’ men gang aft agley”—Robbie Burns.

1.1. Moral and Epistemic Luck

Good people usually do good things; bad people usually do bad things. Good intentions usually lead to good consequences; bad intentions usually lead to bad consequences. But the unusual sometimes happens. What do we do when luck interferes with our best-laid plans? What influence, if any, ought luck to have on our assessments of responsibility? Adequate answers to those questions must accommodate two apparently contradictory impulses. We feel that people should not be held accountable for accidents. The very notion of responsibility seems linked to purposive as opposed to accidental actions. We sometimes also feel that people should be held accountable for the things they cause, even if they do not cause them deliberately. The conflict between these impulses poses a threat to the theoretical coherence of our conception of responsibility, and it poses a practical challenge to our ability to respond consistently when the unusual happens. In this thesis, I shall outline the problems that luck poses for assessing moral and epistemic responsibility and offer a critique of the standard responses to those problems. I shall then tender an evolutionary explanation for the conflict of intuitions we experience in the face of the influence of luck on responsibility. That explanation will, in turn, ground a solution to the problems of moral and epistemic luck.

Where the good or bad consequences of an action are not foreseen, controlled, or intended by the agent in doing that action, those consequences are lucky or luckless as the case may be.¹ Purposive action seems to delimit the arena of human agency, and evaluating humans is a practice that seems only to legitimately take place within that arena. Yet we do evaluate people for the chance occurrences they instigate. We judge one

another and ourselves insensitive when we fail to accurately predict the hurtful effect of a blunt statement, or untrustworthy when we break a long-forgotten promise without intending to do so. Are these judgments out of bounds and, hence, illegitimate? To answer that we are only accountable for the reasonably foreseeable consequences of deliberate actions is to clear-cut the jungle of responsibility. To reply that we are accountable for unpredicted, unintended, and uncontrollable events is to afforest the whole world with that jungle. While human accountability sometimes seems to outstrip the arena of human agency, we should be loath to so broaden the arena that nothing is out of bounds. We require a reasoned way to negotiate the boundaries of responsibility in the face of luck. The problems of moral and epistemic luck present limiting cases where the conflict between jungle and clear-cut displays a paradox within the notion of reasonably apportioned responsibility.

To be lucky or unlucky, an unpredicted, uncontrolled or unintended event must also have some normative repercussions for the agent. When something beyond the agent's ken or control affects that agent positively, that agent is lucky; when something beyond the agent's ken or control affects that agent negatively, that agent is unlucky. Since human lives are affected daily by circumstances that agents cannot predict, cannot control and do not intend, it seems obvious that good luck and bad luck are ubiquitous. Few would deny that luck, as characterized, is rampant. People win lotteries; they are struck by lightning. Less dramatic cases of chance exert their influence over far more agents far more often. There is, however, an aspect of human lives that has conventionally been held immune from luck: our standings as responsible agents. The problems of moral luck and epistemic luck pivot upon the supposed immunity of agents'

standings from luck and their all-too-evident vulnerability to vicissitudes wrought by fortune.

1.2. Moral Luck: The Phenomena

Bernard Williams (1993) and Thomas Nagel (1993) first highlighted the phenomena now known as moral luck. Williams' central claim is that luck infects the life projects by which agents attain their moral standings. He illustrates his point by reference to the painter Gauguin who abandoned his wife and children to pursue his art unburdened by familial responsibilities. For Williams, whether Gauguin's action is justified depends upon whether Gauguin becomes a successful painter. "If he fails...then he did the wrong thing" (23, 1993). However, whether Gauguin becomes a successful painter depends on factors which are both unknown to him at the time of his decision, and beyond his control. If he fails to become a great painter, even for reasons that are beyond his control, then the life for which he left his family will have proven beyond his means to attain. Had Gauguin suffered a fatal naval accident en route to Tahiti, his extrinsic bad luck would have foreclosed the possibility of his decision being justified. The drowning Gauguin "does not and never will know whether he was wrong" (40). Bad luck that is intrinsic to Gauguin's project, such as his unfortunate lack of artistic talent, would render his decision unjustified. If, for reasons he cannot control, Gauguin is unable to muster painterly resources, his decision will have been proven wrong. From the standpoint of his failed life, Gauguin will be unable to justify having left his family, for the only thing that could have made his decision morally tenable is a successful life. Williams' analysis of Gauguin's moral luck alights upon a tragic feature of human experiences. The successful life project can remain unattained for reasons beyond the agent's ability to control, or even to apprehend, yet the agent's moral standing may depend on it.

Thomas Nagel also primes our intuitions about moral luck through examples. Of two negligent drivers, the lucky driver makes it home without incident while the unlucky driver hits and kills a child who runs onto the road. Of two attempted rescuers, the lucky rescuer succeeds while the unlucky rescuer drops the intended rescuee twelve storeys. By hypothesis, the behaviour that is under the control of the lucky and the unlucky individuals in each case is the same. Both drivers acted with disregard for the safety of pedestrians. Both would-be rescuers acted with supererogatory verve. According to the moral intuition that people are only responsible for the deliberate outcomes of their purposive actions, the unlucky driver is no more blameworthy than the lucky driver, and the unlucky would-be rescuer is no less praiseworthy than the lucky rescuer. In these cases, Nagel insists, we oppose the egalitarian verdicts enjoined by the intuition. Instead, we blame the lucky driver for driving negligently, but we blame the unlucky driver for driving negligently *and* for the child's death. We might praise the unsuccessful rescuer for trying to rescue someone, but we praise the successful rescuer for trying to rescue someone *and* for succeeding. The discrepancy between the assessments of these individuals is attributable only to the fact that the one was lucky and the other was not. The additional moral censure that is foisted upon unfortunates is responsive to their bad moral luck. The additional moral praise that is piled upon the fortunates is responsive to their good moral luck.

Nagel defines moral luck as follows:

Where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment, it can be called moral luck. (1979, 64)

Nagel argues that moral luck uncovers a paradox in morality: we must not, but inevitably do, hold people responsible for things that transcend their control. We do seem committed to excusing agents for what lies outside of their control. Why can't we just say that our inclinations to blame the unlucky driver more than the lucky one, and praise the lucky rescuer more than the unlucky one, are wrong? Nagel believes that to do so would jeopardize most of our moral judgments. Although we may be convinced that moral judgments which hinge on luck are irrational, we continue to make such judgments, and "our ordinary moral attitudes would be unrecognizable without [them]" (66).

Williams voices the same complaint. He says that "it would be a kind of insanity never to experience sentiments of this kind... and it would be an insane concept of rationality which insisted that a rational person never would" (Williams, 1981, 44). It seems, therefore, that even after reflection, the clash of intuitions remains. While we insist that luck ought not to make a moral difference, when we examine more closely the conditions under which we make moral assessments, we will find that luck makes a difference at every pass.

1.3. What is Moral Luck?

Moral luck exists if agents' moral standings are determined, at least in part, by luck. The problem of moral luck is most perspicuously cast as a clash of two intuitions, both of which seem central to morality. On the one hand, we believe that moral responsibility only attaches to purposive controlled actions and their predicted and intended consequences. On the other hand, we hold people responsible for their non-purposive actions, and for the unintended and unforeseen consequences of both purposive and non-purposive actions.

1.4. Moral Luck: The Varieties

Nagel outlines four ways that factors beyond an agent's control affect her moral standing, her moral responsibility, or both. Moral assessments of ourselves, and of others, are influenced by the following varieties of moral luck:

- 1) Resultant Luck, or "luck in the way one's actions and projects turn out."
- 2) Circumstantial Luck, or luck in "the kind of problems and situations one faces."
- 3) Constitutive Luck, or luck in "the kind of person [one is, as is evidenced by one's deliberate actions,] inclinations, capacities, and temperament."
- 4) Causal luck, or "luck in how one is determined by antecedent circumstances."

(Nagel, 60)

The cases so far considered are instances of resultant luck. Gauguin was lucky that his project of becoming a successful artist turned out well. The negligent driver and the would-be rescuer were unlucky that their actions turned out so awfully. Resultant moral luck turns on the idea that the actual outcomes of one's actions and projects are essential to our moral evaluation of those actions and of the agents who perform them, even when those actual outcomes are accidental.

The second kind of moral luck, circumstantial moral luck, concerns whether or not one actually performs a particular action. To exhibit circumstantial moral luck, Judith Jarvis Thomson contrasts Judge Actual with Judge Counterfactual. Both judges would accept a bribe if offered. Judge Actual is offered a bribe by an accused, and he accepts. As a result, the moral records of Judge Actual and Judge Counterfactual differ. Judge Actual is guilty of having taken a bribe whereas Judge Counterfactual is not. This difference in their moral standings is, by hypothesis, entirely attributable to something over which neither exercised control. It is only because of "the luck of the courthouse

draw” (Thomson, 1989, 207) that it was Judge Actual and not Judge Counterfactual who was offered and accepted the bribe. Nagel’s example takes two would-be Nazi collaborators, one who remained in Germany and actually became a Nazi collaborator, and another who “left Germany for business reasons in 1930” (59). The actual collaborator failed the moral test he was forced by circumstances to take, and his moral standing reflects that failure. Had he remained in Germany, the would-be collaborator would have behaved just as the actual collaborator did, and would warrant the same reproach. To what do we owe the large difference in the moral standings of the equally malicious actual and would-be collaborator? The former had circumstantial bad luck, while the latter had circumstantial good luck.

Constitutive moral luck is luck in the kind of person one is. “[P]eople are morally condemned for... [some] qualities, and esteemed for others equally beyond control of the will: they are assessed for what they are *like*” (Nagel, 65). By and large, the genetic, environmental, societal, and historical factors that make each of us the persons we are are not under our control. Although one’s character may also be the product of earlier choices, and future choices may cause changes to one’s character, we are formed, at least in part, by fortune. Moreover, one’s temperament and inclinations, hallmarks of one’s character, are not generally subject to the will. As Nagel puts it, for reasons he has neither willed nor caused, “a person may be greedy, envious, cowardly, cold, ungenerous, vain or conceited.... 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” (64). It is the Nazi collaborator’s bad constitutive luck that he is malicious, easily manipulated, and prone to corruption by propaganda. And it is his circumstantial bad luck that he was in Germany during the Third Reich.

In 1972, W.T. Greenough and F. R. Volkman demonstrated that “rats given a rich environment of toys and exercise gear and opportunities for vigorous exploration had measurably more neural connections, and larger brains, than rats raised in a bare, restrictive environment” (274). Though comparable studies on young humans are unfeasible, evidence from wide-ranging cases demonstrate that the type of person one is is grossly affected by factors in one’s early environment.² Take two agents whose characters were formed by a combination of factors, some in their control, some beyond their control. Suppose that their characters are identical in all respects that are under their control. The factors that are beyond their control have wrought a difference between the two agents’ characters. One is envious, while the other is not. (The biblical story of Cain and Abel provides a nice illustration of constitutive moral luck resulting in Cain’s envy.) If we assess these persons differently, as we surely do, we assess them differently for reasons that are beyond their control. This difference in our appraisals of the individuals seems perfectly just while we emphasize the difference in their characters. The difference between Cain and Abel is morally significant, and our assessments of the two individuals should reflect that difference. However, when we ascribe the difference in their characters to the whims of fortune (or God in the story) rather than to anything over which our agents wielded control, the difference in the agents’ moral standings seems undeserved. A similar verdict comes from looking at the prospects for becoming virtuous among the severely oppressed. It is not her fault that a woman who is told she is debased, shown that she is inferior and made to act so, becomes cowardly, unwholesome, or subservient. That she is cowardly, unwholesome or subservient is, nevertheless, something to which our moral judgments ought to be responsive. There seems to be no straightforward means

to reconcile these two judgments, and there seems to be no straightforward rationale for upholding one judgment and dispensing with the other.³

Causal luck is luck in one's antecedent conditions. Causal luck is the centrepiece of the problem of freewill and determinism. If one is caused to do something, then one cannot refrain from doing that thing. If one cannot help but do something, one cannot be held accountable for doing that thing (pace compatibilism). Nagel asks, "how can one be responsible for the stripped-down acts of the will itself, if *they* are the product of antecedent circumstances outside of the will's control?" (66). If we retain the control condition, it seems that one cannot.

Luck imperils our practices of undertaking and attributing responsibility. By undermining the coherence, the legitimacy and the practicability of those practices, moral luck threatens a central and essential aspect of human existence. Our self-worth, our most intimate relationships, as well as much less intimate ones, and our life projects all depend crucially upon practices of ascribing, attributing and undertaking responsibility, praising and blaming, reproaching and commending, and asking, receiving and granting forgiveness. The problem that luck poses to those practices is not merely metaphysical or theoretical. It is a real practical problem, and the harm it threatens is local, immediate and devastating.

Before turning, in Chapter 2, to proposed resolutions to the problem of luck and responsibility, consider the type and degree of threat that luck poses to our self-conceptions. Not only does luck threaten to undermine those of our practices surrounding moral responsibility, it poses an equal and equally devastating threat to practices surrounding our epistemic responsibility. Undertaking, ascribing and attributing responsibility for veridical claims, trusting and relying on others and on oneself,

distributing, hoarding and silencing discoveries, teaching and learning, are essential to humans' ability to navigate our physical and social and moral environments. Without those practices and the knowledge in which they issue, human lives and livelihoods would be impossible, or at least impossibly difficult. The threat that luck poses to the coherence of epistemic norms and the practicability of epistemic endeavours undermines the very presumption of human epistemic agency upon which our interdependent lives and livelihoods depend.

1.5. Epistemic Luck

The phenomena exhibiting moral luck turn on the fact that one can do something morally correct or incorrect by accident. The varieties of moral luck are distinguished from one another by the way factors beyond the agent's control contribute to morally significant results. Just as one can *do* right or wrong by accident, so one can *be* right or wrong by accident. Epistemic luck comes to the fore when we consider the respects in which knowledge can be the result of features over which the knower does not exert control. To call a piece of knowledge lucky or unlucky presupposes two uncontroversial claims. First, it presupposes that knowledge is something that it is good to have and bad to lack. Second, it presupposes that possessing knowledge is something over which knowers have some control. The key role of accurate prediction in human affairs inheres in our stake in the future. Without some sort of cognitive control over the future, we humans would be a sorry lot. The need for true beliefs naturally leads us to seek out others who are liable to tell us the truth. We are accountable to ourselves and to one another for what we purport to know. The dissemination of false information can carry grounds for reproach, just as the dissemination of true information can carry grounds for admiration. The question that epistemic luck poses is whether we can reasonably be praised for specifying future

occurrences without exerting cognitive control over such specification. That is, epistemic luck pits the actual consequences of inquiry against the demand that we hold agents responsible only for what they control. That control is exerted by the methods through which we acquire and justify information. However, whether an agent knows something or not depends on factors other than the methods of inquiring and justifying that can be presumed under her control. True beliefs can result from mere accident, lucky guesswork and pure chance. In "Prediction" Nicholas Rescher pits the merits of controlled inquiry, against those of lucky guesses. He states, "only *rational* prediction that is based on grounds whose merits are discernable prior to the event is of epistemological interest: predictions whose merits are discernable only after the fact are useless" (713).

Where knowledge is concerned, the agent is generally thought to wield control over the justification of her beliefs and over her methods of inquiry. It, therefore, seems natural to praise people for being diligent inquirers, and for making sure their beliefs are justified, and to blame them for being sloppy inquirers and holding whatever beliefs they choose, with disregard for their justificatory status. But we do not always attach our regard for epistemic agents to those things that are under their control. Lucky guessers are praised; unlucky guessers condemned. Moreover, the lucky guesser acquires the reputation of a good epistemic agent, while the unlucky guesser gains the reputation of a bad epistemic agent.

1.6. Epistemic Luck: The Phenomena

When good reasoning produces true beliefs there is little difficulty doling out epistemic praise, and when bad reasoning produces false beliefs there is little difficulty ascribing epistemic blame. Matters are not so straightforward when good reasoning generates false beliefs, or when bad reasoning results in true beliefs.

Take two persons, A and B, both of whom have justified their beliefs to the utmost of their abilities. A's belief is false, but B's belief is true. To employ a controversial example first, suppose A believes that the bottle from which he is drinking contains water. Suppose that A has reason to fear being poisoned, and has undertaken stringent measures to ensure that the bottle reading 'purified water', which he has purchased from a reputable vendor, has not been tampered with. He has a justified belief that the bottle from which he is drinking contains water. Now, suppose B also believes he is drinking water, and, also for fear of being poisoned, has undertaken stringent measures to ensure that his belief is justified. In fact, suppose that A and B are identical in every respect except the truth-value of their beliefs. A's belief is true. B's belief is false. The difference in truth-value between A and B's belief is entirely beyond either's ability to control. B has perhaps been unwittingly teleported to 'twin earth' and is drinking 'twin water.' What matters for the example is whether there is a difference between A and B that is relevant to their assessment as epistemic agents. A has knowledge while B does not. Does this difference matter for assessing A and B's responsibility as epistemic agents?

The answer to this question is controversial, and not just because the example alludes to semantic externalism. If we say that A is more praiseworthy than B, we make that judgement according to something that is beyond A and B's control. If we say that A and B are identical in terms of epistemic appraisal, then we make the judgement that knowledge is not more important, for epistemic responsibility, than is justified belief. Both answers are controversial. To treat A and B equally seems wrong to those wedded to the idea that truth matters for assessing epistemic responsibility. Those wedded to the idea that epistemic responsibility can be determined only by what is within agents' control will

insist that, though truth might matter for other reasons, it cannot matter for agents' responsibilities.

Take two agents, A and B, neither of whom justify their beliefs. Despite it having been drawn to their attention that they lack the expertise to judge accurately on the matter, both A and B persists in holding, without justification, the belief that the tree before her is a beech. A's belief is false. She was looking at an elm. B's belief is true. She was looking at a beech. We may suppose that A and B are identical in every respect, except the truth-value of their beliefs. What matters for the example is whether there is a difference between A and B that is relevant to their assessment as epistemic agents. Neither A nor B has knowledge. However, A has a true belief, while B has a false belief. Does this difference matter for assessing A and B's responsibility as epistemic agents? The positive answer that A is marginally more praiseworthy than B invites dismay among those who insist that the agents are the same in terms of all matters that they controlled. The negative answer that A and B should be judged same in terms of their epistemic responsibility invites consternation among those who insist that truth-value is responsibility-relevant independently of justification.

Resultant epistemic luck is at the root of the controversy. So far, however, the controversy seems rather trivial. Who cares whether A is right about her water? Are we really going to blame someone for mistaking an elm tree for a beech tree? The problem gains its pertinence alongside the importance of getting it right. Consider the controversy of intuitions in an academic setting. Student C submits a well-written environmental philosophy paper that presents a valid argument in favour of killing half (minus one) of the world's human population. Student D submits a well-written environmental philosophy paper that presents a valid argument inferring from the fact that it enjoins

killing half (minus one) of the world's human population that we should be chary of adopting straightforward utilitarianism in environmental contexts. Suppose that C and D warrant the same amount of credit in terms of justification (and originality, reasoning, diction, and so on). The only mark that differentiates C from D is the truth of their premises and conclusions. It seems that D warrants more credit for her claim than C does, simply because D's conclusion is true (or likely to be true) while C's conclusion is false (or unlikely to be true). At the same time, it seems as though C and D should be counted equally epistemically responsible. After all, both did what they could to ensure the veracity of their beliefs. The problem of according epistemic praise or blame seems even more important if we vary the example somewhat, and suppose that C and D are ethical consultants for the Department of Natural Resources. Of course, standards of justification will increase alongside increases in the importance of an epistemic claim. But if we maintain that both agents exercised the same standards of care in arriving at their conclusions, and tailor examples so that the only difference is the truth-value of the belief got thereby, the problem of resultant epistemic luck seems real, it seems important, and it seems to threaten norms of epistemic appraisal.

The same conclusions about the immediacy of the problem of resultant epistemic luck follow when we contrast unjustified true belief with unjustified false belief. A pair of cases from the history of science highlights the role that fortune plays in the acquisition of knowledge, and in the veneration of, or disdain for, epistemic agents. Both agents in these cases inquired with prejudice, failed to justify their beliefs, and falsified or misreported data to support their favoured theories. The agents were equally guilty of epistemic irresponsibility. As things turned out, the theory one supported was true, while the theory the other supported was false. Neither agent had knowledge. Both propounded

unjustified beliefs. Despite the fact that both agents are guilty of intellectual wrongdoing, the agent who supported the true theory is lauded as an intellectual hero, while the agent who supported the false theory is reviled as an intellectual villain.

One of the most infamous cases in the history of science is the Austrian biologist Paul Kammerer's alleged proof of Lamarck's contention that newly acquired characteristics can be inherited. Kammerer undertook to prove Lamarckism by experimenting with Midwife toads. As the Midwife toad breeds on land, it lacks nuptial pads, "rough patches on the hands with which the males of other species grasp the slippery back of the female while mating in water" (William Broad and Nicholas Wade, 1982, 183). Kammerer's midwife toads were mated in water for several generations. Kammerer reported that the toads' descendants were born with the nuptial pads characteristic of other species of toads, seeming to confirm Lamarckism. Unfortunately for Kammerer, insuperable doubt was cast upon the results he reported. In 1926, Kinsley Noble from the American Museum of Natural History examined the sole remaining descendant. Astonishingly, the toad did not have nuptial pads: "the black coloring that marked the pad was nothing but India ink" (184). The Midwife toad case marks one of the most deplored episodes of scientific fraud, and its architect, Paul Kammerer, took his life as a consequence of his public humiliation.

One of the most celebrated episodes in science is Sir Arthur Stanley Eddington's dramatic proof of Einstein's General Theory of Relativity. A consequence of the General Theory of Relativity is that light does not travel only in straight lines, but is bent by gravity. The theory of relativity predicts that light bent by the sun's gravitational field will demonstrate an angular displacement roughly twice that predicted by Newtonian physics. Convinced of the theory's soundness, Eddington took advantage of a solar

eclipse in 1919 to measure the influence of the sun's magnetic field upon the light from stars near the sun. Photographic plates were set up in Sobral, Brazil and in Principe, off the coast of West Africa. According to relativity, the sun's gravity should displace the light from nearby stars by 1.7 seconds of arc. According to Newtonian physics, the light from nearby stars should be displaced by 0.8 second (1/3600 of a degree) of arc. The results of the experiments were mixed at best.

On Collins and Pinch's analysis, "one of the Sobral instruments supported the Newtonian theory, while the other leaned towards Einstein's prediction for his own theory" (48). The two plates from the Principe expedition were the worst of all. Although they offered some confirmation of Einstein's prediction, Eddington calculated their results "using a complex technique that *assumed* a value for the gravitational effect" (48).⁴ In his official report of the experiments, Eddington flatly omitted the Brazilian results that contravened Einstein's prediction. He reported only the results from the sets of photographic plates that corroborated Einstein's prediction. Despite the ambiguous nature of the observational results, and despite the suspect manner in which those results were calculated and reported, only a few months after Eddington reported his data, the Astronomer Royal proclaimed that the observations had decisively confirmed Einstein's theory.

Both Kammerer and Eddington were unscrupulous epistemic agents. Both were convinced of the theory for which they sought confirmation, and that conviction led to their unprincipled epistemic activity. The fact that Einstein was correct does not excuse Eddington's methods of inquiry, just as, had Lamarck been correct, that fact would not excuse Kammerer. Why was Eddington revered and Kammerer reviled? Certainly political connections had something to do with it. However, no amount of political clout

would have elevated Eddington to the status he attained had he been wrong. The truth of Eddington's favoured theory and the falsity of Kammerer's may be construed, for present purposes, to have been the distinguishing factor.

The question is whether that distinction conforms to intuitions about epistemic responsibility. Do they warrant the same responsibility? When we emphasise the behaviour that was under their control, it seems they do. Eddington and Kammerer were both irresponsible in their methods of inquiry, were wilfully blind to countervailing evidence, and both propounded unjustified beliefs. The only difference between the agents is the truth-value of the theory each advocated. This is something neither agent controlled. According to the intuition that people are only responsible for the things over which they exert control, Kammerer's purposive misrepresentation of data in support of Lamarkism is no more blameworthy than Eddington's deliberate misrepresentation of data in support of the theory of relativity. If we limit our judgments of agents to factors they control, then we feel no inclination to judge Eddington and Kammerer differently. However, when we emphasise the truth of Eddington's conclusions and the falsity of Kammerer's, we seem to oppose the egalitarian verdict. Instead, though we will continue to blame Eddington for having failed to inquire judiciously and having failed to adequately justify his beliefs, we blame Kammerer for having failed to inquire judiciously and failed to justify his beliefs *and* for being wrong.

Those who insist that epistemic responsibility corresponds only to what agents control will argue that the reason for the difference in Eddington and Kammerer's epistemic standings is that we are inclined to infer that Eddington *knew* from the fact that his belief was true. Learning that Eddington was not justified overthrows the inclination to count him a better epistemic agent than Kammerer. This apology for the temptation to

count truth-value, which agents manifestly do not control, is tempting. It likely accounts for some of the diffidence we feel, from our historical standpoint, about praising Eddington at all after being duped into believing him to be an upright, even superlative, epistemic agent. However, it does not erase the temptation to count truth-values as important to epistemic responsibility. Consider the story of Jim Watson, Francis Crick along with the uncredited Rosalind Franklin's race against Linus Pauling or Maurice Wilkins to discover the structure of DNA.⁵ Suppose that, worried that they would be beaten by their competitors, Watson and Crick submitted their paper outlining the double helix as the structure of DNA to *Nature* before looking at Franklin's x-ray photographs of DNA which confirmed their theory. Watson and Crick, as they submitted their paper, would know that they were doing something epistemically blameworthy. They would also know that they would have done something ever so much worse if their proposal turned out to be false. Not only would they be more likely to be caught out, they would have disseminated false information. It seems, therefore, that the intuition to count truth-value, which agents do not control, in the determination of epistemic responsibility cannot be so easily dismissed.

Luck influences all three traditional criteria for knowledge. An agent may be lucky with regard to whether her belief is true or false. The way the world is is generally not up to the epistemic agent. Whether or not an individual holds a particular belief may be a matter of luck. Broad and Wade argue that the socialism prevalent in Austrian politics contributed to the appeal of Lamarckism during Kammerer's time, while "the Social Darwinists in England and the United States cited [Darwinism] as support for conservative, laissez-faire policies" (Broad and Wade, 1982, 181). Had Kammerer been reared in England or the United States rather than in Austria, he may not have believed

that newly acquired traits can be inherited. That Eddington believed Einstein's theory of relativity might, similarly, be attributable to factors over which Eddington wielded little control. If we assess epistemic agents by reference only to those beliefs whose origins they are responsible for, then it seems that we cannot praise Eddington for his beliefs and we cannot blame Kammerer for his. On this interpretation, Kammerer was unlucky to have believed in Lamarckism, while Eddington was lucky to have believed in Einstein's theory of relativity.

Finally, whether an agent holds a justified belief may outstrip her control. According to the view that one can only be held responsible for that which one controls, the locus of epistemic appraisal is justification. Justification is the chief means agents employ to ensure that they have good rather than bad beliefs. The responsible epistemic agent tailors her acceptance and rejection of beliefs according to their justificational standing. She rejects, rescinds, or otherwise fails to endorse those propositions for which she has no grounds or for which she has countervailing evidence, and accepts, admits or otherwise endorses those beliefs for which she has reasons.⁶ When luck interferes with an agent's ability to parse her beliefs into good and bad by reference to the reasons for or against them, the result is that the agent, through no fault of her own, cannot be fully responsible for her beliefs. In "Logical Luck" Roy Sorensen notes how semantic externalism can contribute to agents' lacking control over the validity of their inferences. He states "If 'meaning ain't in the head', then no internal check will ensure that I am using terms uniformly. I might try to apply *modus ponens* and fail" (320). Circumstances beyond an agent's control can prevent her from using valid inference forms in her justifications. A less contentious example might be the experience of women who suffered sexual advances in the work place prior to the articulation of the concept of

sexual harassment.⁷ The self-loathing, guilt and shame induced by sexual harassment was a wrong that defied propositional articulation. Women so suffering were precluded, try though they might, from justifying their responses to the wrongs they underwent. An agent may fail to be justified in her belief, despite having striven to hold only justified beliefs, or an agent may be justified in her belief, despite having exercised no virtuous epistemic aplomb. And yet, we hold individuals responsible for their justifications even when luck has played a part.

Following Nagel's definition of moral luck, we can define epistemic luck as follows:

Where a significant aspect of what someone [believes] depends on factors beyond his control, yet we continue to treat him in that respect as an object of [epistemic] judgment, it can be called [epistemic] luck. (1979, 58)

Epistemic luck seems problematic for just the same reasons that moral luck is problematic. We care deeply about having knowledge, and we care deeply about whether epistemic agents' reputations are warranted. By showing that luck makes an epistemic difference, epistemic luck uncovers a paradox in our thinking about knowledge: whether or not an agent has knowledge both must not be, and inevitably is, a matter of luck.

Although we may be convinced that it is irrational to praise people for their accidental knowledge and to blame people for their unavoidable lack of knowledge, our actual practices of assessing epistemic agents seems inextricably tied to factors over which they do not exert control. Whether their beliefs are true or false matters independently of the procedures (or lack thereof) they undertook to ensure the correctness of their beliefs.

1.7. What is Epistemic Luck?

Epistemic luck exists if agents' epistemic standings are determined, at least in part, by luck. The problem of epistemic luck can be cast as a clash between two intuitions that are central to epistemology. On the one hand, we believe that epistemic responsibility only attaches to the active measures that agents can undertake to ensure that their beliefs are correct. This intuition is emphasized by Scott Sturgeon and Dretske and Enç, among legion others. The former claims "Luck does not contribute to our possession of knowledge" (Sturgeon, 1993, 161). The latter assert, "[A] belief, to qualify as knowledge, must have no admixture of accidentality in its correspondence with the facts" (Dretske and Enç, 1984, 517). On the other hand, we hold people responsible for having beliefs, even when those beliefs are caused by factors that they could not avoid. We hold people responsible for the truth or falsity of their beliefs, even when that truth or falsity is not up to them. And we hold people responsible for justifying their beliefs, even when the means for such justification lie beyond the purview of their ken or control. When it comes to our appraisals of epistemic agents, the best that they can do is often not good enough. We hold them to a standard of responsibility that includes things that they cannot control.

1.8. Epistemic Luck: The Varieties

Daniel Statman outlines the parallel between moral and epistemic luck as follows:

[W]ith regard to any act *a* of an agent *P*, *P* might be lucky in two fundamental ways, one concerning the agent, and the other concerning the act:

(1m) Luck in the causes or circumstances that bring *P* to do *a* or to be such a person who wants to do *a*.

(2m) Luck in the destiny, so to speak, of the act; whether it turns out to be morally right or wrong.

Corresponding to situational and resultant luck in morality, we find the following two sorts of luck in epistemology:

(1e) Luck in the causes or circumstances that bring some subject *S* to believe *p*, or to be such a person who believes *p*.

(2e) Luck with regard to *p* being true or false. (1991, 148)

The same distinctions Nagel notices among moral luck phenomena can be employed to parse epistemic luck phenomena.

- 1) Resultant Luck is luck in whether or not one's belief turns out to be correct.
- 2) Circumstantial Luck is luck in the circumstances in which one collects evidence and forms beliefs
- 3) Constitutive Luck is luck in the kind of epistemic agent one is.
- 4) Causal Luck is luck in how one's beliefs are determined by antecedent circumstances.

The cases of the student-cum-ethical consultants, Eddington and Kammerer, and the hypothetical Watson and Crick case exemplify resultant epistemic luck. Each features an agent who was lucky that his or her belief turned out to be true in contrast with her counterpart, whose belief turned out to be false. Resultant epistemic luck turns on the idea that the actual outcomes of one's inquiry are essential to our epistemic evaluation of that inquiry, and of the agents who so inquire, even when those actual outcomes are accidental.

The second kind of epistemic luck, circumstantial epistemic luck, concerns the degree to which fortune or misfortune influences the formulation and justification of particular beliefs. The different temporal and evidentiary situations in which epistemic agents find themselves and the different epistemic problems with which they are

confronted with obvious influence over the particular beliefs agents hold, and the manner by which they come to hold them. To take an uncontroversial example of factors beyond one's control exerting influence over the truth or falsity of one's beliefs, someone born in Mesopotamia in 600 BC would likely hold the false belief that the Earth is flat, while someone born in Mesopotamia in 2004 AD would insist that the Earth is, in fact, roughly spherical. The one has a false belief, while the other has a true belief. When one is born is obviously not something within the agent's control. To the extent that we blame people for their false beliefs and praise them for their true beliefs, we blame and praise them, at least in part, for things that they could not control.

Circumstantial epistemic luck includes luck in the particular situations and problems that one must face as well as luck in one's circumstances, more broadly construed. As Daniel Statman notes, it is one's bad circumstantial epistemic luck to be "in circumstances that are misleading, e.g. presenting optical illusions" (149). It can also be one's good epistemic luck to be in a favourable evidentiary position. The fable in which an apple falling from the tree under which he sat spurs Sir Isaac Newton to formulate his conception of gravity as a mechanically explicable force (rather than an occult force or an essential property of bodies) places Newton in a providential epistemic position.⁸ One can imagine Newton's luckless counterpart who, suffering from an allergy to pollen, remained indoors. Never having been in the auspicious position of having an apple fall on his head, Newton's counterpart never contributes to the advancement of experimental physics. We must surely insist upon a distinction in the epistemic standings of Newton and his hapless counterpart. However, insofar as we do so, the different epistemic standings that we attribute to them are grounded in circumstances over which

the agents wielded no control. It follows that our appraisal of epistemic agents is based, at least in part, on circumstantial epistemic luck.

Any treatment of epistemic luck would be remiss to exclude the most famous examples of epistemic luck. Before turning to constitutive and causal luck, let us briefly visit Gettier cases. In 1963, Edmund Gettier noticed that luck affects whether or not persons' justified true beliefs count as knowledge and, hence, affects persons' epistemic standings. Gettier's examples provide compelling reason to reject the standard definition of knowledge as justified true belief; a belief might be both justified and true yet fail to qualify as knowledge. In Gettier cases, luck plays a central role in determining whether the epistemic agent in question has met the criteria for knowledge.

Take Smith and Jones, both of whom have applied for a job. Gettier asks us to suppose that Smith has strong evidence for the proposition: "(d) Jones is the man who will get the job and Jones has ten coins in his pocket" (142). Smith legitimately infers from this that "(e) the man who will get the job has ten coins in his pocket" (142). Thus, Smith believes (e) and he is justified in believing (e). Suppose further that (e) is true. However, unbeknownst to Smith, Smith also has ten coins in his pocket, and it is Smith, rather than Jones, who will get the job. Smith's belief in (e) is justified, and it is true, and yet it seems one should balk at saying that Smith knows that (e).

Gettier examples show that even when agents have fulfilled their epistemic responsibility to justify their beliefs, they do not exercise control over whether they have knowledge. Poor Smith does not know because his belief and the justification thereof were *lucky*. To see that this is so, contrast Gettier's example with the following case. Trent and King have both applied for a position. Trent has strong evidence for the proposition: "(d) King is the man who will get the job and King has ten coins in his

pocket" (142). Trent legitimately infers from this that "(e) the man who will get the job has ten coins in his pocket" (142). Thus, Trent believes (e) and he is justified in believing (e). Suppose further that (e) is true. The difference between this case and Gettier's is that King *is* the man who will get the job. Trent's belief in (e) is justified, and it is true, and Trent knows that (e). The explanation for the fact that Trent has knowledge whereas Smith does not is that it is only by wild blind luck that it so happens that Smith's justified belief is true, and his justification is adequate. By contrast, no intermediary of wildly good epistemic luck was required for Trent's belief to be both justified and true. Given his evidence, it is perfectly natural that Trent's belief that the man who will get the job has ten coins in his pocket is true. Given his evidence, it is very bizarre that Smith's belief that the man who will get the job has ten coins in his pocket is true.

Linda Zagzebski (1994 and 1996) has diagnosed Gettier examples as involving double luck. There is the bad luck that would ordinarily preclude the agent's justified belief from being true. In addition, there is the good luck that ensures that the agent's belief is true, despite the foregoing bad luck. On my view, Gettier cases involve the operation of two (or more) types of epistemic luck operating at cross-purposes. Smith, for example, profits from good resultant luck that counterbalances the bad circumstantial luck that would otherwise sever the truth connection between his justification and the truth of his belief.

Constitutive moral luck is "luck in the kind of person one is, where this is not a question of what one deliberately does, but of one's inclinations, capacities and temperament" (Nagel, 60). One's inclinations, capacities and temperament also wield palpable influence on the type of epistemic agent one is. Whether one is a precipitous reasoner, quick to jump to unfounded conclusions or a meticulous thinker who is loath to

accept a proposition unless it has been amply substantiated by all available evidence is, to a large extent, a matter of forces over which one does not exert control. If one is the former type of reasoner, one suffers from bad constitutive epistemic luck. If one is the latter type of reasoner, one profits from good epistemic luck. The availability of education and the freedom from oppression as well as the endowments of natural talents and aptitudes are essential to the formulation of good epistemic agents. However, as Claudia Card makes clear, these epistemic goods are apportioned by natural as well as unnatural lottery.

To see the influence of constitutive luck on the beliefs one holds, contrast Kammerer with a hypothetical biologist who would have espoused Lamarckism just as unjustifiably as Kammerer did, had his epistemic character been informed by an extremely oppressive intellectual climate. This agent is, instead, reared under a liberal intellectual regime in which diligence in justifying beliefs is paramount among epistemic virtues. This agent is lucky in the intellectual climate in which his epistemic character is formed. The difference between this agent and Kammerer is attributable only to the fact that the one was lucky enough to have had his methods of belief acquisition formed by a liberal rather than an oppressive social, political and intellectual climate. However, when and where one is born, and under which social-political and intellectual climates one is reared, is hardly up to the agent. What type of reasoner one is and, hence, the care or neglect one shows in generating and justifying beliefs is caused at least in part by one's intellectual milieu. Insofar as the type of intellectual agent one is is a matter over which agents do not exert control, credit or blame for the type of intellectual agent one is constitutes credit or blame for what one does not control. If we restrict our appraisal of

agents to those things that they control, then we must accord Kammerer and his counterpart equal epistemic standings.

Causal epistemic luck is luck in the causal antecedents of one's beliefs. A theory of causal epistemic luck would enumerate the causes that determine one's beliefs. In its extreme version, a theory of causal epistemic luck "might be reductionist, and argue that beliefs are reducible to some kind of bodily activity, which is wholly determined by physical and chemical laws" (Statman, 149). Meting out praise and blame would seem irrational on such a conception of epistemic luck. If one is caused to believe something, then one cannot refrain from believing that thing. If one cannot help believing that thing, one cannot be held responsible for believing that thing. If we uphold the control condition, it seems that we cannot praise or blame agents for any of their beliefs, be they attained by judicious inquiry or by flippant guesswork. According to Nagel, when we offer causal explanations for the actions of agents, we lose our grip on their agency. Notice, however, that the straightforward inference from the fact that agents are caused to do, or believe, something, to the claim that they cannot be responsible for those actions or beliefs, presupposes incompatibilism.

1.9. Moral and Epistemic Luck

Moral and epistemic luck are often intertwined, and the problems they present for the coherence of responsibility individually are compounded and confounded by their conjunction. One's bad moral luck may be the consequence of one's bad epistemic luck. One's bad epistemic luck may be the consequence of one's bad moral luck. For example, one may fail, despite every attempt at justification, to know that one's well-intentioned act will have a harmful effect. One may accidentally bring about a harmful result that engenders an unjustified but true belief. Unpropitious evidentiary circumstances may

lead to unpropitious moral circumstances, and vice versa. One's epistemic faults, though neither willed nor controlled, may lead one to moral faults, equally beyond the purview of one's will. One's moral failings, though not of one's making, may result in epistemic failings. For reasons of perspicacity, I will not address the many permutations and combinations of moral and epistemic luck. Real world examples are rarely so straightforward and uncomplicated.⁹ Nevertheless, the general claims made throughout about the types of moral and epistemic luck will apply to permutation and combination thereof.¹⁰

Each of the four variants of causal and epistemic luck threatens to destroy our deepest and most cherished conceptions of ourselves as moral and epistemic agents. Let us now turn to proposed resolutions to the threat that luck poses for moral and epistemic luck.

Chapter 2. Standard Approaches to Moral Luck

If two cases are indistinguishable from the subject's standpoint but differ in external result because of chance factors beyond the control of the agent(s) (where the result involves harm to one or more others), this will constitute bad moral luck iff

i) The moral statuses of the cases differ

and

ii) One agent does better or worse morally in virtue of the chance factors.

The standard approaches can be parsed into rough categories by reference to their stance on the following inconsistent triad:

1. Our stable moral intuitions about moral praiseworthiness and moral blameworthiness arise from norms that are coherent, rational, and fair.
2. According to our moral norms, persons can never be morally praiseworthy or morally blameworthy for things that are beyond their control.
3. According to our moral norms, persons can sometimes be morally praiseworthy or morally blameworthy for things that are beyond their control.

The positions in the literature divide roughly into the following categories:

A. The Skeptical Approach upholds 2 and 3, but denies 1: One can only be responsible for things over which one exerts control. One may be held responsible for things over which one does not exert control. Therefore, moral norms are incoherent.

B. The internalist approach accepts 1 and 2, but denies 3: Moral norms are coherent. One can only be responsible for things over which one exerts control. Therefore, it is not the case that one may be held responsible for things over which one does not exert control.

C. The Externalist approach accepts 1 and 3, but denies 2: Moral norms are coherent.

One may be held responsible for things over which one does not exert control. Therefore,

it is not the case that one can only be held responsible for things over which one exerts control.

2.1. The Skeptical Approach

The skeptical approach despairs of any solution to the quandary luck poses for moral norms. Skeptics believe both that morality, as we know it, cannot be free from luck, and that morality must be free from luck. They maintain that the existence of moral luck exposes a paradox in traditional moral reasoning: moral norms insist that we must and must not count luck in assigning moral responsibility to agents. It is, accordingly, impossible to coherently ascribe moral responsibility. Both Williams and Nagel, among others, take this view.

Williams insists that moral luck is real and that it undermines the universal scope and overarching value of morality. In so doing, moral luck threatens to substantially and irrevocably alter our conception of morality. The view that morality fairly allots agents' moral standings according to their responsible agency is fundamental to our conception of morality. Moral luck, however, erodes our concept of responsible agency. Resultant, circumstantial, constitutive, and causal luck show that "one's history as an agent is a web in which anything that is the product of the will is surrounded and held up and partly formed by things that are not" (44). The idea that morality is fair is intimately linked with the idea that moral desert is meted out according to what agents have done, rather than according to what has happened to agents. It seems unfair to be held responsible for what is beyond our control. But given moral luck, this seems inevitable. Williams claims that the problem of moral luck shows either "that responsible agency is a fairly superficial concept, ...or else that it is not a superficial concept, but that it cannot ultimately be purified" (45). If responsible agency is a superficial concept, then morality, which trades

in this superficial concept, must be conceded superficial. If morality insists upon pure responsible agency, and such agency does not exist, then it must be conceded that morality has no legitimate realm of application. On either option, we must repeal something central to our conception of morality. Williams concludes as follows:

Scepticism about the freedom of morality from luck...will leave us with *a* concept of morality, but one less important, certainly, than ours is usually taken to be; and that will not be ours, since one thing that is particularly important about ours is how important it is taken to be. (Williams, 54)

Nagel's conclusion is similarly dismal. On his view, the conflict between the control condition and assessments of agents and actions according to luck illuminates a paradox implicit in moral reasoning. "A person can be morally responsible only for what he does; but what he does results from a great deal that he does not do; therefore he is not morally responsible for what he is and is not responsible for. (This is not a contradiction, but it is a paradox.)" (66). Moral luck, for Nagel, challenges the possibility of meaningful ethical discourse.

Neither Williams nor Nagel believes there to be any tenable solution to the problem of moral luck. We must accept that morality is internally conflicted. Our moral norms say that we are not to hold persons accountable for what lies outside their control and that we are to hold persons accountable for what lies outside their control. Since both moral norms are ineluctable, the idea that morality is rational must give way. Not only must we concede that morality is irrational in the face of moral luck, we must also concede that morality is fundamentally unfair. If morality is susceptible to luck, then at least some persons will have moral standings that they do not deserve, and the option of being moral will not be open to everyone, or at least not open to everyone equally.

2.2. The Internalist Approach

Internalism is the second approach to the problem of moral luck. Internalism about moral responsibility limits responsibility to those things internal to the agent over which she can reasonably be assumed to exert control. Perhaps the paradigmatic internalist was Immanuel Kant. For Kant, luck has no legitimate role to play in the ascription of moral worth to actions or agents. There is simply no such thing as moral luck. No matter how tempting, it is a mistake to praise or blame someone for something that is beyond her control. The intuitions that drive us to do so must be jettisoned so that our practices of assessing agents are brought into accord with the dictates of morality. Internalism, following Internalist approaches to morality, restricts agents' moral worth to an area over which agents have control: their will. There is no such thing as *moral* luck because all that matters for morality is a good will.

The good will is not good because of what it effects or accomplishes or because of its adequacy to achieve some proposed end: it is good only because of its willing, i.e., it is good of itself. And, regarded for itself, it is to be esteemed incomparably higher than anything which could be brought about by it in favor of any inclination or even the sum total of all inclinations. Even if it should happen that, by a particularly unfortunate fate or by the niggardly provision of a stepmotherly nature, this will should be wholly lacking in power to accomplish its purpose, and if even the greatest effort should not avail it to achieve anything of its end, and if there remained only the good will (not as mere wish but as the summoning of all the means in our power), it would sparkle like a jewel in its own right, as something that had its full worth in itself. Usefulness or fruitlessness can neither diminish nor augment its worth. (1979, 54)

In this passage, Kant categorically denies moral luck. (To be specific, he denies resultant moral luck.) One may accept Kant's repudiation of resultant luck without conceding that morality is immune from luck. It might be the case that all that matters for moral worth is having a good will. However, whether one is presented with opportunities within which to formulate a good will may be a matter of circumstantial moral luck. Moreover, whether or not one has a good will may be a matter of constitutive moral luck. The Internalist denial of resultant moral luck, if successful, still leaves the agent's moral worth at the mercy of fortune in the circumstances she faces, in the character she has, and in the causal antecedents of her will.

It is a bit misleading to present Kant's conception of morality as a response to the problem of moral luck. In fact, Kant offers one of the two intuitions whose conflict is at issue in the problem of moral luck: the intuition that one can only be held morally accountable for what lies within the purview of one's will. Kant denounces those intuitions that endorse moral luck. To invoke the internalist conception of moral agency is, thus, simply to uphold the intuition that morality must be free from luck, and to deny the validity of those intuitions which tell us that luck does, sometimes, make a moral difference. To accept the internalist's approach is to dismiss the problem outright. However, this dismissal is too quick. We do hold persons responsible for the actual consequences of their wills, and we often think that we ought to. Even if we think that we ought not to do so, since we do not seem to be able to stop doing so, the internalist approach, by itself, does not resolve the conflict among our moral intuitions. To constitute a solution, the internalist's intuition that morality cannot countenance the influence of luck must be married to an explanation of why we believe that there is such a thing as moral luck, and a protocol for how to get rid of such beliefs. This is a project to

which many recent authors have devoted considerable effort. The project is to deny that apparent instances of moral luck are genuinely moral instances of luck.

2.2.1. Internalism and the Epistemic Apology

The most common approach to the problem of moral luck conjoins the internalist intuition that luck cannot make a moral difference to the contention that our temptation to count luck as significant has its roots in our epistemic fallibility. If the only difference between two individuals (and/or their actions) is attributable to luck, then there is no moral difference between the two individuals (and/or their actions). If we were perfect epistemic agents with perfect epistemic access, we would never be tempted to reproach or applaud persons for things that they do not control. The driving insight for this approach is that we are never subject to moral luck, although epistemic fallibility sometimes makes us think we are.

Brian Rosebury, Norvin Richards, Nicholas Rescher, and Michael J. Zimmerman each elaborate versions of the epistemic approach.¹¹ Brian Rosebury argues that apparent cases of moral luck can be adequately explained by human epistemic fallibility.

Rosebury argues that the phenomenon of moral luck does not demonstrate that moral responsibility is sometimes keyed to outcomes determined by fortune. Instead, he insists that the problem of moral luck can be dissolved by an account of moral responsibility, which acknowledges the intellectual finitude with which agents address moral quandaries.

According to Rosebury, when an agent's actions have terrible unforeseen consequences, his reflective awareness of his epistemic fallibility will cause him to believe that "he may have been negligent in some respect not yet clear to him" (1999, 515). Moreover, because other agents can only problematically distinguish between unforeseeable consequences and consequences which could have been foreseen, had

appropriate epistemic care been deployed, unlucky moral agents are judged responsible for the outcomes of their actions. If, however, the consequences of the acts in question are genuinely not foreseeable by the agent, then both self-judgments and other-judgments are in error. Rosebury maintains that the retrospective moral judgment attendant to apparent resultant moral luck is misplaced. He suggests that we guard against unjust judgments of moral responsibility by restricting such judgments to considerations available to agents at the time of their deliberations. So long as agents exercise due epistemic care, they can see that they are not morally responsible for unforeseeable outcomes of their actions and, therefore, ought not to be judged morally responsible for those outcomes.

Rosebury recasts the problem of moral luck in terms of the epistemic problem of “distinguishing between decision making vitiated by avoidably and unavoidably insufficient knowledge” (500). He, thus, suggests a means to re-establish consistency between the intuition that luck cannot play a determinative role in assessments of moral responsibility and the fact that we do praise, blame, and judge agents according to factors which are beyond their control. When we do the latter, we commit the error of viewing moral agents as infallible epistemic agents.

Norvin Richards articulates the relationship between epistemic fallibility and the erroneous supposition that luck makes a moral difference as follows:

A culprit may...be lucky or unlucky in how clear his deserts are. And we must allow his luck to make a difference in how we treat him, if we are not to change our practice into one in which we pretend omniscience. But in doing so we do not contradict the contention that his luck cannot affect what his deserts *are*. So, we

do not act paradoxically, as Nagel contends, but only reflect our epistemic shortcomings, and the agent's good or bad fortune in those. (169)

Rescher states the epistemic thesis as follows:

We have no choice—given the difficulties of epistemic access—to forming our judgments of persons on the basis of what we observe them to say and do. But all that is mere evidence when the moral appraisal of human agency is at issue. The “morally lucky” villain is not *morally* lucky (by hypothesis, he is a villain); he is lucky only in that his reprehensible nature is not *disclosed*. The difference is not moral but epistemic; it is a matter of not being found out, of “getting away with it.” It is precisely because both one's *opportunities* for morally relevant action and ...the *actual consequences* of one's acts lie beyond one's own control that they are not determinants of one's position in the eyes of morality. (156)

Richards and Rescher discount Nagel and Williams' claim that luck can influence one's moral standing. Luck plays a role in our moral appraisals of agents, but when it does so, the luck is epistemic and not moral. The actual consequences of persons' wills, and the actual opportunities that arise to display their wills, are important for moral evaluation of agents only insofar as those consequences and circumstances are genuinely indicative of agents' characters. The actions agents perform and the circumstances that arise for such performances are, according to the epistemic approach, the evidentiary basis for moral judgments. But such evidence can be good or bad evidence.

In the ordinary course of things, good consequences stem from a good will, and bad consequences stem from a bad will. Hence, in the ordinary course of events, we can infer an agent's good will from the good consequences of her acts, and we can infer an agent's bad will from the bad consequences of her acts. When there is a mismatch

between an agent's will and the outcomes of her actions, our ordinary inferences are illegitimate. Rescher states, "the moral bearing of luck enters in because its operation can deflect the evidential linkage by disrupting the normal, standard import of actions" (157). The normal standard import of actions is their connection with the actor's character. When luck intervenes between an agent's intentions and the outcome of her actions, we cannot judge her on the basis of her actions. In so-called cases of resultant moral luck, we must ignore the consequences of an agent's actions, for they are irrelevant to the agent's moral standing. We must, therefore, treat the negligent driver who has run over a child and the negligent driver who has not killed anyone as moral equals; if one deserves moral condemnation, then so does the other. We must also regard the would-be rescuer who drops his intended rescuee twelve stories just as we do the successful rescuer. If the latter warrants our praise, then the former does as well.

Plausible though it is, the epistemic approach is bound to meet with steadfast resistance at this point. Even when we have perfect epistemic access, as in the hypothetical cases of the negligent drivers and the would-be rescuers, we are loath to treat the two individuals as moral equals. Rescher tries to soften this resistance by invoking non-moral (and non-epistemic) grounds for differentiating between good and bad consequences. We may, he says, distinguish successful good intentions from unsuccessful good intentions "on *nonmoral* grounds; for example to reward only *successful* rescues or to punish only *realized* transgressions as a matter of social policy *pour encourager les autres* (159).... But from the specifically moral point of view, it is *desert* that is determinative and fate is ultimately—and mercifully—irrelevant" (164).

The claim that human epistemic fallibility is intimately related to instances of resultant moral luck is, indeed, tempting. Proponents of the epistemic approach believe

that once we admit the fallibility of human predictive capacities, we can realign our moral intuitions to accord with the claim that morality must be immune from luck. The epistemic approach purports to solve the internal conflict which moral reasoning faces by dismissing a large facet of distinctively moral feelings. This approach does not deny that we *feel* that agents who inadvertently cause terrible things ought to be held morally responsible, even when oneself is the agent in question. In cases where an agent, through no fault of her own, causes an unfortunate occurrence, the agent may continue to feel what Williams calls agent regret. However, proponents of the epistemic approach insist that “the concept of moral responsibility loses its coherence if this residual unhappiness is listed among ‘moral’ sentiments” (Rosebury, 517). The approach, therefore, dismisses feelings of agent-regret and moral liability, as well as the moral judgments that they give rise to, as non-moral, or as not representing the world accurately. It is, rather, supposed that they are “psychological facts...[that] merely show that our emotional life is less analytically structured and differentiated than our rational judgments” (Rosebury, 514). The problem with this view is that the distinctively moral tenor of the emotional feelings and judgments that the epistemic approach dubs illogical is not so easily dismissed.

Michael Zimmerman and Judith Jarvis Thomson also advocate the epistemic approach. However, they are dissatisfied with the standard attempts to explain why, even when we possess all the relevant information about two agents, we regard luck as morally significant. In order to account for why we insist on blaming the unlucky reckless driver more than the lucky reckless driver, Zimmerman distinguishes between “‘P is to blame for more events than P*’ and ‘P is more to blame than P*’” (227). He maintains that an agent may be to blame for more events than another without being more to blame than the other agent. Zimmerman identifies the unlucky negligent driver who accidentally kills a

child as the successful driver. The successful negligent driver “is to blame for more events than the ...[unsuccessful] one—more events serve to indicate that he is to be evaluated negatively on this occasion—but this does not imply that he is to be evaluated negatively to a greater extent than is the... [unsuccessful] driver” (227). Insofar as what happens after one has made a free decision is not up to us, then these events, while perhaps serving as indirect indicators of praise and blame, are strictly dispensable in the assessment of moral responsibility.

On Zimmerman’s conception of the epistemic approach, the problem of moral luck necessitates that we tease out the tangle of moral responsibility with a finer toothed comb. Judith Jarvis Thomson also recognizes that a more fine-grained analysis is required to address the problem of moral luck. She believes that the illusion of moral luck can be dispelled by differentiating between two different conceptions of blame. She contends that there is no such thing as genuinely moral luck, and attributes our tendencies to believe that there is such a thing as moral luck to epistemic luck. She explains the influence of epistemic luck on our moral sentiments by outlining two conceptions of blame. Her distinction will both help to clarify the epistemic approach and better enable an understanding of its problems.

Thomson insists that luck cannot make a moral difference. Although she concedes that we do judge the unlucky driver more harshly than the equally careless lucky driver, she claims that “when we take seriously that mere luck *is* what makes the difference, then we do not feel this difference in degree of disapprobation” (213). On Thomson’s view, we can better understand practices of blaming agents for things for which they are not to blame if we distinguish two types of blame. She defines the first type of blame, blame-1, as follows: “A person P is to blame for doing such and such, where P’s doing such and

such is a wrongful act or omission for which P has no adequate excuse” (201). Blame-2 states that “A person P is to a greater or lesser to blame for doing (or being) such and such where his doing (or being) the such and such is unwelcome, just in case P’s doing (or being) the such and such is stronger or weaker reason to think P is a bad person” (202). Although we operate with both conceptions of blame, Thomson believes that the two types of blame are connected. A person fails to have an adequate excuse for being a certain way or doing a certain thing just in case her being that way or doing that thing is a reason for thinking her a worse person. She outlines the connection between blame-1 and blame-2 as follows: “A person P is to blame for doing the such and such —first notion of blame—if and only if P is in some measure to blame for doing the such and such under the second notion of blame” (211). On Thomson’s view, a person is to blame for a wrongful act for which she has no excuse if and only if her performing that wrongful act is a stronger or weaker reason to think she is a bad person.

Like Rosebury, Rescher, and Norvin, Thomson diagnoses our habit of blaming persons for things that they do not control as resulting from our epistemic fallibility. Like Zimmerman, she recognizes that there is a conception of blame that attaches to the consequences of our actions (be they foreseen or unforeseen). She views the type of blame that attaches to such consequences as only legitimate if the consequence is genuinely indicative of a bad character.

In restricting the scope of moral accountability to outcomes that an agent can reasonably be expected to foresee, the epistemic approach leaves no room for rational assessment of an agent as morally responsible for things which she has done through no fault of her own. In the following passage, Williams illustrates the point. He maintains that it is self-deceptive to

say that your feelings about something done involuntarily, or as the lesser of two evils, are to be understood as regret, a nonmoral feeling. This turns our attention away from an important dimension of ethical experience, which lies in the distinction simply between what one has done and what one has not done. That can be as important as the distinction between the voluntary and the non-voluntary. (1985, 177)

It is this second notion of ethical experience which makes reasonable the demands that morality places upon agents to make amends, where possible, to those harmed by the unhappy consequences of their actions. By insisting that self-reproach, guilt, and remorse are nonmoral, the epistemic approach misclassifies restitution and reparation as nonmoral activities. This approach also makes unintelligible the moral reproach that rightfully attaches to those agents who shirk their responsibilities. Why? Because a moral scheme that refuses to admit that responsibilities transcend control cannot make sense of the approbation that attaches to those agents who fulfill their duty to make amends where possible. This suggests that it would be *immoral* to do as Kant and advocates of the epistemic approach suggest and simply view feelings of remorse and their third personal corollaries, praise and blame, as self-deceptive or misplaced. For those who deny the reality of moral luck, the responsibilities that unlucky moral agents incur cannot be moral responsibilities. Yet we find the agent who feels no compunction to fulfill such responsibilities not merely immoral, but morally alien. The problem of moral luck calls for an earthly approach to moral agency.

2.3. The Externalist Approach

Externalists about moral responsibility believe that factors external to the agent do affect the agent's responsibility. Externalists insist that doing the right thing matters for moral responsibility. Margaret Urban Walker and Claudia Card are champions of and harbingers for moral externalism. In "Moral Luck and Impure Agency," Walker analyzes the problem of moral luck from the earthly perspective of a participant in moral practices. Walker assesses the incompatibility of moral luck with the control condition, and maintains that one must eschew the control condition as an inappropriate and unrealistic constraint on morality. She argues that accepting the "regimentation of responsibility embodied in the control condition represents an alteration of our common life far more drastic than may first be supposed" (17). Moral luck is, according to Walker, something necessarily involved in human morality, and, hence, central to any adequate understanding of morality. Accordingly, Walker claims that "moral luck is real and not paradoxical, even if wishfully distorted or simply inadequate views of the human condition make it appear so" (15).

Walker's criticism is clear: the idea that we can hold agents responsible only for what lies within the purview of their will is both conceptually problematic and morally inadequate. Walker contrasts the Kantian idea of a pure agent with Williams's idea of impure luck-ridden agents. A pure agent is self-constituted and liable only for pure acts of the will. Pure agency is "neither diluted by nor implicated in the vagaries of causality at all, or at least not causality external to the agent's will, itself understood as a causal power" (244). By contrast, the impure agent is at least partially constituted by factors over which she has no control, and may be liable for things that outrun her control. Following Williams, Walker defines impure agency as "agency situated within the causal order in

such ways as to be variably conditioned by and conditioning parts of that order” (243). Such agents can be morally responsible for things they did not intend and for outcomes over which they have no control.

Walker believes that the idea of pure agency is conceptually problematic because the agent’s will is, itself, the product of innumerable factors over which the agent possesses no control. Moreover, the help or hindrance of causal factors external to the agent are essential to the exercise of the agent’s will. Her criticism echoes Nagel’s remark that if we take an agent and subtract everything that is not within that agent’s control we are left with almost nothing. This conceptual problem is subsidiary for Walker. Walker believes that the intuition that one can be responsible only for what one controls conflicts with our moral experience in ways that provide strong reasons for rejecting that intuition. She argues that the idea of pure agency correspondent to the control condition is morally inadequate because 1) it cannot accommodate our moral experience and 2) we would not want to share a world with pure agents. Such pure agents would be immoral (or at least amoral) by our lights.

For Walker, moral luck and its correspondent impure conception of agency are deeply entrenched in our conception of morality and our moral practices. In support of this claim, Walker argues that at least some distinctively moral responses are responses to moral luck. We assess agents as worthy of moral blame or praise, at least in part, by reference to their forbearance against bad moral luck or their gratitude for good moral luck. Recognizing one’s impurity as a moral agent constitutes a particularly moral understanding of oneself and of those impure agents with whom one shares a moral sphere of action. Walker outlines three virtues whose existence depends upon our

vulnerability to luck. These virtues, integrity, grace, and lucidity, enable appropriate moral responses to circumstantial and resultant luck.

The first virtue of impure agency is integrity. Whole persons who are not willing to compromise their moral commitments in the face of trying circumstances have integrity. Integrity is the capacity to “keep or make whole what might otherwise suffer deformity, collapse, or desertion: a coherent and responsible moral posture” (242). It is precisely when persons’ actions have unforeseen morally significant consequences or when circumstances display persons’ willingness to make good on their moral commitments that they have the opportunity to display their integrity. The second virtue of impure agency, grace, is appropriate where integrity, though necessary, seems insufficient in the face of moral luck. The unlucky negligent driver who has accidentally killed a child seems an appropriate case. It is incumbent upon the driver who has accidentally killed a child to retain his moral centre. Where no reparative work seems possible, the driver may display grace by his “non-aggrandized daily ‘living with’ [the accident] unsupported by fantasies of overcoming or restitution” (242). Lucidity is a “reasonable grasp of the nature and seriousness of one’s morally unlucky plight and a cogent and sensitive estimate of repairs and self-correction in point” (243). These virtues, Walker insists, are genuinely moral virtues whose very existence depends upon our vulnerability to luck. These virtues are “impossible to capture fully without reference to the vicissitudes of moral luck” (241-2). If integrity, grace and lucidity are moral virtues, which they seem to be, then we must reject the control condition and its attendant pure conception of agency because they preclude the possibility of these actual virtues.

What is most notable about Walker’s analysis is her suggestion that a pure agent, an agent whose moral responsibility extends only to those results over which she has

control, would be immoral when embedded in a society of agents whose actions are subject to whims of fortune. She states, “pure agents may not be depended upon, much less required, to assume a share of the ongoing and massive human work of caring, healing, restoring, and cleaning-up on which each separate life and the collective one depend” (25). By contrast, impure agents, possessed of the appropriate virtues, are those agents upon whom the very possibility of human inter-dependability rests. She states, “The truth of moral luck that the rational, responsive moral agent is expected to grasp is that *responsibilities outrun control*” (241). The fact that we are individually subject to luck, and the fact that we often depend upon others for our ability to respond to circumstances in moral ways, makes us morally responsible for acts which are beyond our power to avoid, and results which are beyond our power to either predict or control. These responsibilities stem not from the autonomy of moral agents, but from their interdependence. She concludes, “In a world where we need so much from each other so often, acceptance of our impurity is not the worst we can do” (26).

Walker rejects the strong intuitions that underlie the view that moral agents can be judged only according to their purposive actions.¹² Walker undermines the intuition that one can only be held responsible for the intended and controlled consequences of one’s purposive actions by showing its unpalatable consequences. If we accept the control condition, then we cannot accommodate such virtues as integrity, grace and lucidity; and if we accept the control condition, then we must exonerate agents for insensitivity and undependability which seem immoral. These consequences of the control condition are unacceptable. Hence, the control condition from which they flow ought to be rejected. Walker argues that moral luck is not at odds with our theoretical conceptions of moral value and responsibility; it is essential to them. If our theoretical pictures of moral value

and responsibility do not accommodate moral luck, then the “acceptability of those pictures needs defending” (Walker, 1993, 238).

Externalism about moral responsibility, laudable though it is, goes too far in rejecting the control condition outright. Though Walker is correct to charge that internalism excludes too much from the arena of moral responsibility, her own view includes too much and is insufficiently sensitive to real variations in the responsibilities that accrue to agents. Under the internalist assumption that we ought to hold individuals responsible only for the intended results of their purposive actions, some individuals will be held more responsible for their worse actions and their worse intentions than others. If we get rid of the control condition on moral responsibility, as Walker suggests, it seems that we lose our grip on the notion of desert. Without the control condition, we are left with no principled manner to distinguish accidents from purposive actions.

Walker’s vision of a moral order without the control condition enjoins us to play with the cards we are dealt. Take the negligent driver who is so unfortunate as to accidentally kill a child. His insistence that it was accidental will not exonerate him. The child who is born to abusive parents, and whose prospects for moral growth are stunted has no recourse. If he ends up a suspicious, hardened, violent man, incapable of caring relationships, it is just his foul moral luck. This seems manifestly unfair, yet Walker’s externalism does not explain this unfairness. The control condition enables us to make sense of holding the unfortunate driver less responsible than someone who has run over a child deliberately, and holding the survivor of abuse less accountable for his bad character than someone whose bad character can be traced to his deliberate choices. Because it renounces the control condition, externalism is left without principled means to account for the differences between deliberate controlled wrongdoing and accidental wrongdoing.

Walker would likely respond that the lack of principled means does not imply a lack of any means for differentiating moral responsibility for purposive harms from moral responsibility for accidental harms. On her view, there is an inexact match between control and accountability. Persons can be accountable for things they do not control. Walker's view, thus, does not preclude the possibility of a relationship between control and accountability. She, therefore, may be open to a conception of responsibility that sees a central role for control in differentiating between the kinds of responsibility that purposive harms and benefits invite and the kinds of responsibility that non-purposive harms and benefits invite.

She would, therefore, likely insist that judgments about responsibility depend on the particularities of each case. The type of harm and its severity, along with the relation between the harmer and the harmed, will matter for the type of and degree of responsibility the harmer bears. Surely, however, whether or not the harm was done purposively will play a central role in determining "evaluations of a full repertoire of perceptions, judgments, expectations, responses, attitudes and demands with respect to ourselves and others in the matter of our conduct, its meaning, and its impact" (Walker, 1991, 238).

On the broader focus that Walker enjoins, people can clearly be seen to be responsible for things that they do not control. In particular, they can take responsibility for X, and that does not require having controlled X. However, from this broad focus, the clear difference between the responsibilities that accrue to someone who has done something by accident and those that accrue to someone who has done something on purpose are blurred, and, in consequence, cannot adequately be accounted for.

Once we narrow the focus to the praise and blame, or credit and discredit, and accountability senses of responsibility the fact that there is a broader conception of responsibility that does not invoke the control condition no longer seems to do away with the problem that moral luck. In narrowing the focus of responsibility undertakings and attributions to discredit, blame and accountability considerations, I do not intend to marginalize the care-taking conceptions of responsibility. The credit, praise, and accountability conception of responsibility does not appear to be prior to, or more important than, care-taking conceptions of responsibility. To the contrary, I see both care-taking and credit/discredit senses of responsibility as functioning together to maintain cohesive and collaborative groups. The emphasis on the backwards-looking, credit/discredit, sense of responsibility is an artefact of the problem that luck presents for that conception of responsibility. Moreover, the problem that luck presents for the credit/discredit, praise/blame, sense of responsibility does not appear to be solved by moving to the care-taking sense of responsibility.

What, then, would Walker say about the difference between controlled and uncontrolled consequences appears to make for attributions and undertakings of credit/discredit responsibility? Walker is right to be chary of views and positions that purport to stem from 'our' intuitions or 'the' moral agent.¹³ Such views tend to occlude significant differences between moral agents and their intuitions. They also tend to occlude the social identity of the person propounding the view. Nevertheless, it is incumbent on any view that purports to have some purchase on practices of attributing and undertaking responsibility to address one of the central means employed to determine responsibility. Practices of fairly blaming and praising persons are surely essential to our conception of morality. Insofar as the control condition is central to those practices, a

theoretical picture of moral value and responsibility that eschews the control condition is inadequate. Thus, while Walker has shown the depth of the intuition that luck makes a moral difference, she has not demonstrated that our moral practices can be made sense of without the control condition. The role of luck in our morality remains deeply problematic.

Like Walker, Claudia Card is leery of the Internalist assumption that persons can be held responsible only for what they control. On that view, “the same basic character development is accessible to all” (1996, 4). However, Card insists that this idea is contravened by empirical evidence. On her view, “being inside socially constructed ethnic and color categories becomes part of our [constitutive and circumstantial] moral luck” (9). The different starting points of individuals do seem to affect their abilities to become good as well as their opportunities to perform acts of good will. Card’s focus is on the types of moral luck that infect women and other oppressed groups, and her goal is to outline the implications for moral reasoning that can be drawn from oppressed persons’ unlucky starting points.

Card takes constitutive moral luck as an empirical given. What Card finds problematic is the existence of socially constructed asymmetries in the luck with which different groups must deal. If we neglect to weigh the constraints that the oppressed must overcome in their pursuit of virtues, then the oppressed will receive univocally worse moral score sheets than those lucky enough to live without such constraints. By contrast, if we give credence to the idea that luck can and does make a moral difference to the characters that agents are able to attain, we are in a better position to correct for luck in our assessments of individuals. She states, “Luck is involved both in the motivation to take responsibility and in our ability to carry through. Where that seems unfair, we may

be able to take the unfairness into account, morally, in some of our evaluations” (1996, 27). By noticing that luck makes a moral difference, we can work to overcome the differences that luck makes. We will not be able to overcome all such differences. To believe that we could do so would be to commit the error of viewing agents as immune from luck, which they manifestly are not. However, there appear to be some differences in the luck to which persons are subject that can be lessened, if not erased.

The upshot of Card’s thesis is that oppressed individuals face larger burdens of luck than do non-oppressed individuals, and that this difference in their vulnerability to moral luck is grossly unjust. Although one cannot efface the influence of luck on agents’ moral standings, one can work to erode the type of bad moral luck that attends being born to a society that constrains oppressed persons’ acquisition of virtuous characters and their ability to perform virtuous actions.

Card clearly views moral luck as a real phenomenon. Things over which they have no control affect persons’ moral standings. However, she also seems committed to the view that luck ought not to make a difference to persons’ moral standings. The main thrust of her argument for recognizing moral luck seems to be that our recognition is invaluable to redressing the differences in persons’ moral standings that are due to moral luck. Card’s motivation for erasing the differences that luck makes to persons’ moral standings is that such differences are unfair. She, therefore, seems to rely on the Internalist intuition that luck ought not to make a moral difference after all. Are Card’s commitments coherent? How can she reconcile the view that we ought, where we can, subtract the influence of luck from evaluations of agents with a conception of agency wherein luck is inextricable? The answer to this question lies in her conception of moral responsibility.

Like Walker, Card argues “appreciating the impact of luck on our lives can add depth to our understanding of responsibility and increase our sense of morality’s importance” (1996, 21). The problem of moral luck concentrates on what Card calls the “backward-looking, or credit,” (1996, 22) conception of responsibility. She distinguishes this backward-looking conception of responsibility from forward-looking conceptions of responsibility. According to the former, responsibility is a predicate of agents. An agent is or is not responsible for her character, or her circumstances, or her will, or her acts and the consequences thereof. According to the forward-looking conception of responsibility, responsibility is a virtue to which agents should aspire.

If we look at the things agents can be responsible for through the prism of variegated impure agency, we can see that responsibility outstrips the control condition. Card outlines four dimensions of responsibility:

(1) the administrative or managerial sense of responsibility—undertaking to size up and organize possibilities comprehensively, deciding which should be realized and how; (2) the accountability sense of responsibility—agreeing to answer or account for something, or finding that one should be answerable, and then doing so; (3) the care-taking sense of responsibility—committing oneself to stand behind something, to back it, support it, make it good (or make good on one’s failure to do so), and following through; (4) the credit sense of responsibility—owning up to having been the (morally) relevant cause of something’s happening or not happening, taking the credit (or blame) for it. (28)

The second and last of these dimensions of moral responsibility have been the focus of the problem of moral luck. According to the control condition, one can only be held accountable for something if one can be legitimately credited or discredited for that thing.

In rejecting the control condition, Card and Walker reject the view that one can only be responsible for things over which one deliberately exercises causal control. They insist that different dimensions of responsibility enable one to be responsible for things that one did not control. One need not be in control of the consequences of one's act to take responsibility for those consequences. (Although it is noteworthy that one must be in causal control of whether one takes responsibility for the outcomes of one's acts in order to be praised or blamed for doing so.)

Card's contention that there are different dimensions of moral responsibility seems eminently plausible. There is, however, an ambiguity in the position that Card and Walker advocate which has not yet been resolved. They both insist that one need not be in control of something to be responsible for that thing. But in light of Card's edifying distinction among different senses of responsibility, the question of which types of responsibility transcend control naturally arises. Do Card and Walker believe that the credit and discredit sense of responsibility can be legitimately ascribed in the absence of the control condition, or do they maintain that, in addition to the credit and discredit conception of moral responsibility for which the control condition is essential, there are other conceptions of moral responsibility whose legitimate imputation does not presume the control condition? Another option is that neither the forward-looking nor backward-looking senses of responsibility require the control condition. On any of the three interpretations, it becomes clear that Walker and Card's dismissal of the control condition does not satisfactorily resolve the problem of moral luck.

If we take Walker and Card's rejection of the control condition at face value, we must count the blameworthiness of the unlucky negligent driver unmitigated by his unluckiness. According to this intuition, the "logic of blame sometimes permits censure

of the helpless" (Sorensen, 334). The fact that killing a child was something that he did not intend or control will, on this interpretation, not mitigate his blameworthiness. While this seems unjust, Walker and Card may insist that such injustice is a small price to pay for a moral order that is responsive to our vulnerability to luck. We can dispel the illusion that we are free from moral luck, but the price of doing so is that we must live with the realization that morality is fundamentally unfair.

Externalists might respond that construing morality as responsive to luck only seems unfair because we have not relinquished the intuition that one can only be praised or blamed for things that one controls. As soon as one rejects that intuition, one no longer has grounds for complaining that a lucky moral order is an unfair moral order. Although this argument has some merit, it does not aptly characterize Walker and Card's approaches. Walker and Card do wish to mitigate the unfairness of moral luck. Walker's virtues of impure agency and Card's forward-looking conceptions of responsibility seem designed specifically to reintroduce fairness. One cannot have integrity without being in control of one's moral self. Being in control of one's moral self is just what it means to be a person of integrity. Similarly, one cannot *take* responsibility for something accidentally. One must be in control to take responsibility for something.

In response to this, externalists would likely claim that, just as the control condition is not necessary for imputations of responsibility in cases of resultant moral luck, so the control condition is not necessary for imputations of responsibility where agents display, or fail to display, integrity. Similarly, one can be responsible for taking responsibility even where one does not control whether or not one takes responsibility. That is, externalists would seek to reject the control condition at every level where

responsibility is to be assessed. There is, however, room in their views for a role for control.

Despite the interpretive havoc that it wreaks, I believe the most plausible and charitable interpretation of Walker and Card's position sees both the backward and forward looking conceptions of responsibility as rejecting the control condition, but nevertheless leaving room for a central role for control. On this interpretation, the unlucky negligent driver and the lucky negligent driver would be equally blameworthy because their controlled contributions were equal. But in addition to being blameworthy for having driven negligently, Card believes that it is incumbent upon the unlucky negligent driver to *take* responsibility for having killed a child. The unlucky negligent driver can then be assessed as praiseworthy or blameworthy by reference to whether he takes responsibility for having killed a child. And in addition to being blameworthy for having driven negligently, Walker believes that it is incumbent upon the unlucky negligent driver to maintain his integrity and accept his inability to make reparations with grace. However, the unlucky negligent driver can only legitimately be morally appraised for taking responsibility for the unforeseen and uncontrolled consequence of his actions if taking responsibility for those consequences is something over which he has control. Moreover, the unlucky negligent driver can only legitimately be praised or blamed for maintaining his integrity and gracefully accepting the terrible thing he has done if his integrity and grace are things he controls. On this interpretation, Walker and Card have not dismissed the control condition. They have displaced it.

We may accept Walker and Card's dismissal of the control condition at face value. In that case, we must interpret them as doing irrevocable damage to our theoretical pictures of moral value and responsibility. This damage seems at least as unpalatable as

the damage Walker argues is a consequence of upholding the control condition. On this interpretation, Walker and Card have failed to show that moral luck is not problematic. We may also view Walker and Card as embracing the control condition in the form of taking responsibility and displaying integrity. Insofar as they attach moral value to, and apportion moral praise and blame for, taking responsibility and displaying integrity, the problem of moral luck remains.

Card is surely correct to note that there are different dimensions of responsibility. The question that presents itself, however, is whether or not each conception of responsibility that Card outlines bears the same problematic relationship with the control condition as does the credit-blame sense of responsibility. There is an ambiguity in Card's address of moral luck. While she agrees with Walker that responsibility outruns control, it is unclear whether she believes that one can always be praised or blamed for something that one does not control. Card is surely correct to claim that one can take responsibility for an event one has not controlled. In that sense, responsibility does outrun control. However, it seems as though one cannot be praised or blamed for taking responsibility for an event if one is not in control of whether or not one takes responsibility for that event.

While Card has illustrated that our conceptions of responsibility are not limited to the credit-blame conception of responsibility, it seems as though valuing agents for whether or not they are responsible in the broader senses that Card outlines itself depends fundamentally upon the control condition. If the latter is the case, then Card's investigation of moral luck has shed considerable light on our conceptions of moral responsibility. She has not presented a resolution to the problem of moral luck. Card might wish to deny that this is the case. She might want to uphold the view that one can,

as a matter of fact, always be praised or blamed for things that outrun one's control.

However, if she does this, then she, like Walker, has dismissed an apparently fundamental aspect of our moral systems.

2.4. Summary of Problems with the Standard Approaches to Moral Luck

1. The sceptical view undermines the coherence of our moral norms and leaves the paradox unresolved. It is, therefore, a position of last resort.
2. The Internalist view irrevocably damages our actual practices of moral assessments. It does not adequately deal with our intuition to count cases of moral luck. The Internalist approach simply denies the force of the consequentialist intuition that we must count the results of actions. Insofar as the intuition remains even after reflection, to deny the intuition does not solve the paradox. Moreover, as Margaret Urban Walker argues, the change effectuated by restricting moral ascriptions to the sphere of controlled action, even if possible, would not be desirable for moral reasons. We would be less moral were we luck-free pure Kantian agents. The Internalist approach also fails to adequately address the regress problem. One's intentions, as well as one's character, may be possessed because of luck. As Martha Nussbaum notes, "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" (Nussbaum, 96).
3. The epistemic response misclassifies genuinely moral judgments as judgments of mere epistemic malfunction. Moreover, (as I shall establish below) there is a parallel problem with epistemic luck. Therefore, diagnosing the problem of luck as an epistemic problem simply pushes the requirement that we deal with luck back a step.¹⁴
4. The externalist approach enjoined by Walker and Card denies the force of the deontological intuition that we must count only the intentions with which actions are

performed. Insofar as that intuition remains after reflection, simply denying it fails to solve the paradox exemplified by cases of moral luck. The externalist approach is also open to the regress problem. Advocates of this approach do not have a reasoned means for stopping the influence of luck at any particular point. Moreover, in denying the control condition, the externalist approach leaves us without means to distinguish paradigmatic instances of praiseworthy and blameworthy behaviour from paradigmatic instances of accidents.

Card's identification of forward-looking, senses of responsibility as well as backward-looking senses of responsibility, is intriguing and seems correct. Nevertheless, the problem of moral luck is centrally concerned with backwards looking senses of responsibility. According to the consequentialist intuition, an agent's failure to *take* responsibility for a morally salient consequence she has wrought (even if accidentally) is something for which that agent is (backward-looking) responsible.

None of the standard approaches to the problem of moral luck is adequate.

Chapter 3. Standard Approaches to Epistemic Luck

If two cases are indistinguishable from the subjective standpoint but differ in external result because of chance factors beyond the control of the agent(s), this will constitute epistemic luck iff

i) The epistemic statuses of the cases differ

and

ii) One agent does worse epistemically in virtue of the chance factors.

Standard approaches to the problem of epistemic luck can be parsed into rough categories by reference to their stance on the following inconsistent triad:

1. Our stable epistemic intuitions about epistemic praiseworthiness and epistemic blameworthiness arise from norms that are coherent, rational, and fair.
2. According to our epistemic norms, persons can *never* be epistemically praiseworthy or epistemically blameworthy for things that are beyond their control.
3. According to our epistemic norms, persons can *sometimes* be epistemically praiseworthy or morally blameworthy for things that are beyond their control.

The standard approaches to the problem of epistemic luck are as follows:

- A. The skeptical approach upholds 2 and 3, but denies 1: Skeptics believe that luck must not but does make a moral difference, and are accordingly skeptical about the coherence, rationality, and fairness of epistemic norms.
- B. The Internalist Approach upholds 1 and 2, but denies 3: Internalists believe in the rationality, coherence and fairness of epistemic norms, and are skeptics about epistemic luck.
- C. The Externalist approach upholds 1 and 3, but denies 2: Externalists believe both in the rationality, coherence and fairness of epistemic norms and in epistemic luck.

3.1. The Skeptical Approach

There is something dreadfully unsettling about the idea that knowledge can be lucky and the closely related idea that people's epistemic standings are not deserved seems unfair. Despite this, we seem to attribute knowledge to individuals whose beliefs are true by luck, and we seem to assess people's epistemic agency, at least in part, according to things that they do not control. The skeptical approach despairs of any solution to the quandary luck poses for epistemic norms. Skeptics believe both that epistemic inquiry, as we now know it, cannot be free from luck, and that epistemic inquiry must be free from luck. They maintain that the existence of epistemic luck exposes a paradox in our concept of knowledge: epistemic norms insist that we must not count luck in assigning epistemic responsibility to agents, and yet we sometimes do count luck in assigning epistemic responsibility to agents. It is, accordingly, impossible to coherently ascribe epistemic responsibility.

Skepticism about the compatibility of epistemic norms and practices insists that epistemic luck is real and that it undermines the value of knowledge. Epistemic luck, thus, threatens to substantially and irrevocably alter our conception of knowledge. Part of what we mean when we attribute knowledge to people is that they are justified in believing as they do. One must be responsible for one's belief if one is to have knowledge, and one exercises that responsibility by ensuring that one's beliefs are justified. However, we almost invariably hold people responsible for the falsity of their justified beliefs, or for the truth of unjustified beliefs. While we insist that people are only accountable for what they control, we praise and blame them for things that they do not control. Moreover, such practices of praising and blaming seem ineluctable. Because our actual practices of assessing agents diverge so radically from the norms of epistemic

appraisal, the skeptic despairs of coherently accounting for knowledge. Our epistemic norms say that we are not to hold people accountable for what lies outside their control and that we are to hold people accountable for what lies outside their control.

The skeptical implications of epistemic luck are most salient when we consider resultant epistemic luck. We allow resultant epistemic luck to influence our assessments of agents when we praise or blame them for the fortuitous truth-value of their beliefs. If S has a justified belief that p, and S* has a justified belief that q, and yet, despite the fact that both S and S* have done everything in their power to ensure that their beliefs are true, q turns out to be false while p is true, then S has knowledge and S* does not. S is praised for the uncontrolled consequence of believing as she does; that is, she is praised for having a true belief. Meanwhile, S* is blamed for the uncontrolled consequence of believing as she does; that is, she is blamed for having a false belief. The difference between justified belief and knowledge is a difference that outstrips the agent's control. If we hold agents accountable only for what they control, then we should hold the justified belief to be just as praiseworthy as knowledge, and the justified believer to be just as praiseworthy as the knower. But we do not. In our appraisal of epistemic agents, the truth of a justified belief is a significant—indeed almost ineluctable—vector of appraisal. We praise people for knowledge, not just for justified belief. And we blame people for having false beliefs, not just for having unjustified beliefs.

The traditional conception of knowledge includes the requirement that an individual be justified in her belief. However, epistemic luck phenomena demonstrate that we credit people for false beliefs and praise them for true beliefs, whether they are justified or not. This conflict among norms of epistemic appraisal is problematic. If we allow resultant luck to play a role in our epistemic appraisals, then knowledge becomes

nothing more than true belief since justification is unnecessary. If, however, we refuse to sanction the influence of luck in our epistemic appraisals, knowledge seems to become nothing more than justified belief since truth is unnecessary. On either option, we must concede something that is central to our conception of knowledge.

The consequence of allowing luck to enter our appraisal of agents is dire: our traditional conception of knowledge seems unattainable. The conflict among our intuitions about knowledge comes to the fore once we realise that possessing knowledge is not something over which agents wield control. How can we reconcile the contention that knowledge must be something agents control, as evidenced by our commitment to justification as a necessary criterion for knowledge, and our praise of knowledge over and above mere justified belief? The skeptical approach licenses no such reconciliation. Moreover, advocates of this approach believe that we cannot relinquish either intuition without doing irreparable damage to our conception of knowledge.

However, it is not only resultant epistemic luck that jeopardizes our conception of knowledge. Circumstantial, constitutive and causal epistemic luck also undermine our attributions of knowledge to agents, and our conception of knowledge. If one takes seriously the many ways that luck influences one's epistemic agency, the very notion of epistemic agency loses its coherence.

The skeptical implications of epistemic luck follow from what Mylan Engel Jr. calls "the incompatibility thesis" (1992, 62). The incompatibility thesis simply contends that luck and knowledge are incompatible. According to that thesis, if a belief is somehow coincidentally or fortuitously true, it does not count as knowledge; and if a belief is to count as knowledge, it cannot coincidentally or fortuitously true. Given the undeniable role played by factors beyond the agent's control in the acquisition of beliefs,

the contention that lucky beliefs cannot be knowledge challenges the very possibility of knowledge, as it is normally construed. If, however, we allow lucky beliefs to count as knowledge, then attributing knowledge to someone ceases to have its standard normative force. Justification, which is presumed to be under the agent's control, marks the distinction between mere true belief and knowledge. Why not simply dismiss those intuitions that press us to count luck in our assessment of knowledge and knowers? This is precisely the strategy enjoined by the second approach to the problem of epistemic luck.

3.2. The Internalist Approach

Epistemic internalists insist that agents are responsible only for their justifications. A belief is justified just in case there is some combination of internal states—typically other beliefs and experiences—that provide support for that belief. According to epistemic internalists, an agent's true belief only counts as knowledge if it is justified, and that justification is internally accessible to the agent.¹⁵ As Hilary Kornblith describes it, internalism is responsive to “the laudable goal...to form beliefs in a way that manifests integrity...to form beliefs which fully satisfy [one's] own subjective standards” (2001, 6). The internalist's response to epistemic luck limits the realm of epistemic appraisal to an arena the agent controls: whether or not her belief meets her internal standards of justification. According to the internalist, if an agent has done all she can do to ensure that her beliefs are true, then she does not deserve reproach if her beliefs turn out to be false. Since all that agents can be expected to control is whether or not their beliefs are internally justified, an agent whose belief is justified but false deserves as much acclaim as an agent whose belief is justified and true. An agent who has not met her own standards for accepting or rejecting her own beliefs warrants reproach, even if her belief

turns out to be true. Since all that agents can be expected to control is whether or not their beliefs are internally justified, an agent whose belief is unjustified but true deserves as much reproach as an agent whose belief is unjustified but false.

Andrew Latus puts the internalist's response to the problem of epistemic luck as follows: "A person should not, says the internalist, form a belief unless he has in his possession good reasons for that belief. Otherwise, the person is simply lucky that his belief turns out to be true" (2003, 163). To accept internalism as a response to epistemic luck is to dismiss the temptation to count the results of inquiry in one's evaluation of epistemic agents. According to the internalist, justification, which must be within the agent's control, is the only appropriate locus of epistemic appraisal.

The driving intuition behind internalism as a response to the problem of epistemic luck is to restrict the locus of epistemic appraisal to an arena which is supposed to be secure from luck: one's justified belief. If we take this internalist response to epistemic luck at face value, it claims either that we can be praised or blamed for knowledge, but that knowledge reduces to mere justified belief, or it insists that, in addition to being justified belief, knowledge must also be true, but that we cannot be praised or blamed for having knowledge. If the only epistemically praiseworthy behaviour is trying to achieve true beliefs, trying to achieve true beliefs and having true beliefs are equally praiseworthy. If we follow the former option only justified belief can be praiseworthy. 'Knowledge' is an empty accolade. If we follow the latter option, there is knowledge, but possessing it does not accrue to our credit. Neither option does justice to our intuitions about knowledge. In fact, analysing the internalist response to epistemic luck only brings the conflict among our epistemic norms into sharper focus.

The first problem with the internalist resolution to the problem of moral luck is that it does not appear to solve the problem. In effect, the internalist affirms one of the intuitions whose conflict gives rise to the problem of epistemic luck. Internalists insist that agents are only to be held responsible for the justification of their beliefs. While this position seems plausible when considered alone, its ramifications for our conception of knowledge are both perturbing and pervasive.

What internalism requires, if it is to be construed as a response to the problem of epistemic luck, is some explanation for why it is that we praise people for their unjustified true beliefs, and why we blame people for their justified false beliefs. That explanation seems readily available. When we praise people for true beliefs, we take the truth of a belief as evidence for its justification, and when we blame people for false beliefs, we take the falsity of a belief as evidence for its lack of justification. This explanation might well be psychologically accurate. However, far from substantiating the view that our epistemic practices accord with our epistemic norms, this explanation identifies the reason why we do not count only justified belief in our appraisal of epistemic agents: we care deeply about truth, and our epistemic norms would be unrecognizable without that esteem.

There is also another problem with the internalist's position. The internalist strategy is to restrict the scope of epistemic appraisals to justifications because those, at least, are under the agent's control. However, justification is something over which an agent might fail to have control. The strategy of restricting epistemic appraisal to the agent's justification presumes that the agent's justification is itself something over which the agent exerts control and, hence, about which praise and blame are appropriately apportioned. In light of circumstantial, constitutive and causal luck, this presumption

seems unfounded. Where the justificatory standing of our beliefs arises from circumstances beyond our control, restricting appraisal of agents to whether or not their beliefs are justified flouts the very intuition it was invoked to protect. For the internalist, if we do not control our beliefs' justification, it does not seem appropriate to praise or blame us for that justificatory standing. Not only does internalism do violence to the intuition that truth matters for responsibility, it may not be a viable strategy for guarding against epistemic luck.

Take the following example: Smith believes that Jones owns a Ford. He is quite justified in this belief. For the many years Smith has known Jones, Jones has owned a Ford, and Smith has just seen Jones driving a Ford. Now, unbeknownst to Smith, Jones is merely pretending to own a Ford. One day, as he is walking in the rain, Smith's umbrella is stolen. Smith ducks into the nearby auto rental place to avoid getting wet. There, he stumbles upon Jones returning the Ford he has been renting. As a result, Smith is no longer justified in believing that Jones owns a Ford. Now, Smith*, Smith's counterpart, does not have the misfortune of having his umbrella stolen. Thus, Smith* does not come to see Jones returning the Ford he has been renting, and Smith* continues to be justified in believing that Jones owns a Ford. While Smith is not justified in believing that Jones owns a Ford, Smith* is justified in believing that Jones owns a Ford, and the difference in their justifications is attributable only to luck.

It seems as though epistemic luck poses the same problem for justification as it does for knowledge. Either we limit praise or blame for justification to what an agent controls, or we praise or blame agents for whether or not their beliefs are justified, whether or not they wield control over that justification. If we opt for the former, then the area of legitimate epistemic appraisal seems diminishingly small. In light of the many

external factors influencing both belief and justification that are outside the agent's control, that which can legitimately be attributed to the agent seems very little. The latter option is not viable for those wishing to restrict responsibility to things agents control. If we opt for the latter, then we will be praising epistemic agents for things that they do not control. If we are willing to do so for justification, why not also do so for knowledge?

3.3. The Externalist Approach

The third approach to epistemic luck suggests that knowledge, justification, and beliefs warrant praise or blame, even though not under agents' control. Externalism about epistemic responsibility has been largely untouched. However, the intuition to count things beyond agents' control in the determination of agents' knowledge, for which they are deemed at least *prima facie* responsible, finds articulation in justificatory externalism. Externalist theories of justification reject the traditional internalist contention that an agent's justification for her beliefs consists of marshalling internally accessible beliefs in support of those beliefs. While the internalist seeks to meet the deontological concern to ensure that agents form their beliefs responsibly, the externalist conception of justification is responsive to the laudable objective to form beliefs in ways that have some objective purchase on the truth. According to externalism, the justificational status of a belief is a function of the process by which that belief was formed.¹⁶ External justification is a causal or nomological relation between the process that gives rise to a belief and the truth of that belief. A consequence of the externalist conception of justification is that the agent need not be in control of her belief's justification—she need not even be aware of it—for her belief to qualify as knowledge. For example, in his "Reliabilism: What is Justified Belief?" Alvin Goldman advances the following externalist account of justification: "the justificational status of a belief is a function of the reliability of the

process or processes that cause it, where...reliability consists in the tendency of a process to produce beliefs that are true rather than false" (1999, 311). Reliability is a conditional relationship between inputs and outputs of a process such that (in general) if the inputs are true, then the outputs will also be true.

Peter Unger's externalism (1968) is forthright in its eschewal of accidental knowledge. He proposes an account of knowledge in which S knows that p if and only if p is true, S believes p, and it is no accident that S believes p. If a belief results from a process whose function is to produce true beliefs, then the fact that the belief is true is not an accident. On Alvin Goldman's diagnosis, standard internalist accounts of justification fail to exclude accidental knowledge because they neglect to identify the causal processes that underlie belief acquisition. He states, "a strategy for defeating a noncausal principle of justifiedness is to find a case in which the principle's antecedent is satisfied but the belief is caused by some faulty belief-forming process" (1986, 311).

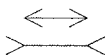
Goldman's reliabilism seeks to exclude accidental knowledge by requiring that putative knowledge claims be produced by processes that tend, on the whole, to be truth-productive. Externalists believe that no matter how good an internal justification one might have for a belief, that belief may only turn out to be true by a remarkable coincidence. The externalist justificatory requirement is meant to preclude such accidentally true beliefs from counting as knowledge. In order to preclude accidental beliefs from counting as knowledge, externalist conceptions of justification insist that any beliefs that qualify as knowledge must derive from processes geared towards attaining true beliefs.

Externalism does not deny that luck exerts influence on agents' beliefs. However, a belief that is lucky does not count as knowledge. If, by chance, a true belief originates

from a process that normally produces false beliefs, that belief will not qualify as justified, and, hence, is not a suitable candidate for knowledge. Similarly, if an ordinarily truth conducive process accidentally produces a false belief, that belief cannot qualify as knowledge. The externalist strategy for dealing with epistemic luck is to uphold the intuition that what matters for our appraisal of epistemic agents is their tendency to produce true beliefs rather than false beliefs.

One of the chief motivations for justificatory externalism is its promise to exclude resultant epistemic luck from an account of knowledge. Justificatory externalism posits a relationship between justification and truth such that a justified belief cannot be true by accident. However, externalists count justification that agents do not control as the justificatory criterion on knowledge. Justificatory externalism, therefore, admits factors beyond agents' control in the determination of knowledge. The additional question invoked by considerations of epistemic luck is whether agents can be responsible for knowledge, if they are not in control of the justificational component of knowledge.

Most incarnations of externalism reject the intuition that one must be in control of one's belief for that belief to count as knowledge. Paradigm reliable belief acquisition processes include "standard perceptual processes" (Goldman, 312)¹⁷ which, by and large, cause true beliefs rather than false ones. On Goldman's view, when Jones has a good look at a mountain goat that is 30 yards away, Jones knows that there is a mountain goat. However, Jones does not control his visual mechanism. The causal production of beliefs by perceptual mechanisms, like other causal processes, is not something generally held to be within the control of the agent. Take, as example, the simple visual experience elicited

by the following figure:  .¹⁸ Try as we might, we cannot make ourselves see

the two lines as the same length. The very fact that visual experience is something over which agents wield little control is what makes it a paradigm of reliabilist justification. The foregoing illusion notwithstanding, perceptual mechanisms usually cause true beliefs, and the agent's inability to control what she sees is a virtue that precludes the likelihood of false beliefs induced by wishful seeing.

Suppose that under normal conditions someone looks out her window and sees rain pelting the sidewalk. On the face of it, her perceptual evidence that it is raining constitutes a justification for believing that it is raining. However, it is difficult to see how her justification is within her control. Moreover, suppose that it is lucky that our agent happened to look out her window while it was raining; but for circumstances beyond her control, she would not have formed the justified belief that it is raining. If she is not responsible for the origination of her belief, or for whether or not it is justified, it seems as though we cannot praise or blame her for having a justified belief after all.

Externalists do not believe that justification is something over which agents wield control. Justification, rather, is an external relation between inputs and outputs such that an output belief is likely to be true if appropriately induced. Thus, although externalism accords well with the epistemic intuition to count the truth or falsity of an agent's beliefs in our assessment of whether that agent has knowledge, it flies in the face of the contention that knowledge is something for which agents are responsible. One of the major complaints against externalist conceptions of justification is that they seem to accord no responsibility to the agent for her belief. As Laurence Bonjour puts it, an externalist theory of knowledge:

holds that the conditions required for knowledge in addition to true belief concern facts about the relation between belief and world which need not (and normally

will not) be within the cognitive grasp of the believer; thus such a believer can have knowledge without having any reason for thinking that his belief is true (or even likely to be true). The contrast is with an internalist theory according to which it is a requirement for knowledge that the believer have a grasp of such a reason. (1987, 297)

On Bonjour's view, externalist justification does not suffice for knowledge. To substantiate his point, he offers the following as a counterexample to externalist accounts of justification:

Samantha believes herself to have the power of clairvoyance, though she has no reason for or against this belief. One day, she comes to believe, for no apparent reason, that the President is in New York City. She maintains this belief, appealing to her alleged clairvoyance power, even though she is aware of a massive amount of apparently cogent evidence...indicating that the President is at that time in Washington D.C. Now the President is in fact in New York City.... Moreover, Samantha does in fact have a completely reliable clairvoyant power..., and her belief about the President did result from the operation of that power. (1985, 38)

In "A Critique of Externalism" Keith Lehrer also faults externalism for allowing the counterintuitive result that a given cognizer who fails to have any reason to trust her belief would be justified. Lehrer insists that an agent whose belief is reliably formed would fail to have knowledge on the grounds that she lacks epistemic access to the reliability of her belief acquisition procedure. On his account, "a person totally ignorant of the external factors connecting her belief with truth might be ignorant of the truth of her belief as a result" (1999, 312). Lehrer believes that the central tenet of externalism

that some causal relation to the external world suffices to convert true belief to knowledge conflates knowledge with mere information. He, therefore, argues that the externalist conception of justification is better conceived as providing “accounts of the possession of information rather than of the attainment of knowledge” (312).

The internalist insists that “One must have some way of knowing that the information is correct” (Lehrer, 313). Knowing whether or not our beliefs are justified is implicit in the normative account of epistemology. Moreover, such access is constitutive of the normative force of justified, as opposed to unjustified, beliefs. Although externalism meets the epistemic intuition that the truth of an agent’s beliefs should count in our appraisal of those agents, it does so at the cost of making justification something for which the agent is not responsible. Thus, for the externalist, an agent may know something without being responsible for her knowledge. If we praise epistemic agents for knowing, construed externally, then we praise them for something that they do not control.¹⁹

If neither internalism nor externalism can adequately accommodate both our reactions to epistemic luck and norms of epistemic appraisal, it seems as though skepticism about the coherence of epistemic norms is appropriate. Before turning to skepticism, a brief foray into the promise of internalist-externalist hybridism about justification seems appropriate. Internalism is too lenient, allowing, as in Gettier examples, beliefs that have, given the circumstances, no anterior likelihood of being true to count as justified. Internalism is also too stringent, since it disallows as knowledge cases, such as perceptual knowledge, that appear to count as knowledge. Externalism is too stringent, allowing beliefs that have, from the subject’s perspective, no likelihood of being true to count as justified. Externalism is also too lenient, since it allows as

knowledge cases, such as Samantha's clairvoyant belief, that do not appear to be knowledge. What, then, of a view that combines internalist and externalist constraints on justification? What hopes do hybrid views have for solving the clash of intuitions about persons' responsibilities for their beliefs?

The trend towards hybrid views of epistemic justification supports the contention that neither internalism nor externalism suffices for epistemic justification. Externalists make allowances for internalism. Some, for example, require that agents' have epistemic access to the reliability of the processes whereby their beliefs are formed. Others move from the reliability of processes to agent reliabilism.²⁰ Internalists make allowances for externalism, requiring, for example, that the very process of adducing reasons in favour of a belief tends, on the whole, to be a process that generates true rather than false beliefs.

In *Problems of Knowledge: A Critical Introduction to Epistemology*, Michael Williams articulates a contextualist approach to justification in which justificatory criteria meet both internal and external constraints of a belief's being justified. For Williams, standards for justifying a knowledge claim are fixed by the context. Knowledge is accordingly fixed by context. In one context, A could truly say 'S knows that P', while at the same time, B, in a different context, could truly say 'S does not know that P'.

For Williams, the demand for justification depends on context. Not just every challenge is reasonable. Justification exhibits a 'default and challenge' structure. One is entitled to beliefs by default so long as there are no reasons to the contrary. In order to meet the challenge that the 'default and challenge' model of justification is purely dialectical, Williams includes the following five externalist constraints on justification:

- 1) The intelligibility of semantic constraint (unless we are right about many things, what we are talking about is not even clear).

- 2) Methodological constraints (some default entitlements are necessary to determine the direction of inquiry).
- 3) Dialectical constraints (within the specific context of inquiry).
- 4) Economic constraints (if the costs of error are high, the standards are raised).
- 5) Situational constraints (facts about the actual situation).

Williams also invokes epistemic responsibility (personal justification) and evidential justification (adequate grounding) and thinks both are requisite for justification, though both can be default-justified. It doesn't suffice that we form beliefs reliably; we also must know that we form beliefs reliably. This reliability doesn't, for Williams, justify our first-order beliefs. It is, rather, a precondition on our having beliefs at all. Beliefs are possible only if embedded in practices of inference and argument. But we know when to draw inferences from our beliefs and when to withdraw our knowledge claims when we know something about our reliability.

There seems nothing objectionable, and much to be praised, about Williams' careful articulation of a contextualist approach to justificatory internalism-externalism. Some sort of concession to the demands of its competitor theory of justification seems now to be nearly ubiquitous among externalists and internalists alike. Williams' employment of contextualism to mediate beliefs requiring internal and external justification is intriguing and it seems to hold some promise to serve the function for which it was conceived: it provides reasons to refuse to engage in dialectic with sceptics.²¹

Does the introduction of hybrid views of justification solve the problem epistemic luck poses for norms of epistemic responsibility? Though the requirement that agents be internally and externally justified in their beliefs solves many of the problems of

justificatory internalism and of justificatory externalism, the problem of epistemic luck appears to remain unaltered by the introduction of any conception of justification short of infallibilist justification. So long as a belief that is justified, even a belief that meets the requirements of a contextualized hybrid conception of justification, may nonetheless be false, there remains room for the sort of mismatch between a person being justified and a person being right. A and B have the same internal and external justification for P in context C. A is right. B is wrong. Do we praise A more than B or not? Conjoining internalism and externalism and adding contextualism doesn't seem to make a difference to the problem or its solution.²²

If our epistemic norms are incoherent, knowledge seems beyond our grasp. Although there is a difference between skepticism about the coherence of epistemic norms and epistemic practices on the one hand, and skepticism about the possibility of knowledge on the other, the former seems to lead inexorably to the latter. The former variety of skepticism asserts two things. First, it asserts that knowing that *p* cannot be a matter of luck. Second, it asserts that we attribute knowledge that *p*, at least in part, because of luck. The problem of epistemic luck pulls in two directions. Either knowledge is not as we traditionally construe it; that is, knowledge can be accidental, and epistemic standings can be meted out according to luck, or knowledge is as we traditionally conceive of it; that is, knowledge cannot be accidental, and epistemic standings must be meted out only according to the purposive methods of inquiring and justifying that its possessors have undertaken. The problem with the former option is that justification is an essential component of our conception of knowledge. If someone believes something by accident, it may qualify as true belief, but it fails to qualify as knowledge. The problem with the latter option is that it flies in the face of our actual practices of attributing

knowledge to individuals, and of our actual practices of assessing epistemic agents. Were we to abide by the strictures of the traditional conception of knowledge, the phenomena of epistemic luck would imply that vanishingly few beliefs qualify as knowledge.

In addition to skepticism about the coherence of epistemic norms in the face of luck, luck poses a problem for the very possibility of knowledge. There is considerable agreement about the jeopardy that luck poses for the coherence of epistemic norms.

Mylan Engel Jr. (1992), Andrew Latus (2003), and Daniel Statman (1993) to name a few recent contributors to the field of epistemic luck, agree that the influence of luck on belief, truth, and justification engenders skepticism about the possibility of knowledge. Daniel Statman, for example, believes that luck undermines the possibility of knowledge in at least two ways. He states:

First, if my belief in **p** is determined by factors over which I have no control—biological, physiological and others—it seems that doubt as to the truth of **p** is unavoidable. There always exists the fear that these factors mislead us, that they are means in the hands of a Cartesian demon. And if what I believe to be true is forced upon me by such a demon, I can never claim to have knowledge. Second, since justified belief does not guarantee truth, there is always a possible gap between what one believes to be true, relying on the best evidence available to one, and what is really true (1991, 152).

There is an intimate connection between the influence of luck on knowledge and skepticism about the possibility of knowledge. This is because global skepticism about the possibility of knowledge turns on the introduction of luck. René Descartes' evil genius hypothesis requires that we admit the possibility that our beliefs are systematically misleading. If it turns out that we are being manipulated by Descartes' evil genius, then

we are epistemically unlucky. We fail to have knowledge, despite even our best efforts, because our beliefs fail to be true. If, however, it turns out that we are in the real world, and our beliefs are not wholly false, we still do not have knowledge, because it is a matter of luck that our beliefs are true. After all, we could have been systematically deceived by an evil demon.²³

In "Is Epistemic Luck Compatible with Knowledge?" Mylan Engel, Jr. argues that the view that luck and knowledge are incompatible entails skepticism. On Engel's diagnosis, the incompatibility thesis leads directly to skepticism about the possibility of knowledge. Engel, therefore, insists that "anyone who wants to reject scepticism [must] face the logical music and reject the incompatibility thesis as well" (1992, 59). Engel maintains that, coupled with the incompatibility thesis, both externalism and internalism lead directly to skepticism. He defines internalism as "a theory of justification which maintains that epistemic justifiedness is exclusively a function of the cognizer's internal states...e.g. belief states, memory states, perceptual states, etc." (60). Externalism "maintains that epistemic justifiedness is (at least) partly a function of external features to which the cognizer lacks cognitive access, e.g. the actual reliability of the process producing the belief in question" (60). Engel argues that internalist theories of justification lack a "truth-connection." In order for justification to be conceptually connected with truth, "it must be the case that for every possible world W, if conditions C make person S's belief justified in W, then conditions C make it probable that B is true in W" (60).²⁴ The problem with internalist theories of justification is that "for any set of entirely internally specifiable conditions CI, there will always be a possible world, call it WD, where an evil demon has seen to it that we possess the requisite internal states and, hence, satisfy CI, even though all of our contingent beliefs are false" (60-61). From the

fact that internalism fails to provide a truth connection, Engel concludes that it is always a matter of luck when internally justified beliefs turn out to be true. Since internalist epistemologies cannot eliminate the influence of luck on having a true belief, if luck and knowledge are incompatible, then “no internalist epistemology can give rise to knowledge, and so, every internalist epistemology results in skepticism” (62).

Externalist epistemologies fare no better. While it is not always a matter of luck when an externally justified belief turns out to be true, it is always a matter of luck when a belief turns out to be externally justified. This is because externalist theories make justification a function of external conditions to which the agent lacks epistemic access, and over which the agent fails to wield control. Engel concludes: “Since no externalist epistemology can eliminate luck’s role in acquiring E-justified [externally justified] true beliefs, if the incompatibility thesis is true, it follows that no externalist epistemology can yield knowledge, and therefore, that all externalist epistemologies lead to skepticism” (64).

In “Moral and Epistemic Luck” Andrew Latus voices substantial agreement with Engel. He believes that we should despair of coming up with an account of propositional knowledge that does not countenance certain lucky guesses as knowledge. He states, “for any set of conditions we might care to propose, it is always possible to conceive of a case in which those conditions hold as a matter of luck” (2003, 162). If luck is incompatible with knowledge, and luck cannot be eliminated, it follows that there can be no such thing as knowledge. Latus echoes Engel’s claim that the antagonism between epistemic internalists and externalists is fomented by their adherence to the incompatibility thesis. On his view, “the debate between internalists and externalists has essentially consisted of

one side accusing the other of proposing accounts of knowledge that allow too much room for luck" (163).

The solution to the impasse that Engel proposes and Latus endorses is to deny the incompatibility thesis. In so doing, Engel and Latus move beyond current internalist approaches to epistemic responsibility, since these assert that luck is incompatible with knowledge. Engel distinguishes between veritic luck, which he believes to be incompatible with knowledge, and evidential luck, which he believes to be compatible with knowledge. An agent has veritic luck if "given her evidential situation, it is simply a matter of luck that her belief turns out to be true" (67). An agent has evidential luck if "she is lucky to be in the evidential situation she is in but that, given her evidential situation, it is not a matter of luck that her belief is true" (67). To substantiate his point that veritic luck is incompatible with knowledge, he asks us to consider the case of Dylan the euchre player. After shuffling the deck of cards, the euchre dealer asks Dylan which card he believes to be on top. Dylan forms the belief that the jack of hearts is the top card, and the top card just happens to be the jack of hearts. Given his meagre evidence, Dylan is very lucky that his belief is true. Dylan plainly fails to have knowledge in this case. The reason he fails to have knowledge, as Engel well notes, is that Dylan's belief that the top card is a jack of hearts is not justified. Standard cases of veritic luck turn on lack of justification (either internally or externally justified beliefs). Engel's example of evidentiary luck features Nadine who is lucky to be in the evidentiary circumstances that give rise to her belief. Nadine looks out the window of her study and, upon seeing rain and lightning, and hearing thunder, forms the true belief that it is storming outside. Now Nadine almost never uses her study. She usually spends her afternoons in a windowless, soundproof carrel in the library. It is by mere happenstance that Nadine is in such a good

evidential situation, rather than the poor evidential situation she would have been in, had she followed her normal routine. Engel believes that Nadine has knowledge, despite the influence of luck on her formulation of the belief that it is storming outside. He concludes, "Nadine is evidentially lucky regarding the storm, she nevertheless knows that it is storming. Therefore, evidential luck is not incompatible with knowledge" (69).

If we endorse the view that knowledge cannot depend on luck, a person will not be able to have knowledge of anything that she believes due to evidence that results from her situation, if that situation is somehow lucky. Latus concludes from Engel's example that "Clearly, not all sorts of luck are incompatible with knowledge" (165). The moral that Engel draws from the presumed compatibility of evidential luck with knowledge is that an appropriate epistemology must include some externalism.

Externalist epistemologies sanction inadvertent knowledge. One may fail to have knowledge despite all the evidence available to oneself. Insofar as failure to have knowledge redounds to one's reputation as an epistemic agent, this seems unfair. In "How Internal Can You Get?" Hilary Kornblith outlines the implications of the view that we may not be justified in believing ourselves to be justified.

The realization that there is a gap between what an agent is justified in believing and what she is justified in believing that she is justified in believing may cause us to distance the concept of justification from that of blame. Alternatively, it may force us to recognize that we legitimately may be blamed for things for which, by our lights, we should not be blamed. (123)²⁵

It is important to note, however, that Nadine's evidential luck does not preclude her belief that it is storming being justified. In fact, she has solid perceptual evidence for her belief that it is raining. Engel's Nadine example does not include resultant epistemic

luck, for her belief is true. Rather, Nadine's case turns on circumstantial epistemic luck. Engel assumes from Nadine's case that circumstantial epistemic luck is compatible with knowledge. That assumption merits investigation.

To see that Nadine profits from circumstantial epistemic luck, consider her counterpart, Ned. Ned, much like Nadine, spends most of his afternoons at his carrel deep in the bowels of the library. On the stormy day in question, Ned remains at his subterranean post throughout the duration of the storm. Consequently, Ned does not come to believe that it is storming out, while Nadine does come to believe that it is storming out. Now, had Ned whimsically decided to vary his routine and work by a window as Nadine had, Ned, too, would know that it is storming. Nadine's epistemic record reflects a true belief, while Ned's does not. The difference in their epistemic standings is fairly innocuous. Nevertheless, there is a difference, and that difference is entirely attributable to circumstantial epistemic luck.

What implications can we draw from Engel's claim that circumstantial epistemic luck is compatible with knowledge? One can accept the influence of luck on one's circumstances, and still be said to know, just as one can accept the influence of luck on one's circumstances, and still be said to have done the right thing. Just as some forms of circumstantial moral luck are compatible with morally correct action, so some forms of circumstantial epistemic luck are compatible with knowledge. To see that this is so, consider the difference between an individual who is beset by bad circumstantial moral luck and one who is not. Suppose that Nadine, upon sitting at the window of her study, notices an elderly woman struggling against the driving rain, and promptly offers the woman shelter for the duration of the storm. Ned would also have helped the woman, had he but been in a position to do so. Nadine is lucky not to have been at her carrel, and she

has performed a morally correct act. Being circumstantially lucky and performing moral acts are not incompatible. Now the difference between Nadine's moral standing and Ned's moral standing is slight. If, however, we vary both the moral example and the epistemic example, the importance of circumstantial epistemic luck comes to the fore. The threat that circumstantial epistemic luck poses is not to the status of one's belief as knowledge, but to one's status as a knower.

The standard fare of epistemologists, which features subjects who may or may not know that it is storming, or that there is a mountain goat 30 yards away, or that Jones owns a Ford, does not do justice to the importance of knowledge, nor does it do justice to the importance of epistemic standings. The difference between the epistemic standing of Nadine, who knows that it is storming, and Ned, who does not know that it is storming, fails to reflect the importance of people's epistemic standings. It is only when one considers more significant knowledge claims, and the accretion of such claims over a career, that the influence of luck on one's epistemic standing seems unfair.

3.4. Summary of the Problems with the Standard Approaches to Epistemic Luck

1. The skeptical approach undermines the coherence of our epistemic norms and leaves the paradox unresolved. It is, therefore, a position of last resort.
2. The Internalist view irrevocably damages our actual practices of epistemic assessments. It does not adequately deal with our intuition to count cases of epistemic luck. The internalist approach simply denies the force of the consequentialist intuition that we must count the results of inquiry in our assessments of epistemic agents. Insofar as the intuition to do so remains even after reflection, to deny the intuition does not solve the paradox. The internalist approach also fails to adequately address what we may call

the regress problem: One's justifications, as well as one's epistemic character, may also be possessed because of luck.

3. The externalist approach denies the force of the deontological intuition that one must have access to the justificational status of one's belief in order to have knowledge. Insofar as that intuition remains after reflection, simply denying it fails to solve the paradox exemplified by cases of epistemic luck. Moreover, in denying the contention that internally accessible justification is essential to knowledge, the externalist leaves us without means to distinguish paradigmatic instances of praiseworthy and blameworthy epistemic behaviour from mere guesswork.

A hybrid justificatory schema that includes both internalist and externalist criteria remedies many of the problems that each theory of justification faces by itself. Such views are, therefore, to be preferred over justificatory internalism and justificatory externalism. However, such hybrid views do not, by themselves, solve the problems that luck pose for epistemic responsibility. This is because one may be both internally and externally justified in believing something, and nevertheless be wrong.

It seems, therefore, that none of the standard approaches to the problem of epistemic luck is adequate.

Chapter 4. Proposed Resolution to the Problems of Moral and Epistemic Luck

The foregoing analysis of moral and epistemic luck brings into focus two apparently ineluctable but contrary dictates that govern attributions of responsibility. The first avers that one can only be held responsible for things over which one exerts control. The second insists that one may be held responsible for things over which one does not exert control. One cannot coherently espouse both dictates, and yet neither seems eliminable. As I have argued, the problem with standard solutions to the problems of moral and epistemic luck is that in seeking to eliminate one of the two requirements, they thereby fail to do justice to our actual practices of attributing responsibility to ourselves and to others. The resultant conceptions of responsibility are not satisfactory. One might, on these grounds, be tempted to join Nagel and Williams in their skepticism about the coherence of responsibility. If our conception of responsibility commits us to contrary positions, then it seems that one has no choice but to concede that our conception of responsibility is incoherent. The onus of this chapter and the next is to offer a framework that reconciles the intuitions that underlie these contrary claims and thereby forestall the necessity of acquiescing to skepticism in the face of moral and epistemic luck. I do so by offering a hybrid account of responsibility in which both intuitions that underlie the contrary dictates are necessary for legitimate attributions of responsibility.

To preview: responsibility is an evolved set of practices whose function is to coordinate the behaviours of intelligent, reflective, interdependent, norm-governed animals. That is, the notion of responsibility circumscribes a set of practices whereby persons gauge others', and prove their own, worthiness for cooperative endeavours and communal living. We hold others and make ourselves responsible for the good and bad things that we do and believe so as to ensure that our cooperative endeavours—

paramount among which is living—are coordinated. In order to serve this function, our practices of attributing and accepting responsibility must be responsive to what we actually do and believe, and they must be responsive to the reasons for which we do and believe the things we do and believe.

The problems of epistemic and moral luck pit the objective demand that we actually get things right—that we have true beliefs and perform good deeds—against the subjective requirement that we do everything within our power to ensure that we have true beliefs and perform good deeds. These competing epistemic and moral conditions on responsibility each have an evolutionary rationale. The objective demand that our actions generate the right result (namely true beliefs and morally correct actions) is responsive to the fact that our survival depends on having true beliefs and on doing the right things. There is, thus, a straightforward rationale for the intuitions that press us to count correct results as essential in assessments of moral and epistemic agents. The subjective demand that our actions stem from the right reasons (namely responsible inquiring and justifying procedures²⁶ and morally good intentions) exists because we are interdependent organisms who need a basis upon which to predict whether our cohorts will perform the right actions and arrive at the right beliefs. It follows that the reasons for which actions are performed are vital in assessments of moral and epistemic agents. The apparent paradoxes exemplified by moral luck phenomena and epistemic luck phenomena reflect the competition among these normative demands. Given the rationales for each, it can be seen that both demands are appropriate for beings like us. Moreover, given this diagnosis of the problems, it becomes obvious that the competition among the objective and subjective requirements on responsibility is the exception to the rule of their confluence. Trying with all your might to do the right thing and to arrive at true beliefs usually leads

you to do the right thing and believe truly; doing the right thing and having true beliefs usually follows from trying to do the right thing and trying to obtain true beliefs.

The first intuition centres on the importance of the subject's control. The dictate that one may only be held responsible for things over which one exerts control identifies the agent's subjective stance as the appropriate stance from which to evaluate her responsibility. The intuition to praise or blame persons only for what they control—their intentions or their inquiring and justifying procedures—is responsive to this subjective condition on responsibility. The second intuition centres on the importance of the event or state of affairs that the agent brings about. The dictate that one may be held responsible for things over which one does not exert control identifies the objective facts about what the agent did or believed as the appropriate measure for evaluating the agent's responsibility. The intuition to praise or blame persons for acts or beliefs that they do not control is responsive to this objective condition on responsibility.

My proposal is straightforward. I contend that both the subjective and objective requirements are necessary for legitimate attributions of responsibility. To be genuinely responsible for doing or knowing something, the agent whose doing or knowing is in question must somehow be in control of what she does or knows, and the agent whose doing or knowing is in question must have actually done or known the thing in question. Neither merely intending to do something nor merely doing something without intending to do that thing suffices for being fully responsible for that thing. Similarly, neither merely justifying one's belief that P, nor merely truly believing that P suffices for being fully responsible for knowing that P. I shall call the view that responsibility requires both that the agent fulfill her subjective obligations to do the right thing or have true beliefs

and that the agent actually do the right thing and have true beliefs the hybrid account of responsibility.²⁷

Very roughly, on this hybrid account, an agent is responsible for something if she 1. tries to bring that thing about and 2. succeeds in doing so. For example, agent A is responsible for consequence C just in case 1. A intends to bring about C, and 2. A brings C about. A is, thus, responsible for a successful intentional act. And an agent A is responsible for knowing that P just in case 1. A justifies her belief that P, and 2. P is true. A is, thus, responsible for a justified true belief. This (as-yet skeletal) hybrid account of responsibility seeks to accommodate the reasons that underlie the contrary intuitions we feel pressed to espouse when faced with cases of moral and epistemic luck. This conception of responsibility applies to paradigmatic instances of ordinary doings and believings. Agency involves being able to bring about an end that is specified by one's reasons. Epistemic agency involves bringing about true beliefs by inquiring judiciously. Moral agency involves bringing about correct results by intending to do so. Paradigmatic agency applies to cases where luck does not intervene between an agent's subjective striving to bring about a state of affairs and the objective fact that state of affairs is brought about. And it applies where luck does not intervene between an agent's subjective striving to attain true beliefs and the objective fact that she attains true beliefs. By emphasizing the subjective and objective conditions on legitimate attributions of full responsibility under ordinary conditions, the hybrid account of responsibility possesses the wherewithal to diagnose the problems of moral and epistemic luck. And once this diagnosis is at hand, the appropriate remedy will become clear.

The problem of moral luck, that chance both must not and does affect the moral status of individuals, and of epistemic luck, the view that chance both must not and does

affect the epistemic status of individuals, have parallel logical structures. Cases of moral and epistemic luck arise when only one of the two necessary conditions on paradigmatic responsibility is met.

In this chapter I address the conflict exhibited in cases of moral and epistemic luck from an unabashedly naturalistic perspective. A naturalistic perspective is one that finds responsibility to be a part of the empirical world, and, thus, to be open to understanding through empirical investigation. By examining how our practices of attributing responsibility evolved, and what function they play, I hope to provide a better understanding of what responsibility is. In particular, I shall argue that the dictates which we seem pressed to espouse in the face of moral luck are not genuinely reflective of the conditions on responsibility. Once one has in view the function of responsibility, it becomes clear that we can accommodate the intuitions that guide our avowal of contrary claims without requiring that we believe both that one can only be held responsible for things over which one exerts control and that one may be held responsible for things over which one does not exert control.

4.1. The Hybrid Solution and Evolution

Although I begin by offering a naturalistic explanation of the origins and function of responsibility, the solution to the problems of moral and epistemic luck I tender does not depend, logically, on that explanation. The speculative evolutionary account to follow is not essential to the credibility of the hybrid solution to the problems of moral and epistemic luck. That solution is independently defensible given the basic intuitions regarding responsibility that underlie the solution and that fit the numerous examples discussed in preceding chapters (and more besides). Central to the hybrid conception of responsibility is the vision of interdependent shared lives which explains the need and

purpose of moral and epistemic norms. Our interdependence cements the requirement on each of us to fulfil our objective and subjective epistemic and moral duties. Moreover, the dual understanding of morally significant action and epistemically significant belief that follows in Chapter 5 provides theoretical underpinnings for the hybrid solution on offer.

Nevertheless, it is natural to ask if it is possible to explain, in a way that is consistent with a naturalist perspective on the origin of responsibility, why we have a conception of responsibility in the first place, indeed a hybrid conception of responsibility, and thus to explain the function of responsibility as an adaptive product of biological and cultural evolution. This explanation provides a philosophically illuminating understanding of moral and epistemic norms. It will help us to understand in what respects our actual moral and epistemic practices are social manifestations of our most basic epistemic and moral predicament. The explanation to follow, though speculative, is consistent with the most current theory of evolution. This explanation reinforces the credibility of the solution offered, and adds depth to its significance, since it implies that the fundamental intuitions upon which the solution is based are not only independently credible but also are rooted in the origin and function of morality given our limitations as interdependent and vulnerable social animals.²⁸

4.2. The Evolution of Normative Control

On my diagnosis, both epistemic luck and moral luck phenomena arise because of the evolutionary conditions in which holding people accountable arise. Our practices of holding one another responsible for our intended consequences serves dual functions: by being responsive to consequences of agents' actions, they enforce norms, and by fixing the intentions with which agents' actions are performed, they facilitate prediction of who

among our would-be cooperators is liable to conform to those norms and who is liable to violate them.

Practices of punishment and reward have the function of coordinating the behaviours of groups who engage in such practices. Our practice of attributing responsibility to others and taking responsibility for ourselves is a sophisticated extension of the evolved practices of punishment and reward. I shall here investigate the crucial step from the primitive practices of normative control manifested by chimpanzees, wolves, and orca (to name a few social mammals that exhibit such control) to our conception of responsibility. An account of evolutionary origins of our practices of holding one another responsible for violations of, or conformity with, moral and epistemic norms obviously requires significant scaffolding.

The first bit of scaffolding demanded is an account of the evolution of moral and epistemic norms.²⁹ Fortunately, much of that scaffolding is on offer. Here I sketch a very brief overview of some conclusions from a large and complex literature on the evolution of norms. Allan Gibbard, among others, suggests that “systems of normative control in human beings...are adapted to achieve interpersonal coordination (1990, 64). The practice of punishing agents who fail to cooperate with their group members (and rewarding those who do cooperate) is a mechanism for solving cooperative impasses that allows groups to enjoy the benefits of cooperative activity. Such normative control is characterized by a group in which there is “a partly learned tendency to behave in a pattern beneficial to the group and members of the group reinforce that tendency” (Campbell and Woodrow, 2003, 354). Boyd and Richerson (1992) show that normative control is likely to arise wherever the cost of punishment is relatively low. This can be virtually guaranteed whenever there emerges a practice of punishing those who don’t

punish. Holding people accountable for violating moral norms helps ensure conformity with those norms because it ensures that such conformity serves the self-interest of the would-be violator. Holding people accountable for violating epistemic norms helps ensure conformity with those norms because such conformity serves the self-interest of the would-be violator.

Norms have their roots in the benefits of coordinated activity. Moral norms have their roots in relative benefits that accrue to groups of altruists as compared with groups of selfish individuals. Evolutionary game theoretic models demonstrate that, under certain conditions, the strategy of sacrificing one's immediate individual welfare to the long-term welfare of one's group is an evolutionarily stable strategy.³⁰ Epistemic norms arise because group benefit from true beliefs depends on the individual members of the group having true beliefs. Because group survival and flourishing depends significantly on the information accumulated by its individual members, it is unsurprising that practices of punishing people for false beliefs, and rewarding them for true beliefs should have arisen.³¹ Epistemic norms have their roots in the benefits that accrue to groups who rely on one another's epistemic accuracy in cooperative activity that requires true beliefs. The costs of false beliefs are high, and holding people accountable for the consequences of their inquiry helps to ensure that people inquire judiciously.

4.3. From Normative Control to Responsibility

The onus of this section is to explain how our practices of holding one another and ourselves responsible for what we do and what we believe evolved from more primitive practices of normative enforcement. More specifically, I seek to explicate, in evolutionary terms, the crucial step from systems of normative control which centre around punishment and reward for objectively accessible overt behaviour to systems of

normative control which take as their central tenet *justified* responses to behaviour. What justifies primitive normative control? We can offer an external justification for such practices in terms of the betterment of the group that attends coordinating their behaviours. But such a justification leaves us cold when we attend to the particular norms that are thereby justified. And when we attend to the particular instances of punishment and reward for overt behaviour, we feel as though it is unfair to punish or praise agents for non-deliberate acts.

In *The Sources of Normativity*, Christine Korsgaard articulates a similar worry about the sufficiency of evolutionary justifications for norms of responsibility. Korsgaard allows that something like the foregoing evolutionary theory of responsibility might be explanatorily adequate. But she insists that the evolutionary explanation for the origination of moral norms does not justify the claims that morality makes on us. She asks, “why after all should the preservation of the species count so much more than the happiness of the individuals in it?” (996,15). She claims that the evolutionary theory violates the ‘transparency’ condition. The transparency condition states that a normative theory “must be one that allows us to act in the full light of knowledge of what morality is and why we are susceptible to its influences, and at the same time to believe that our actions are justified and make sense” (17). And on her view, the fact that normative control increases the fitness of animals who engage in normative control provides no incentive for agents to abide by the systems of normative control that are diffused within their group. Finding out that the source of one’s moral motivation is moral instinct does not justify that motivation for that individual. She states, “given the strength of the moral instinct, you would find yourself overwhelmed with the urge to do what morality

demands even though you think that the reason for doing it is inadequate....you might be moved by the instinct even though you don't upon reflection endorse its claims" (15).

Korsgaard believes that normativity derives its authority from the ability to reflect on and endorse our motivations to conform to norms.

A lower animal's attention is fixed on the world. Its perceptions are its beliefs and its desires are its will. It is engaged in conscious activities, but it is not conscious *of* them. That is, they are not the objects of its attention. But we human animals turn our attention on to our perceptions and desires themselves, on to our own mental activities, and we are conscious *of* them. That is why we think *about* them.

(92-93)

Humans have the morally and epistemically significant point of view that we have because of the way we evolved as a particular kind of social creature. Happily, humans have evolved to a point from which we can evaluate our own normative practices, and decide whether and how to change those practices for the better. Once we are able to take our own mental states as objects of reflection, we become able to ask the questions: "Is this perception a *reason* to believe?" and "Is this desire a *reason* to act?" How does this ability arise?

The ability to exert control over our selves, over our actions and beliefs, and the special character of human consciousness, its reflective nature, emerged simultaneously in the evolution of our species. The evolution of human consciousness and of human self-control are evolved mechanisms for individual agents to thrive in groups in which primitive normative control had evolved. According to Nietzsche³² the development of normative control among groups necessitated the development of self-control and self-consciousness among those animals who occupy the lower echelons in the hierarchy of

the group. The capacities for self-control and self-consciousness are, for Nietzsche, the responses of our progenitors to being punished. Nietzsche believes that an intelligent willful animal with the instinct to dominate, when punished by others, turns its instinct to dominate inward, and learns to dominate itself. And in that way reflective distance and the autonomy that goes with it came into being. Korsgaard captures Nietzsche's argument as follows:

The only way to prevent an aggressive animal from behaving aggressively is to punish it—that is, to inflict pain on it. And this is not like the pain that arises naturally in the life of an animal, when nature and circumstance threaten its physical identity, and it finds it must revolt against the threat. This pain is imposed from without, when the animal is acting naturally, that is, doing what is made imperative by the identity it naturally has. So this pain, the pain of punishment, forces the animal to revolt against its own identity. And then for the first time it says to itself: I should be different from what I am, and it experiences guilt. And that is the origin of normative thought. (158)

I think that Nietzsche was right to suggest that the capacity for self-control and self-consciousness are evolved responses to social circumstances in which animals punish and reward one another for their violation of or conformity with norms. And I think that Korsgaard is right to suggest that our ability to consciously reflect on the reasons for our desires and beliefs, and either endorse or reject those reasons as justified reasons, is central to particularly human normativity. But I think that Nietzsche was wrong about the details of the evolution of genuine normativity. In particular, I think that his approach suggests that self-consciousness and self-control themselves serve no evolutionary function. By contrast, I insist that they serve the all-important function of enabling the

sort of internal and first personal justification of particular norms that Korsgaard finds lacking in evolutionary explanations of normativity.

Why does an *explanation* for the origins of our two dimensions of epistemic and moral appraisal constitute a *justification* for the continued employment of those two dimensions of epistemic and moral appraisal? What is taken to be the rationale for practices of attributing and undertaking responsibility is the role those practices play in enabling cooperation and collaboration among humans. The evolutionary explanation offered here shows how the two dimensions of responsibility fulfill that normative function. In so doing, the explanation for the two dimensions of responsibility also constitutes a justification for the two dimensions of responsibility.

Moral and epistemic norms are externally justified. They are good things to have for beings like us. But they must also be internally justified. That is, we must be able to endorse the reasons for following particular norms in particular instances. Primitive normative responses are keyed to overt behaviour. In primitive practices of normative enforcement—punishment and reward among Capuchins for example—one cannot expect there to be a separation between the consequences that an organism brings about and the mechanism by means of which the organism brought it about. The subjective and the objective aspects, or the motivational and behavioural components, of a capuchin's activity are conjoined. And her counterparts' punishments or rewards are responses to a motivational-behavioural complex. In *Good Natured*, Franz deWaal provides evidence that practices of punishment and reward among capuchins is unintentional. By contrast, chimpanzees seem to be conscious, at least to some degree, of what they are doing when they punish or reward their counterparts. This suggests that practices of punishment and reward first evolved as mechanisms to coordinate group activities, and that mechanisms

of conscious recognition and motivation were a later adaptation which served to greatly enhance the reliability and efficiency of those practices in coordinating cooperative activities.

In "Why Moore's Open Question is Open: The Evolution of Moral Supervenience," Richmond Campbell and I advance an evolutionary explanation for how humans evolved genuine normative practices from the more primitive systems of normative control which presumably existed among our forefathers and foremothers. There is a large step between the kind of normative control which can be expected to have arisen among our progenitors, and the sophisticated, linguistic, intentional, and dialectical practices which fall under the rubric of our concept of responsibility. The crucial step in that evolutionary argument is the requirement that agents come to be able to identify what prompts their responses to circumstances in a way that not only identifies the reason for acting as a causal originator, but also employs that reason as a justification for the response in question. The crucial difference between human normative practices and those of other social animals is that we provide justifying reasons.

Interdependent animals whose behaviour is norm-governed face a quandary: they must simultaneously meet the incompatible demands that they be predictable to one another and that they be responsive to changes in their social and physical environments. To be predictable, animals' normative motivations must be transparent and intractable. However, to be socially and environmentally adaptive, animals must be docile. They must conform their motivational structures to those of their group to maintain group membership, and they must be able to change their norm-governed behaviours to suit new circumstances. By making explicit the circumstances to which an animal's norm-governed behaviour is responsive, language provides a means to resolve the conflicting

drives towards stubbornness and towards docility. Once animals have the wherewithal to indicate the reasons for their responses, the aptness of either stubbornness or docility in the face of new situations can thereafter be socially negotiated.

Practices of attributing and taking responsibility are sophisticated extensions of more primitive systems of punishment and reward. The evolutionary function of those practices is to enable the coordination of behaviours. To blame someone takes the force of a challenge to that person's commitment to the norm she has violated. The ability to distinguish between an agent's violation of a norm and an agent's having meant to violate a norm enables agents to demonstrate their commitment to or rejection of norms. Citing the reasons for which one has responded in particular ways can be exculpatory because such agents thereby display their espousal of the norm in question.

How could this practice of identifying the reasons for one's behaviour have evolved? In "How to be both Docile and Autonomous: The Drive to Make Reasons Explicit"³³ Richmond Campbell and I examine the evolutionary origins of our capacity to distinguish between the reasons for which an agent does what she does, and what the agent actually does. That is, we provide an evolutionary explanation for our capacity to differentiate the agent's subjective motivation in acting from an objective characterization of the agent's act. The hypothesis that we defend is that the ability to make explicit the reasons for which we act evolved as a social mechanism to resolve normative uncertainty. In the following, I summarize the argument we present there.

A number of elements must be in place in order to account for the mechanism for making reasons explicit. The first of these is normative expression. Normative expression is simply approval or disapproval directed at reason-response pairs. We think that there are good reasons for which normative expression would have evolved. For example,

normative expression would replace physical punishment for violation of norms, and thereby reduce the costs of maintaining normative control. Notice that for normative expression to operate, animals must be able to represent norms.

Docility and autonomy are also evolved necessary ingredients for the proper functioning of norms. Docility is the tendency of animals to conform to the norms of their group. This is necessary simply for the coordination of normative behaviours. A dearth of docility will prompt challenges from group members. Autonomy is the tendency of animals to consistently respond to situations. This is necessary for the coordination of normative behaviours, and it is required for group members to predict one another. A dearth of autonomy will prompt challenges from group members. An animal who is insufficiently docile or insufficiently autonomous will be a poor bet for cooperative behaviours.

We think that for animals capable of normative expression, and limited linguistic aptitudes, the pressures exerted by the requirements of docility and autonomy can straightforwardly result in a pattern of challenge and response. We think that the step is small in evolutionary terms, but the results are tremendous.

The challenge, the expression of disapproval towards overt behaviour has, in this pattern, the force of “why” directed at the appearance of insufficiently docile or insufficiently autonomous behaviour. The appropriate response to such a challenge is to articulate the norm that the individual is following. Notice that as an answer to the question “why did you violate the norm?” an articulation of the norm will be the challenged individual’s justification to her challenger. She will see her behaviour as justified by the norm, and if the challenger agrees that it is, then the behaviour itself will be justified.

I will use a food sharing norm exhibited among chimpanzees to demonstrate how the pattern of challenge and response can re-establish normative certainty where the boundaries of norms are fuzzy.³⁴ The food sharing norm states that when food is neither so scarce nor so abundant as to make food-sharing useless, share food. When the boundaries of the food sharing norm are fuzzy, normative uncertainty arises. For example, when food is somewhere between moderate scarcity and scarcity—represented here by two leaves—some animals try to share the food while others will hoard it. When person A sees two leaves, she shares. When person B sees two leaves, she does not share. Animals will not be able to reap the benefits of coordinating their food sharing behaviours.

The pattern of challenge and response would arise in this circumstance because, to A, B seems not to be docile. B seems to be violating the norm “whenever three leaves, share.” A will, thus, challenge B’s behaviour. In reply, B can try to convince A that she is being docile after all by articulating the norm she is following. Thus, B can state the group norm “whenever one leaf, don’t share” as a justification for her behaviour “two leaves, don’t share.”

The transition from primitive normative control to reasons-governed normative control is a transition from systems of punishment and reward, which are keyed to a complex of motivational and behavioural elements, to our systems of attributing responsibility according to the agent’s explicitly articulated reasons for behaving, as well as what they actually do.

Richmond Campbell’s insight that humans’ evolved ability to inhibit actions they are motivated to perform is a necessary precursor to the evolution of genuine reasons-sensitive normativity also provides an explanation for how responsibility can attach to

motivations divorced from their usual manifestation in objective behaviour, and to objective behaviour detached from motivations. The ability to inhibit one's motor cortex from manifesting behaviour is a later adaptation on the reasoning-acting mechanism of our predecessors. In his "Two Faces of Moral Judgment" Richmond Campbell cites the function of muscle spindles in support of the view that "a process in the brain that can be identified as a single event or unified series of events can be hooked up to other processes that result in the original process performing functions that are manifestly different from or even opposed to each other" (33).³⁵

The now famous experiments by Benjamin Libet suggest that voluntary control over behaviours takes the form of inhibiting messages sent from one's primary motor cortex to one's spinal nerve cells. The brain begins its discriminative and evaluative work as stimuli are received, and it responds to its evaluations just as seamlessly. The ability to consciously reflect on those evaluations and the appropriateness of the responses in which they issue is a later adaptation resulting from our being subject to punishment and reward for those responses.³⁶ Libet's experiments provide evidence that suggests that our original motivational-behavioural complex was originally one process which was later (in evolutionary terms) divided into a motivational function and an acting function, with the capacity to inhibit behaviour one is motivated to perform.

Given our interdependence, humans are responsible for 1) trying our best to get things right, both morally and epistemically, and 2) getting things right, both morally and epistemically. Without actually getting things right, we would not survive. That is, the benefit that attaches to holding persons responsible for what they actually cause and believe is meted out according to actually getting things right. Human beings have a fundamental need for true beliefs. Without them, we would surely perish. As Edward

Craig claims in his *Knowledge in the State of Nature*, this basic need for truth drives us to seek out reliable others, so as to pool our epistemic resources.³⁷ Without the objective sense in which one is actually following the moral and epistemic norms of one's group, there would be no coordinated activity, and hence no evolutionary benefits to normative control in a group. Without the ability to tell one another that one was trying to get things right, groups of humans would not be able to negotiate the fuzzy boundaries of norms, resolve normative conflict, or respond in coordinated ways to new features of their physical and social environments.

In order to perform the evolutionary function of coordinating behaviours, practices of attributing and taking responsibility must be keyed to what persons actually do. However, in order to perform the evolutionary function of coordinating the behaviours of agents who are able to inhibit their motivations (and therefore for whom the proliferation of norms can be expected)³⁸ practices of attributing and taking responsibility must be keyed to the reasons for which persons do what they do.

4.4. Normative Control and the Objective Condition on Responsibility

The demand that our actions generate true beliefs and morally correct actions is responsive to the fact that group survival depends on its members having true beliefs and doing the right things. Holding one another accountable for the events that we cause and the beliefs we disseminate has the evolutionary benefit of ensuring that group members conform to the epistemic and moral norms of their group.

Whether one has cooperated in the past provides an inductive basis upon which to gauge whether one will cooperate in the future. Thus, the structure of normative enforcement—holding one another responsible for the states of affairs and beliefs that we actually bring about—serves as a mechanism for coordinating group behaviours and as a

mechanism for predicting which group members will be apt for future coordinative behaviours and group living.³⁹ But notice that there are significant costs attached to the inductive evidence required to predict which members will be good bets for future cooperative activity. Waiting to see who does evil things, or who disseminates false information, carries with it the costs of failed collaborative enterprises.

4.5. Normative Control and the Subjective Condition on Responsibility

With the transparency of motivations (or better yet, the ability to communicate one's intentions and justifications—see Richmond Campbell and my paper for how that might have arisen) comes the ability to predict what individuals intend to do.⁴⁰ We can then use what they are motivated to do as a means of deciding which group members are poor bets for cooperative epistemic and moral activity without the costs incurred by their actual failure.

In a very interesting paper on the social dimensions of epistemic practices, Miranda Fricker outlines the difference between an agent's credibility and her rational authority. Rational authority attaches to agents who are good at arriving at true beliefs, and who are, therefore, trustworthy.⁴¹ Credibility attaches to agents who are able to display their epistemic competence to their epistemic collaborators. Someone who has a poor track record of distinguishing poisonous red berries from non-poisonous red berries will be a poor bet for epistemic cooperation because she lacks rational authority and credibility. However, Fricker notes, someone who has recently learned the secret for distinguishing poisonous from non-poisonous berries will not gain credibility until her track record has been reformed. Thus, she may possess rational authority without being credible.⁴² If, however, that agent is able to make clear to her counterparts the reasons for her epistemic behaviour, by citing the distinguishing characteristics of poisonous berries,

she will incur others' trust. The use of reasons to supplement mere track records enables a rapid transition from being a non-credible to being a credible agent.

When someone defects, the costs can be enormous. Individual lives and group livelihood are at stake. Where one's life, or the livelihood of the group is at stake, a means for discovering which individuals will cooperate and which will defect *before* the opportunity for defection arises would have enormous benefits. Given the (not too controversial) assumption that intentions can be made transparent, an individual's intentions offer a mechanism for parsing future cooperators from future defectors. There is, thus, an evolutionary rationale for parsing an agent's actions into a subjective component responsive to the agent's reasons and an objective component responsive to the truth-value of the agent's beliefs and the moral value of the states of affairs she brings about.

In the ordinary course of events, good intentions produce good actions, and bad intentions produce bad actions. Someone's intentions, therefore, provide a basis for predicting what that person will do. We care deeply about consequences, and that is why we care deeply about intentions. Intentions lead to actions, and the ability to discern someone's intentions enables one to predict that person's actions.

I contend that our concept of responsibility is an essentially normative concept that attaches to successful intentional actions and justified true beliefs. Both intentions and consequences, justificational standing and truth-value, are necessary for our concept of responsibility to apply. When a purposive action does not produce its intended consequences, our notion of responsibility does not apply (or does not apply without qualification) to the actor. One can only be less than fully responsible for an intention if

the intended consequence does not follow, and one can only be less than fully responsible for a consequence if that consequence is not intended.

Our concept of responsibility is responsive to two evolutionary drives: to ensure that individuals get the right results, and to ensure that groups get the right results by allowing cooperators to gauge one another's aptness for cooperation. So long as one another's intentions are transparent, the causal link between intentions and the consequences they generate form an inductive, predictive package we employ to assess our counterparts' adequacy for cooperative endeavours and communal living.

4.6. Responsibility and Luck

Situations involving unfulfilled intentions or accidental action, justified false belief or unjustified false belief, pose a genuine problem to ordinary means of assessing responsibility. Cases that exemplify moral and epistemic luck fail to be sufficient for attributing responsibility because they meet only one of the necessary criteria for attributing responsibility. A bad consequence alone, without the corresponding intention to bring about that bad consequence, is not sufficient for attributing responsibility. Similarly, a bad intention alone, without a corresponding bad consequence, is not sufficient for attributing responsibility. A person's intentions generally lead to consequences. They are indicative of the future harms or benefits that she will bring about. But our concept of responsibility is not merely a mechanism by which to predict future consequences. It is also a mechanism we employ to enforce patterns of behaviour that accrue to group-benefit. The consequences that a person brings about are indicative of her willingness to conform to group norms. But our concept of responsibility is not merely a mechanism to enforce behaviours. It is also a mechanism to gauge agents' aptness for cooperative endeavours.

4.7. Hybrid Epistemology

The dispute between internalists and externalists about justification centres on whether or not epistemic agents must have cognitive access to the evidentiary standing of their beliefs in order for those beliefs to qualify as knowledge. While internalists maintain that the justificatory component of knowledge must be available to agents, externalists do not. The motivation underlying externalism to connect knowledge with what we believe to be true in normal worlds has the virtue of being able to divert skepticism about whether our beliefs are isomorphic with the external world. By insisting that agents can know about the external world without knowing that their beliefs qualify as knowledge, externalism can circumvent the external world sceptic's challenge to demonstrate that we have knowledge about the external world.

This achievement is not without sacrifice. In securing immunity from sceptical threats to the very possibility knowledge, externalism concedes a vital feature of epistemic achievement. It repudiates the requirement that evidence for a target belief be reflectively accessible to the epistemic agent. The internalist refrain that externalists have accounted for the possession of information, but not knowledge, is captured by a reluctance, even among externalists, to ascribe knowledge to agents whose true beliefs are got by means that are reliable in normal worlds, but who lack any access to the fact that the mechanism that produce those beliefs are reliable. Two examples from the literature underline this point.

Consider first Alvin Plantinga's example of an uncommon brain lesion that causes its victim to believe that he has a brain lesion.

S suffers from this disorder and accordingly believes that he suffers from a brain lesion. Add that he has no evidence at all for this belief: no symptoms of which he

is aware, no testimony on the part of physicians or other expert witnesses, nothing.... Then the relevant mechanism will certainly be reliable; but the resulting belief—that he has a brain lesion—will have little by way of warrant for S. (Plantinga 1993a, p.199)

A second example concerns the uncommon but reliable ability of chicken sexers to parse male from female chicks. Agents who possess this ability are noteworthy in their lack of cognitive access to the way they come to believe that a chick is male or female, and the reliability of their ability to so know.

According to internalists, agents with the uncommon brain lesion and chicken sexers fail to meet a necessary requirement for knowledge: they lack access to the evidentiary grounds of their beliefs. Though externalism is able to obviate such veritic luck by tying the truth or falsity of beliefs to their generation by reliable mechanisms, they are unable to obviate what Duncan Pritchard has called reflective luck. Reflective luck differentiates between agents whose beliefs are not luckily true, given the information available to them, and agents whose beliefs are luckily true, given the information available to them. Reflective luck, thus, concerns the standing of an agent's belief as justified for the agent. Pritchard defines reflective luck as follows:

For all agents, ϕ [sic], the truth of an agent's belief in a proposition, ϕ , is reflectively lucky if, and only if, the agent's belief that ϕ is true in the actual world, but, in nearly all possible worlds consistent with what the agent is able to know by reflection alone, were the agent to form a belief that ϕ , that belief would be false.

Reflective luck can be seen by comparing chicken sexers who lack access to how they sex chickens and do not know that they are reliable chicken sexers with enlightened chicken sexers. Enlightened chicken sexers know that their ability to parse male from female chicks redounds to an acute sense of smell. They, moreover, realize that their sense of smell is highly reliable at differentiating male and female chicks. The unenlightened chicken sexer is lucky to have a true belief in a respect that the enlightened chicken sexer is not. The former has no knowledge reflectively accessible to her to indicate that her belief is reliable in normal worlds.

Reflective luck also differentiates between an unenlightened brain lesion sufferer and an enlightened brain lesion sufferer. Insofar as the latter has strong grounds for believing that his belief that he has a brain lesion stems from a reliable source, his belief that he has a brain lesion is not lucky. By contrast, the agent who does not have any reason to believe that his belief that he has a brain lesion results from a mechanism that is reliable in the actual world, and across nearby possible worlds, is lucky in his belief that he has a brain lesion.

In *Socializing Evolutionary Epistemology*, I defend a hybrid conception of epistemic justification.⁴³ I argue that internal dialogical justification is a reliable truth-producing mechanism for groups, and that external reliabilist justification for primitive beliefs carries internalist justification in the form of evidence for those mechanisms' truth-conduciveness. Internalist justification is an irreducibly social process that involves complex dialogical interactions among members of an epistemic community. Although externalism is appropriate for beliefs arising from primitive belief acquisition procedures such as perceptual beliefs and enumerative inductive beliefs, more complex beliefs arising from more advanced belief acquisition procedures such as the testimony of others,

deduction, abduction and Bayesian induction must be justified inter-subjectively to qualify as knowledge. Notice, however, that the externalist conception of justification which cites the reliability of the source of one's beliefs serves as justification at least in part because we are enlightened about the reliability of perception and induction. Were we unaware that these belief acquisition mechanisms have 'better than coin tossing chances of being true'⁴⁴ in conditions similar to those in which the mechanisms were naturally selected, citing the provenance of perceptual beliefs would not suffice to confer justification.

When agents espouse beliefs they present themselves as able to provide reasons for the acceptability of those beliefs, and meet intelligible challenges to those beliefs. Justification involves giving and receiving observational and inductive data, subjecting those data to criticism, and negotiating inter-subjective agreement about those beliefs. Given agents' membership in epistemic communities, merely citing the origins of a belief does not suffice to justify that belief intersubjectively. The internalist conception of justification is grounded in our membership in interdependent epistemic communities. The responsibilities placed upon agents in virtue of that membership are not just to attain true beliefs by reliable means. It is also to conduct themselves as responsible epistemic agents. That involves monitoring their belief acquisition mechanisms in such a way as to wield indirect control over the truth-to-falsity ratio of their beliefs.

4.8. Hybrid Morality

The dispute between internalists and externalists about moral responsibility concerns the acceptability of counting factors external to the agent in determining the agent's moral standing. For externalists, because we are individually subject to luck, and we depend on others for our ability to respond to circumstances in moral ways, we are responsible for

things that are beyond our ability to either predict or control. These objective responsibilities stem not from the autonomy of moral agents, but from our interdependence.

The driving insight behind moral externalism connects moral evaluations to our actual practices of attributing and undertaking responsibility in a world that is not of our making. It, thus, connects moral norms with the social realities to which they are responsive. This feat, however, is not without costs.

In securing immunity from sceptical threats to the very coherence of moral norms, moral externalists cede a vital feature of moral achievement. They renounce the requirement that the agent's reasons for an action be central to evaluating that agent. The internalist exhortation that control is essential to moral appraisal is captured by a reluctance, even among externalists, to ascribe full moral responsibility to agents who mean to cause something bad, but cause something good instead or vice versa. Two examples underline this point.

S is a well-trained Emergency Medical Technician who administers CPR to a child. Unbeknownst to S, the child suffers a rare and undetectable heart condition. The child dies because of S's attempt to save her life. T is a murderous child abuser who beats a prone child. Unbeknownst to T, the child suffers from a rare and undetectable heart condition. The child lives because of T's attempt to murder her.

The hybrid view has the virtues of internalism and externalism without their faults. Epistemic externalists insist that agents are responsible for the truth-to-falsity ratios of their actual beliefs. However, what makes them responsible for the truth or falsity of beliefs is that other agents who depend on them and upon whom they depend hold them responsible. Moral externalists insist that agents are responsible for the good-to-bad

ratios of their actual behaviours. However, what makes them responsible for the goodness or badness of their actions is that other agents who depend on them and upon whom they depend hold them responsible.

Epistemic internalists insist that agents are responsible for the reasons to which their beliefs are responsive. What makes them responsible for their reasons for believing is that they constitute the means of achieving control over the truth and falsity of their beliefs. Others depend upon those reasons, both for evaluating the beliefs and for evaluating the agent. Moral internalists insist that agents are responsible for the reasons to which their behaviours are responsive. What makes them responsible for their reasons for acting is that they constitute the means of achieving control over the goodness and badness of their behaviours. Others depend on those reasons, both for evaluating the actions and for evaluating the agent.

4.9. A Skeptical Solution or a Straight Solution?

Intentions and consequences, taken together, serve as the best indicator of an individual's commitment to norms and aptness for future cooperative endeavours. But the two criteria for responsibility are responsive to different evolutionary functions. The causal criterion, that the individual whose responsibility is at issue caused the consequence, is responsive to the evolutionary demand to have behaviours under normative control. The intention criterion, that the individual whose responsibility is at issue intended the consequence, is responsive to the evolutionary demand to have predictable group members without the cost of waiting and seeing whether they will bring about good or bad consequences. Because the two criteria, intention and consequence, usually coincide we have one concept, responsibility, to fulfill both roles. Our notion of responsibility attaches to a composite of intention and consequence.

Evolutionary theory has provided an explanation for how our concept of responsibility arises. But hasn't it just shown that, in cases of moral and epistemic luck, our concept of responsibility simply breaks down? If it doesn't apply, then haven't we just proven that the skeptic is correct and that our moral and epistemic norms are incoherent in the face of moral luck? On my view, the fact that the concept of responsibility serves two functions can illuminate a decision procedure for when those functions come into conflict. In particular, if holding people accountable for the consequences of their actions serves the purpose of ensuring that those individuals do the right thing in the future, and thereby serves the purpose of coordinating behaviours so that groups can reap the rewards of cooperative endeavours, when people do bad things accidentally, we still ought to hold them "consequence-accountable." Meanwhile, if holding individuals accountable for their intentions serves the purpose of determining whether they will be good future cooperators, when people have bad intentions which fail to have their attendant bad consequences, we still ought to hold them "intention-accountable."

An analogy will clarify my position on the twin functions of responsibility, and the promise of appropriately fitting our concept of responsibility to the cases where those functions diverge. In "Pushmi-Pullyu Representations" Ruth Millikan identifies a type of representation that has "both a descriptive and a directive function, yet they are not equivalent to the mere conjunction of a pure descriptive representation and a pure directive representation" (145). Directive representations have the proper function of "guiding the mechanisms that use it so that they produce its satisfaction condition" (149). Directive representations, like GEM Anscombe's (1957) grocery list, serve their function when they make the world fit its words. Descriptive representations serve their function

when its words fit the world. Millikan provides the food call of a hen to its brood as an example of a pushmi-pullyu representation. She states, "The call is directive, saying something like, 'Come here now and eat!' But it is also a condition for proper performance of the call that there be food when the hen calls, so the call is also descriptive, saying something like, 'Here's food now'" (151). Pushmi-pullyu representations are not merely conjunctions of a descriptive plus a directive representation. They perform both functions at the same time. They are more primitive than either purely directive or purely descriptive representations. "Representations that tell only what the case is have no ultimate utility unless they combine with representations of goals, and, of course, representations that tell what to do have no utility unless they can combine with representations of facts" (152).

The dual functions of attributions of responsibility for actions are served when an individual successfully brings about an intended consequence. When a consequence is not intended, responsibility, in its two-dimensions, does not apply, and when an intention does not bring about its intended consequence, responsibility, in its two-dimensions, does not apply. The analog is as follows: When the hen's brood does not come, the directive function of the brood call fails, and if there is no food present, the descriptive function of the food call fails. Just as pushmi-pullyu representations can be divided into their directive and descriptive functions, so responsibility can be divided into an objective function and a subjective function. The concept of responsibility, thus, circumscribes two functions: 1. an objective function that identifies individuals as apt cooperators by reference to what they have done and 2. a subjective function that identifies individuals as apt cooperators by reference to the reasons for which they did what they did. Although by itself, neither function serves the function of paradigmatic attributions of

responsibility, they nevertheless serve a valuable role; and neither performs roles we can afford to overlook.

Like any evolved trait, the practice of attributing responsibility to one another and ourselves has its appropriate domain of application. It fulfils the function for which it was selected most of the time. Just as one would not deny that eyesight is a reliable means of negotiating the physical environment on grounds that it is not reliable in bizarre circumstances that interfere with its proper functioning (such as when a mirror is interposed between a perceiver and a candle which seems to her to be 10 metres away), so one should not deny that responsibility is a reliable means of navigating our social environment on grounds that it is not reliable in bizarre circumstances which interfere with its proper functioning (such as when a child runs into the road in front of one's vehicle thereby interfering with the ordinary functioning of one's motivational and behavioural system). The point, here, is that placing practices of responsibility in a natural framework enables one to see that the existence of some counterexamples to its proper functioning do not thereby make the concept itself illegitimate. However, I do not intend to argue that contexts that feature luck belong (straightforwardly) to the class of contexts in which practices of attributing responsibility fail in their proper functioning. Rather, I argue, the fact of our vulnerability to luck is instrumental to our development of an explicitly hybrid account of responsibility. Isolating the distinct features of actions and beliefs to which our habits of praising and blaming are responsive better enables the employment of the concept of responsibility for the end for which it was naturally selected: the coordination of our behaviours.

Chapter 5. The Two Necessary Facets of Responsibility

Responsibility has its origins and function in interpersonal dealings. The full-fledged notion of juridical responsibility according to which oneself and others are held accountable for deliberately wrought consequences has the function of enabling cooperative interdependent activity. The consequences that someone brings about provide inductive grounds for predicting what consequences she is likely to bring about in the future. The predictive merit of keeping a moral record of agents' acts undergirds the practice of holding people responsible for the states of affairs and events that they bring about. Fortunately, we have more bases to go on in predicting what people will do than what they have done in the past. Equally important to accurately predicting what people will do are the reasons for which they bring about particular states of affairs or events, and the reasons they marshal in support of a particular course of action. Inferring someone's future acts from her reasons for performing an act enables us to adjudicate between prospective cooperators by assessing their dependability before the success of real cooperative endeavours are at issue.

In the preceding chapter, I argue that our evolved practices of assigning and undertaking responsibility comprise a deontological dimension, responsive to the reasons for which acts are undertaken, and a consequential aspect, responsive to the states of affairs or events brought about. However, the preceding sketch of the evolutionary origins and function of responsibility needs some fleshing out. To be objectively and subjectively responsible for an act, a person must have done the act in question, and done it on purpose. What is now required is a principled means of distinguishing successful intentional acts, for which attributing responsibility is unproblematic, from human actions for which attributions of responsibility are vexed. Examples of moral luck bring to light

two categories of human behaviour for which we feel praise or blame is appropriate but which fail to meet the requirements for full-fledged attributions of responsibility.⁴⁵ First, agents whose acts bring about unintended consequences that are morally significant are subject to appraisal on the bare ground that they caused the event or state of affairs in question. Second, agents who intend to bring about a particular event or state of affairs but who are forestalled from doing so by factors external to the normal functioning of their deliberating and acting mechanisms are appropriately subject to appraisal on the bare ground that they intended to bring about the event or state of affairs.

The most salient difference between accidental causation and miscarried intentions, on the one side, and successful intentional causation, on the other, centres on the actor's lack of control. In cases of accidental causation the agent lacks control over the event or state of affairs that she (in some sense) brings about. Factors external to her deliberating and acting mechanism bring about the state of affairs for which the agent is thereafter deemed responsible, in some sense of that term. In cases of miscarried or aborted intentions, factors external to the agent's deliberating and acting mechanism prevent the agent's bringing about the state of affairs or event she intended to bring about. In such cases, the agent is still responsible, in some sense of that term, for the aborted intention.

Examples of moral luck provide reason to believe that persons who unintentionally cause morally odious events are responsible, in some sense of that term, for having done so. What underlies the remorse that the unlucky negligent driver feels for what she has done? What licenses the recrimination that would follow were the unlucky driver to fail to feel remorse for what she has done? Moral luck examples also highlight

the fact that persons who intend but fail to cause morally repugnant events are responsible, in some sense of that term, for so intending.

An analysis of the necessary and sufficient conditions for paradigmatic cases of responsibility ascription, and the relevant differences between those paradigmatic cases and the more problematic cases of accidental causation and aborted intentions will provide the conceptual wherewithal to appropriately deal with the ways that luck influences responsibility. In providing an analysis of paradigmatic responsibility in terms of necessary and sufficient conditions, I do not deny, but, rather, hope to shed some light on the relationship between central and outlying cases of responsibility. Central cases will non-problematically meet the requisite criteria for full-fledged responsibility; less central cases will meet conditions that resemble, to varying degrees, those criteria. The appropriateness of attributing responsibility in such cases will be a matter of the degree to which the criteria met resemble those met in more central cases.

The main impetus, however, for the diagnosis of responsibility is to illuminate the relationship between cases involving moral luck, such as accidental causation and aborted intentions, and cases of full-fledged responsibility. Where luck plays a determinative role, the appropriateness of attributing responsibility is not a matter of degree. In meeting only one of the criteria for full-fledged responsibility, cases of moral luck activate distinctively moral responses both on the part of the agent and on the part of those who judge her. Outlining the criteria that activate such responses and their relationship to the criteria for full-fledged responsibility will explain and validate moral responses to our impure agency. In so doing, I hope to settle a question that threatens the very coherence of norms of moral appraisal.

In the following, I shall endeavour to provide an analysis of responsibility which answers three questions: (1) What are the conditions on successful intentional acts that legitimate attributions of responsibility for those acts? (2) What are the conditions on intending to bring about states of affairs or events, without actually doing so, that legitimate attributions of responsibility for those intentions? And (3) What are the conditions on bringing about states of affairs or events unintentionally that legitimate attributions of responsibility for producing those states of affairs or events? The outcome of this analysis will, then, be applied to the various cases of moral luck and epistemic luck.

5.1. Responsibility and Successful Intentional Acts

What are the conditions on successful intentional acts that legitimate attributions of responsibility for those acts? Informed by the free will and determinism debate, philosophical literature on responsibility has been dominated by a concern with freedom. All standard positions in that debate, hard determinism, libertarianism, and compatibilism, maintain that if a person is to be legitimately held responsible for an act, then that person must have performed the act in question freely. This rare point of agreement among disputants in the free will and determinism debate is precisely what is called into question by the problems of moral and epistemic luck. Moral and epistemic luck phenomena pit our entrenched practices of holding people responsible for things they do not do freely, because they do not control them, against a deep commitment to holding people responsible only for what they do freely. This tension is complicated by differing views about what I will call the negative condition on freedom. In fact, the fulcrum of the dispute among the three principal positions on free will and determinism is the negative condition on freedom: If one is causally determined to perform an act then one cannot be

said to have freely performed that act. Hard determinists and libertarians uphold this negative condition, maintaining that freedom from causal determinism is a necessary condition for free acts, while compatibilists insist that one can be both free to do something and causally determined to do that thing.

Important as the central dispute over the negative condition is, mere freedom from causal determinism is a hopelessly impoverished condition for free acts. For compatibilists, freedom from causal determinism is neither necessary nor sufficient for free acts. They, therefore, endeavour to outline positive conditions for free acts that are compatible with those acts being causally determined. And, while incompatibilists believe freedom from causal determinism to be necessary for free acts, it is clearly not sufficient. Random acts are, in Richard Taylor's words, "erratic and jerking... without any rhyme or reason" (1999, 343). In order to act freely, persons must also have some power to choose what they do. Both compatibilists and incompatibilists, if they are to propound a tenable view of action, require a positive account of the conditions of free agency. This positive account must enable us to attribute actions to agents in a way that supports the claim of responsibility.

When we are interested in free actions, we are particularly, if not exclusively interested in rational actions; actions performed for reasons. We want to show that those actions are free. Whether or not the mechanism or process that issues in an act is reasons-responsive, therefore, seems a reasonable starting point for a positive condition for that act being free.⁴⁶

In *Responsibility and Control*, Martin Fischer and Mark Ravizza investigate the conditions for legitimate attributions of responsibility. They contend that control is the freedom-relevant condition for responsibility.⁴⁷ Although their discussion focuses on the

compatibility of responsibility with causal determinism, it is their claim that the kind of control necessary for freedom can be captured in terms of the reasons-responsiveness of those mechanisms that result in acts which interests me here.

For Fischer and Ravizza, what distinguishes agents who are responsible for their actions from agents who are not is that the former exhibit a kind of control over their actions, while the latter do not. For reasons expounded by Harry Frankfurt (1969), Fischer and Ravizza deny that the sense of control an agent possesses if she is able to do otherwise is requisite for responsibility. Frankfurt's examples feature agents who are responsible for their actions despite the fact that they could not have done other than what they actually do. Suppose Paul learns of a dearth of his blood type at the local blood bank, and, therefore, decides to donate blood. Unbeknownst to Paul, his neighbour, Mary, has implanted a device in his brain. Should Paul's resolve weaken, Mary would activate the device and make Paul donate blood. Paul's resolve does not waver, and he donates blood. It seems clear that Paul is responsible for his act; we can justly praise Paul for donating blood, even though he could not have done otherwise. It follows that being able to do otherwise cannot be necessary for legitimate attributions of responsibility.

Part of the reason that Paul's act is praiseworthy is that he seems to have been in control of what he did. If Paul had changed his mind, and Mary had activated the device and made Paul donate blood, then Paul would not have been in control of his actions. And he clearly would not have been responsible for donating blood. Fischer and Ravizza hone in on the sense of control that Paul exhibits as the ground of our attributing responsibility to him. They call that sense of control "guidance control" in contrast to "regulative control," which involves the ability to have done otherwise. They elucidate the difference between guidance and regulative control by reference to Sally's control

over her act: turning her car to the right. Sally exhibits guidance control while driving when she decides to turn her car to the right, and guides her car right on the basis of that intention. She plainly controls her car's movements in a way that pedestrians and her passengers do not. Sally exhibits regulative control when, in addition to being in guidance control of turning her car to the right, she could have done otherwise. That is, Sally is able to form the intention to turn her car to the left rather than the right, and were she to do so, she would turn the car to the left. Fischer and Ravizza believe that guidance control suffices as the control requirement necessary for responsibility. What, then, are the criteria for guidance control such that persons who have guidance control over their actions are responsible for those actions? Fischer and Ravizza identify two aspects of guidance control that, together, differentiate those actions for which agents are responsible from those for which they are not responsible. An agent exhibits guidance control over an action insofar as 1) the mechanism that actually issues in the action is her own, and 2) that mechanism is reasons-responsive.

The human deliberative mechanism is responsive to reasons when it operates in normal, unimpaired way. The mechanism whereby an agent brings about a result is reasons-responsive when three conditions are met. First, there must be a tight fit between the reasons there are for performing an act, and the reasons the agent has for performing the act. Second, there must be a tight fit between the agent's reasons for performing the act and his or her choice to perform the act. And, third, there must be a tight fit between the agent's choice to perform the act and the agent's performance. Fischer and Ravizza flesh out the notion of a tight fit in counterfactual terms. If an actually operative kind of mechanism, K, issues in an act, it is reasons-responsive where "if K were to operate and

there were sufficient reason to do otherwise, the agent would *recognize* the sufficient reason to do otherwise, and thus *choose* to do otherwise and *do* otherwise” (41).

The deliberative and performative mechanism that actually results in Paul’s action, donating blood, would be strongly reasons responsive if: (1) Paul is able to apprehend what reasons there are for donating blood, and those reasons are modularly motivational; that is, what reasons there are become Paul’s reasons.⁴⁸ Thus, when Paul learns that there is a shortage of his blood type at the local blood bank, that fact, together with his other beliefs and motivations, become his reasons for wanting to donate blood. (2) Paul would choose to donate blood, given his reasons for wanting to donate blood. (3) Paul would donate blood if he chose to donate blood.

Now, Fisher and Ravizza believe that reason-responsiveness, as outlined, is too strong a requirement for moral responsibility, and hence cannot be necessary for moral responsibility. To see why, consider the case where Paul does as he does, but where, were there a sufficient reason for him not to do so, say his phobia of needles, Paul would have staunchly carried out his blood donation. There would have been a sufficient reason for Paul to refrain from donating blood, Paul would have recognized that sufficient reason, but he would have donated blood all the same. Here, Paul is responsible for donating blood, although the mechanism whereby he does so is not strongly reasons-responsive; if there were a sufficient reason for Paul not to choose to donate blood, he would nevertheless choose to do so.

Paul may also fail to be aware of, or fail to lucidly appraise, what reasons there are for not donating blood while still being responsible for his act. If, for example, Paul’s information was incorrect, and there were no shortage of his blood type at the local blood bank, Paul’s deliberative mechanism would fail to be reason-responsive. But Paul would

still be responsible for donating blood. Moreover, if, for reasons not beyond Paul's ken Paul's child would be psychologically scarred by witnessing her father donating blood Paul's failure to understand the effect of his action would be a failure in his responsiveness to what reasons there are. We would not, I think, forbear from calling Paul responsible for donating blood on those grounds.

The third respect in which an agent's deliberative mechanism may fail to be reasons-responsive without impugning that agent's responsibility is the case where, were the agent to choose to perform otherwise, the agent would be unable to do so. Such cases conform to the format of Frankfurt's examples. In the alternative sequence, Paul is unable to translate his decision not to donate blood into the 'act' of not donating blood because his neighbour, Mary, would activate the device she implanted in his brain and make him donate blood.

An agent's deliberative action-producing mechanism may be insufficiently receptive to reasons, insufficiently reactive to reasons, or it may fail, in the counterfactual sequence, to produce the chosen act, without undermining the agent's responsibility. Fischer and Ravizza, therefore, set out to redeem reason-responsiveness by weakening the counterfactual fit of the agent's reason-receptivity, reason-reactivity, and her ability to translate her reason-reactive choice into action. They hold that the weak reasons-responsiveness of the agent's deliberative act-producing mechanism is the guidance-control requirement that is necessary for legitimately attributing responsibility to an agent.

Weak reasons-responsiveness is satisfied when there is some possible world in which there is a "sufficient reason to do otherwise, the agent recognizes this reason, and the agent does otherwise" (44). Accordingly, agents may act against some good reasons

without thereby excusing themselves from responsibility for their acts. However, in order to exhibit guidance control, an agent's deliberative acting mechanism must be responsive to some sufficient reasons. Paul would fail to exhibit guidance control over his actual act of donating blood if (1) there were no possible scenario in which he would be receptive to the reasons for him not to donate blood; or (2) there were no possible scenario in which Paul would choose according to the reasons for him not to donate blood; or (3) there were no possible scenario in which Paul would act according to his choice.

Suppose Paul does just as he does, but if he were apprised of the fact that he is a hemophiliac, and would surely die were he to donate blood, his failure to be receptive to such a reason, or any other sufficient reason, would indicate that he lacks guidance control over his action. Similarly, if in every possible world, he would consider his imminent death to be a sufficient reason not to donate blood, but choose to donate blood all the same, this would indicate that he lacks guidance control. Finally, if in every possible world, Paul were to choose not to donate blood, but he would be unable to forestall his donation of blood, he would likewise lack guidance control.

This weak reasons-responsiveness, however, is too weak. Firstly, it allows for bizarre patterns of responsiveness to reasons. That is, it allows for Paul to be responsive to his death as a sufficient reason not to donate blood, but fail to be responsive to his excruciatingly painful death. Secondly, as Michael Hymers notes weak reasons-responsiveness "will almost always be met because it requires only that it be logically possible that the agent should modify her behaviour in response to the recognition of some sufficient reason."⁴⁹ Two further refinements of guidance control are in order. The first refinement restricts the scope of counterfactual responsiveness. Moderate reasons

responsiveness applies to ‘sufficiently close possible worlds’ rather than all possible worlds. The second refinement restricts the pattern of counterfactual responsiveness.

To restrict the pattern of counterfactual responsiveness, moderate reasons responsiveness strengthens the requirements set out by weak reasons-responsiveness by invoking two sub-capacities. In David Zimmerman’s terms, moderate reasons-responsiveness comprises “a robust sub-capacity for reasons-*receptivity*, i.e., the capacity ‘regularly’ to *recognize* those factual considerations which count as reasons for action, and a somewhat less robust sub-capacity for reasons-*reactivity*, i.e., the capacity to *transform* such recognition into action” (2002, 204). Accordingly, 1) the structure of an agent’s receptivity to what reasons there are must exhibit “a pattern of actual and hypothetical recognition of reasons that is understandable and minimally grounded in reality” (Fischer and Ravizza, 76). So, for example, if Paul’s deliberative mechanism were receptive to the reason of his death by hemophilia, he ought also to be receptive to the reason of his death by contracting AIDS. Barring some reasonable explanation, we would think that something has gone dreadfully awry were Paul receptive to his death but not his excruciatingly painful death, as a sufficient reason not to donate blood. Moreover, 2) Given the understandable pattern of reasons that agents take up as their reasons, in actual and hypothetical cases, agents are to act according to that pattern.

Publicly shareable norms regulate both the moral appraisal of an agent’s actions and the account of her actions that she provides to others. Guidance control is exhibited by agents whose deliberative and acting mechanism is receptive to an understandable pattern of reasons that is grounded in reality, is reactive to those reasons, and acts according to its reaction to those reasons. For the time being, I intend to employ guidance control as the control condition on successful acts for which agents can reasonably be assessed as

responsible. That is, to the extent that we demand that the acts for which agents are fully responsible be their reasons-responsive acts, Fischer and Ravizza's account of what it means to be reasons-responsive offers a useful starting point from which to assess the hybrid notion of responsibility I shall articulate. To anticipate, I shall argue that an agent A is fully responsible for producing a state of affairs or event (S/E), briefly causing S/E, just in case

- 1) A is a participating member of a community in which causing S/E is norm governed.
- 2) The mechanism whereby A's bodily movement causes S/E is reasons-responsive and reasons-reactive: there are reasons for causing S/E, and A takes up those reasons, according to an understandable pattern, as her reasons to cause S/E, and given her reasons to cause S/E, A chooses to cause S/E.
- 3) A's bodily movement is causally efficacious in bringing about S/E: A's bodily movement makes the difference between the ordinary course of events and S/E.

Both the reasons-responsiveness and the reasons-reactiveness of A's deliberative-acting mechanism meet the counterfactual conditions that there is some sufficiently close possible world in which, if A's deliberative-acting mechanism were to operate, and there were a sufficient reason for A not to cause S/E, and that reason fits into an understandable pattern of actual and hypothetical reasons-responsiveness, A would take up that reason, and so A would not choose to cause S/E, and would not cause S/E.

5.2. Responsibility and Unfulfilled Intentions

What are the conditions on intending to bring about states of affairs or events, without actually doing so, that legitimate attributions of responsibility for that intention? When an agent's deliberative mechanism is reasons-responsive and reasons-reactive, but fails to

manifest the result chosen according to those reasons, the agent may nevertheless be held accountable for intending to bring about the state of affairs or event specified by the reasons to which she is reactive and responsive. That is, when the agent is able to lucidly apprehend and respond to a reasonable pattern of reasons, and chooses to act on the basis of that pattern, that agent can legitimately be held accountable for her intention, even when the agent fails to bring about the state of affairs or event specified in that intention. A is morally accountable for intending to bring about a norm-governed state of affairs or event (S/E) when:

- 1) A is a participating member of a community in which causing S/E is norm-governed.⁵⁰
- 2) Causing S/E is reasons-responsive and reasons-reactive for A: there are reasons for causing S/E, and A takes up those reasons, according to an understandable pattern, as her reasons to cause S/E, and, given her reasons to cause S/E, A chooses to cause S/E.
- 3) Given her choice to cause S/E, A tries, but fails to cause S/E.

Both the reasons-responsiveness and the reasons-reactiveness of A's deliberative-acting mechanism meet the counterfactual conditions that there is some sufficiently close possible world in which, if the deliberative-acting mechanism were to operate, and there is sufficient reason for A not to cause S/E, and that reason fits into an understandable pattern of actual and hypothetical reasons-responsiveness, A takes up that reason, and so A does not choose to cause S/E.

What differentiates this case from a case in which responsibility for the result S/E is non-problematically attributed to A is that A fails to bring about S/E. A's deliberative and acting mechanism, though responsive and reactive to A's reasons for causing S/E,

fails to be causally efficacious in making it the case that S/E. So long as A's failure to cause S/E is beyond her control, she is accountable for intending to cause S/E. Such failure may take two forms: A's deliberative and acting mechanism produces bodily behaviour in pursuit of S/E which fails to cause S/E, or A's deliberative and acting mechanism fails to manifest bodily behaviour in pursuit of S/E.⁵¹

How does being responsive and reactive to reasons for bringing about a state of affairs or event, without actually succeeding to bring about that state of affairs or event, legitimate particularly moral responses? There is little doubt that we feel tempted to hold persons who do so responsible for their intentions. In fact, driven by the intuition to hold individuals responsible for their intentions, the internalist about moral responsibility places all praiseworthiness and blameworthiness under the behest of agents' intentions. A more earthly perspective on our practices of ascribing and undertaking responsibility understands the rationale for holding persons to account, morally, for their miscarried intentions in terms of the insight this provides into what those persons are liable to do in the future. The moral verdict that persons are legitimately blameworthy or praiseworthy on grounds of their intentions alone has its "legal counterpart in the punishment of abortive attempts to commit crimes" (Hart and Honoré, 59). The analogy between legal repercussions for attempted and completed crimes can be further pressed to highlight the real difference, in terms of the moral responsibility agents bear, between their unfulfilled and their fulfilled intentions. Because they indicate agents' future actions and their unfitness for interdependent living, genuine but unfulfilled evil intentions warrant indignation, reproach, and censure. Fulfilled evil intentions warrant all this as well as indignation, reproach, and censure for actually doing the intended evil.

5.3. Responsibility and Accidents

5.3.1. What Justifies Holding Persons Responsible for their Unintentional Acts?

What are the conditions on bringing about states of affairs or events unintentionally that legitimate attributions of responsibility for those states of affairs or events? I have already provided a sketch of the rationale for holding agents responsible for the states of affairs or events that they bring about involuntarily: One is responsible for what one actually does—even if what one actually does is not within one's control—because what one actually does provides an inductive indicator of one's aptness for future cooperative endeavours. A history of bad but unintentional results indicates either an incomplete grasp of the norm or an underlying ineptitude at causing intended states of affairs.⁵² One can redress the appearance that one has an incomplete grasp of norms, or that one is particularly inept, by identifying the state of affairs as accidental. Nevertheless, the very fact that causing a norm-governed state of affairs accidentally demands this redress indicates that it is important for attributing and undertaking responsibilities.

Normative enforcement—the practice of holding agents responsible for the things they actually do—has the evolutionary function of coordinating cooperative group behaviours. In order to serve this function, group members must have some means to teach and learn the normative structures of their group, they must have some means to gauge one another's commitment to the group's norms, and they must have some means to gauge one another's aptness for particular cooperative endeavours. The practice of attributing and accepting responsibility for the consequences of persons' actual behaviour provides agents with the wherewithal to teach norms, to display their commitment to norms and to demonstrate their trustworthiness as cooperators. Since the function of holding persons responsible is to coordinate group behaviours, if someone's actual

behaviour is out of synch with those of her group, that person is not a good bet for coordinated cooperative living. The question of whether the agent intends to coordinate her behaviour with that of her cohorts is a secondary question. To perform its evolutionary function, our practices of attributing responsibility must be responsive to the bare fact that her behaviours tend not to coordinate with those of her group members. If someone consistently defects on her would-be cooperators, even if doing so is something she does not control, that person is not likely to be a dependable participant in cooperative endeavours. And if someone consistently violates the group's norms, that person is not a good bet for group membership.

It is a consequence of the hybrid view on offer that the Kantian claim that 'ought' implies 'can' is not always true, and, therefore, it is not the case that 'cannot' implies 'not ought.' Although, for our moral norms to have evolved, it must be the case that 'ought' almost always implies 'can', there are exceptions. My contention is that the Kantian tenet that we have an obligation, moral or epistemic, to do or believe x only if doing or believing x is something over which we have voluntary control is wrong. There exist some obligations for which the inability to comply with the obligation does not defeat the fact that one is obliged. For example, a legal obligation to repay a debt is not defeated by being unable to do so.⁵³ Moral and epistemic obligation to the other members of one's epistemic and moral community to do the right thing and believe the right thing, I think, are on a par with such legal obligation. One's actual inability to fulfill one's obligation does not defeat that obligation.

Peter van Inwagen claims that "almost all philosophers agree that a necessary condition for holding an agent responsible for an act is believing that the agent *could have* refrained from performing the act" (1975, p. 189). I dissent from this near unanimity.

Being responsible for an act or a belief does not imply that one could have failed to perform that act or maintained that belief. What is essential to determining whether an agent is responsible is not whether she could have done or believed other than what she actually did or believed. Rather, to determine whether someone is responsible is to determine whether the agent performed the act or maintained the belief, and the reasons for which the agent performed the act or affirmed the belief.

In *Elbow Room*, and *Freedom Evolves* Daniel Dennett argues against the Kantian tenet that 'ought' implies 'can.' For Dennett, when we attribute responsibility, "it is seldom that we even *seem* to care whether or not a person could have done otherwise" (133).

"Here I stand," Luther said. "I can do no other." Luther claimed that he could do no other, that his conscience made it *impossible* for him to recant. He might, of course, have been wrong, or have been deliberately overstating the truth. But even if he was—perhaps especially if he was—his declaration is testimony to the fact that we simply do not exempt someone from blame or praise for an act because we think he could do no other. Whatever Luther was doing, he was not trying to duck responsibility. (133)

What we care about when assessing whether or not someone is responsible for what he or she did are that she did what she did, and the reasons that prompted the act, not whether he or she could have done otherwise.

Agents, according to the foregoing, are responsible, at least sometimes, for things that they do not control. When someone intends to do something evil, but is debarred from doing so for reasons she neither controls nor endorses, that agent's moral standing is assailable on grounds of the evil intention alone. Although the fact that she has not

brought about the state of affairs specified by her reasons for acting and consequent proximal intention to act is beyond the agent's control, the intention itself is, at least sometimes, within the agent's guidance control. It is responsive to an understandable pattern of reasons that the agent takes up, and chooses to act in accordance with. Thus, at least where intentions that are in character for the agent are concerned, attributions of responsibility do not violate the internalist's injunction against holding persons responsible for things that outstrip their control. That injunction is violated by practices that hold agents responsible for states of affairs or events that they bring about accidentally. What grounds such practices? Externalists such as Walker and Card embed forward-looking conceptions of responsibility, such as an agent's taking responsibility for harms she causes accidentally, in the broader social practice of maintaining caring relationships. Though I commend their approach, the question I wish to address involves the relationship between the agent and the event or state of affairs that she causes such that it is incumbent upon her to take responsibility. What is it about an agent causing harm accidentally that makes it reasonable to expect her to make reparations, and that licenses moral censure and moral reproach if she fails to do so? What grounds the practice of holding an agent responsible for her failure to take responsibility for having unwittingly caused harm? What governs such responses is causing a morally significant state of affairs or event, and what justifies such responses is the role that they play in cementing relations of interdependence upon which individual lives and the communities of which they are members depend.

Like externalists, I view moral responsibility as essentially social. It is from the perspective of interdependent agents whose very agency is contingent upon a world that is not of their making that the demand for agents to be responsible for their accidental

actions makes sense. Because relations of interdependence and trust necessary to human thriving are cemented by practices of undertaking and ascribing responsibility, indications of agents' unreliability, even their inadvertent unreliability, matter for moral responsibility. They require the redress afforded by practices of undertaking and ascribing moral responsibility. Having accidentally caused harm, thus, demands that the agent discharge the moral onus to display her continued commitment to the moral norms of her moral community.

We can judge from the fact that apology and redress are required from the accidental wrongdoer that she is morally responsible, in a non-trivial sense, for the accidental harms she causes. This conclusion is supported by the intuition that the unlucky negligent driver warrants more blame than the lucky negligent driver. It is also supported by the existence of distinctly moral virtues that Margaret Walker outlines for coping with our vulnerability to luck. These externalist insights into responsibility show that when an agent's causing a state of affairs is norm-governed, this invites moral responsibility even when such causal efficacy is divorced from any intention to perpetuate the state of affairs in question. Responsibilities for their accidental harms accrue to agents qua participants in moral communities. The linchpin that distinguishes lucky from unlucky drivers is the attributability of a morally significant event, a child's death, to the unlucky driver's accidental causal efficacy. Accordingly, it is reasonable to suppose that responsibility for causing a morally significant S/E applies not just to reasons-responsive and reasons-reactive behaviours that have unintended consequences, but also to completely unintentional behaviours that have (now obviously) unintended morally significant consequences.

Once we take seriously the externalist intuitions that press us to blame the unlucky negligent driver more than the lucky negligent driver, to praise the successful rescuer more than the unsuccessful rescuer, and that compel us to applaud the exercise of grace, integrity, and lucidity in the face of accidentally caused harms, these externalist intuitions about moral responsibility can be appropriately categorized as responsive to agents' causal efficacy. In light of this, the externalist intuition to count factors beyond agents' control in assessing their moral responsibility can be seen to apply just as legitimately to causing completely unintentional harm as they do to accidentally causing harm by means of intentional behaviour.

Though the conclusion is controversial, there are good reasons to believe that persons are morally responsible for unintentionally caused states of affairs/events. So long as they are active participants in the moral community in which S/E is norm governed, bringing about S/E, even if wholly unintentionally, carries with it particularly moral obligations. Those obligations—to make reparations where possible, to feel remorse, to apologize, to exhibit Walker's virtues of impure agency, and so on—are possible and necessary because the agent is the cause of the state of affairs or event.

Unlike skeptics such as Nagel and Williams, unlike internalists such as Kant and Rosebury, and unlike externalists such as Walker and Card, I recognize no salient difference between the case where someone acts with the intention of bringing about one thing and brings about another wholly unintended thing, and someone who brings about the second thing wholly unintentionally. Once one divorces the causal role played by agents and the intentions with which they cause what they cause, the causal role, alone, can be assessed as a dimension of responsibility.

If I am cautiously and attentively driving my well-maintained car and suffer an unforeseeable seizure which causes me to run over and kill a child, I bear responsibility for causing that child's death. Being the cause of that child's death makes a moral demand upon me at least in the following senses: Firstly, if I don't feel deep remorse, there is something seriously wrong with me as a member of moral community. Walker's vilification of a pure agent's reaction to having unintentionally caused something awful seems as appropriate here as it would be in the case where an agent causes the death of a child in pursuit of another intentional end. Secondly, I have an obligation to make reparations where possible, to express my deep and abiding regret (and that it was unintentional!), and to try to alleviate the burden of grief my completely accidental bodily movement produced.

As I have argued, what grounds these moral obligations is that I caused the child's death. Remorse, self-recrimination, and guilt are appropriate moral sentiments in such instances, and their absence would be strong grounds for jettisoning an agent who failed to experience such sentiments from the moral community. This is not because an absence of remorse indicates that I did it on purpose after all (to forestall an objection from the internalist camp). It is, rather, because having killed a child is, by itself, morally significant. If an agent is morally responsible for causing a child's death, why does moral blame not seem appropriate? We rely on the agent, qua moral agent, to blame herself. If she fails to do so, though, she's an appropriate target for moral censure.

Ascribing responsibility to agents for harms caused by their bare bodily movements is admittedly controversial. In particular, staunch internalists about moral responsibility will be sanguine about rejecting outright, or explaining away, the intuition to hold persons to account for accidents. Nevertheless, given their relevant similarities,

those who find themselves tempted to endorse the externalist intuition in cases of moral luck, who find themselves blaming the unlucky negligent driver more than his lucky counterpart, or praising the successful rescuer more than her unlucky counterpart, should feel equal pressure to endorse the externalist intuition when the unintentional harm or benefit was brought about by completely unintentional behaviour. Not only is the claim that responsibility accrues to agents for their accidents consistent with intuitions about cases of moral luck in the literature, holding persons morally responsible for accidents is an ineluctable facet of broader quotidian practices of attributing and undertaking responsibility. If I trip and fall into someone or sneeze on someone's dinner plate, I am responsible for having done so. It is, therefore, incumbent upon me to apologize for my unintentional act. My failure to respond appropriately to my having jarred or disgusted someone might not invite moral outrage, but it would invite ill will. Finally, as I will argue below, the independent plausibility of holding persons to account morally for their bare bodily movements is justified by the fact that doing so supports and is supported by a social conception of responsibility that features interdependent agents living in a world that is not of their making.

5.3.2. How do we Attribute Unintentional Acts to Agents?

Having established that moral responsibility for accidents is defensible there remains the conceptual problem of identifying the criterion of causal agency to which such responsibility attributions are legitimate responses. Praise or blame of an agent for the events she accidentally causes is grounded in practices of gauging that agent's worthiness for cooperative endeavours and communal living. Those practices, of course, depend upon our ability to tell whether an event or state of affairs has been brought about by an agent. However, the question of how to identify an agent's act as the cause of an event or

state of affairs is problematic. The problem is that as soon as we begin to speak of an agent's causing an event accidentally, we lose track of the agency according to which we ordinarily attribute responsibility.⁵⁴ What is required is a means of describing unintentional human causation, states of affairs or events that are caused by humans, though *completely unintentionally*. Morally significant states of affairs or events that are caused by humans' 'bare bodily movement' seems an appropriate candidate.⁵⁵ Though, as I am at pains to emphasize, events caused by a human's bare bodily movement lack something central to human agency; they are not caused by the agent's deliberative-acting mechanisms in accordance with the agents' apprehension and adoption of reasons and consequent choices.

Standard means of attributing events or states of affairs to agents are intertwined with a theory of action that deems purposiveness as the central criterion by which we make such attributions. For example, according to the causal theory of action expounded by Alfred Mele, "a certain event occurring at *t* is Norm's raising his hand at *t*—an action—partly in virtue of its having been produced 'in the right way' by certain mental items (or their neural realizations)" (2000, 288). The target mental item (or its neural realizations) is Norm's intention to raise his hand. For Mele, agents' actions are like money or sunburns. If a piece of paper is not produced in the right way by the treasury department, no matter how much it might seem like money, it will not be; if a red and blistering patch of skin is not caused by the sun, no matter how much it resembles a sunburn, it is not a sunburn. To be an agent's act *A*, for Mele, is to be produced in the right sort of way; it is to be causally initiated by the agent's proximal intention to bring *A* about. Suppose the act in question is to throw a rock currently held in the agent's hand. The acquisition of a proximal intention to throw the rock would initiate a command signal

to the agent's motor cortex where it would be transformed into specific fine-grained executive signals to specific muscles. The signals would command the muscles to move in particular ways, the completion of which would ensure satisfaction of the intention to throw the rock.

On this account of acting, it would be impossible for an agent to act accidentally; to perform act A is to be caused by one's intention to perform A. This is not to say that Mele's conception of action precludes the possibility of accidents. Rather, it is to say that, for Mele, if something is an accident, it cannot properly be described as an agent's act.

This conception of action contrasts with a broader notion of action disparagingly described by Mele as the sense of acting according to which even "acids, waves and wind act" (290). But humans do 'act' in the same sense as acids waves and wind act. They do so when they act inadvertently. It is this second sense of human action, action divorced from intention, that is required to support the attribution of responsibility for accidentally caused states of affairs or events. A broader view of action which includes unintended human bodily behaviour is espoused by George Wilson. For Wilson, all human bodily movements are acts; when a human moves, a human acts. He lists the following bodily movements among acts of agents: "A man performs a convulsive and spasmodic movement when he clutches and cannot loose a live electric wire, and someone undergoing an epileptic seizure may perform a series of wild and wholly uncontrollable movements" (1999, 49). Wilson does not deny the importance of intentional actions. For an agent to intentionally or consciously perform an act, *b*, "the mechanisms of the agent's bodily control, as exercised in the performance of *b*, [must be] systematically and selectively responsive to the agent's perception of her environment" (1989, 146).

However, Wilson insists that, even in the absence of responsiveness to reasons, human bodily movement constitute acts. Sentient or conscious actions simply form a subset of the larger category of human actions comprising all human bodily movements.

In order to justify the intuition that merely causing an event is something to which our practices of attributing responsibility is responsive, we require an acceptable account of what it is for a person to act inadvertently. Following Wilson, we may identify human acts with human bodily movement. Intentional bodily movement, which is reasons-responsive and reasons-reactive bodily movement, forms a subclass of the larger class of bodily movement. Unintentional bodily movement is simply bodily movement that is neither reasons-responsive nor reasons-reactive.

This view of action severs the conceptual entanglement of agent causation with agent intention in a way that enables identifying an agent as the cause of a state of affairs or event without imputing to that agent any intention to bring about that state of affairs or event. The problem is that, in severing the connection between human causation and reasons for acting, this view of agency seems to also sever the connection between human causation and responsibility for what is caused. As his examples make clear, bare human bodily movement is as lacking in intentionality and control as are the 'acts' of inanimate forces. How, then, can such acts support attributions of responsibility to humans while refusing to sanction similar attributions of responsibility to inanimate forces? The conception of what it is for a person to cause an event, in the same way as inanimate objects cause events, must be commensurate to the fact that humans are held accountable for the occurrences they eventuate while inanimate objects are not.

5.3.3. Responsibility and Bare Bodily Movement

How can we identify an agent's bare bodily movement as the cause of an event or state of affairs? Before addressing the reasons for which we might legitimately hold humans to account for what they cause by their bare bodily movements, we must answer the question of how to isolate *the cause* of an event or state of affairs. Leaving aside Humean worries about reifying causal forces, identifying the cause of an event or state of affairs is, by itself, problematic. Any investigation into the cause of a state of affairs or event will show that innumerable conditions must exist for that event or state of affairs to come about. How can we isolate one of those conditions as the cause of the ensuing event or state of affairs? In *Causation in the Law*, H. L. A. Hart and A. M. Honoré pose this quandary as follows: "Our causal generalizations inform us only that an occurrence of a given kind regularly follows when a complex set of conditions is satisfied. So when we identify single events as causes it appears that we choose one element from such a set, although each of the members of the set is equally required for the production of the effect" (16). Another way to articulate this problem is to employ J. L. Mackie's (1974) conception of a cause as an 'INUS' condition for an effect. On Mackie's view, a cause is an *insufficient* but *non-redundant* part of an *unnecessary* but *sufficient* condition for an effect. Take the causal claim "if this match is struck, it will ignite." Striking the match, by itself, is not a sufficient condition for it to ignite. Other conditions, such as there being oxygen in the room, that the match not be wet, together with the condition that the match is struck, form the sufficient condition for the match to ignite. Striking the match is a necessary part of the conjunction of conditions that form the sufficient condition for the match to ignite.

For Hart and Honoré, the process of answering how, in the absence of the agent's intention to bring about a state of affairs, we identify the agent as the responsibility-relevant cause of the state of affairs in question must begin by answering the broader question of how to distinguish between the cause of an occurrence and mere background conditions. Then when a human act, conceived broadly so as to include all human bodily movement, is the cause of an event, rather than a mere condition necessary for the event, we shall say that the human was the cause of the event.

Hart and Honoré identify two contrasts as the fulcrum of the distinction between the cause of an event and mere conditions that are necessary for that event. The first is the contrast between what is normal and what is abnormal. The second is the contrast between voluntary human action and all other conditions. When we seek to identify the cause of a particular occurrence, in contrast with cases where we seek to explain the causal underpinnings of types of occurrences, it is the abnormality of the occurrence that compels our inquiry. For example, when we seek to identify the cause of a house fire, we do so because it is abnormal for houses to burst aflame. The difference between normal and abnormal enables a ready distinction between mere conditions for the particular occurrence and the cause of that occurrence. Those factors that are present in both normal and abnormal cases are mere conditions; they are "present as part of the usual state or mode of operation of the thing under inquiry" (32). Thus, the presence of oxygen, while a necessary condition for the house fire, is a mere condition. The cause, by contrast, is what makes the difference between the house remaining in its normal state and it burning. When human activity makes the difference between things going on as normal and abnormal occurrences, then that human activity is the cause of the abnormal occurrence in question.

As Hart and Honoré note, what is normal and what is abnormal is relative to the context of inquiry in two ways. First, the presence of oxygen, though normal, may be identified as the cause of a particular occurrence, as when the proper functioning of a dangerous chemical experiment requires the exclusion of oxygen. When requisite precautions to exclude oxygen are not undertaken, the presence of oxygen is not a feature common both to the ensuing disaster and the normal functioning of the experiment. In such contexts, the presence of oxygen may properly be identified as the cause, not a mere condition, of the ensuing fire. Second, the interests of the inquirer may serve to distinguish what is the cause and what is merely a condition in one and the same case. Hart and Honoré's example features a man who suffers from stomach ulcers. His wife identifies eating parsnips as the cause of his indigestion, while his doctor identifies his ulcerated condition as the cause of his indigestion, and his ingestion of parsnips as a mere condition. Hart and Honoré maintain that this is no indication that the distinction between cause and condition is arbitrary. Rather, it is an artefact of the reason for which the inquiry is undertaken. What differentiates the doctor's diagnosis from the wife's is that the doctor seeks to answer "What gave this man indigestion when other men do not get it?"; for him, what the man ate...is a mere condition—part of the normal conditions of most men's lives" (34). By contrast, the man's wife seeks to answer 'What gave this man indigestion when this man does not usually get it?'; for her, that the man has an ulcer is a mere condition—part of the normal condition of that man's life.

Non-voluntary human acts can be causes and can be distinguished from mere conditions according to whether the involuntary act 'makes a difference' in the ordinary course of events. "What we call 'accidents' are often, though not always, occurrences of

an untoward or surprising character which have as their explanation, and so as their cause, a human action which is done unintentionally or involuntarily” (Hart and Honoré, 39).

The second contrast that differentiates causes from mere conditions is that between human voluntary action and all other conditions (including involuntary human actions). The relation between human voluntary acts and causal inquiries is noteworthy in two ways. First, suppose we find a lit match was the cause of a house fire. If we find, further, that an arsonist deliberately lighted said match, we identify the arsonist’s action as the cause of the fire. While we do not thereby deny the lighted match as the cause of the fire, we view it as the mere means by which the arsonist produced the fire. A deliberate human act is “something *to* which we trace the cause though intervening causes of other kinds” (41). Second, suppose we find that the arsonist’s greed caused him or her to light the match, we do not (ordinarily) speak of the arsonist’s greed as the cause of the fire. Although we do say that greed caused the arsonist’s act, we do not, in Hart and Honoré’s words, “trace the cause through the deliberate act” (40).

In the moral judgments of ordinary life, we have occasion to blame persons when they cause harm, and to demand that they compensate those to whom they have caused harm. There is little dispute about whether or not occurrences caused by voluntary human acts support attributions of responsibility for those occurrences. The more contentious question at issue here is whether occurrences caused by involuntary human acts, by their bare bodily movement, support attributions of responsibility for those occurrences. Following Hart and Honoré, we may identify an involuntary human act as the cause of an event when it makes the difference between the ordinary functioning of particular processes and the eventuation of an abnormal occurrence whose origination demands explanation. Someone may cause something involuntarily in two ways. First,

she might act voluntarily intending to bring about event or state of affairs S/E, but external factors interfere such that her act causes a different event or state of affairs S*/E* instead (where S*/E* is under normative control in her group). When an agent's deliberative mechanism is reasons-responsive and reasons-reactive but fails to occasion the consequence chosen according to those reasons and occasions instead another consequence, we may say that the agent caused the actual consequence involuntarily. Sally's decision to throw a snowball at a signpost might be reasons-responsive and reasons-reactive. She is then in control of her act. If, however, a child runs in front of the signpost, and Sally's snowball hits the child rather than the signpost, we shall say that Sally involuntarily caused the child to be hit by the snowball despite the fact that her act, throwing the snowball, was voluntary.⁵⁶ Although Sally was in control of throwing the snowball, she was not in control of the snowball hitting the child.

The second way that an agent may act involuntarily is when the agent's bare bodily movement causes an event or state of affairs that is under normative control in the agent's group. Such an act is neither reasons-responsive nor reasons-reactive. Thus, suppose Sally is holding a snowball when a gust of frigid wind causes her to shiver uncontrollably. Her arm jerks up, and her fingers release the snowball. The snowball flies through the air and hits a child. In this case, even though her act was not under her control, Sally caused the child to be hit by the snowball. In both instances, the mechanism that actually causes the child to be hit by the snowball outstrips Sally's control. There are, thus, two ways in which an agent may unintentionally cause an event or state of affairs. First, the agent's bodily movement, which makes the difference between the ordinary functioning of a process and the occurrence of the event or the state of affairs, is not under the agent's guidance control. That is, there is some sufficiently

close possible world where, if the same mechanism were to operate and there were sufficient reason to do otherwise, the agent would *not recognize* the sufficient reason to do otherwise, or would *not choose* to do otherwise for those reasons, or would *not do* otherwise. Second, the agent's bodily movement, which makes the difference between the ordinary functioning of a process and the occurrence of the event or the state of affairs, is under the agent's guidance control, but the event or state of affairs caused by the act is not the event or state of affairs specified in the agent's choice to act according to particular reasons. The agent's bodily movement was under her guidance control, but what it engendered was not.

We are now in a position to outline the conditions on legitimately holding persons responsible for the events or states of affairs they cause unintentionally. An agent A is consequence-responsible for state of affairs or event S/E just in case

1. A is a participating member of a community for whom causing S/E is norm-governed⁵⁷
2. A's bodily movement is causally efficacious in bringing about S/E: A's bodily movement made the difference between the ordinary course of events and S/E.

There remains, however, the question about what grounds attributions of responsibility to humans for accidentally bringing about a norm-governed state of affairs or event in a way that excludes inanimate objects or forces from inviting moral responsibility for unintentionally bringing about norm-governed states of affairs or events.

5.3.4. Human Causation and Inanimate Causation

What differentiates an unintentional causation by agents from inanimate causation such that the former, but not the latter, can legitimately be held accountable for what is caused? To say that someone is the appropriate subject of moral appraisal just because she caused something requires a principled distinction between a state of affairs brought about by an agent and a state of affairs brought about by a non-agent. We hold agents accountable for the occurrences they cause, but we do not (or at least not legitimately) hold inanimate objects or forces accountable for the occurrences they cause. On the view of acting with which we are operating, human 'agency' seems no different from the 'agency' of inanimate objects or forces in causing events or states of affairs. And yet we hold the one responsible and not the other. What, if anything, licenses attributions of responsibility for the occurrences caused by involuntary human acts? Any positive answer to this question ought to fund an explanation for why we do not attribute responsibility to non-human causes of events or states of affairs. I will argue that what differentiates occurrences wrought by humans from occurrences wrought by other causes is that humans are able to take responsibility for their acts. That is, humans are members of groups whose conformity to a pattern of normative engagement with one another necessitates that they see their actions from the perspective of participants in that network of norms. It is in virtue of their active membership in groups whose interaction is characterized by holding one another accountable for what they do that our practices of holding persons responsible for the events or states of affairs that they cause are legitimated.

Is it appropriate to hold persons responsible for events that they have caused by accident just because they are capable of providing moral justifications for their behaviour? Does the agent-regret that accompanies having caused harm by accident

constitute arrogation of responsibility, or are such reactions legitimate? Are such responses, and the responses that they invite, sustained by the requirement that agents make their commitment to moral norms intelligible to the others in their shared community? One way to answer these questions is to look at what differentiates excuses for certain behaviours from exemptions from moral assessments. While excuses mitigate responsibility for a certain class of actions, exemptions make it inappropriate to hold the agent in question morally accountable for the consequences he or she brings about. In "Freedom and Resentment," P. F. Strawson distinguishes two types of exemptions from full-fledged moral responsibility for the consequences one causes. The first are temporary impediments to the agent's ability to employ his or her powers of reflective self-control. "[H]ypnotism, extreme stress or physical deprivation and the short-term effect of certain drugs" (155), interfere with the reasons-responsiveness of an agent's actions. However, they do not undermine her status as a participant in the normative practices of her group. The agent is a participant in the normative practices of her group insofar as she is able to take responsibility for the events she causes or intends to cause or both, and is liable to be held responsible for the events she causes or intends to cause or both. It is in virtue of her continued membership in a group characterized by responsibility attributions and responsibility acceptances that it is incumbent upon the agent who violates norms, even if accidentally, to rectify the damage (if any) done to others and to rectify the damage done to her own moral record. Citing such conditions as lack of control over one's violation of a norm is a way of expiating the wrong one has done; a way of expunging the black marks on one's moral record. By contrast, more systematic and persistent impediments to an agent's ability to exercise reflective self-control jeopardize the agent's status as a participant in practices of accepting and

attributing responsibility for violating norms. They thereby jeopardize the agent's aptness for cooperative endeavours and communal living. As Strawson notes, "[I]nsanity or mental illness, extreme youth, psychopathy, and the effects of systematic behavior control or conditioning" (155) render the agent's general bodily behaviour immune from assessment along the ordinary lines of moral appraisal.

The additional requirement on human unintentional causation that distinguishes it from inanimate causation, and which therefore supports the attribution of responsibility to persons for their accidentally caused states of affairs or events is the person's active participation in a moral community. To be a responsible moral agent is to possess the ability to understand and be moved by moral reasons. To treat an agent as responsible for the events or states of affairs that she brings about inadvertently is to treat that agent as an active and reactive member of norm-governed group. Gary Watson argues that "the social status of one who is worthy to give a moral account of her conduct (by seeking to justify it, excuse it, admit fault, or the like) is also that of one worthy to disclose her valuational system through her actions" (2003, 78). Regarding oneself as morally responsible is a precondition on an agent's being morally responsible; that is, it is a precondition on her entering moral dialogue about her actions. For Watson, we are responsible because and when our conduct is self-disclosing. As John Dewey put it, "that conduct is ourselves objectified in actions."⁵⁸ On this view, excusing, mitigating and exempting conditions indicate circumstances that prevent or qualify self-disclosure. The general trait of being morally responsible involves being worthy of a certain social standing, that of "an eligible participant in the various kinds of moral exchange, such as offering reasons, seeking excuse, begging forgiveness, and so forth.... We might understand the conditions of moral responsibility as conditions of intelligible moral

address” (Watson, 79). The appropriateness of holding someone responsible for the states of affairs she brings about, either deliberately or not, will depend on the appropriateness of holding the person to the normative demand of responding suitably to her relationship to the event or state of affairs.

Our judgments of culpability are related to whether the individuals in question are persons we are justified in including in the moral community. And our blame takes the form of a demand that those persons make reparations for what they have caused, and thereby demonstrate their possession of traits which make them liable for cooperation rather than for defection. When we hold persons responsible for events or states of affairs that they actually cause—inadvertently or advertently—we mark those events or states of affairs down on their moral records. The mark is not indelible. In fact, by blaming persons, we enable them to defend or excuse their behaviour. On this view, locutions such as “It was an accident,” “I could not have done otherwise,” and “It was out of my control” are not expressions of conditions that exempt their utterers from being held accountable for the events that they bring about. To the contrary, they are the first steps (and sometimes the last step) in a dialogue that begins with realizing that one is causally responsible for the event brought about, and whose end is to demonstrate one’s commitment to the norm one has violated, and thereby demonstrate one’s aptness for cooperative endeavours and communal living. Offering lack of control over the event one has caused is a way to repair the damage that being causally responsible for an event has done to one’s moral reputation. The appropriateness of apologizing for inadvertent acts underwrites this claim. We are sorry, and appropriately so, for things we cause accidentally. As Paul Benson argues, moral response is an essential part of morality. When we make moral judgments, we expect the agent judged to respond. “It is a

normative expectation concerning standards of response for one who is fairly held responsible” (2000, 79).

5.4. Responsibility, Justification, and Truth-Value

5.4.1. Doxastic Responsibilism and Control

The idea that knowledge, justification and belief are things for which agents can, and should be, held responsible is not new. Descartes, for example, maintains the notion of duty or obligation plays a central role in the whole doxastic enterprise. Descartes claims that:

If I abstain from giving my judgment on any thing when I do not perceive it with sufficient clearness and distinctness, it is plain that I do not act rightly.... even though I judge according to truth, this comes about only by chance, and I do not escape the blame of misusing my freedom. (1955, 176)

Despite its venerable history, the project to subsume epistemic responsibility under the same rubric as moral responsibility faces an immediate and trenchant objection: doxastic states, unlike acts, do not seem to be under the purview of the will. While we seem free to choose to turn right rather than left, to raise a hand or not, we do not seem free to choose to believe that it is raining rather than that it is not, that today is Friday and not Monday.⁵⁹

5.4.2. Indirect Doxastic Voluntarism and Epistemic Justification

In “Concepts of Epistemic Justification” William Alston argues that deontological conceptions of epistemic responsibility are untenable because, by and large, doxastic phenomena outstrip voluntary control.⁶⁰ If one lacks control over which beliefs one assents to, “then it is futile to discuss whether I am *permitted* to believe that p at t or whether I would be *irresponsible* in choosing to believe that p at t” (85). On his view, we

do not get to choose which beliefs we hold. Rather, beliefs are pressed upon us. Which propositions we believe is not a matter over which epistemic agents generally exert control.

The claim that belief states are impervious to determination by the will, therefore, requires some attention. Are our beliefs under voluntary control? To presume that they never are seems false for it commits us to the view that where we find ourselves believing something for which we have no justification, we have no recourse but to continue believing that thing. We can weigh the evidence for our beliefs, and, at least much of the time, can rescind beliefs for which we find there are no grounds.

Though we may lack direct control over which beliefs to hold, we enjoy indirect voluntary control over our beliefs. Thus, one may fulfill one's subjective epistemic obligations by ensuring that one has undergone inquiring procedures that can reasonably be expected to deliver true beliefs rather than false ones.⁶¹ One can thereby ensure that one is fulfilling one's overriding epistemic obligation to have true beliefs rather than false ones without thereby committing oneself to the implausible view that one can choose to believe whatever one wants to believe independently of the evidence that supports or undermines such belief.⁶²

This shift from direct voluntarism of beliefs as the control relevant condition on responsible acts, to indirect voluntarism as the control relevant condition on responsible beliefs, is echoed in the emphasis that pundits on epistemic luck have placed on the degree of control that agents have over the justification of their beliefs, rather than on the degree of control that agents have over whether or not they believe a particular proposition. Accordingly, the division between externalists, who licence attributions of responsibility to agents for beliefs that outstrip their control, and internalists, who deem

only what is within the agent's ken and control as warrantably inviting epistemic approval or disapproval, centres on the legitimacy of holding agents' beliefs to be justified when that justification is not internally accessible, and hence not controllable, to the agents.

In "A Critique of Internalism" Alvin Plantinga distinguishes between objective and subjective duty or rightness. He states "you are guilty or blameworthy if you fail to do your *subjective* duty, but not necessarily guilty for failing to do your *objective* duty" (1999, 62). He insists that blameworthiness attaches only to the dereliction of subjective duties. "You are guilty, or to blame, or properly subject to censure only if, as we say, you *knowingly* flout your duty" (362). But we do not ordinarily accept ignorance as an excuse for dereliction of one's objective duties. Plantinga claims that this is because there is a connection between objective and subjective duties, and that connection is forged by "our nature." Whether one fulfills one's subjective epistemic duties is generally up to the agent. By contrast, the fulfillment of one's objective epistemic duties is met with serious impediments.

On the naturalistic hybrid account, the means of fulfilling one's objective epistemic duty to have true beliefs is by means of fulfilling one's subjective epistemic duty to try to have true beliefs. This striving to have true beliefs takes the form of employing belief acquisition and belief affirmation procedures that one has grounds for believing have a high probability of generating true beliefs. This approach, however, leaves the final determination of whether one has fulfilled one's objective epistemic duty up to determination by factors beyond one's ken or control. To fulfill one's objective epistemic duty is to regulate one's belief according to its truth or falsity. As Plantinga notes, "Merely *trying* to regulate it thus is not sufficient; I must *succeed* in so doing if I am not

to be blameworthy” (363). The objective requirement that one’s beliefs be true, I have argued, stems from the interdependence of epistemic agents.

The internalist requirement on epistemic responsibility meets the demand that an agent be epistemically responsible to herself.⁶³ It does not, however, meet the demand that an agent be epistemically responsible to others. To be epistemically responsible to others, one’s objective epistemic duty is not just to regulate one’s beliefs so that one only affirms those beliefs that are supported by the evidence available to one. It is, rather, that one only affirms those beliefs that are true. Or that one only affirms those beliefs that are founded on the reasons that there are for believing it, rather than merely those she has ascertained. This is why outlining reasons is an intersubjective—rather than merely intrasubjective—activity.

This internalist approach to epistemic luck makes justification a matter of the knower’s normative situation. It places the knower’s belief that *p* in the context of her or his community’s intellectual norms, standards, obligations, and duties. This puts the internalist constraint that one must have epistemic access to the grounds of one’s beliefs to be justified in those beliefs as following from a deontological conception of what it means to be justified. According to internalists, “So long as I am not violating any intellectual duties I am ‘in the clear’ in believing that *p*, whatever my chances for truth” (Alston, 2001, 81). By contrast, for externalists, so long as one’s belief is formed by a process or mechanism that usually generates true beliefs in the actual world and in close possible worlds, “it is of no concern how well one is carrying out intellectual duties” (81).

Under the present conception of responsibility, responsibility is a mechanism—a set of practices—whereby persons coordinate their epistemic and moral behaviours. Particular epistemic norms and moral norms are the mechanisms whereby this

coordination of behaviours is ensured. Practices of holding one another responsible for violating norms serve their participants as means to gauge their own and others' aptness for cooperative endeavours and communal living, which, in turn, ensures that participants' normative behaviours are coordinated. As Lorraine Code notes, "Much of epistemic life is concerned with determining the credibility of others and with establishing and maintaining one's own credibility" (75). The overarching aim of practices of attributing epistemic responsibility to one another is to coordinate behaviours. To meet that objective, those practices must not only track agents' justifications, it must also track the truth-values of their beliefs.

5.5. Full-fledged Epistemic Responsibility

Full-fledged epistemic responsibility corresponds to two dimensions of assessment: 1) the subjective internal requirement that agents exert (indirect) control over the truth or falsity of their beliefs, and 2) the objective external requirement on the truth or falsity of agent's beliefs. To be subjectively and objectively responsible for a true belief, a person must believe the claim in question, and believe it with justification. To be fully responsible for a false belief, a person must believe the proposition in question, and believe it without justification. Application of the foregoing analysis of the conditions on moral responsibility to their epistemic analogues is relatively straightforward. I shall briefly outline the necessary and sufficient conditions for justified true beliefs and unjustified false beliefs. What is required of actual agents in actual epistemic circumstances will vary from context to context. The present concern is to characterize prototypical cases with a view to developing a philosophically illuminating concept of epistemic responsibility. That concept will serve as contrast class for more philosophically and practically problematic instances of epistemic responsibility. I shall then identify the

relevant differences between such paradigmatic cases for epistemic responsibility and more problematic instances. That analysis will provide the conceptual apparatus for appropriately dealing with the ways that luck influences epistemic responsibility.

The analysis of epistemic responsibility will answer three questions: 1) What are the conditions on justified true beliefs and unjustified false beliefs that legitimate attributions of epistemic responsibility for those beliefs? 2) What are the conditions on justifying beliefs where this fails to yield true beliefs, or failing to justify questionable beliefs where this generates true beliefs, that legitimate attributions of responsibility for the presence or absence of justificational activity alone? And 3) What are the conditions on having a true, or false, belief without its accompanying justificational standing that legitimate attributions of epistemic responsibility for holding a true or false belief?

Unproblematic paradigmatic cases of responsibility for beliefs involve two dimensions: the reasons-responsiveness and reasons-reactivity of the mechanism whereby the agent comes to affirm the belief, and the truth-value of the belief. What is now required is a principled means of distinguishing justified true beliefs and unjustified false beliefs from human epistemic activity for which attributions of responsibility are more contentious. Examples of epistemic luck bring to light two variants on epistemic behaviour for which we feel praise or blame is, in some sense, appropriate but which fail to meet the requirements for paradigmatic epistemic responsibility. Both cases involve a mismatch between the justificational standing of an agent's belief and that belief's truth-value. First, agents whose diligent inquiring and conscientious justifying fail to generate true beliefs are subject to epistemic appraisal on the bare ground that they believe the false proposition in question. Second, such agents are, at the same time, epistemically praiseworthy for their diligent inquiring and appropriate justifying. Similarly, agents who

form true beliefs capriciously, or believe a true proposition for wrong reasons are unreliable. They, therefore, invite epistemic blame for their mismanagement of evidence or their failure to justify questionable beliefs. Despite this, they also invite epistemic praise simply for having true beliefs.

The most salient difference between paradigms for epistemic responsibility in which a belief's justificational standing and truth-value accord, and more problematic cases of epistemic responsibility attributions, involves lack of control on the part of the agent. That is, such cases involve the breakdown of the indirect control agents exert over the truth or falsity of their beliefs. The means agents employ to control their beliefs are inquiring and justifying procedures. The intermediary of luck, in epistemic contexts, takes the form of a mismatch between the justificational standing of a belief and its truth-value.

There are almost as many theories of justification as there are epistemologists. These range from strict internalism, which admits only the cognizer's reasons as putatively justificatory to strict externalism, which views the reliability of the cognizer's belief producing mechanism, independently of any grounds the cognizer might have for believing that mechanism to be reliable, as sufficient for justification. As noted in Chapter 2, I, along with the preponderance of contemporary contributors to the literature on internalism and externalism, prefer a hybrid construal of justification that invokes internalist and externalist constraints on justification. However, in order not to become entrenched in the independent question about the legitimacy of putative conditions on justification, I propose to leave the details of the justificatory criterion on epistemic responsibility largely untouched. So long as that justificatory criterion does not guarantee the truth of the belief justified thereby, it should qualify as the justificatory criterion on

epistemic responsibility. The point regarding epistemic responsibility raised by considerations of epistemic luck is whether or not the justificatory criterion, the criterion over which agents exert control, is the only dimension according to which cognizers are deemed epistemically responsible. While internalists will insist that it is, externalists will insist that it is not. The hybrid view about the conditions for full-fledged epistemic responsibility sees both justificational standing and truth-value as dimensions according to which epistemic responsibility is properly assessed.

An analogue of guidance control for justification is the first dimension according to which agents are deemed responsible for their beliefs. An agent exhibits guidance control over her justification insofar as 1) the doxastic mechanism that actually issues in the belief is the agent's own, and 2) that mechanism is reasons-responsiveness. Reasons-responsiveness of agents' belief-producing mechanisms is constituted by two conditions. First, there must be a moderate fit between what reasons there are for believing a proposition and the agent's reasons for believing that proposition. Second, there must be a moderate fit between the agent's reasons for believing a proposition and her acceptance or rejection of that proposition. Because there is little difference between an agent's acceptance or rejection of a proposition, given her reasons for or against that proposition, and her belief in that proposition given her acceptance or rejection of the reasons for that proposition, condition 3) for moral reasons-responsiveness becomes redundant in epistemic contexts.

Counterfactually, agent A's actually operating belief-producing mechanism, B, is moderately reasons responsive where if B were to operate and there were some possible world in which there is a sufficient reason to believe otherwise, the agent would recognize that sufficient reason to believe otherwise, and thus have reason to believe

otherwise, and would accordingly believe otherwise.⁶⁴ In addition, the reasons to which A responds must be patterned understandably, in actual and hypothetical cases, and they must foment belief. The doxastic mechanism whereby A comes to believe P is responsive, according to an understandable pattern, to the reasons there are for that proposition, it is reactive to those reasons, and the consequent belief accords with the agent's reasons. To satisfy externalists about justification, the mechanism itself must be reliable at producing true beliefs. An agent is paradigmatically responsible for believing a proposition if her belief is responsive, according to an understandable pattern, to the reasons there are for that proposition, her belief is reactive to those reasons, and the consequent belief accords with the agent's reasons for believing that proposition.

An agent A is paradigmatically responsible for her true or false belief that P where

- 1) A is a participating member of a community for whom, in the particular context of inquiry, P is at least minimally significant.
- 2) P is reasons-responsive and reasons-reactive for A: there are reasons for believing that P (not believing that P), A takes up those reasons, according to an understandable pattern, as her reasons to believe P (not to believe that P), and, given her reasons to believe that P (not believe that P), A believes P (does not believe that P).
- 3) A's doxastic mechanism caused her belief that P (lack of belief that P).
- 4) The truth-value of P matches the agent's reasons to believe that P (not to believe that P), i.e., if P is true, A is justified and if P is false, A is not justified.

The reasons-reactive criterion on full responsibility for false beliefs might, at first blush, invite skepticism on grounds that a belief in a proposition that is reactive to the reasons that there are for a proposition would preclude that belief's falsity. To

circumvent such incipient skepticism, consider the reason 'it reflects positively on me' as a reason for a particular proposition to which a believer would react according to an understandable pattern. The egoistic cognizer would endorse such a reason as her own, while someone with a less inflated self-conception would not.

5.6. Epistemic Responsibility and Justification

Subjective responsibility attends the active justifying and inquiring procedures agents undertake. They may be responsible for holding justified beliefs that turn out to be false. They may also be responsible for holding unjustified beliefs that turn out to be true. What is important for this dimension of epistemic responsibility are agents' active undertakings to exert what control they can over the truth value of their beliefs. An agent A is subjectively responsible for her inquiring and justifying procedures about proposition P despite the contravening truth-value of P where

- 1) A is a participating member of a community for whom P is at least minimally significant.
- 2) P is reasons-responsive and reasons-reactive for A: there are reasons for believing that P, A takes up those reasons, according to an understandable pattern, as her reasons to believe P, and, given her reasons to believe P, A believes P.
- 3) A's doxastic mechanism caused her belief that P.

What justifies holding persons responsible for their inquiring and justifying procedures, independently of the truth-value of the belief that is thereby generated? Indefatigable inquiring, astute reasoning, and active justification of disputable putative beliefs on the part of a cognizer provide dialogical indications of future true beliefs. Even when they fail, in particular circumstances, to result in true beliefs, the processes themselves indicate the agent's aptness for cooperative epistemic endeavours. Similarly,

cursory inquiring, slovenly reasoning and a failure to justify disputable putative beliefs on the part of a cognizer provide dialogical grounds for believing that the agent will form false beliefs in the future. Even when, in particular circumstances, such processes yield true beliefs, the processes themselves indicate the cognizer's inaptness for cooperative epistemic endeavours.

5.7. Epistemic Responsibility and Truth-Value

An agent is responsible, in some sense, for believing a proposition if that proposition is true or false, but in nearby normal worlds, its truth or falsity is not responsive to the kind of mechanism whereby the agent's belief in the proposition was actually generated.

Agent A is consequence-accountable for believing P where

1. A is a participating member of a community for whom P is at least minimally significant.⁶⁵
2. A's doxastic mechanism caused her belief that P.

What justifies holding persons responsible for the truth-values of their beliefs, independently of whether or not they match the means by which such beliefs are ordinarily generated? Because true beliefs are the overarching aim of epistemic activity, an agent's holding a true belief is, by itself, grounds for epistemic commendation; her holding a false belief is, by itself, grounds for epistemic reproach. The truth-values of one's beliefs are nodes of epistemic approval and disapproval because, for interdependent agents, individual lives and collective livelihoods depend upon members' having true rather than false beliefs. What validates practices of holding humans to account for the truth-values of their beliefs, and invalidates practices of holding thermometers to account for their veridical or non-veridical readings is that persons, unlike thermometers, are

members of epistemic communities whose interdependence makes reasonable the demand that their beliefs be true.

Chapter 6. Applying Hybrid Responsibility to Moral Luck

The ubiquitous influence of luck on human agency brings to light the fact that we lack a reasoned means to apportion the weight of responsibility. Because theorists have overlooked the dual nature of responsibility, they have been forced to insist on one of two implausible conceptions of responsibility. Externalists declare that responsibility does not entail control. This approach does not do adequate justice to the intuition that control is an essential aspect of responsibility. By rejecting the control condition, externalists forfeit the wherewithal to distinguish between things for which persons are clearly and unquestionably responsible from those things for which responsibility attribution is less clearly appropriate or not at all appropriate. Although proponents of this approach are right to wish to expand the horizon of things for which persons are responsible beyond the horizon of the things over which they exert control, there must be some natural means to differentiate things for which one is paradigmatically responsible from things for which one is not. The volitional causal agency of the executor in bringing about an event or state of affairs seems a natural means by which to forge such a demarcation. Though its motivating insights and its aims are admirable, externalism about moral responsibility is simply too broad.

Internalism about moral responsibility fails because it is too narrow. Internalists insist that one is responsible only for what one controls. This approach does not do adequate justice to the intuition that what we do is at least as important a measure of human agency as what we do on purpose. Because particularly human agency involves being responsible to others, the apparatus whereby responsibility is meted out must be reactive to those things we do, even when we do not control them; and it must also be

responsive to those things we intend to do, even when our intentions are not brought to fruition.

6.1. Hybrid Responsibility and Resultant Moral Luck

Standard examples in the resultant moral luck literature hold fixed what is within agents' control, their motivations, intentions, and choices, and vary the outcome of the agents' deliberative behaviour. In Nagel's drunk drivers example of resultant moral luck, Sam and Perry have the same intentions, the same motivations, and make the same choice. However, Sam's action—driving drunk—has disastrous consequences, while Perry's action—driving drunk—does not. Similarly, the would-be rescuer and the actual rescuer have the same intention: to save a child from a burning building. The outcomes of their actions vary dramatically; one drops the child to her death, the other saves her. These cases illustrate that, contrary to the deep-seated belief that one can only be held responsible for what one controls, evaluations of moral agents take something that is beyond agents' control to be determinative, at least in part, of agents' moral responsibility.⁶⁶

The various positions in the literature concern the legitimacy of that practice. To refresh, internalism, the majoritarian view, insists that such practices are illegitimate as they conflict with the tenet that one is responsible only for what one controls. Externalists such as Walker and Card claim that such practices are legitimate, as we can be responsible for things that outstrip our control, and that it is the tenet that one can only be responsible for what one controls that is illegitimate. The sceptical approach, advanced by Nagel and Williams, asserts that the conflict between practice and dictum indicates that norms of moral responsibility are incoherent.

Adopting the stance of participants on our actual practices involving attributions of responsibility shifts the focus away from the conflict among internalist and externalist intuitions. The naturalist hybridism that results from this shift in focus asserts that having control over what one does is a necessary condition for paradigmatic moral responsibility, and that the consequences one brings about are a necessary condition for paradigmatic moral responsibility. My approach, like Nagel and Williams', seeks to accommodate the intuitions that drive both internalist and externalist approaches to the problem of moral luck. However, unlike Nagel's and Williams', my diagnosis accommodates both intuitions without acquiescing to the sceptical conclusion that norms of moral responsibility are incoherent.

The solution to the problem of resultant luck is to underscore the importance of both the deliberative and causally efficacious aspects of ordinary human agency to the concept of responsibility. Both the result that an agent causes and the intentions with which an agent acts are essential to evaluating that agent's aptness for cooperative endeavours. In unproblematic, paradigmatic cases of responsibility attribution, we hold agents responsible for their deliberately caused consequences. Once it is clear that full responsibility for results attaches to successful intentional acts—deliberate acts whose intended consequences are achieved by means of the acts—the diagnosis of moral luck phenomena becomes transparent.

6.1.1. Example 1: Good Resultant Moral Luck

Resultant moral luck phenomena trade upon the presence of one, but not the other, of organically concomitant dimensions of quotidian human agency: the intentions and causal efficacy of the actor. Comparing a successful and a would-be rescuer illustrates this point. In the middle of one night, Alice is torn from her sleep by screams for help. She rushes

outside to see her neighbour's child clinging desperately to the windowsill of her burning house. Alice speedily hauls her ladder to the windowsill and climbs it. Alice pulls the child into her arms, and carries the child safely back down the ladder. Alice's counterpart Bob is not so fortunate. After being awoken by screams, Bob erects and climbs his ladder to save his neighbour's child from a terrifying and painful death. Just as he reaches out his arms to gather the child to safety, a gust of wind blows the child from Bob's reach. The child's life is dashed on the pavement below him.

Alice and Bob have the same motivations and intentions, and both undertake acts to fulfil the intention to save a child's life. Their acts are identical up to the point where factors beyond their control intercede in the causal efficacy of their intentional acts. The difference between their actions is the presence for Bob, and the absence for Alice, of an uncontrollable gust of wind.

When we evaluate the cases individually, there seems to be no problem determining that Alice is morally praiseworthy. There is also no problem in judging Bob to be morally praiseworthy. Bob intended to save the child, and did everything that he could in fulfillment of that intention. But though there was a will, there was no way. The phenomenon of moral luck comes into light when Alice's praiseworthiness is compared to Bob's praiseworthiness. We are tempted, and I believe rightly so, to count Alice as more praiseworthy than Bob. And, yet, what made the difference between Alice's successful rescue and Bob's failed rescue is not something over which either wielded control. The end result of their acts, attributable to luck though it is, is nevertheless something that makes a difference in our evaluation of Bob and Alice's moral responsibility. Alice is objectively and subjectively responsible for saving the child. Bob is not.

To highlight the diagnosis of resultant moral luck phenomena as pivoting on the presence of only one of two internally related aspects of human agency, intentions and causal efficacy in producing results, contrast Alice and Bob with Carman, the accidental rescuer, and Dave, the non-rescuer. Carman and Dave are avid fishers who set out one morning with the intention of catching fish. Carman and Dave have the same motivations and intentions, and both undertake acts in the fulfillment of their intention to catch a large fish. Their acts are identical up to the point where factors beyond their control intercede in the consequences of their intentional acts. Carman casts her lure, feels what she believes to be a fish bite at the end of her rod, and reels in her fishing line. To her surprise, she pulls a spluttering child, and not a fish, onto shore. Dave, however, is not so fortunate. Just at the point where his counterpart Carman hooks the child with his lure, a gust of wind blows Dave's fishing line out of range of the child. The child's lungs fill with water and she drowns. Here, again, the presence or absence of the gust of wind is something neither Carman nor Dave could have controlled.

There is little difficulty assessing Carman and Dave. Carman saved a child, and her having done so is morally significant, if not precisely praiseworthy. Dave did not save a child, and, therefore, warrants no praise. Comparing Carman with Dave introduces the phenomenon of moral luck. We are tempted, and I believe rightly so, to count Carman as more praiseworthy than Dave. And, yet, what made the difference between Carman's accidental rescue and Dave's uneventful fishing expedition is not something over which either wielded control. The difference in the results of their acts, attributable to luck though it is, is nevertheless something that makes a difference in our evaluation of Carman and Dave. Carman is consequence responsible for saving the child. Dave is not responsible for saving the child.

We now have at hand all the permutations of the two facets of ordinary human agency to which responsibility attributions are responsive. Alice is fully responsible for saving the child. She fulfills the objective criterion on resultant responsibility of having caused a morally significant state of affairs, and she fulfills the subjective criterion on responsibility of acting in response to her reasons for saving the child. Though Bob is responsible for intending to save the child, he cannot be counted responsible for saving the child. Bob fulfills only the subjective criterion on resultant responsibility of having reasons to save the child and having acted in response to those reasons. Bob is only morally judgeable along the intentional dimension of moral responsibility. Carman is causally responsible for saving the child. She fulfills the objective criterion on resultant responsibility of being the isolable cause of a normatively significant state of affairs. However, she is not responsible for saving the child, because she does not fulfill the subjective criterion of having done so in response to her reasons for doing so. She is only objectively responsible for causing the child to be saved. Dave is not responsible for saving the child, as he lacks both necessary components of fully responsible behaviour.

Resultant moral luck is usually highlighted by contrasting two intentional acts with different results. The same difficulty in applying our ordinary concept of responsibility occurs when we hold fixed the results, but vary the intentions with which agents act. Comparing Alice, who is responsible for saving the child with Carman who is only consequence-accountable, or objectively responsible, for saving the child illuminates the uncontroversial point that the intentions with which an agent undertakes actions are essential to legitimate responsibility-attribution. Although both caused the child to be saved, Alice's responsibility is more thoroughgoing than Carman's because Alice intended to save the child, while Carman did not. The same point comes to the fore when

Bob's case is contrasted with Dave's. Although neither Bob nor Dave saves a child, there is a respect in which Bob is responsible and Dave is not. That something more is intending to save a child, and undertaking to bring about the end specified in his intention.

6.1.2. Example 2: Bad Resultant Moral Luck

The problem of resultant moral luck arises when only one of the two elements necessary for attributing paradigmatic responsibility is present. One is being the salient cause of a state of affairs or event. The other is intending, but failing, to bring about a state of affairs or event that is under normative control. Each is a key indicator of one's aptness for communal living, and each is, therefore, something to which our concept of responsibility is responsive. To underscore this point, consider the same distribution of intentions and results as considered above, but where the event or state of affairs that is brought about is objectively bad, rather than objectively good.

Alice deliberately and successfully murders a child. Bob tries to murder a child but fails for reasons beyond his control. Carman accidentally kills a child. Dave neither tries to nor succeeds at killing a child.

6.1.3. Analysis of Hybrid Resultant Moral Luck

A frightful fact about our lives is that they are fraught with fortune, both good and ill. Responsibility theory has tended to discount the degree to which even our most cherished pursuits are beholden to factors over which we have no control. Once we rescind the view of agency according to which agents are in control of everything that they do and believe, a theory of responsibility that accommodates both deliberate and non-deliberate acts becomes necessary.

According to this more realistic construal of human agency, resultant moral luck phenomena result from the detachment of the two facets of ordinary human agency whose concurrence invites attribution of full responsibility. However, even disjoined from its counterpart, each of the facets provides grounds for predicting agents' generation of future results. The two dimensions of human agency to which resultant responsibility corresponds are 1) the agent's reasons for the result in question and 2) the agent's causing the result in question. The agent's reasons provide inter-subjectively accessible dialogical grounds for predicting what the agent will do in the future. The agent's causal efficacy provides objectively accessible inductive grounds for predicting what the agent will do in the future.⁶⁷

When luck does not interfere with human agency, the same motivations and intentions produce the same results. When things beyond agents' control sever the standard link between similar intentions and similar results, our ordinary conception of responsibility fails to apply. As many philosophers who advocate rejecting the control condition on moral responsibility rightly underline, nothing that we do is entirely up to us.⁶⁸ In virtue of the plain fact that we live in a world that is not of our making, everything that we do is determined, at least in part, by factors that we do not control. When luck interferes with the ordinary functioning of agency that fact becomes salient for the attribution of responsibility. Moral luck arises when factors beyond our control make a difference to the usual concurrence of similar intentions with similar results. Their contribution problematizes ascriptions of responsibility. Mere contingency does not throw our conception of responsibility into disarray. It is when contingency intercedes in the normal production of humanly wrought results—behaviours and beliefs—that our ordinary concept of responsibility lacks applicability.

Legitimate attributions of full moral responsibility for results is two-dimensional. The first of these dimensions circumscribes what the agent controls. The degree to which the mechanism whereby the agent seeks to bring about a particular result is reasons-responsive and reasons-reactive is a determinant of the agent's subjective responsibility. The second dimension to which responsibility properly attaches is the agent's causal efficacy. Whether the agent's bodily movement made the difference between the ordinary course of events and the state of affairs or event whose responsibility is at issue is a determinant of the agent's objective responsibility. Because we live in a world that is not of our making, these two dimensions of moral responsibility are imperfectly correlated.

Moral responsibility is an instrument of assessment that enables cooperative interdependent living. To fulfill that role, practices of attributing and undertaking moral responsibility must be responsive to both dimensions of human agency. By tracking the morally significant states of affairs and events people cause, whether deliberately or not, moral responsibility enables us to distinguish future cooperators from future defectors. It enables normative control and it enables teaching and learning behaviours that are under normative control. Having a state of affairs or event that is morally significant offers inductive evidence that one will wreak similar consequences in the future. The states of affairs and events that an agent causes are not, as the internalist claims, only valuable insofar as they betoken the agent's intentions. They are intrinsically demonstrative of an agent's worthiness for cooperative endeavours. We do not only care about what people mean to do. We care about what people do.⁶⁹

By tracking the persons' reasons for acting, whether successful or not, moral responsibility enables us to distinguish accidental from deliberate defection from norms,

it enables propositional articulation of norms so as to teach norms and so as to socially negotiate fuzzy normative boundaries, and it enables us to distinguish future would-be cooperators from future would-be defectors before the costs of their defection are incurred. Having intended to bring about a state of affairs or event that is morally significant serves as a direct inductive basis on which to predict future intentions, and, hence, as an indirect basis on which to predict future events or states of affairs.

Responsibility attaches to the ordinary course of actions where consequences follow the intentions with which deliberate human actions are performed. Unintended consequences, however, abound, as do unfulfilled intentions. When this occurs, I have argued, we are faced with the problem of moral luck. We are faced with conflicting intuitions, because that to which our intuitions about responsibility are responsive is conflicted. When we reflect on what it means to hold one another and ourselves responsible, we find that we are unwilling to rescind our commitment to either of the necessary conditions on moral culpability.

Unproblematic cases of resultant responsibility correspond to acts whose perpetrators deliberately and voluntarily brought about a particular state of affairs or event, where that state of affairs or event is under normative control in the community in which the perpetrator is an active participant. There are two problematic cases of responsibility attribution for results. The first is where an agent brings about a particular state of affairs or event, but did not control the mechanism whereby the state of affairs or event came to be. This class includes good and bad luck, unavoidable mistakes, accidents, sheer happenstance, and so on. These are cases where an agent's bare bodily movement makes the difference between the ordinary course of events and the particular state of affairs or event whose responsibility is at issue. In such cases, agents do not meet both

objective and subjective criteria for responsibility. They are, nevertheless, consequence accountable.

Responsibility is not univocal. As Claudia Card points out, there are many different senses of responsibility. However, once one looks at the function that practices of attributing responsibility plays in our evolutionary history and in our actual day to day lives, it becomes apparent that a common thread is interwoven through the fabric of responsibility-ascriptions and responsibility-undertakings. That common thread is the employment that such usages have in determining one's standing in one's moral and epistemic community.

Why is the common role played by responsibility attributions and responsibility undertakings so important? Surely, one might object, the manifold roles responsibility plays are important enough that it is their diversity which deserves to be emphasized, rather than their uniformity. There is some merit to this objection. To simply reduce the various uses of responsibility to the role that they play in cementing epistemic and moral standings would unjustly blur important distinctions. Such is not my intent. To the contrary, it is the very diversity of roles that responsibility plays that cements interpersonal relations and thereby enables cooperative living and inquiring.

The moral that I wish to draw is that the backwards looking juridical mode of attributing responsibility to oneself and others—the one that gives rise most dramatically to the problems of moral and epistemic luck—is part of a temporally extended nexus of interpersonal relations that surround and uphold the ability for persons to interrelate. If one's bodily movement is causally determinative of an ethically salient event or state of affairs, and one's act is reasons-responsive and reasons-reactive, then one is fully responsible for that ethically salient event—in the juridical mode of responsibility.

Yolanda, who is envious of her co-worker Zach's recent promotion, spreads a vicious and false rumour about Zach. Zach is so humiliated by the slander that he resigns from his position. Because her malicious act was causally efficacious and responsive to her reasons, Yolanda is responsible for slandering Zach. However, the fact that she is fully responsible for slandering Zach does *not*, on my view, eternally doom Yolanda to an abject moral record. There are remedies. Yolanda can take responsibility for her slander. Taking responsibility for what she has done, which involves recognising her moral wrong, attempting to undo it, and undertaking not to behave similarly in the future, performs the role of re-establishing Yolanda's standing in her own eyes and in the eyes of her moral community.

This, unfortunately, cannot always be done. Had Zach committed suicide as a result of Yolanda's malice, her wrong could not be undone. Though taking responsibility would alleviate some of the burden of her wrongful act, Yolanda would have to live with the consequences of her moral error. There are some actions that leave indelible blemishes upon one's moral record, and their indelibility can attach to accidents just as easily as to deliberate acts.⁷⁰ None the less, once one sees the juridical mode of responsibility as one of a variety of mechanisms we employ to enable cooperative living and thinking, finding out that one can be, and sometimes is, responsible for things that lie beyond the purview of one's will no longer has quite the sting of unfairness that it seems to at first.

I have argued that practices of ascribing responsibility to people for causing states of affairs or events that are under normative control performs the evolutionarily significant role of ensuring conformity to norms within groups. It does so by keeping track of who has done wrong (or right), and who will do wrong (or right), by tracking

persons' successful intentional acts. This analysis of the purpose and function of responsibility attributions enables a diagnosis of resultant moral luck phenomena in terms of the presence of only one of the two components to which responsibility attributions are normally a response: either having merely caused a state of affairs or event that is under normative control, or having merely intended to do so. A decision procedure follows straightforwardly: if the former, then one is consequence-accountable; if the latter, then one is intention-accountable.

I have, moreover, argued that holding people consequence-accountable, and holding people intention-accountable are explained and justified by the value such practices have for predicting who among our prospective collaborators is apt for cooperative endeavours and communal living. Holding people consequence-accountable is invaluable because it makes clear and transparent the normative structure of the group, and it serves as an inductive basis on which to base predictions about future causal behaviour. Holding people intentions-accountable is invaluable because it provides an inductive basis on which to predict future intentions, from which it is legitimate to infer results specified in the intention. Holding people objectively and subjectively responsible for successful intentional acts is invaluable because it provides the strongest indication of future successful intentional acts.⁷¹

6.2. Hybrid Responsibility and Circumstantial Moral Luck

When agents who are identical in terms of character, temperament, and disposition face different situations, the intercession of luck threatens to undermine the legitimacy of moral assessments of persons for their successful intentional acts. Circumstantial moral luck has the same formal structure as resultant moral luck. The difference inheres in where circumstantial luck interferes with ordinary human agency. Circumstantial luck

phenomena occur when luck interferes with the relation of regular production between a person's character and her successful intentional acts.

Adopting the naturalistic perspective on our practices of ascribing and accepting moral responsibility enables us to view the relationship between a person's character and her intentional acts as one that admits of imperfect correlation without thereby defeating the claim that what one does deliberately and successfully is at least partially constitutive of who one is. Just as responsibilities for results could not have evolved if good intentions did not, as a rule, produce bad results and vice versa, so our practices of attributing responsibility for successful intentional acts would not have evolved if a person of good character could always perform bad intentional acts, or a person of bad character could always perform good successful intentional acts.

It will be worthwhile to briefly review the reasons that circumstantial moral luck poses a problem. Circumstantial moral luck phenomena jeopardize the view that people can only be responsible for what lies within their control, when what lies within their control, now, is a successful intentional act. The problem of circumstantial moral luck is not the problem of gauging responsibility for having brought about particular circumstances. If factors beyond the agent's control have contributed to those circumstances, then the problem of how to appropriately apportion responsibility for the circumstances in question falls under the auspices of resultant moral luck. The problem of circumstantial moral luck, rather, is that we do not all face the same circumstances.

Circumstantial moral luck phenomena arise because even where we have the means to assess an agent as fully responsible for perpetrating a successful intentional act whose result is norm-governed, we are left with the problem of evaluating that person's responsibility relative to a counterpart who would have performed the same successful

intentional act, had she been in the same situation. The circumstances that mark the difference between two otherwise identical agents is attributable only to the mutability of factors that lie beyond the agents' control. The intuition that responsibility ought to be immune from luck, therefore, compels the conclusion that the two agents are equally morally responsible, even though one agent is fully responsible for an act that the other agent is not responsible for.

In some respects circumstantial luck seems more pernicious to the coherence of the concept of moral responsibility than resultant luck. The world is rife with unfulfilled intentions and with accidents. However, most norm-governed behaviour is both purposive and successful. States of affairs that are brought about by mere bodily behaviour, behaviour unaccompanied by reason-responsiveness, are anomalous. They are *accidents*. Abortive reasons-responsive mechanisms are similarly anomalous. Luck in the circumstances, situations and problems that people face is eminently more pervasive. Much though we might wish that it were not the case, problems and morally fraught circumstances arise whether we will them to or not.⁷² If we admit that circumstantial moral luck exists, as I contend we should, then very much of what we are responsible for doing is attributable to luck.

Notice, however, that circumstantial luck does not undermine moral responsibility for what we do. That is the purview of resultant luck. Circumstantial moral luck undermines the view that persons who perform morally significant acts are responsible for those acts, while their counterparts who would have done the same are not. So long as the same verdict of moral responsibility applies to the actual and counterfactual agent, the internalist intuition that moral responsibility ought to be immune from the influence of luck will have been satisfied. Satisfying that intuition, however, flies in the face of the

equally tempting and well-reasoned externalist intuition that presses us to count persons' actual performance of successful intentional acts and discount persons' would-be performance of successful intentional acts in our determination of persons' responsibility. The second intuition presses us to count what the agent has done in the actual sequence as determinative of what she is responsible for. The hybrid view of responsibility holds persons accountable for those reasons-responsive actions that they perform, so long as those actions are responsive not merely to those persons' reasons, given particular circumstances, but also to those persons' deep characters.

Examples of circumstantial luck hold fixed the agents' characters, dispositions, and temperaments, and vary the situations and problems that the agents face. The reason that circumstantial moral luck poses a problem is that persons' moral responsibility hinges upon what they do. And what they do hinges upon the problems and situations that they face. By showing that agents' successful intentional acts are contingent upon factors that they do not control, circumstantial moral luck demonstrates that even those successful intentional acts for which we ascribe full resultant responsibility are sometimes attributable to factors over which the agent does not wield control. There is, on the view I espouse, a clear difference between what resultant moral luck undercuts and what circumstantial moral luck undercuts. While resultant moral luck jeopardizes our conception of agency by demonstrating that luck can interfere between what someone intends to do and what that person does, circumstantial moral luck jeopardizes our conception of agency by demonstrating that luck can interfere between the sort of person one is, and what one intentionally does.

The predictive warrant of acts that are in character, as compared with out-of-character acts, is veiled by the format of examples that explicitly postulate the agents'

characters. The standard examples of circumstantial luck are, therefore, not the most propitious for applying the naturalistic hybrid conception of responsibility. The reason for this is that the examples postulate a static character about which everything is known and vary the problems and circumstances that character faces. This is done to highlight the unfairness of otherwise equal agents having different moral standings because of things that lie beyond their control. More realistic examples that feature the small temptations that people fail to, or manage to, overcome, better exemplify the rationale for holding persons accountable for what they do when they face moral tests. On the naturalistic view of moral responsibility, the circumstances and problems people face matter for their culpability not only because they are demonstrative of the type of person the actor is, but because they are formative of the type of person the actor will be. Let me explain.

Persons form proximal intentions to act and act on them in particular circumstances. The reasons to which one's successful intentional acts are responsive become entrenched by being employed in particular circumstances.

6.2.1. Example 1: Bad Circumstantial Moral Luck

Consider Edith and Frank who are identical in terms of character, temperament and disposition. Edith happens upon a dropped wallet containing its owner's identification, credit, and bankcards, and ten dollars. Edith, swayed by the inducement of unearned money, pockets the ten dollars. Had Frank found the wallet, he would also have stolen the ten dollars. Sheer luck in circumstances makes the difference between Edith and Frank such that Frank has done nothing wrong, and Edith has voluntarily and successfully done something that is morally impugnable, and therefore something that is morally imputable to her.

Edith and Frank are equal in terms of character. That equality is spelled out by the fact that Frank would have behaved exactly as Edith behaved, and for the same reasons. It is only because Edith faced a moral test that Frank did not face that Edith and Frank differ in moral responsibility. Luck intervened. Employing the apparatus already in place, we can ascribe full resultant responsibility to Edith. Edith's theft was reasons-responsive and reasons-reactive, and Edith's bodily movement made the difference between the theft and the ordinary course of events. What then of Frank? He is not responsible for intending to steal ten dollars. Circumstances did not arise to which that intention would have been reasons-responsive for Frank. Nevertheless, it is clear that Frank warrants moral reproach. Frank is, after all, the kind of person who would steal.

Edith and Frank are identical in all respects up to the difference in their circumstances. They have the same pattern of beliefs, desires, motivations, the same pattern of responses to external reasons, the same vices and virtues, the same character, temperament, and disposition. Their common status as untrustworthy, avaricious, and thieving persons ought to be reflected in our moral evaluation of the agents. The fact that Edith committed a successful intentional act that is morally odious, while Frank did not, also ought to make a difference in our evaluation of Edith and Frank. How are we to judge Edith and Frank? We can neither hold them equally morally responsible nor can we dismiss the respects in which they bear equal moral responsibility. Though equal in terms of character, Edith is fully responsible for a successful intentional act while Frank is not. The hybrid view enjoins counting Edith objectively accountable for being objectively and subjectively responsible for a successful intentional act. It enjoins holding Frank subjectively responsible for being the kind of person who would do what Edith did.

Circumstantial moral luck phenomena pivot on the presence of only one of two ordinarily concurrent aspects of human agency: character and successful intentional acts. Practices of holding persons responsible for their successful intentional actions correspond to the usual case in which a person's successful intentional actions disclose her character. However, because we live in a world that is not of our making, the correlation between someone's character and her successful intentional acts is not perfect. For reasons she does not control, someone can perform a successful intentional act that is out of character. For reasons she does not control, someone can fail to perform a successful intentional act that would be in character. When we evaluate Edith and Frank, we evaluate them by reference to their fulfilment of the two dimensions of responsibility for successful intentional acts. Edith exhibits both dimensions of responsibility. Frank exhibits only the character-responsive dimension of responsibility for successful intentional acts.

The present diagnosis views circumstantial luck in terms of the fulfilment of only one of the two dimensions of responsibility for successful intentional acts: that the successful intentional act be done and that it be character responsive. To underscore this point, compare Actual Edith and Counterfactual Frank with the case of Gertrude, the unwilling thief and Horace, the non-thief, who are identical in terms of character. They are honest persons for whom stealing would be out of character. Gertrude, however, is having an extraordinarily bad evening. She has just found out that she immediately and desperately requires an expensive medical intervention. As poor Gertrude is walking, she finds a wallet containing the usual complement of identification and credit and bankcards, and a ten-dollar bill. Gertrude pockets the ten dollars and turns the wallet in to the local post-office. Though it is out of character, in light of the circumstances Gertrude faces,

she forms the intention to take the ten dollars, and does so. Had Horace been in Gertrude's circumstance, he, too, would have stolen the ten dollars. However, as circumstantial luck has it, Horace does not face the harrowing circumstance that Gertrude faces. As a result, Horace does not intentionally steal ten dollars.

6.2.2. Example 2: Good Circumstantial Moral Luck

Consider another example in which circumstances beyond their control make a moral difference between the responsibility of two agents with similar characters, temperaments, and dispositions. Edith and Frank are trustworthy persons. Edith, while walking to work one morning, happens upon a lost wallet. She picks it up, locates the owner's address, and returns the wallet, money intact. Had Frank been faced with the same circumstances, Frank would also have returned the lost wallet to its rightful proprietor. However, Frank is not faced with such circumstances. Accordingly, Frank does not intentionally return a lost wallet. Edith is clearly responsible for returning the wallet. She successfully performed an intentional act that flowed from her character. Frank is not responsible for intentionally returning a wallet. However, that the circumstances arose is a matter over which neither Edith nor Frank wields control. It seems clear that there is a difference in the moral responsibility that we can ascribe to Edith and Frank. Even though the difference is attributable solely to luck in the circumstances they face, Edith is responsible in a respect that Frank is not.

A brief word is necessary here to outline the relationship between responsibility at the level of results, and responsibility at the level of circumstances. As I have argued above, full responsibility at the level of results is legitimately attributed for those acts that are reasons-sensitive. When only one of the two criteria that ordinarily coincide for human agency is present, the hybrid view enjoins imputation of responsibility either for

the reasons-sensitivity of the mechanism that ordinarily produces action, or for the production of the action. The problem of circumstantial moral luck underscores the need for an additional requirement on responsibility. That is, although it suffices for attributing responsibility for results that the agent performed the successful intentional act, there is another dimension of responsibility that is not satisfied by merely identifying the agent as responsible for the result in question. The additional question of whether the agent's successful intentional act is in character refines and amplifies the predictive merit of resultant-responsibility ascriptions. Acts that are in character provide a better predictive measure for judging agents' future actions than acts that are out of character.

However, acts that are out of character also provide a straightforward inductive ground for predicting what an agent will do in future similar circumstances. Just as what one does unintentionally provides an inductive basis for predicting what one will do in future similar instances, and, hence, ought to be reflected in the agent's moral record as consequence-accountability, so what one does intentionally, but out of character, provides an inductive basis upon which to predict the agent's future behaviours, and, hence, ought to be reflected in that agent's moral status.

6.2.3. Analysis of Hybrid Circumstantial Moral Luck

Our vulnerability to circumstances and problems that are not of our making or choosing delimits the predictive merit of knowing persons' characters. In acting from a desire, motive, intention, etc., that is alien to one's character, one fails to conform to the criteria requisite for attribution of both dimensions of responsibility. In the foregoing, I have offered a rough characterization of the conditions that ground attributions of responsibility over and above the reasons-responsiveness and causal efficacy of the actor's deliberative and acting mechanism in terms of the character-responsiveness of the

reasons for which the agent acts. Loosely, the character-responsiveness of a reason consists in the actor's actual or counterfactual reflective endorsement of that reason as a reason for choosing to act in a particular way. Character-responsiveness may be a property of internal or of external reasons. For example, the agent may actually or counterfactually reflectively endorse her own avarice as a reason for stealing ten dollars. In so doing, she authenticates the intention to steal ten dollars as legitimately flowing from her character. An agent may also authenticate an external condition on acting as being a reason for acting that is in character. So, for example, Edith's construal of the existence of a found wallet as a reason to steal its contents reflects her character. Edith, thus, reflectively endorses the existence of the circumstances that present opportunities for expressing her character. Inauthentic intentional acts, by contrast, are only generalizable to other virtually identical circumstances.

When the determining reason for acting is instilled by the happenstance of untoward circumstance or responsive to parts of the agent's character that she eschews and is actively attempting to eradicate, the grounds for inferring future similar acts are less firm. In such circumstances, a person's character does not uniquely determine a course of action. Moreover, when there is no set character that dictates that the agent go on in the same way as she has done before, there are few reliable grounds for predicting what the agent will do. When a situation arises that is significantly outside of the normal range of reasons to which an agent's deliberating and acting mechanism is keyed the result is an intentional act that is out of character. In such cases, an agent's ensuing action may be out of character simply because the agent has not yet developed a stable pattern of response to such situations. A situation might also arise that activates more than one of the reasons to which the agent's deliberating and acting mechanism is keyed. A

widespread case of circumstantial luck is the case of moral dilemmas. In such circumstances, an agent must choose between two evils. Thus, although the agent does not act voluntarily on either horn of the dilemma, the successful intentional act she undertakes is appropriately subject to moral assessment both by the agent herself, in terms of agent-regret and self-reproach, and by the agent's community, in terms of blame and a demand that she make up, somehow, for the wrong she perpetrates, however out of character.

Circumstantial luck phenomena highlight the grounds for holding persons responsible for their intentions. Authentic intentions, intentions that flow from the agent's character, are useful for predicting what an agent will do, because of their internal relation to attendant consequences, and because they illuminate the actor's character. The authenticity of intentions licenses holding persons responsible for their aborted intentions, because those intentions are demonstrative of the agent's nature, and hence of what intentions she or he is likely to form in future relevantly similar scenarios.

Inauthentic intentions and the intentions of agents who are genuinely conflicted between two courses of action, as well as the intentions of agents who have no set character and hence no set of intentions that would be in character, are each grounds upon which to predict future actions and evaluate character. They are, however, not as solid grounds as those afforded by authentic acts.

6.3. Hybrid Responsibility and Constitutive Moral Luck

Constitutive moral luck is luck in the kind of person one is. It is luck, good and bad, in having the particular character, dispositions, and temperament that one has. The problem of constitutive moral luck is particularly trenchant because we have no more basic level of evaluation to appeal to so as to ground attributions of responsibility than the agent's

character.⁷³ When one's character is itself attributed to luck, rather than to the agent, it seems we have no grounds for justified moral reproach for bad character, or for moral approval for good character. This quandary stems from the conjunction of mutually incompatible intuitions about responsibility. The first is that one can only be responsible for what one controls. Given the bare fact that who one is is, at least in part, beyond one's control, it follows that to the degree that who one is is not within one's control, one is not responsible for who one is. The second intuition at play in constitutive luck is the view that one is responsible for who one is, irrespective of one's provenance. Neither intuition is satisfactory. The first fails because it accords insufficient weight to the intuitions underlying the second, while the second fails because it accords insufficient weight to the underlying merits of the first.

6.3.1. Is Constitutive Moral Luck a Real Problem?

Before outlining a response to the problem that constitutive luck poses to the coherence of moral evaluation, a brief analysis of the concept of constitutive luck is in order. A number of commentators have insisted that constitutive luck is incoherent. There is a marked difference between resultant and circumstantial luck, which concern what agents do, and constitutive luck, which concerns what agents are. Daniel Statman, Susan Hurley and Nicholas Rescher maintain that the idea of constitutive luck posits a property-less person to whom luck adds properties. Hurley, for example, writes that constitutive luck requires us to think "in terms of each of us having an equal chance, from some relevant perspective, of having any particular constitution, hence of being any particular person" (1993). Rescher claims that "with people, there is no antecedent, identity bereft individual who draws the lot at issue with a particular endowment. One has to be there to be lucky" (1995, 30-31). Meanwhile, Statman asserts that "luck necessarily presupposes

the existence of some subject who is affected by it" (1993, 12). On Rescher's view, the idea of factors being within or beyond one's control does not apply to personality traits. Subsuming an agent's constitution under the rubric of things over which she has or lacks control is, for Rescher, to commit a category mistake.

In his introduction to *Moral Luck* Daniel Statman appeals to Harry Frankfurt's view of responsibility and character to object to the very idea of constitutive luck on grounds that we do not generally control our characters. On Frankfurt's view, the question of an agent's responsibility for his or her character has not to do with the origination of that character, but with whether the agent has "taken responsibility for his characteristics. It concerns whether the dispositions at issue...are characteristics with which he identifies and which he thus by his own will incorporates into himself as constitutive of what he is" (1988, 171-172). There is nothing incoherent about the idea that the abilities, capacities and temperament we have are things that are beyond our control. Indeed, it sometimes seems as though these things are, to a large extent, beyond individuals' ability to control.⁷⁴ The problem, however, is about how to attribute responsibility for those things when they are admitted beyond agents' control. It, moreover, seems as though we do attribute responsibility for character. In fact, there seems no more basic level of moral or epistemic appraisal.

Andrew Latus identifies Statman, Rescher and Hurley's objection to the idea of constitutive luck as follows: "there is no chance *you* could have been constituted differently, because anything that was constituted differently would not be you" (471). Latus's response is "the way we think about chance is not in terms of the chance of *you* being constituted differently but in terms of the chance of *a person* being constituted that way" (472). Although Latus's point here seems to capture some of what we consider

when we view the possibility of constitutive luck, there seems to be more at work than merely the likelihood of any particular individual possessing the character trait that is under normative appraisal. My contention is that one can limit the scope of personal identity to close possible worlds. There are possible worlds in which a person born to my parents would not be me. There are, however, closer possible worlds in which "I" designates me, and the "I" has characteristics, temperament, and dispositions dissimilar to my own. Thus, while it makes sense to say, "I could be less cowardly," it does not make sense to say "I could be my sister."

I sympathize with the worry that these authors share. Accordingly, in what follows I suggest curtailing the scope of constitutive moral luck. Nevertheless, I believe that there is sense to be made of the idea that matters over which one lacks control are at least partially determinative of who one is. By limiting constitutive luck to factors that are clearly within the scope of an agent's personal identity, one can legitimately arbitrate between possibilities in which an agent exists, but lacks a particular character trait, temperamental feature, set of inclinations or dispositions that she possesses in the actual world, and cases where luck is credited to a non-existent subject. Not only can some sense be made of the metaphysics of constitutive luck, it seems to be a genuine problem not satisfactorily dismissible simply by asserting that there must be an antecedently extant person to whom luck can be attributed.

We are often faced with cases in which we can trace the history of a person's rotten disposition to a particular factor over which that person did not wield control. (Those so fortunate as not to often face such cases can easily imagine them.) In facing these cases, we fall into a now familiar conflict between the urge to exculpate the person on grounds that the rotten disposition was not deliberately wrought and excoriate the

person because of the rotten disposition. The hybrid approach to moral responsibility explains why this conflict arises, and it offers a reasoned decision procedure for what to do when it does. Given the delimitation of the scope of constitutive luck to real, or close to real, problems, and the prospects of a solution to those problems, constitutive luck, complicated though it is, seems worthy of address.

Statman echoes Judith Thomson's discernment that "axiological concepts, such as 'good' and 'bad', admirable and so on, are not threatened by luck in the same way as are the concepts of responsibility and justification.... If David is arrogant, a bully, a coward, and full of envy, then...it would be crazy to think that, because his having these features is (supposedly) a matter of luck, he is, therefore, not a bad person" (2). We can, I think, agree with Thomson's appraisal of David as a bad person. The question that constitutive moral luck raises is not whether David is a bad person, but whether his being so is something that warrants our holding him responsible. That question has to do with the trait's origination and voluntarism.

The central problem with constitutive luck is that who an agent is is attributable, at least in part, to factors that are outside the agent's control, and yet the agent is worthy of moral judgment on that ground. According to the internalist intuition that one may be held responsible only for what one controls, insofar as who one is lies beyond the scope of one's control, one cannot be legitimately held responsible for being the person that one is. The problem with this approach is that it accords insufficient weight to our actual practices which hold persons responsible for their characters, independently of the provenance of those characters. The contrasting externalist intuition holds that one may be responsible for things over which one does not control. The verdict, according to this second intuition, is to hold persons responsible for who they are, independently of the

provenance of those characters. The problem with this approach is that it fails to accord sufficient weight to the countervailing intuition that control over who one is is essential to legitimate attributions of responsibility.

The naturalist view of constitutive moral luck tries to do justice to the intuitions underlying both approaches. On this view, the exercise of partial control over the kind of person one is is necessary for judging an one's responsibility for being the kind of being that one is. Nonetheless, those causal contributions to who one is that are beyond one's control are also necessary for judging agents' moral responsibility. That one is the way that one is is at least as important a determinant of moral responsibility as that one is the way one is on purpose.

6.3.2. Resultant and Circumstantial Moral Luck as Constitutive Moral Luck

As we have seen above, in terms of resultant moral luck, the morally significant states of affairs one brings about, even if not on purpose, are at least partially determinative of who one is. Similarly, in terms of circumstantial moral luck, those problems and situations that one faces are at least partially determinative of the kind of agent that one is. Humans are not static creatures. And we are, much though we might wish it were otherwise, not at liberty to simply create ourselves exactly and only as we wish to be. The first two ways in which matters beyond our control affect who we are is when we do, or fail to do, things by accident, and when we face problems, situations and circumstances that are not of our choosing. Resultant and circumstantial moral luck phenomena are two ways in which who we are is not up to us; the accidents we wreak or that befall us, as well as the circumstances, problems and situations we face, are also two major instances of constitutive moral luck.

In the foregoing, I present the hybrid solution to resultant and circumstantial moral luck. However, the results one brings about and the circumstances one faces constitute two classes of factors that exert influence over the development of agents' moral characters, dispositions and temperaments. The salient difference between these two categories of moral luck and the third category is that the prior two categories posit agents who possess moral characters, temperaments, and dispositions. By contrast, constitutive moral luck investigates external influences over the achievement of moral agency. Insofar as resultant moral luck and circumstantial moral luck can and do affect what kind of characters their victims become, they are also forms of constitutive moral luck.

A second respect in which the data considered in addressing resultant moral luck are important is where the state of affairs or event that is caused by the agent is a change in that selfsame agent's character. Agents can and sometimes do deliberately and successfully better themselves.⁷⁵ Insofar as being a better person, being less quick to anger, more courageous, or forthright, or caring etc. is a consequence of their successful intentional acts, persons may be counted fully responsible for such consequences. When we accidentally better ourselves, we may be counted consequence-accountable for doing so. Meanwhile, in the tragic case where we do everything within our power to better ourselves, but fail for reasons beyond our control, we may still be counted intentions-accountable for doing so.

Part of the rationale for holding persons of otherwise identical character who have different moral histories to have different moral standings is that what one does in light of problems that one faces changes one's character, sometimes deeply. Because of the entrenchment of behavioural patterns, Alice, who stole money from a found wallet, is

likely to be different in terms of proclivity to do so in the future than her counterpart, Bob, who would have stolen, but did not steal, money from a found wallet had the circumstances arisen. The kinds of problems one faces, and the situations and circumstances one finds oneself in are, therefore, important determinants of the kind of person one is. They are, therefore, essential to evaluating agents' responsibility for being the kinds of persons that they are.

In addition to how the results we cause and the circumstances we face shape us as moral beings, luck in the different starting points from which each of us begins constitutes a potent example of how our moral status is influenced by factors we do not control, and yet for which we invite moral responsibility. The difference in starting points includes variation in genetic and other early developmental resources that are beyond agents' control, but which exert influence over the character, temperament and disposition that eventually ensue. This third category of constitutive moral luck also includes such differences as the social and environmental atmosphere in which an agent develops. Poverty or prosperity, oppression or liberty, war or peace, jeopardy or safety, cruelty or kindness, even drought or abundance exert influence over the moral development of agents. Moreover, particularly when they occur early in an agent's history, circumstances of trauma, abuse, and oppression are flagrant examples of factors beyond their victims' control that undermine the assumption of equality regarding persons' capacity for moral agency.

6.3.3. Voluntariness and Origination

Traditionally, the question of an agent's responsibility for an outcome that is norm-governed in that agent's society is addressed by answering whether the agent originated the outcome in question, and whether that origination is voluntary. A further question

about responsibility arises in the context of the consistency of agents' voluntarily originated acts with their deep characters.

On the view I espouse, the external causes that determine an agent's character are good predictors of that agent's aptness for cooperative activity. They are, therefore, salient for the attribution of responsibility to that agent. We blame an agent more for being a bad person when she is afforded all the opportunities for being a good person. We blame an agent less for being a bad person when she is afforded few or none of the opportunities for being a good person. Acceding that luck influences character formation does not necessitate the conclusion that all praise and blame is to be meted out to factors external to the agent. The hybrid view enables the division of two ordinarily concurrent factors that ordinarily comprise human characters: external factors such as genetic and other early developmental resources, and familial, social and physical environment on the one hand, and internal factors such as reasons, dispositions, character and temperament on the other.

The important and problematic cases of good or bad luck in being who one is concern really possible alternative histories. I could have been more courageous. I could have been less acquisitive. I could have had more integrity. There but for the adulteration of chance go I. I could have been less caring. I could have been more irritable. I could have been less stalwart. There but for the grace of chance go I.

Examples of constitutive moral luck hold fixed what is outside the scope of an agent's personal identity, and vary, by means over which the agent lacks control, some relevant subset of traits that are within the scope of that agent's personal identity. The biblical example of Cain and Abel, alluded to in Chapter 1, provides a familiar illustration of the influence that factors beyond one's control exert on the development of agents'

moral characters, temperaments, and dispositions. But for God's preference of Abel's offering, Cain would not have become the envious and eventually murderous agent that he became.

6.3.4. Example 1: Bad Constitutive Moral Luck

My first example of constitutive moral luck involves identical twins, Ingrid and Jill, whose common genetic, embryonic, familial, physical, and social influences eliminate complicating variance in their starting points so as to better isolate the cause of a particular character trait, temperament, or disposition that one possesses but the other does not. Ingrid and Jill are born to a compassionate familial and social environment that seeks to imbue Ingrid and Jill with standards of moral rectitude. Both Ingrid and Jill demonstrate initial indications of becoming upstanding participants in their moral community. Then luck intercedes. When she is six, an old man abducts Jill. It takes three hours for her frantic parents to locate her. As a result of that terrifying interlude, Jill develops a deep-seated antipathy towards strangers. Despite her acceptance of the moral reasons in favour of compassion, Jill develops into an unsympathetic person. Jill's character is so structured that her deliberative and acting mechanism does not view the quotidian plights of others as a reason to act. It is important for the example that Jill's lack of compassion does not result from her acceptance of the reasons, such as they are, for being uncompassionate. She fully affirms the grounds for altruism. However, for reasons she cannot control, she is not compassionate.

By contrast, Ingrid harbours no underlying coldness towards strangers. Though initially merely injunctions, the reasons for altruism elucidated by her parents are taken up, approved of and accepted and employed by Ingrid as her own. Henceforth, Ingrid's reasoning and acting mechanism recognizes the quotidian plights of others as a reason for

acting. As a result, Ingrid undertakes successful intentional acts that have as their outcome the selfless betterment of others. When Ingrid faces problems or circumstances that easily enable her to benefit others, she does so. Doing so is in character for Ingrid.

Contrasting Ingrid with Jill brings to light the variance between the two agents in terms of their status as moral agents. Ingrid fulfills both the objective criterion on responsibility for her character trait of being compassionate, and the subjective criterion on responsibility for her character trait of reflectively affirming the reasons for altruism and actively participating in the fostering of that character trait. Jill's cold-hearted character trait is externally inculcated, and she does not authenticate it. Thus, Jill does not fulfill the subjective criterion on constitutive responsibility of authentication and attendant active or passive partaking in the formation her character trait. However, Jill does fulfill the objective criterion on constitutive responsibility: Jill is not a compassionate person. Although her untoward character trait was unlucky, Jill's callousness toward strangers is none the less a moral fault. A system of moral reckoning that overlooks the difference between Ingrid and Jill would omit information that is necessary to determining Jill's aptness for cooperative endeavours.

By the same token, an approach to moral responsibility that refuses to note the salient difference between Jill's externally inculcated accidental callousness and someone for whom being unsympathetic is a voluntary undertaking would likewise erroneously overlook an important facet of persons' responsibility for their characters, temperaments and dispositions. To see that this is so, contrast the case of Ingrid and Jill with that of Kim and Laurel. Kim and Laurel share starting points with Ingrid and Jill. Though inculcated with their parents' reasons for becoming compassionate agents, both Kim and Laurel repudiate those reasons. Kim and Laurel witness a neighbouring child's

dismissive attitude toward another, and emulate the neighbour's behaviour. Both children evaluate and rescind the reasons for altruism that had been inculcated in them. In their stead, Kim and Laurel espouse stringent individualism. There is, however, a difference between Kim and Laurel. Though she upholds reasons against being compassionate, and attempts to inculcate cold-heartedness in herself, Kim is unsuccessful at developing the character trait. Laurel, by contrast, successfully develops cold-heartedness as a character trait. It seems clear that there is a difference in the degree and type of moral responsibility that we can legitimately attribute to Kim and Laurel. Although both persons approve of callousness, the responsibility that Laurel bears for her character trait outweighs that Kim bears for the simple reason that Kim is not callous. While she is subjectively responsible for trying to become callous and for upholding the reasons for callousness, Kim fails to qualify for the objective criterion on constitutive responsibility of being callous. Kim accordingly warrants less moral reproach for her character flaw than Laurel does for hers.

Finally, compare Ingrid and Jill with Laurel and Kim. Insofar as being compassionate is concerned, Ingrid is objectively and subjectively responsible. She fulfills both the subjective requirement of affirming the reasons for compassion and undertaking to mould her character to those reasons and the objective criterion of being compassionate in the actual world. Jill is subjectively, but not objectively, responsible for being compassionate. This is because she identifies with the moral norm of being compassionate and would conform her character accordingly were she able to. Kim does meet the objective requirement of being compassionate. However, because she does not endorse, and in fact repudiates the reasons for compassion, Kim does not meet the subjective requirement on being responsible for her character trait. Laurel is not at all responsible for being compassionate. She fulfills neither the objective requirement of

being compassionate nor the subjective requirement of upholding or identifying with the reasons for being compassionate.

6.3.5. Example 2: Good Constitutive Moral Luck

Larry and Mark, also identical twins, have very different starting points from Ingrid and Jill's. Larry and Mark are born to parents whose individualism runs deep. They proceed, by emphasising the import of individual autonomy and scorning those who need help from others, to inculcate in their sons a callous attitude towards the struggles of others. At the age of six, Mark becomes lost for a period of three hours. During that time, a stranger comforts him. As a result of this incident, Mark develops a deep-seated kindness towards strangers. This character trait persists. When opportunities for doing so arise, Mark will lend what aid he can. By contrast, Mark's twin Larry suffers no help from strangers, and provides none. He is unsympathetic to others' ills, and treats them with coldness.

It should be clear that Mark is worthy of moral approbation for his character trait, while Larry warrants opprobrium for his. This is so despite the fact that Mark's good character trait was attributable to luck. According to the hybrid view of moral responsibility, the bare fact that Mark possesses his good character trait is worthy of moral commendation, while the bare fact that Larry possesses his bad character trait is worthy of moral censure. This is because, as outlined above, being the way one is is the archetypal indicator for aptness for cooperative endeavours.

Now contrast Mark's case with Ned's. Ned shares Mark and Larry's upbringing. Although he is inculcated with the same reasons for disdaining the needs of others, and has the same examples of behaviour as Mark and Larry, Ned repudiates the reasons for callousness towards others that have been instilled in him. The inducement to do so,

suppose, stems from witnessing a neighbouring child's kindness towards another. Instead of callousness, Ned deliberately adopts the grounds for empathising with the travails of others, and proceeds to develop sympathy as a character trait.

6.3.6. Analysis of Hybrid Constitutive Moral Luck

Just being that way is as important a determinant of responsibility for a character trait as being that way for reasons that you uphold as your own. By the same token, upholding being a certain way, but being unable to be that way constitutes some ground for approbation or opprobrium, depending upon the character trait, temperament or disposition. According to the hybrid view of responsibility, having a character trait for reasons one accepts and upholds corresponds to being fully responsible for the character trait. However, merely having the character trait for reasons one does not authenticate, or even for reasons one rails against, is nevertheless a ground for attributing responsibility to that agent for that character trait. Similarly, upholding and authenticating reasons for a character trait that one does not possess is also a ground for attributing responsibility to that agent for that character trait.

6.3.7. Rejecting Voluntariness and Origination

Robert Adams, like Claudia Card, Martha Nussbaum, and Walker, argues against the view that persons are only morally responsible for what lies within their voluntary control. His focus, unlike theirs, is not on actions, but rather on character traits. Adams argues that "many involuntary states of mind are objects of ethical appraisal and censure in their own right, and that trying very hard is not all that is morally demanded of us in this area. We ought not only to *try* to have good motives and other good states of mind rather than bad ones; we ought to *have* good ones and not bad ones" (1985, 12). The voluntary provenance of one's good or bad character traits is not necessary for one to be

responsible for them. Having a bad character trait, disposition or temperament is by itself a reason for an ethical indictment. Despite such indictment, there is a moral injunction to try to improve oneself. Trying to improve one's character, temperament and disposition warrants moral approbation, even though success is not guaranteed. The hybrid view of moral responsibility explains this approbation in terms of an agent's being intentions-accountable, though not causally efficacious, and hence not objectively and subjectively responsible. The struggle against an unwanted character trait, disposition or temperamental predilection involves *taking responsibility* for the flaw. To echo Adams and Card, if you take responsibility for it, you also do not see it as something that just happens to you. You see your callous approach as one that you yourself are taking, not voluntarily, but none the less really, to the plight of others. Recognizing that one is callous, for example, and that callousness is a moral flaw, redefines the relation between one's deliberative and acting mechanism and the object of that acting mechanism, in this case, one's own character.

One's moral flaws can be involuntary and externally imbedded in one's character. They are none the less one's own. There simply is no more basic criterion on which to base moral evaluation of an agent than the agent's authenticated character, temperament and dispositions. Being a bad person who affirms the grounds of his or her badness warrants moral condemnation. Being a good person who affirms the grounds of his or her goodness warrants moral commendation. Being a bad person, even when one does not authenticate one's reasons for being so, warrants moral reproach. Not as much as the former case, but some nonetheless. Being a good person, even when one does not authenticate one's reasons for being so, warrants moral approval. An agent's character, temperament and dispositions are excellent predictive measures of an agent's aptness for

cooperative living. Even when the reasons for the temperament are not upheld, the mere fact of their possession of those traits is strong grounds for holding them responsible. Moreover, holding fast to the reasons for character trait, disposition, or temperament is, by itself, ample grounds for moral approbation or disapprobation.

The criteria of reasons-responsiveness and causally efficacious origination for actions, and for character-responsiveness and successful intentional actions for circumstances, are each important because they serve either as dialogical or inductive grounds for discerning an agent's aptness for cooperative interdependent living. One's character laid bare provides an accurate measure of one's aptness for cooperative living. The degree to which one upholds the reasons, be they causal or cognitive, for one's character traits is an important predictor of those character traits' persistence, and also of the likelihood that those character traits will manifest in successful intentional actions.

6.3.8. Constitutive Responsibility and Agent Responsibility

In "Involuntary Sins" Robert Adams sketches a criterion by which to distinguish those involuntary states for which agents are appropriately held responsible from those for which they are not appropriately held responsible. Among the latter he counts pneumonia, blindness, mental retardation, and a broken arm. Among the former he counts states of mind that fall in many psychological categories: desires, emotions, attitudes, beliefs etc... but each of which has an intentional object. Adam's criterion is that agents can only be morally responsible for those states of mind that are "responding, consciously or unconsciously, to data that are rich enough to allow adequate ethical appreciation of the state's intentional object, and the intentional object's place in the fabric of personal relationships" (1985, 26). There are two things I wish to emphasise about Adams' confessedly sketchy criterion. The first is that ratification of the causal,

emotive, and cognitive reasons for one's own character traits need not be conscious. That the reasons be genuinely the agent's reasons suffices for responsibility for those reasons to be appropriately attributed to the agent. The reasons must arise in the agent from data sufficiently rich to enable an agent to gauge the reasons' effects upon other people. This constraint is placed upon a moral agent simply in virtue of her membership in a moral community.

The second claim Adams articulates is that the criteria of sole causal origination and direct voluntary control are not appropriate criteria when what is at issue is responsibility for one's own character traits, dispositions, and temperament. Rather, causal contribution and self-authentication determine responsibility. Where, then, does the agent's own responsibility for her character come into play? People are, simply in virtue of being members of a moral community, answerable for their characters, temperaments and dispositions. They can be reasons-accountable by accepting and approving their characters, dispositions and temperaments even when they do not create those traits. They can be 'consequence' accountable for undertaking to uphold those sets of reasons of which they approve, and undertaking to change those of which they disapprove, even when they do not have direct voluntary control over those traits. Insofar as their reasons-responsive undertakings are successful, they are objectively and subjectively responsible for the character traits that result. Responsibility for being the beings that they are is, thus, something accessible to agents once they are participating members of their moral society.

Unless agents ratify their character traits, they are only consequence-accountable for those traits. Agents are fully responsible for having the character traits that they have and endorse *even though* that element of control that agents have over their character

traits—the counterfactual reasons-responsiveness of those traits—is dependent upon factors that the agents did not create. There are two points that I outline which militate against the seeming unfairness of people being fully responsible for their endorsed traits despite the fact that their endorsement is itself a product of externally inculcated reasons that they begin with. The first is that the distribution of moral resources is itself unfair. The theory reflects this unfairness. The second is that responsibility for character is vested in us by the community of agents with whom we are interdependent. The demand that agents be sole authors of their characters so as to be responsible for those characters is, as Statman points out, metaphysically problematic. It demands that we construe agency from some extensionless, propertyless point so as to attribute characteristics had by luck. Not only would meeting the demand be metaphysically problematic in terms of self-manufacture, the demand itself stems from a conception of responsibility explicitly denied by the participatory perspective on practices I enjoin. On my view, an agent is responsible for her character traits not in virtue of her being a self-made autonomous being, but in virtue of her interdependence with others. It is others' demand, given that interdependence, that we be answerable for our character traits that constitutes our responsibility for those character traits. It is not that we have made ourselves as we are, much less is it that we have voluntarily made ourselves as we are.

One's moral flaws can be involuntary and externally imbedded in one's character. They are nonetheless one's own. There simply is no more basic criterion on which to base moral evaluation of an agent than the agent's authenticated character, temperament and dispositions. Being a bad person who affirms the grounds of his or her badness warrants moral condemnation. Being a good person who affirms the grounds of his or her goodness warrants moral commendation.

6.3.9. Constitutive Responsibility and Societal Responsibility

If you genuinely welcome your moral fault, then were it the case that you could create the world, including yourself, then you would create yourself to have that fault. In that sense, you are in control of your being the kind of agent that you are. Recall that being in control of what kind of person you are means being responsive to your reasons for being that way. It means being responsive to your reasons such that there is some sufficiently close possible world in which, if there is a sufficient reason for you to be otherwise, you would take up that reason, according to an understandable pattern of reasons, respond to that reason, and so make yourself otherwise. Your being actually unkind, for example, fulfills the objective element. Your endorsing the reasons for being unkind fulfills the subjective requirement. You are exactly as you would have created yourself to be. There seems not to be any luck involved that would undercut attributions of responsibility for that character trait. Or is there?

The deep social shaping of character complicates most questions of moral responsibility for character traits. One might object to the foregoing solution to constitutive moral luck that it has missed its mark. Claudia Card, for example, devotes the bulk of her reasoning about moral luck to the kind of luck that oppressed persons face as contrasted with socially empowered persons. From Card's standpoint, the real and immediate problem of constitutive moral luck is luck in different persons' and different groups' starting points. Oppressed persons and groups are not afforded the same opportunities for achieving moral and epistemic virtues. Moreover, this kind of constitutive moral luck is exacerbated by the endorsement of character traits inculcated under oppressive regimes. Thomas Hill Jr.'s Deferential Wife example is apt. She is "...utterly devoted to serving her husband [, and does] ...not to form her own interests,

values, and ideals; and, when she does, she counts them as less important than her husband's.... She just believes that the proper role for a woman is to serve her family" (1979, 135). On Hill's analysis, the woman displays morally objectionable lack of self-respect. This 'servility' has epistemic as well as moral dimensions. It is, on Hill's analysis, a failing to understand her moral rights.

Marilyn Friedman (1985) insists that the deferential wife's failing, though having epistemic dimensions, is a failure to exhibit whole personhood; it is a failure of integrity. A person who "simply acts uncritically to satisfy [her preferences]... fails to achieve whole moral personhood; she lacks moral integrity. It matters not whether her preferences are, in an important sense, originally hers or originally the preferences of other persons to whom she devotes her life" (7). The importance of the example is the supposition that the deferential wife believes that she should be deferential; she endorses her deference. Given the preponderance of social pressures in favour of deference among wives, this would not be surprising. Moreover, if we gerrymander the example by placing it in a slightly more remote historical context, we may suppose that there do not exist any reasons that might persuade her to critically assess her deference. There would, then, also be no reasons available to the wife against which her servility is deemed an epistemic failing.⁷⁶

Does the deferential wife suffer from bad constitutive moral and epistemic luck? Not according to the present diagnosis. She is subjectively and objectively responsible for being servile. Nevertheless, it seems that she is unlucky, in some sense of that term, to be objectively and subjectively responsible for being servile. This example demonstrates that the difference between someone who critically reflects on and endorses her character and someone who does not endorse her character leaves aside an important

respect in which luck influences character. Different persons' and groups' starting points makes it more difficult for some, impossible for others, to achieve virtuous moral characters. There is a sense in which even though an agent is exactly as she would make herself if she could have made herself, those character traits, motivations and dispositions to which her being servile are responsive are externally inculcated and therefore unlucky. Even though the deferential wife welcomes her deference, she is unlucky to be deferential. It is still her bad luck to be deferential in the first place, since that is not in her control. This second sense of constitutive moral luck does not concern agents' responsibility for their own character traits. It concerns others agents' responsibility for a particular agent's particular character trait.

There is a bootstrapping problem that affects paradigm cases of constitutive responsibility. When characterizing deference as being in the wife's control, as being counterfactually responsive to her reasons, the reasons cited to which her deference is counterfactually responsive are themselves externally inculcated. The reasons that support her deference, such that her deference is reasons-responsive, are reasons instilled by her parents, siblings, immediate social environment, husband and family, and so on. Now that, it seems, is at the root of the worry that even those persons who qualify for paradigmatic constitutive responsibility, those persons who are exactly as they would make themselves, are lucky to be as they are. The wife of Hill's example is servile for her own reasons, but those reasons antedate her agency. Agents do not create the incipient character traits that are instilled in them by social and environmental forces and from which their self-approval stems. In that sense, the reasons to which agents' having particular character traits are responsive are beyond their control.

The subject of responsibility for character invites equivocation.⁷⁷ The reason is that an agent's character is both a) the thing to which responsibility is attributed and b) the thing about which responsibility is attributable. The question of constitutive moral luck that I have dealt with above, in terms of the agent's objective and subjective responsibility for a trait, is when the thing about which responsibility is attributable, the agent's character, is attributed *to the agent*. The two criteria for attributing responsibility for b) to a) are 1) that the agent possesses the trait and 2) that she ratifies the reasons for the character trait in question. (Note that, taken alone, the non-ratified possession of the character trait, or the ratification of the character trait without its possession, constitutes a necessary, but insufficient, ground for paradigmatic responsibility attribution for said character trait to the agent. For 1), the agent is objectively responsible for the character trait. For the 2), the agent is subjectively responsible for the character trait. If she fulfills 1) and 2), she is paradigmatically responsible for possessing a particular trait.

The second issue of responsibility for an agent's character traits, dispositions and temperament emphasises the external springs of the agent's character traits. In this case, appropriate attribution of responsibility for the reasons, be they causal, affective, or cognitive, for an agent's character traits is to something outside of the agent. When deliberate and causally efficacious actions on the part of family and society result in a set of well-ordered deep reasons, family and society can be fully responsible for the moral upbringing of agents. That is, persons other than A can be held responsible for A's character trait.

If agents are the product of biological, environmental, social, and historical forces, it is difficult to see how praise or blame could rationally attach to them. It would seem, rather, that any moral praise or blame would attach to the causal antecedents of particular

agent's characters. A view of moral personhood as created and constituted by social contexts need not eliminate the agent autonomy upon which reasonable assessments of praise and blame depend.⁷⁸ A perspective that views moral agents as constituted by biological, environmental familial and social factors allows room for the important insight that one's moral environment contributes to one's moral makeup.⁷⁹

One of the advantages of a system that recognizes the existence of constitutive moral luck is that it enables a more fair evaluation of moral agents by recognizing, and accommodating, their different starting points. An approach to morality that shuns constitutive moral luck and apportions responsibility to agents only for what they control would be less apt to understand, much less deal with, the real variations in moral character, disposition, and temperament that follow from factors wholly external to the particular agent whose moral agency is at issue. Recognizing the ubiquity of luck helps to alleviate the unfairness of moral judgements between persons of different familial, economic, and societal background. The background condition for the development of agency is non-agency. An infant or a child is expressly not an agent; it is something upon which external factors exert causal influence over what will, if things go well, become an agent. Before becoming an agent, luck is ubiquitous; *everything* lies beyond one's control.

As elaborated in the foregoing, practices of attributing and undertaking responsibility is something that is done among participating members of moral communities. The ordinary trajectory from non-participating infant to fully participating member of one's moral community is complex. Nevertheless, under ordinary conditions of development a child moves from being a non-participant to being a participant in her moral community. The ordinary does not always occur. Matters beyond a putative

agent's ken or control can interfere with the very development of agency, and thereby preclude full membership in moral community. Recognizing different causal contributors to the development of moral agency, and recognizing that luck can seriously impair or even preclude moral agency, leaves room for a more forgiving approach to the morally impaired. It is incumbent upon persons to take responsibility for their histories. However, persons' histories vary. The difficulties those oppressed by socially proscribed iniquities face in developing and maintaining integrity are particularly trying. The insight that follows from acknowledging constitutive moral luck enjoins what seems fair: tailoring moral judgement to the specifics of agents' histories.⁸⁰

Examining moral luck through the prism of oppression highlights the inequality of the burdens agents face. Miranda Fricker (1998 and 2003), Susan Babbitt (1996), Victoria Davion (1991), María Lugones (1990), and Amelie Rorty and David Wong (1990), in addition to Claudia Card and Margaret Walker, along with almost innumerable others, emphasise the respects in which oppression curbs, or renders more difficult, persons' abilities to achieve epistemic and moral virtuosity. Moreover, because the construct of normal or ordinary agency with which much of moral and epistemic theory concerns itself is with what is normal or ordinary for those with the most social power, those voiced members of society, the oppressed often find themselves precluded by fiat from moral and epistemic normalcy.⁸¹ Theorizing about moral luck and epistemic luck is, unfortunately, no exception to this rule. However, dealing with issues of agency and control as it does, it is uniquely positioned to place different life histories of disparate control within a framework that explicitly studies, rather than ignoring or discounting, the degree to which factors beyond agents' control exert influence over their agency. Those oppressed by class, race, ethnicity, gender, and disability face the trials of moral and

epistemic luck more often, and the burdens such luck places on them are heavier. Being oppressed delimits one's arena of action. It forecloses oppressed agents' control over whether or not their intentions are successfully borne out. The oppressed, therefore, shoulder an unfair portion of resultant luck. The frustrations of circumstantial luck are also endemic to oppressed lives. Because the oppressed are more likely to be placed in double binds, forced to make decisions between evils, they are also more likely to form proximal intentions that they would not endorse were they not in such trying circumstances. Perhaps most tragically, oppressed persons face a disproportionately large encumbrance in the form of constitutive luck. Oppressed persons are uniquely situated such that many epistemic and moral virtues are more difficult to attain, if not impossible, much though they wish and strive to do so. Self-esteem, confidence, and even integrity, for example, are very difficult to achieve whilst being formed under a social regime that continually denigrates one's worth. Even more, oppressed persons are vulnerable to bad constitutive luck in their starting points. They may be possessed of character traits that they endorse, but their life histories so situated as to preclude even the ability to be dissatisfied with an unbearable character trait, disposition or temperament in question.

María Lugones' work suggests, in this connection, that there might be virtues of oppressed agency.⁸² 'Virtues of oppressed agency' are those traits that would enable oppressed persons to forebear against the bad constitutive, circumstantial, and resultant luck that characterises their positions as denigrated persons. Subordinated persons face very real obstacles to achieving character traits, dispositions, and temperaments that they themselves endorse. At the same time, they are more often than their empowered counterparts in the unenviable position of possessing character-traits, temperaments and dispositions that they neither endorse nor approve.⁸³ Those suffering from resultant luck

bear moderated responsibility for the consequences they cause, and for their unfulfilled intentions. Those suffering from circumstantial luck bear mitigated responsibility for successful intentional acts that are out of character, and for their characters, though unexpressed in particular circumstances. And those suffering from constitutive luck invite less than full responsibility for inauthentic character traits, temperaments and disposition, as well as for their sincere striving for traits they are unable to inculcate in themselves. Persons who suffer from bad circumstantial luck in starting points can be both objectively and subjectively responsible for their character traits, temperaments and dispositions. However, according to the hybrid view, the role played by the oppressors, or other socially empowered groups, forming those character traits, temperaments and dispositions, makes them responsible as well. If the role played was deliberate, they warrant subjective and objective responsibility. If the role played was not deliberate, but was nonetheless causally efficacious in producing the character trait, then they warrant consequence accountability.

Even so cursory an overview makes clear the unequal share of luck that oppressed persons endure. The hybrid view enables the identification of an agent's responsibility, moderated though it is by its determination by luck, for things she causes, things she intends, and things she is. In so doing, it isolates the particular relationships oppressed persons have to their own bad (and sometimes good) moral luck. Unlike Immanuel Kant, Brian Rosebury and Nicholas Rescher's internalist solution to the problem of luck, the hybrid approach recognizes the fact that morality doles out unfair burdens. Unlike Margaret Walker's externalist solution to the problem of luck, the hybrid view sees moral and epistemic luck as a genuine problem. Unlike Claudia Card's externalism, the hybrid view places constitutive luck within the broader context of cases where responsibilities

outstrip control. Card does not overlook the fact that those suffering from bad constitutive luck bear some responsibilities. She claims "...it is morally problematic for beneficiaries of oppression to hold its victims responsible for bad conduct, its victims have responsibilities of their own to peers and descendants" (1996, 41). The hybrid solution identifies the underpinnings of such forward-looking responsibilities. By upholding the importance of control to responsibility, the hybrid view identifies the lack of control persons wield over their acts, their intentions, and their characters as problematic. It, therefore, provides the means to say that it is wrong that oppressed persons have less control over their acts, intentions and characters. Moreover, it enables the attribution of responsibility, one dimensional or two-dimensional, to those who cause, either directly or indirectly by participating in systemic discrimination, uneven moral and epistemic playing fields. In so doing, the hybrid view is better placed to make demands on the socially empowered to redress the harms wrought by social and environmental iniquities.

6.4. Hybrid Responsibility and Causal Moral Luck

Causal moral luck occurs when one is lucky, or unlucky, in how one is determined by antecedent circumstances, and yet we continue to treat that person in that respect as an object of moral judgment. Causal moral luck threatens to undermine the very possibility of legitimately holding persons responsible. By highlighting the apparent conflict between a person's action being caused and it being under her control, causal moral luck jeopardizes even those actions for which responsibility attribution is unproblematic. Causal moral luck threatens to subsume everything that agents do under the rubric of moral luck. Accordingly, if agents' character-responsive successful intentional actions are caused, then it seems as though agents are not in control of their intentions, what they do

in response to various circumstances, or whether or not they actively partake in the formation of their moral characters.

6.4.1. Causal Moral Luck and the Problem of Freewill and Determinism

The problem of causal moral luck resurrects the hoary metaphysical problem of free will and determinism. The problem of free will and determinism concerns whether responsibility is compatible with universal causal determinism, the thesis that everything is caused. The threat of causal moral luck is that everything, even those things over which the agent seems to wield control and for which responsibility attribution is unproblematic, becomes a species of moral luck. If even our intentional actions are caused, then they are not under our control, and it would seem irrational to hold us responsible for them. Couched in the terms of the freewill and determinism debate, causal moral luck is the problem of being caused to do something and being responsible for that thing. This raises the question of the relationship between freedom from causal antecedents and those central intentional actions to which both dimensions of responsibility correspond.

As is well known, the conceptual territory is logically exhausted by two positions: Compatibilism and Incompatibilism. Compatibilists believe that agents can be both caused to do something and responsible for that thing. Incompatibilists believe that agents cannot be both caused to do something and responsible for that thing. Incompatibilists, in turn, divide into libertarians and hard determinists. Libertarians believe that at least some of our actions are free from causal determination and, hence, at least some of our actions are appropriate targets for responsibility. Hard determinists believe that all of our actions are causally determined, and, hence, that none of our actions legitimately invite responsibility.

6.4.2. Compatibilism and Causal Moral Luck

If one is a compatibilist, then causal moral luck presents no problem beyond the problems already solved. Causal moral luck is luck in how one is determined by causal antecedents that are beyond the agent's control but which compel an agent to action. Causal *luck*, accordingly, is nothing over and above the following varieties of luck: Resultant luck, Circumstantial Luck and Constitutive Luck.

If one is a compatibilist, the remaining causal antecedents of agents for which we treat them as objects of moral judgment do not constitute luck. The agent's reasons for bringing about a state of affairs or event which cause her to bring about the state of affairs or event are not lucky causal antecedents. The agent's character traits, temperament and dispositions, which generate the agent's reasons for bringing about a state of affairs, which in turn generate the agent's action in accord with her reasons for bringing about the state of affairs or event in question, are not lucky causal antecedents. And the agent's identification with the reasons, causal, emotive and cognitive, that underlie and comprise her character traits, dispositions, and temperament, and generate the agent's ordered sets of reasons for acting, are not lucky.

The remaining causal antecedents, such as hypnosis, drug inducement, brain washing, compulsion by force or threat of force, either fall into the category of non-moral luck because persons so caused fail to be appropriate objects of moral judgment, or they fall into one of the categories of moral luck to which the hybrid solution to problems of moral luck applies. If one is a compatibilist, causal luck does not present any problem in addition to what has been presented, diagnosed, and solved by the naturalist hybrid approach to responsibility.

6.4.3. Incompatibilism and Causal Moral Luck

The problem of causal moral luck might at first appear more trying if one is an incompatibilist. This appearance is misleading.

If one is a hard-determinist incompatibilist, then there is no problem of moral luck. If one believes that agents' intentional actions and their characters are caused by antecedent states of affairs or events, and that their being so caused precludes agents from being responsible for them, then there is no problem of causal moral luck. This is because causal moral luck occurs when agents are held responsible for actions, intentions and character that are brought about by causal antecedents. Because there is no holding persons responsible from the hard-determinist's stance, there is no problem of moral luck.

If one is a libertarian incompatibilist, then there is no moral luck beyond the varieties of moral luck considered thus far. Libertarians insist that persons must be free from causal determinism in order to be candidates for responsibility. The targets for freedom from causality are voluntary formulation of intentions for resultant luck, consistency with character for circumstantial luck, and self-authorship for constitutive moral luck. For the libertarian, causal moral luck threatens because it renders agents' intentions, character consistency, and self-authorship as caused, and hence not available for responsibility. However, by rejecting determinism, the libertarian thereby rejects this threat to responsibility.

Though I favour compatibilism, my solution to the problem of moral luck is amenable to a libertarian's construal of freedom as requiring indeterminism. Accordingly, a solution to causal moral luck broaches the topic of whether responsibility, as I have outlined it here, is compatible with universal causal determinism. I see no impediment to an analogous libertarian theory according to which the naturalistic hybrid account of

moral responsibility in its ordinary functioning outlined herein is indeterministic. The libertarianism advocated by Laura Waddell Eckstrom seems a particularly apt candidate.⁸⁴

In “Indeterminist Free Action,” Eckstrom identifies the best freedom-enhancing location for indeterminism as just prior to the agent’s formation of a unique preference about what to do. On her view, in the deliberative process about what to do, a number of considerations about what to do occur to the agent. As the free agent decides what to do, which preference she will form as the resolution of the decision process is causally undetermined by the past. The preference the agent eventually forms has causal antecedents in the considerations that occur to the agent during the deliberative process. However, “the considerations *cause but do not determine* the deliberative outcome” (2001, 148). Considerations, or reasons for or against a course of action, are probabilistic causes of the formation of the agent’s preference for a particular course of action. Although she treats only of necessity for indeterminism in the relationship between the agent’s reasons for a course of action and the consequent activation of the agent’s acting mechanism, analogous indeterminist treatments of character-responsiveness and character authentication are not wholly implausible. The libertarian would then have an indeterministic account of reasons-responsiveness, character-responsiveness and character authentication.⁸⁵

Whether one is a compatibilist, or an incompatibilist, the naturalist hybrid solution to the problem of moral luck has covered all the ground that it is possible to cover.

6.4.4. Freewill and Determinism is a Different Problem from Moral Luck

The problem of free will and determinism does not raise any additional moral luck phenomena. Moreover, the problem of freewill and determinism centres on a different

question from that addressed by moral luck. The rare point of agreement among compatibilists and incompatibilists of both libertarian and hard determinist stripe is their consensus that moral responsibility entails free will. The problem of moral luck, by contrast, is about whether or not moral responsibility entails free will. Moral luck phenomena show that one can be responsible for things that one does not control.

The positions on moral luck include internalism, which maintains that moral responsibility does entail free will, and externalism, which maintains that moral responsibility does not entail free will. Internalists insist that one can only be held responsible for things over which one exerts control. Externalists insist that one can be held responsible for things that outstrip one's control. The hybrid approach marries the intuitions that underlie both internalism and externalism about moral responsibility.

Despite the fact that they address different questions, there are obvious and important relationships between the subject matter in the freewill and determinism debate and the debate about the existence of moral luck, and its implications for the coherence of moral norms. Most libertarians, for example, would without hesitation ally themselves with internalists about moral luck. Standard libertarian insistence upon the necessity of both voluntariness and uncaused origination, or sole authorship, for freedom seems irreconcilable with the view that persons may be held responsible for things that they do not control. They would, thus, likely echo the internalist's dismissal of commendation and reproach for accidental results and miscarried intentions as groundless. However, satisfying as it does the intuitions that underlie the internalist's response to the problem of moral luck, the hybrid approach I advocate should go some lengths to assuaging libertarian animadversions.

It should be clear that from the naturalistic hybrid view of moral responsibility, being caused is no impediment to being responsible. There may therefore be more spontaneous sympathy among standard compatibilists for the externalists' avowal that persons may be responsible for things that outstrip their control. However, the hybrid view rejects the standard compatibilist insistence that one cannot be even partially responsible for actions, reasons, or character traits that are not free.⁸⁶

I think that there are sound reasons to reject the spectre of incompatibilism about responsibility and causal determinism. As should be clear by my insistence that responsibility outstrips reasons-responsive causal control, the demand that responsibility be limited to those beliefs over which the agent wields uncaused sole authorship is far too stringent to accommodate the types and degrees of responsibility that ordinary agents bear for the beliefs that they hold. Moreover, the evolutionary function of responsibility attributions and acceptances as means of determining agents' aptness for participation in cooperative and interdependent living would applaud, rather than deplore, causal relations among agents' reasons-sensitivity and the beliefs they actually hold.

There are important intersections between the problem of moral luck and the problem of free will and determinism. However, they are different problems. The problem of moral luck concerns the coherence of norms of moral appraisal given the ineluctable penchant to hold persons responsible for their accidental actions, and the equally ineluctable commitment to the view that what matters for moral responsibility are non-accidental actions. The problem of freewill and determinism concerns whether one can freely choose things that are antecedently caused.

Not only are they different problems, but the presupposition that responsibility requires freedom upon which incompatibilists and compatibilists concur is called into

question by the problem of moral luck. Moreover, positions in the two debates do not align neatly. One may be an internalist about moral luck and a compatibilist about freedom and determinism. One may be an externalist about moral luck and an incompatibilist about freedom and determinism. Although being a hybridist about moral luck strongly indicates being a compatibilist about freedom and determinism, one can be a hybridist and be either a libertarian or a compatibilist. Though the naturalistic hybrid view offered here has implications for the freedom and determinism debate, those implications are ancillary to the central objective to solve the problem of moral luck.

The burden of this dissertation is to offer a well-reasoned, consistent diagnosis of the practical problem of the apparent incoherence of moral norms while doing justice to the intuitions underlying the combatant positions. The diagnosis offered meets that burden. Moreover, naturalist hybridism carries with it a resolution to the practical problem of where and how to assign moral responsibility when conflicts arise.

Chapter 7. Conclusion

7.1. Hybrid Responsibility and Resultant Epistemic Luck

Just as whether an act is good or bad is generally not up to the actor, so whether a belief is true or false is generally not up to the believer. What *is* generally up to the believer is whether she brings it about that she has a true belief, and the standard means of doing so is by justifying her beliefs; that is, agents control whether they have true or false beliefs by holding justified beliefs and rejecting unjustified beliefs. On the moral side, what is generally up to the actor is whether she brings it about that she does a good deed, and the standard means of doing so is by intending to perform a good deed. In moral and epistemic pursuits, agents customarily exercise control over the results they bring about, whether right actions or true beliefs. Full responsibility for those results attaches to the standard cases in which control is exercised. However, in non-standard cases, where there is a mismatch between the means by which a result is ordinarily wrought, and the result, standard practices of evaluating agents fall into disarray. We are left with only one of the two criteria we ordinarily use to evaluate epistemic agents. However, each criterion, the justificatory standing and the truth-value of the target belief, is by itself a vector for epistemic appraisal. Whether or not beliefs are true makes a difference to an agent's epistemic standing, even where the truth-value of the agent's belief is beyond her control. Similarly, whether or not the agent's belief is justified makes a difference to an agent's epistemic standing, even if the justified belief turns out to be false, or the unjustified belief turns out to be true.

If we follow the intuition underlying epistemic externalism and allow that resultant luck plays a determinative role in the epistemic responsibility of agents, then we thereby erode the distinction between knowledge and true belief. If, however, we follow

the intuition underlying epistemic internalism and refuse to count resultant luck in assessing epistemic agents, then we thereby erode the distinction between knowledge and justified belief for purposes of attributing epistemic responsibility. The sceptical approach asserts that the conflict between practices of holding persons responsible for the truth-values of their beliefs and the dictum that enjoins us to count only what agents control indicates that norms of epistemic responsibility are incoherent.

Epistemic responsibility comprises two dimensions of assessment. Central cases are two-dimensional, involving a justificatory component and a factual component. Holding persons responsible for their beliefs involves assessing whether they did what they could do ensure their beliefs were true and whether or not their beliefs are true. Once it is clear that responsibility attaches to ordinary epistemic agency—where unjustified beliefs are false and justified beliefs are true—the diagnosis of resultant epistemic luck phenomena becomes transparent. When someone exhibits epistemic praiseworthiness along one of these dimensions of assessment, but fails to exhibit its usual concomitant, the hybrid approach enjoins judging them just so: as praiseworthy or blameworthy as dictated by the particulars of each dimension of assessment. In consequence, hybrid epistemic responsibility is far more messy than either internalist or externalist epistemic responsibility. One can be praiseworthy for having a true belief but blameworthy for failing to justify that belief. One can be praiseworthy for justifying a belief but blameworthy for believing something false.

Hybridism accommodates the intuitions that drive both the internalist and externalist approaches to the problem of epistemic luck. The solution to the problem of resultant epistemic luck underscores the importance of both the deliberative and truth-conducive aspects of ordinary human epistemic agency to the concept of epistemic

responsibility. Both the truth-value of the belief that an agent arrives at and the internally accessible inquiring and justifying procedures with which an agent forms her beliefs are essential to evaluating that agent's aptness for cooperative epistemic endeavours.

7.1.1. Example 1: Bad Resultant Epistemic Luck

Resultant epistemic luck holds fixed what is within agents' control, their inquiring and justifying procedures, and varies the outcome of the agents' deliberative behaviour. Take Anne and Bill, who are equal in terms of their access to and weighting of information, their reasoning skills, their proneness to arriving at beliefs, and the justification they deem sufficient for holding particular beliefs under particular circumstances. Anne must establish whether or not there are weapons of mass destruction in country C. Her country's policy towards C will be governed by Anne's determination. Given the importance of her inquiry, Anne carefully solicits and considers all of the data that are available to her. The evidence, C's recent imports of plutonium, C's government's consistent claims to the right to a nuclear weapons program, its refusal to allow members of the United Nations security council entry into its weapons facilities, and satellite images of radioactivity, points to the existence of a nuclear weapons program. Anne carefully judges that the hypothesis that there are nuclear weapons in C renders the data more likely than were there no nuclear weapons in C. Anne, nevertheless, imaginatively considers alternative explanatory hypotheses for the data, and eliminates each as less likely and less explanatory of the data than the conclusion that there are nuclear weapons in C. Anne, therefore, concludes that there are nuclear weapons in C. Anne's justification of her belief cannot be faulted. She has discharged her subjective epistemic duty to ensure that the belief she entertains is true. Anne's diligent inquiring and justifying bears

fruit: C is actively engaged in creating nuclear weapons. Anne knows that there are nuclear weapons in C.

Now contrast Anne with Bill. Bill is identical with Anne in all respects except for the truth-value of their beliefs. Bill discharges his subjective epistemic duty as diligently, imaginatively and adroitly as Anne does. However, Bill's justified belief is false. All available evidence to the contrary, the country in question does not have a nuclear weapons programme.

There seems to be little problem evaluating Anne and Bill individually. Anne warrants that highest of epistemic accolades: Anne possesses knowledge. Bill, too, is epistemically praiseworthy. He did everything that was in his power to ensure that his belief was true. The phenomenon of resultant epistemic luck comes to light when we compare Anne with Bill. Although they are equal in terms of what was within their control, Anne's belief is true while Bill's is false. The difference between the two epistemic agents is, by hypothesis, something over which neither wielded control. Nevertheless, that difference is epistemically significant. Anne knows, while Bill does not. Poor Bill suffers from bad resultant epistemic luck.

According to the hybrid account of epistemic norms, our appraisal of Anne and Bill responds to two concomitant facets of epistemic agency: the agent's subjective epistemic duty to do everything within her control to ensure that she has true beliefs and the agent's objective epistemic duty to have true beliefs. Accordingly, we may legitimately count Anne's belief as knowledge, and accord her subjective and objective responsibility for that article of knowledge. Bill has fulfilled only one of the two necessary criteria for the attribution of epistemic responsibility to agents. Thus, although

he may be counted subjectively responsible for his belief, he may not also be counted objectively responsible for the simple fact that his belief was not true.

7.1.2. Example 2: Good Resultant Epistemic Luck

Comparing Anne and Bill with Cindy and Doug introduces the full gamut of agents' objective and subjective duties, and reveals the degree and type of responsibility that corresponds to each permutation. Cindy and Doug are bad epistemic agents. Although they know that their country's department of foreign relations will determine policy towards country C by their determination, neither exercises the inquiring and justifying diligence that is required of them in forming their beliefs. Both Cindy and Doug base their conclusions on biased inquiry with wilful blindness to countervailing evidence. Doug's belief that there are weapons of mass destruction in C is false. As it happens, C is not developing a nuclear weapons program. Cindy's belief, however, is true.

When we take the agents individually, our epistemic evaluation of Cindy and Doug is not particularly troublesome. Doug invites outrage. He shirked his subjective onus to justify his belief, and his belief was false. Cindy also merits outrage. Just not quite so much as Doug does.

Comparing Doug with Cindy illuminates the different dimensions of epistemic responsibility in the case. In all respects that were under their control Cindy and Doug are epistemic equals. However, the fact that one has a true belief and the other a false belief merits a difference in our epistemic evaluation. This difference is explained by the fact that Cindy has fulfilled the objective requirement on epistemic agents to have true beliefs, while Doug has not. For reasons neither controlled, Cindy has a better epistemic record than Doug does. Cindy profits from good resultant epistemic luck.

We are now in a position to illustrate the hybrid account of epistemic responsibility by reference to the full range of permutations made possible by the imperfect correlation between justificational status and truth-value. The hybrid account enables a reasoned distinction between Anne, who qualifies for full responsibility for her knowledge, and Bill, who is only subjectively responsible for having a justified belief. It enables a reasoned distinction between Cindy, who satisfies the objective epistemic criterion of believing truly and Doug who does not satisfy the epistemic criterion of believing truly. The difference between Anne and Bill and between Cindy and Doug redounds to the truth-value of the agents' beliefs. Despite the fact that each pair of agents is identical in terms of what they control, there is, and should be, a marked difference in our appraisal of the agents who got it right and the agents who did not. This difference can, given our ability to articulate reasons, be erased. Nevertheless, the requirement that explanation for error be provided itself indicates that falsity, even justified falsity, blemishes one's epistemic record.

Resultant epistemic luck is highlighted by contrasting doxastic acts that are indiscernible in terms of what the agents control, their inquiring and justifying procedures, but which vary in terms of truth-value. However, the same difficulty in applying our ordinary concept of epistemic responsibility occurs when we hold fixed the truth-values of agent's beliefs, but vary the justificatory standing of the agents' beliefs. Comparing Anne, who is objectively and subjectively responsible for knowing that there are nuclear weapons in C, with Cindy, whose belief is true, but unjustified, illuminates the uncontroversial point that the justificatory standing of agents' beliefs is ineliminable to legitimate responsibility-attribution. Although both have a true belief that there are nuclear weapons in C, Anne is more praiseworthy than Cindy because Anne's belief is

justified, while Cindy's is not. Cindy was lucky that her belief was true. The same point comes to the fore when Bill's case is contrasted with Doug's. Although neither Bill nor Doug has a true belief, Bill is praiseworthy in a way that Doug is not. Bill has done everything within his power to have a true belief. It is bad luck that his belief was false. No intermediary of luck was necessary for Doug's belief to be false. The hybrid view of epistemic responsibility underscores the important distinctions that arise from different combinations of truth-value and justificatory status.

On the present diagnosis, resultant epistemic luck is a phenomenon of the presence of one but not the other of the ordinarily concurrent aspects of human epistemic agency. Persons who have justified beliefs that are also true are unproblematically counted epistemically praiseworthy. For similar reasons, persons whose beliefs are unjustified and false are unproblematically counted epistemically blameworthy. The conflict among our norms of epistemic appraisal arises when agents exhibit a mismatch between the subjective justificatory criterion of epistemic evaluation and the objective truth-value criterion of epistemic evaluation.

The problem of resultant epistemic luck arises when only one of the two elements necessary for attributing full responsibility for knowledge claims is present without its customary attendant. The truth-value of an agent's beliefs is an important inductive indicator of that agent's aptness for collaborative epistemic endeavours.⁸⁷ Similarly, simply trying, or failing to try, to bring it about that one's beliefs are true is a key indicator of one's aptness for communal living. Justifying beliefs and having true beliefs, even when divorced from its usual concomitant, is something to which our concept of responsibility is responsive.

7.1.3. Analysis of Hybrid Resultant Epistemic Luck

These cases illustrate that evaluations of epistemic agents take something that is beyond agents' control to be determinative, at least in part, of agents' epistemic responsibility.

The hybrid view, as it applies to resultant epistemic luck, affirms the difference between justified belief, and knowledge for attributions of epistemic responsibility. This is neither surprising nor controversial. Most epistemologists count the truth-value of a belief to be both necessary for legitimate attributions of responsibility, and up to the world rather than the agent. Where the hybrid view presents a departure from standard epistemologists' fare is in applying what is implicit in the judgment that there is a difference between justified belief and justified belief that is true to other variants of the ordinarily conjoined aspects of human epistemic agency.

7.2. Hybrid Responsibility and Circumstantial Epistemic Luck

Circumstantial epistemic luck challenges the legitimacy of according responsibility to agents for their epistemic achievements or failures by contrasting agents whose epistemic achievements we applaud or deplore with otherwise identical agents who would have achieved the same epistemic feat, had they been in the same circumstances.

According to the intuition that epistemic responsibility is immune from the influence of luck, a person who has a justified true belief ought to be accorded the same epistemic status as an otherwise identical person who would have justifiedly believed a true proposition, were she in the same circumstance. Both persons might be careful inquirers, and warrant epistemic credit for this trait. The difference between the two agents, if there is one, is attributable only to the fact that one, and not the other, was presented with circumstances that provided the reasons for her to believe as she did.

Because neither agent wields control over which circumstances are faced, the difference

redounds to factors beyond their control. At the first level of responsibility attribution, responsibility for results, there is a clear difference between someone who has a justified true belief and someone who does not. So long as the first agent meets the objective and subjective criteria set out in the hybrid account, she is responsible for her justified true belief. Since the second agent meets neither of the two criteria on which we judge epistemic responsibility for results, the second agent is clearly not responsible for having a justified true belief.

There is, however, a difference between the two agents in terms of their epistemic responsibility. That further difference hinges on circumstantial rather than resultant epistemic luck. Mylan Engel Jr. describes a person profiting from resultant, or “veritic,” epistemic luck as one for whom, given her evidential situation, “it is simply a matter of luck that her belief turns out to be true” (1992, 67). Circumstantial or “evidential” luck, by contrast, attaches not to the fortuitous truth-value of agents’ belief, but to the fortunate or unfortunate circumstances in which the agent’s belief is formed. An agent is circumstantially lucky if she is fortunate to be in a particular evidential situation, but “given her evidential situation, it is *not* a matter of luck that her belief turns out to be true” (67). Engel’s example of good circumstantial luck features Nadine, who usually studies in her carrel in the basement of the library, but quixotically decides to study at a carrel facing a window. Given her circumstance, she is in a position to have positive evidence that it is storming out. She is, thus, in a fortunate evidentiary position vis a vis the proposition ‘it is storming.’ Given her evidentiary position, Nadine does not profit from resultant luck; she profits from circumstantial luck. Moreover, Nadine’s circumstantial luck is compatible with according her full resultant responsibility for knowing that it is storming. Circumstantial epistemic luck does not invariably undermine

knowledge claims. Rather, circumstantial epistemic luck endangers the internalist supposition that persons' epistemic standings are equitably meted out.

7.2.1. Gettier Cases and Epistemic Luck

Some circumstantial epistemic luck, however, does undermine knowledge claims. Take Gilbert Harman's barn case as an example.⁸⁸ John is driving through the countryside, he notices a barn, and forms the belief 'that is a barn.' Now, although the barn that John spies, and about which he forms the belief 'that is a barn,' is in fact a barn, the countryside through which John has been driving has recently been the set of a movie for which a number of barn façades were built. Given John's unpropitious evidential situation, it is lucky that his belief turned out true. He might have just as easily, and just as justifiedly, formed the belief 'that is a barn' about something which was in fact a barn façade. On the naturalistic view, what explains our reluctance to attribute knowledge to John is that the circumstance itself generally leads to false belief. Although John does, under the circumstances, have an internally justified belief, his internal justification does not have a reliable truth connection. Compare John's case with Karen's. Karen, like John, drives through the countryside, spots a barn, and forms the belief 'that is a barn.' The countryside through which Karen has driven is not rife with barn façades. She, therefore, has a justified true belief. Moreover, there seems no reason to deny that Karen has knowledge. In such cases, the circumstances that differentiate between Karen and John is something over which neither wields control. That difference nevertheless makes an epistemic difference. Karen knows while John does not. What grounds the intuition that John does not know? John's justified true belief is like a belief that is justified but false in that we would expect John's normal belief producing mechanism to produce mostly false beliefs in such circumstances.

7.2.2. Example 1: Good Circumstantial Epistemic Luck

Persons' epistemic reputations are built up over time. Nevertheless, circumstantial epistemic luck is most salient when it concerns significant discoveries and significant failures to discover. Louis Pasteur said that luck favours the prepared mind. True, luck favoured *Pasteur's* prepared mind. But what of those prepared minds luck failed to favour? By contrasting an epistemic agent's good fortune with that of an equally well prepared mind, whom luck did not favour, we can see that circumstantial luck undercuts the view that persons' epistemic reputations are always warranted. Pasteur discovered that the fermentation of milk and wine was due to the multiplication of bacteria and other microorganisms. This discovery, in turn, led to the discovery of the role of microorganisms in human and animal disease. Now we rightly applaud Pasteur for his discovery. However, Pasteur's fictional counterpart who has an equally well prepared mind and would have discovered the role of microorganisms in fermentation, had he but been in Pasteur's circumstance, would receive no applause.

Another example of the influence of luck on agents' epistemic reputations takes place in Pasteur's own field of inquiry. Sir Alexander Fleming describes his epochal discovery in 1928 of the antibacterial powers of the mould from which penicillin is derived as a "triumph of accident and shrewd observation."⁸⁹ Although Fleming and Pasteur are right to note that the role circumstantial luck plays in epistemic pursuits is secondary to the role played by observation, experiment and reasoning, that role is nevertheless partially determinative of those agents' epistemic reputations. To see that this is so, consider Fleming's counterpart who worked in a slightly more fastidiously kept laboratory. The only thing that differentiates Fleming from his otherwise identical

counterpart, Lady Gleming, is that the counterpart had no accident upon which her powers of shrewd observation could have been put to use.

Acknowledging the influence of luck on epistemic discoveries commonly repudiates the merit that would otherwise accrue to the agent. Just as acknowledging the influence of luck on persons' economic standing has a levelling effect, so noting its influence on epochal epistemic discoveries often serves to discredit the discoverer's epistemic standing. This is perhaps what accounts for the relish with which tales of lucky discoveries are told and heard. Ordinary agents may entertain the idea that had they but been in Fleming's laboratory, they would have discovered penicillin. By the same token, detractors of science often use chance discoveries to denigrate science as accidental, arbitrary or capricious. In so doing, they reflect the intuition that accidents should not count in our assessment of agents' epistemic reputations. To infer that epistemic agents' lack of control over their discoveries wholly invalidates the merit that would otherwise accrue to them, however, would be too quick. To see that this is so, contrast Fleming not with Gleming, the counterpart who would have inferred the antibacterial properties of mould from the unexpected contamination of staphylococcus by spores of penicillium notatum, but with one who would not have done so.

Hleming is a laboratory technician who lacks the attentive and imaginative mind-set requisite for serious scientific inquiry. Hleming is, moreover, an extroverted and painstaking diarist who meticulously inscribes and circulates a detailed chronicle of his daily observations amongst his cohorts. Hleming notices and records that the introduction of penicillium notatum to a specimen of staphylococcus stopped the growth of the staphylococcus. Hleming has, accordingly, discovered that penicillium notatum is an antibacterial agent. Although Hleming has a justified true belief, and fortuitous

circumstances dictated that this was so, there is a significant difference between Hleming's case and Fleming's. Fleming's discovery, attributable to lucky evidentiary circumstances though it was, was nevertheless attributable to Fleming's epistemic character. We are not surprised to find out that Fleming made a significant discovery. By contrast, given his character, Hleming's inadvertent discovery is surprising.

Finally, compare Hleming with Ileming. Ileming is also a painstaking extroverted diarist who lacks the epistemic character to deliberately and wilfully set upon a course of scientific observation, theory formulation, experimentation, and observation. Ileming, however, haunts a more fastidiously kept laboratory than does Hleming. As a result, she happens across no unexpectedly contaminated staphylococcus sample. Comparing Hleming and Ileming elucidates the respect in which having a justified true belief, even if not in character, redounds to one's epistemic good standing. Our epistemic evaluation of Ileming is dismal, at least as it bears upon her aptness for scientific discovery.

We rightly credit Fleming with a sound epistemic reputation. Although we may rightly applaud Gleming for her intellectual aplomb, we may not credit her with having discovered penicillin. Hleming, whose happening upon the discovery of the antibacterial properties of penicillin did not result from his epistemically virtuous character warrants applause, but not full credit for the discovery. Ileming, who would only have chanced upon the discovery but did not chance warrants no epistemic praise.

Circumstantial luck phenomena occur when luck interferes with the relation of regular production between a person's deep character and her justified true beliefs. Just as the truth value of an agent's belief matters, in and of itself, for the agent's epistemic responsibility and matters because it betokens the agent's reasons-sensitive mechanism

whereby the belief was got, so justified true beliefs matter in and of themselves, in addition to mattering because they normally betoken the agent's epistemic character.

The hybrid view of responsibility holds persons fully accountable for those reasons-responsive beliefs that they hold, so long as those beliefs are responsive not merely to those persons' reasons, given particular circumstances, but are also responsive to those persons' deep characters. Fleming and his counterpart Gleming share epistemic characters, traits and dispositions. They equally meet the subjective criteria for circumstantial responsibility. It is, however, clear that Fleming is and should be held responsible for something for which Gleming simply cannot be held responsible: the discovery of penicillin. In terms of character-responsiveness, Gleming's discovery would have been just as triumphant as Fleming's was, and her consequent epistemic reputation as warranted, circumstances were such that she cannot be credited with the discovery. Her disposition towards evidence, her meticulous mind, and her epistemic character dictate that she be held responsible for being a good epistemic agent. However, she fails to meet the objective criterion on circumstantial responsibility: she is not resultantly responsible for a justified true belief. Circumstances beyond either agent's control dictate that the discovery does not redound to her credit, but to Fleming's.

The same difference between Hleming and Ileming dictates that the former, but not the latter, is results-responsible for a momentous discovery. Hleming meets the objective criterion on circumstantial responsibility: he is resultantly responsible for a justified true belief. Ileming meets neither objective nor subjective criteria. The difference between Fleming and Hleming's epistemic standings and that between Gleming and Ileming's underscores the importance of character responsiveness to the attribution of responsibility for their resultant beliefs to agents.

7.2.3. Example 2: Bad Circumstantial Epistemic Luck

Unpropitious circumstantial luck can also differentiate between the epistemic responsibility that accrues to otherwise identical epistemic agents' standings. According to the hybrid approach on offer here, two conditions are necessary for ordinary attributions of epistemic standings to agents. The first, objective, criterion concerns the actual sequence in which the agent comes to hold a particular belief, including the justificational standing of that belief. The second, subjective, criterion concerns the agent's epistemic character, and whether the belief arrived at in the actual sequence, and the means by which the belief was arrived at, is character responsive. An agent committed to high standards of intellectual rectitude who meets those standards in a particular circumstance may be held fully responsible for the epistemic reputation she or he invites. An agent who commonly and systematically flouts standards of epistemic rectitude, and who, given a particular circumstance believes something false hastily, or with wilful blindness to countervailing evidence, or prejudicially, warrants the epistemic disrepute she thereby invites.

Circumstantial epistemic luck phenomena occur when luck intercedes between the kind of epistemic agent one is, in terms of one's character, temperament and dispositions, and the sort of beliefs one comes to hold. Because agents do not all face the same trials, their epistemic standings reflect differences that are attributable only to the fact that some face and fail epistemic tests that their otherwise equal counterparts would but do not. According to the intuition that one can only be held accountable for what is within one's control, agents who believe, and agents who would believe, ought to have the same epistemic standings. According to the intuition that responsibility outstrips control, what agents do believe counts, while what agents would believe does not. The hybrid view

upholds the intuitions that underlie both approaches. By counting both a person's epistemic character and the beliefs that flow from that character to be determinative of agents' epistemic standings, the hybrid view accommodates the intuition to count beliefs that are responsive to agents' deep epistemic characters as well as those beliefs that they come to hold that do not reflect their standard epistemic practices.

A common case where persons face and fail epistemic tests is when deeply held or highly cherished beliefs are threatened by countervailing evidence. To see that persons' responsibility for their epistemic reputations is keyed to both reasons-sensitive epistemic characters, and to circumstance-sensitive actual sequences, consider the case of Neville Chamberlain. One of the better known historical cases of bad circumstantial epistemic luck befell the British Prime Minister Neville Chamberlain during the 1930's. Following World War I, under the Treaty of Versailles, Germany had been subject to stringent restrictions on the size and activities of its military. Adolf Hitler flouted those restrictions. Under the mistaken belief that it would ensure peace, Chamberlain brokered the Munich Agreement which ceded the Sudetenland to Hitler in exchange for a cessation of Germany's expansionism. Upon his return to England, Chamberlain proclaimed peace: "My good friends, this is the second time in our history that there has come back from Germany... peace with honour. I believe it is peace for our time".⁹⁰ Chamberlain was egregiously mistaken. Hitler promptly annexed the remainder of Czechoslovakia, thereby violating the Munich Pact. Moreover, it was not until Germany invaded Poland, almost a year later, that Chamberlain acquiesced to the failure of his appeasement policy, and Britain and France declared war on Germany.

Chamberlain's epistemic reputation is deservedly tarnished. His belief that he had engendered peace despite weighty evidence to the contrary was, if not an outright case of

self-delusion, at least an example of wishful thinking. According to the criteria set out for resultant epistemic responsibility, Chamberlain warrants ascription of subjective and objective responsibility for failing to know that appeasement would not preclude war. His belief was unjustified, and it was false.

However, when we consider the unenviable and unusually fraught epistemic circumstances under which Chamberlain's belief was formed, some pardon seems appropriate. The prospect of military intervention so soon after the end of 'the war to end all wars' would have been inducement enough to prompt many to bouts of wishful belief. The question that is relevant to the problem of circumstantial epistemic luck is whether or not Chamberlain's unjustified false belief was in character. For the purposes of this example, it will be useful to highlight two scenarios.

According to the first scenario, Chamberlain was an epistemic miscreant, prone to discount reliable evidence, generalize hastily and counter-induce, whose bad epistemic character would predictably lead to unjustified false beliefs in a large variety of circumstances. His epistemic failing in believing that ceding the Sudetenland would stop German expansionism exemplified his usual pattern of discrediting sound evidence. On this scenario, Chamberlain's blemished epistemic standing is fully warranted. Now compare Chamberlain with a fictional counterpart, Dhamberlain. Dhamberlain also boasts a substandard epistemic character. She jumps to conclusions, she mismanages evidence, and she is wilfully blind to anything that contradicts her preferred beliefs. Had Dhamberlain been prime minister in 1938, she would have believed exactly as Chamberlain believed. In evaluating Chamberlain and Dhamberlain, it is clear that both warrant poor, even abysmal, epistemic standings. However, it is also clear that Chamberlain warrants more epistemic dispraise than does Dhamberlain. This is so for the

simple fact that Chamberlain, and not Dhamberlain, unjustifiedly believed something false.

To underscore the importance of the character-responsiveness of unjustified false beliefs to legitimate attributions of responsibility, assume, what is more plausible, that Chamberlain (now Ehamberlain) had a healthy respect for truthfulness and an abiding commitment to holding all and only justified true beliefs. On this scenario, Ehamberlain was simply unsuited to forbear against the particular epistemic test he faced. His profound advocacy of peace led him to erroneously but understandably overrate the evidence in favour of peaceful resolution to German expansionism. Ehamberlain's wishful thinking and attendant false belief represents a blemish on his epistemic record, but it is insufficient reason to think that he is a bad epistemic agent. It is, moreover, not generalizable to other behaviour of similar sorts, unless the highly complex and unusual circumstances that prompted his epistemic failing are present.

We can now compare the responsibility that accrues to Ehamberlain, an agent of good epistemic character who faced and failed an epistemic test, with that which accrues to an agent of equally good epistemic character Fhamberlain who would have believed as Ehamberlain did. Had she faced the unenviable test Ehamberlain faced, she too would have abandoned epistemic pursuits to the blandishments of wishful thinking. Fortunately for Fhamberlain, she did not face the problem Ehamberlain faced, and therefore did not fail in the face of that problem.

When we take the cases individually, there is little trouble evaluating both epistemic agents. Ehamberlain, though ordinarily an upstanding epistemic agent, is fully resultantly responsible for his unjustified false belief. As a result of her circumstantial fortune, Fhamberlain's epistemic standing remains untarnished. Things become

somewhat more vexed when Ehamberlain and Fhamberlain are compared to one another. Both warrant good epistemic standings based upon their common virtues as epistemic agents. However, Ehamberlain is assessable along an additional dimension of responsibility attribution than is Fhamberlain. Even though she would have wilfully discredited evidence that military action was necessary had she been in Ehamberlain's position, the actual sequence remains salient to our appraisal of the two agents. The fact remains that Ehamberlain failed a trial that Fhamberlain did not. Accordingly, his failure warrants inclusion in his epistemic record.

On the current diagnosis, circumstantial epistemic luck phenomena trade upon the presence of one, but not the other, of the ordinarily concurrent factors that ordinarily comprise human epistemic agency: the agent's character on the one hand, and the justificational status and truth-value of the agent's beliefs on the other hand. The predictive advantage to holding persons responsible for fulfilling subjective requirements of holding justified beliefs holds for those beliefs that are responsive to the agent's epistemic character. The above comparisons of Chamberlain to Dhamberlain and Ehamberlain to Fhamberlain highlight the role that the actual sequence plays in attributions of epistemic responsibility.

Circumstantial epistemic luck is usually highlighted by contrasting two otherwise identical agents one who, because she or he faces different problems, situations or circumstances, is fully responsible for a belief claim, while the other does not, and accordingly is not. The same difficulty in applying our ordinary concept of responsibility occurs when we hold fixed the unjustified false belief, but vary the character responsiveness of the agents' justifications. Comparing Chamberlain, who is objectively and subjectively circumstantially responsible for his unjustified false belief to

Ehamberlain, whose unjustified false belief is out of character illuminates the uncontroversial point that the degree to which an agent's beliefs are in character is ineliminable to legitimate responsibility-attribution. The same point comes to the fore when Dhamberlain is contrasted with the second Fhamberlain. Although neither Fhamberlain nor Dhamberlain fails to appropriately weigh evidence against appeasing Hitler, Dhamberlain is responsible for something more than Fhamberlain. That something more is being the sort of person who would, across a broad variety of circumstances, mismanage evidence and thereby arrive at false beliefs.

7.2.4. Analysis of Hybrid Circumstantial Epistemic Luck

As I have argued above, objective and subjective responsibility at the level of results is legitimately attributed to agents whose beliefs have truth-values that accord with their justificatory standing. When only one of the two criteria that ordinarily coincide for human epistemic agency is present, the hybrid view enjoins imputation of responsibility either for the reasons-sensitivity of the mechanism that ordinarily produces belief, or for the truth-value of the produced belief. The problem of circumstantial epistemic luck underscores the need for an additional arena of epistemic responsibility. The additional question of whether the agent's justified true belief is in character refines and amplifies the predictive merit of resultant-responsibility ascriptions. Beliefs that are in character provide a better predictive measure for judging agents' future beliefs than beliefs that are out of character.

However, beliefs that are out of character also provide a straightforward inductive ground for predicting what an agent will believe in future similar circumstances. Just as what one believes without justification provides an inductive basis for predicting what one will believe in future similar instances, and, hence, ought to be reflected in the

agent's record as consequence-accountability, so what one believes with justification, but which is out of character, provides an inductive basis upon which to predict the agent's future epistemic conduct, and, hence, ought to be reflected in that agent's epistemic standing.

7.2.5. Reliabilist and Responsibilist Virtue Epistemologies

The literature on circumstantial epistemic luck is sparse. However, the field of Virtue Epistemology concerns the relationship between epistemic character traits and true beliefs. Virtue Epistemologists divide roughly into reliabilist virtue epistemologies and internalist virtue epistemologies. Of the latter, the responsibilist virtue epistemology, exemplified in the work of Linda Zagzebski and Lorraine Code, and also that of Miranda Fricker, and Vrinda Dalmiya, underscores the import of epistemic agents' character traits to their epistemic responsibility.⁹¹ The following are numbered among the epistemically virtuous character traits: intellectual attentiveness, thoroughness, tenacity, courage, autonomy, humility, and imaginativeness, awareness of one's fallibility, fair-mindedness, open-mindedness, and sensitivity to evidence. The central dispute between reliabilists and responsibilists surrounds the value that attaches to epistemic virtues. Reliabilists maintain that epistemic virtues are valuable only insofar as they are conducive to true beliefs, while responsibilists maintain that epistemic virtues are also valuable in and of themselves.

The relationship between being an epistemically virtuous agent and coming to true beliefs is complex and varied. Nonetheless, there is ample reason to affirm that there is a positive causal relationship between epistemic virtue and true beliefs. In dealing with circumstantial epistemic luck, the relationship between being a virtuous epistemic agent and coming to true beliefs was explored. Although there is merit in being a virtuous

epistemic agent, independently of whether in a particular instance one arrives at true beliefs by means of exercising one's epistemic virtue, epistemic responsibility attaches most firmly when the agent's epistemic virtue results in her true belief, or when her epistemic vice results in her false belief. But before we can usefully discuss why epistemic virtue is valuable, we need to address a prior question: can epistemic virtue, whether valuable in itself or for its results, be itself a matter of luck? And if it were a matter of luck, what would that mean for the possibility of epistemic responsibility and epistemic virtue?

7.3. Hybrid Responsibility and Constitutive Epistemic Luck

The chief virtue of particularly human means of navigating our social and physical environments is the inter-generational transmission of intellectual resources. Our predecessors bequeath each new human not only with the biological endowments necessary to see, hear, smell, touch, and taste the world and the cognitive capacity to retain information, generalize from particulars, and choose among competitive explanations, but also with epistemic virtues or vices, and contentful propositional knowledge.

7.3.1. Is Constitutive Epistemic Luck a Real Problem?

The question that vexes constitutive epistemic luck is the warrant of holding agents responsible for the very possession of their epistemic faults and their epistemic virtues. Constitutive epistemic luck concerns the influence of factors beyond the agent's control which contribute, at least in part, to the agent's status as an epistemic agent and for which the agent is held responsible. As with constitutive moral luck, constitutive epistemic luck must be circumscribed by really possible alternative histories of the particular agent whose epistemic character traits, dispositions and temperament are at issue. What is

central to an agent's nature cannot be due to luck, because it would not be contingent with respect to that agent. Thus, an agent may be lucky to be fair-minded, or attentive, or sensitive to evidence, and unlucky to be biased, inattentive, or insensitive to evidence. However, it is not the case that the agent suffers from bad constitutive epistemic luck because she is not omniscient.

The problem that constitutive epistemic luck poses is to isolate the appropriate locus of epistemic praise or blame for epistemologically significant character traits, dispositions, and temperament. How can agents be praised or blamed for their epistemic characters, temperaments and dispositions when those selfsame characters, temperaments and dispositions are determined by the agents' intellectual heritage?

Competing intuitions dictate contrary answers to this question. According to the internalist intuition that one can be praised or blamed only for things that one controls, agents can be legitimately held responsible for their epistemic characters only insofar as their characters result from voluntary actions on the part of those agents. This approach enjoins holding fixed external determinants of character and attuning judgements about epistemic fitness only to those character traits, dispositions and temperaments, or parts thereof, that are voluntarily self-wrought by the agents. An agent who finds herself with an undesirable epistemic character trait, say closed-mindedness, the origins of which lie clearly outside of her control could not be held responsible for being closed-minded.

The first problem with the internalist approach is its impracticability. Suppose it possible to distinguish deliberately self-formed character traits, temperamental features, and dispositions from those that derive from the agent-in-training's biological and environmental and social influences. What grounds would distinguish involuntary externally imbued traits from voluntary self-imbued traits? In *Choosing Character*,

Jacobs articulates the following grounds: "human activity is voluntary in the sense that it is uncompelled, the agent is doing what she wants to do, and the agent is aware of, and has some measure of control over what she is doing" (2001, 10). Though there are many ways to articulate the difference between voluntary and involuntary, central to that difference is the idea that the voluntary stems from the agent's own appreciation of the reasons for or against a way of being or a course of action. What makes character-trait acquisition voluntary is its conformity with the agent's antecedently established want, desire, value, affirmation, or identification with that character-trait. The internalist approach to constitutive luck which discounts externally inculcated character traits and acknowledges only deliberately self-inculcated character traits thus faces a bootstrapping problem: Those self-forming actions that are voluntary are voluntary because of their relation to the antecedent externally inculcated character traits, temperament, and dispositions that the internalist seeks to omit from judgment about the agent's epistemic status.

The second problem with the internalist's approach to the problem of constitutive epistemic luck is that it fails to accord sufficient weight to the countervailing insight that, independently of their provenance, the epistemologically significant character traits, temperament, and dispositions that one actually possesses can not be eliminated from an accurate evaluation of one's aptness for epistemic pursuits. Independently of whether they are possessed voluntarily or involuntarily, independently of whether they are self-caused or caused by the external factors that constitute one's epistemic inheritance, the epistemic vices and virtues one actually has matter for judging one's status as an epistemic agent. By restricting the scope of character responsibility to voluntarily self-manufactured traits,

the internalist excludes factors that are necessary to accurate evaluation of epistemic agents. Internalism about constitutive epistemic responsibility is too narrow.

The second proposed resolution to the problem of constitutive epistemic luck underscores the importance of the character traits, dispositions and temperament that one actually possesses. On this externalist approach, restricting agents' responsibility to all and only those character traits, temperaments, and dispositions that are voluntarily self-inculcated is not only hopelessly flawed from the standpoint of practicability, but it would be inadvisable from the standpoint of the aims and purposes of epistemic responsibility. From the externalist perspective, responsibility for one's epistemic character, temperament, and dispositions can, does, and should transcend the boundaries of voluntarily self-inculcated traits. Whether one is a diligent, attentive and courageous inquirer or an indolent, inattentive and timorous inquirer matters because one is depended upon, and depends upon others, for well-reasoned and accurate judgments.

There is much to be applauded about the externalist's solution to the problem of constitutive epistemic luck. The externalist solution disregards whether our capacities, temperaments and dispositions are voluntarily self-wrought, emphasizing instead the degree to which responsibilities accrue to agents in virtue of the fact that we are interdependent epistemic agents. As I have argued, what warrants holding others responsible for their epistemic prowess or ineptitude in the first place is our dependence upon them for veridical judgments. It is, therefore, both unsurprising and fitting that what grounds persons' responsibility for their epistemic characters is their membership in a community of interdependent inquirers.

There are, however, two problems that caution against whole-hearted adoption of the externalist solution. Firstly, there is something about the distinction between

voluntarily possessed character traits and involuntarily possessed character traits that is significant for the attribution of epistemic responsibility. Because they discount that distinction, externalists' concept of epistemic responsibility is overly broad and insufficiently variegated. Secondly, the chief virtues cited by externalists in favour of their approach depend upon a distinction that their approach eschews. Externalists such as Claudia Card and Margaret Urban Walker commended (moral) externalism because it enjoins forbearance towards those whose moral failings are foisted upon them by the misfortune of impoverished moral circumstance. Card states "Luck is involved both in the motivation to take responsibility and in our ability to carry through. Where that seems unfair, we may be able to take the unfairness into account...in some of our evaluations" (1996, 27). For Walker, moral luck in the circumstances agents face and in the results of their actions "will matter for our judgments on the degree and type of wrong, and on the appropriateness of blame, its amount and nature" (1991, 18). This attitude, laudable though it is, is not immediately accessible given the strictures of the externalist stance on epistemic responsibility. The call to temper blame that the unlucky would otherwise face demands that there be some responsibility-conferring element that is attributable to voluntarily self-wrought traits as opposed to involuntarily or other-wrought traits.

The hybrid approach to constitutive epistemic luck upholds the merits of the internalist's emphasis on voluntary self-manufacture of character traits, temperament and disposition and the externalist's emphasis on the actual possession of character traits, temperament and disposition. Unproblematic cases of responsibility for a particular character trait involve an actually possessed trait that is or has been consciously or unconsciously accepted, identified with, approved of, or otherwise recognized and authenticated as one's own. Both internalists' insistence on the purposefulness of the trait

and the externalists' insistence upon the possession of the trait are met. Cases of constitutive epistemic luck can now be parsed into familiar varieties. Traits that are possessed by agents but not authorized by the agent satisfy the externalist, but not the internalist, intuition to count the actual possession of a trait as determinative of responsibility. Traits that are authorized by the agent but not possessed by the agent fulfill the internalist, but not the externalist, intuition to count what lies within the agent's control in determining the agent's responsibility. The grounds for upholding the externalist intuition, even when the agent does not authorize her own character trait, is that possession of a trait provides strong inductive, objectively accessible grounds for evaluating that agent's aptness for cooperative epistemic pursuits. The grounds for upholding the internalist intuition, even when the agent does not possess the trait that she approves of, are that the agent's reasons for or against possessing a trait provide strong inter-subjectively accessible grounds for evaluating that agent's aptness for cooperative epistemic pursuits.

7.3.2. Example 1: Bad Constitutive Epistemic Luck

There are two classes of constitutive epistemic luck. The first variety of constitutive epistemic luck deals with situations, problems, and circumstances one faces early enough in one's career as an epistemic agent-to-be that they are formative of one's character. To pinpoint this variety of constitutive epistemic luck, consider a set of quadruplets, Olive, Penelope, and Quincy and Ruth. The Quadruplets share all the external resources necessary to their eventual development into open-minded epistemic agents. Their parents train them to be open-minded by citing the truth-conducive properties of open-mindedness and by being open-minded inquirers. The agents-in training are on their way to developing this epistemically virtuous character trait.

Olive predictably and without untoward intervention develops into an open-minded epistemic agent. Olive affirms the reasons for being open-minded, identifies those reasons as her own, and is open minded. During a critical period before Penelope fully grasps and identifies with the reasons for being open-minded, she listens to a persuasive orator whose reasons to scorn dissenting or merely different points of view are infectious. Penelope adopts the reasons against open-mindedness, and authenticates them as her own reasons. However, Penelope is unable to (or simply does not) actually become a closed-minded epistemic agent. Whenever the opportunity presents itself, she solicits the input of marginalized voices, listens to their verdicts and concerns, and accommodates multiple perspectives.

Quincy, like Penelope, possesses only one of the criteria for ordinary attribution of responsibility for a character trait. Quincy affirms and identifies with the reasons in favour of open-mindedness, but she is not open-minded. While her behavioural patterns are not yet set, Quincy attends a birthday party in which only authoritative conformist views are attended to. It so happens that Quincy's expertise is called upon, while her dissenter's views are quashed. The experience is sufficiently memorable that Quincy thereafter harbours a closed-minded attitude. Finally, there is Ruth who both scorns the reasons for open-mindedness and is not open-minded.

What difference, in terms of their epistemic responsibility, does Penelope's bad constitutive epistemic luck make? What difference does Quincy's bad constitutive epistemic luck make? Because her justification for open-mindedness is internally related to her open-mindedness, Olive is authentically open-minded. That character trait, and the approbation that attends it, is rightfully attributable to Olive both on the objectively accessible grounds of her being open-minded, and on the inter-subjectively accessible

grounds of her affirmation of the justificatory reasons for being that way. Olive is, accordingly, objectively and subjectively responsible for being open-minded. By contrast, both Penelope's and Quincy's position qua authentic open-minded epistemic agents is dubitable. While she possesses the character trait which manifests in justified true beliefs when in propitious evidentiary circumstances, Penelope lacks a central ingredient for being an appropriate object of approval for her open-mindedness. Because she lacks or repudiates the justificatory reasons for being open-minded, Penelope does not merit full responsibility for being open-minded. Penelope satisfies the objectively accessible requirement of possessing an epistemic virtue. However, she lacks the inter-subjectively accessible requirement of reflectively assessing the justificatory reasons for open-mindedness, and approving of and identifying with those reasons. Quincy's case is simply the reverse of Penelope's. While she warrants approval for reflectively assessing her reasons for being open-minded, and approving of and identifying with those reasons, Quincy does not warrant approbation for open-mindedness because she is not open-minded. Ruth, meanwhile, possesses neither the objective requirement of being open-minded nor the subjective requirement of identifying with or upholding open-mindedness. She is, accordingly, not at all responsible for possessing the virtue.

7.3.3. Example 2: Good Constitutive Epistemic Luck

Sam, Thomas, Ursul and Victor share all the resources necessary to their eventual development into closed-minded epistemic agents. By verbal means and by so tailoring their own beliefs, their parents train the quadruplets to scorn the epistemic value of soliciting multiple perspectives and attending to diverse or marginalized voices. They emphasize instead tailoring one's beliefs to those held by venerable authorities. Sam predictably develops into a closed-minded epistemic agent. He affirms the reasons for

closed-mindedness, decisively identifies with his reasons for being closed-minded, and is closed-minded.

Thomas' history as an agent differs from Sam's in a crucial respect. Both are incipient closed-minded agents. Although they can cite the reasons for closed-mindedness, they do not yet manifest closed-mindedness. During the critical period prior to the entrenchment of closed-mindedness in Thomas's cognitive makeup, he is exposed to a persuasive orator who touts the epistemic virtues of open-mindedness. During a critical period during which he was particularly susceptible to the entrenchment of behavioural patterns, Thomas attends a birthday party at which wide ranging and wildly diverse opinions are solicited and considered. Thomas emulates the behaviour he witnesses and over a remarkably short period becomes an open-minded epistemic agent. However, Thomas's character trait of open-mindedness lacks something necessary for the ordinary ascription of responsibility for being open-minded. Thomas is unable to effectively marshal reasons in support of his open-mindedness. Rather, the reasons that Thomas has at his disposal are reasons against open-mindedness and in favour of closed-mindedness. Thomas might even decisively identify with the reasons for closed-mindedness, and scorn those for open-mindedness. He would, accordingly, be unable to justify his open-mindedness.

Enter Ursul and Victor who share Sam and Thomas' genetic, embryonic, familial and broader social and physical environment. Ursul is sufficiently receptive to countervailing influences at this point that he repudiates the reasons for closed-mindedness and adopts the reasons in favour of open-mindedness. When questioned about the relative merits of open-mindedness, Ursul musters numerous arguments each lauding open-mindedness as a creditable epistemic character trait. Ursul even decisively

identifies with the reasons for open-mindedness, and thereby affirms those reasons as his reasons. Ursul, however, is unable to shake the influence his parents have exerted over his character so as to actually become open-minded. Ursul is accordingly internally inconsistent. He affirms and identifies with open-mindedness, but is closed-minded. Though he cannot legitimate being so to others or to himself, Ursul possesses the epistemic character trait of closed-mindedness. When presented with the opportunity to listen to alternative or marginalized perspectives on a subject, Ursul will refrain from listening. He will neither solicit nor give credence to perspectives other than those of established authorities. Finally, Victor manifests neither criterion for being closed minded. He deplores the reasons for closed-mindedness and is not closed-minded.

Contrasting Sam's authentic closed-mindedness with Thomas's inauthentic closed-mindedness underscores the import of character-traits, dispositions, and temperaments being consistently entrenched in the agent by means of relations of justification between the agent's reasons for possessing a particular trait, and the possession of that trait. Sam is objectively and subjectively responsible for being a closed-minded agent, while Thomas warrants only objective responsibility for being closed-minded. Thomas exhibits closed-mindedness in all matters over which he exerts control: he decisively identifies with the reasons for closed-mindedness. However, because he is not actually closed-minded, much though he might like to be, Thomas does not warrant full ascription of responsibility for being closed-minded.

There is a clear difference between Thomas and Ursul on the one side and Sam and Victor on the other in terms of the integration and entrenchment of their character trait. In consequence, it seems more reasonable to hold Sam responsible for his closed-mindedness, ineluctable though its origination was, than it is to hold Thomas responsible

for his open-mindedness or Ursul for his closed-mindedness. Though possessed of and identifying with the epistemically significant reasons for the character trait, Thomas lacks the appropriate relationship between being closed-minded and the reasons in support of closed-mindedness that feature so prominently in assessments of agents' fitness as would-be epistemic partners. Thomas is not closed-minded, though his reasons are for closed-mindedness.

Ursul also lacks the appropriate relationship between being closed-minded and the reasons in support of closed-mindedness that feature so prominently in assessment of agents' fitness as would-be epistemic partners. Ursul is open-minded, though his reasons are for closed-mindedness.

Thomas' closed-mindedness, unauthenticated though it is, is none the less an appropriate measure of Thomas' epistemic aptitude. This can be seen by contrasting Thomas who is closed-minded, albeit inauthentically so, with Victor who manifests neither objective nor subjective criteria for closed-mindedness. Thomas warrants epistemic disapproval for the very possession of the deplorable trait. Victor does not. Meanwhile, Thomas's authentication of the reasons for closed-mindedness is deplorable, but his failure to manifest closed-mindedness is creditable. He is not so discreditable as Sam whose closed-mindedness results from his affirmation of the reasons in favour of closed-mindedness.

One might object that reasons for being ambivalent about attributing responsibility to Thomas for his closed-mindedness concern not whether Thomas' reasons and character trait are appropriately related, but whether, given their relation, Thomas is genuinely closed-minded. Similarly, reasons for being ambivalent about attributing responsibility to Ursul for his closed-mindedness do not stem from the lack of integration of Ursul's

reasons with his character trait. Rather, unwillingness to hold Ursul responsible for being closed-minded concern whether Ursul is genuinely closed-minded. This objection underscores the point that what legitimates attributions of responsibility is the genuine possession of a character trait, rather than the voluntary possession of that trait or the self-originated possession of that trait.

What authenticates agents' responsibility for their character traits, dispositions, and temperaments, is not the source of the reasons but rather the degree of conformity between the character traits they possess and the character traits they authenticate by means of self-reflective assessment in light of the norms of epistemic rectitude of the agent's epistemic community.

7.3.4. Analysis of Hybrid Constitutive Epistemic Luck

Comparing the responsibility that the agents bear for their epistemic characters suggests that the anterior likelihood that each agent-to-be has of developing the particular trait that she or he developed is immaterial to the ascription of responsibility for that trait. The criteria that determine the agent's epistemic responsibility is whether the character trait is possessed and whether the agent affirms the reasons for possessing that trait.

Comparing Victor to Penelope, however, brings to light an additional condition that seems important to the assessment of their responsibility for their character traits.

Penelope's situation regarding the eventual acquisition of the epistemic virtue of open-mindedness is lucky when compared with Victor's. The difference in the developmental resources afforded Penelope and those afforded Victor constitutes a second variety of circumstantial epistemic luck.

Given the preponderance of developmental resources that foster closed-mindedness as an epistemic trait, there is a certain sense in which Victor's achievement

seems praiseworthy over and above that of someone for whom becoming open-minded was easy. Victor's development of open-mindedness is unusual, given the preponderance of developmental resources propelling him in the direction of closed-mindedness. However, it was not 'lucky' in the sense of failing to be responsive to Victor's reasons. Victor has successfully translated his endorsement of open-mindedness into open-mindedness. He has exerted all the control that is available to him as an agent who has not the metaphysically problematic ability to wholly create himself. Penelope, too, has successfully exerted all the control available to her in terms of self-manufacture, as an agent who is a product of her social and physical environment. She exemplifies the open-mindedness she wholeheartedly endorses.

7.3.5. Constitutive Epistemic Luck and Agent Responsibility

What difference in terms of responsibility, then, does the difference between Victor and Penelope's starting points make? Victor's open-mindedness is no less laudable than Penelope's for being less likely to have obtained. Penelope's open-mindedness is no more laudable than Victor's for being more likely to have obtained. It is reasonable to assume that it was much more difficult, given his starting point, for Victor to achieve open-mindedness than it was for Penelope. As Card notes, "Potentialities for becoming responsible may be realized without much self-consciousness in a moderately favourable environment.... What develops without much self-consciousness under moderately favourable conditions may be stunted or damaged by oppression or abuse" (47). The difference in the ease and effort with which the vices and virtues for which agents are fully responsible are acquired does seem to be significant for attributions of epistemic responsibility.

What responsibility do agents bear for their starting points? The developmental resources that distinguish Victor's intellectual inheritance from Penelope's are external to both agents. Neither agent was the salient cause of the developmental resources that constitute their intellectual heritage. Therefore, neither agent can be causally accountable for the developmental resources from which their eventual agency was produced. Neither agent's developmental resources are responses to their reasons. Moreover, the agents' intellectual heritage precedes their authentication or affirmation of that intellectual heritage. Therefore, the question as to whether the difference in their starting points is, by itself, significant to the ascription of responsibility cannot be answered by whether or not the agents identify with their intellectual heritage.

According to the internalist intuition that agents can only be responsible for what they control, agents bear no responsibility for their starting points. According to the externalist, agents do bear responsibility for their starting points. The hybrid view seeks to accommodate the intuition that it seems unfair and unreasonable to blame agents for their impoverished developmental resources. However, it also seeks to accommodate the externalist intuition that simply being as one is warrants attributing and undertaking responsibility. Accordingly, the hybrid view refrains from attributing full responsibility to agents for the developmental resources from which their agency stems. Nevertheless, in order to accommodate the idea that the intellectual resources with which agents begin are important to determining agents' aptness for future epistemic cooperation, the hybrid view holds agents objectively, though not also subjectively, accountable for the starting points from which their trajectory as epistemic agents begins.

In the same way as one is consequence accountable for beliefs that one holds by accident, so one is actual-sequence accountable for character traits, temperament qualities

and dispositions that one inherits because of the particularities of one's history. The foci of this second variant of constitutive epistemic luck are the genetic, biological, physical, and parental, familial, cultural and social determination of agents' epistemic resources. There is some inducement to reject the attribution of responsibility to agents for the families, cultures, and societies into which they are born. It seems wretchedly unfair to hold children responsible for the results of their parents', culture's, or society's failings.

There are two responses that undermine the sting of the charge that the hybrid construal of epistemic responsibility unfairly ascribes responsibilities to agents for facets of their epistemic makeup that they neither willed nor controlled. The first response underlines that the fact that people have different epistemic responsibilities solely in virtue of their different starting points *is* unfair. Some epistemic-agents-in-training are read to nightly, have their questions answered, their curiosity encouraged, are trained to reason for themselves and to attend to the judgments and reasons of different groups, while others are never read to, have their questions unanswered, their curiosity stifled, are discouraged from reasoning for themselves, or attending to judgments and verdicts of different groups. The hybrid view locates a difference in responsibilities here, given a difference in resulting characters and it also recognizes that this difference is unfair.

The fact that the hybrid view acknowledges this unfair distribution of responsibilities, however, does not constitute a reason to reject the hybrid view. To the contrary, the hybrid view highlights the unfairness present in socially, culturally, or familially constituted differences in the responsibilities agents must bear. Groups and individuals subordinated for reasons of class, race, ethnicity, gender and disability are thereby tendered an unequal share of epistemic resources. Moreover, they are more likely, given their situatedness along lines of marginalization, to face larger burdens of

resultant, circumstantial and constitutive epistemic luck. The poverty of educational resources accessible to oppressed persons places them in a position of greater difficulty in acquiring true beliefs. The inaudibility of marginalized voices to empowered ears disqualifies oppressed persons from involvement in discursive justificatory practices. The luxury of being an intrepid inquirer requires self-confidence that is 'uppity' and, therefore, quashed by oppressors.

The fact that it seems unfair to judge epistemic agents on the basis of their intellectual inheritance reflects the existing unfairness in the distribution of epistemic resources. Thus, what at first appears to be a reason to reject the hybrid approach to constitutive epistemic luck in starting points may in fact be a reason to accept it. The view makes clear how unfair it is that persons' epistemic characters, the evidentiary circumstances they face, and the truth or falsity of their beliefs, are governed by the societal positions. By underscoring the degree to which persons' statuses as epistemic agents differ solely on grounds of their differing intellectual inheritance, the hybrid approach provides both the conceptual room and the practical impetus to enact changes that would level the playing field among intellectual agents. The hybrid approach provides reason to tailor judgments about blameworthiness and praiseworthiness for agents' epistemic statuses to the different resources from which their careers as epistemic agents began.

The second response to the charge that the hybrid view enjoins unfairly saddling agents with responsibilities is met by underlining the fact that the agents in question do not bear both objective and subjective responsibility for their intellectual inheritance. Because the hybrid view demands that a trait be responsive to the agent's reasons before full responsibility is vested in the agent, the hybrid view can accommodate the distinction

of the responsibilities that accrue to someone, Sam in the preceding example, in virtue of his impoverished epistemic developmental resources and those that accrue to those whose acts or omissions, deliberately or non-deliberately, causally contributed to the iniquitous distribution of epistemic resources. The hybrid view, therefore, allows blame and credit for an agent's intellectual inheritance to accrue to those to whose reasons that intellectual inheritance is responsive. That is, the hybrid view lays blame at the doorstep of those who are responsible for others' impoverished epistemic starting points. Those who contribute directly, by successful intentional acts that limit persons' access to education, their contribution to forums of discussion, their fulfillment of epistemic virtues, to others' limited epistemic horizons are blameworthy. Those who contribute unintentionally, by upholding the status quo or by contributing to systemic discrimination that limits oppressed persons' epistemic horizons are consequence accountable for doing so.

7.3.6. Taking Responsibility for Constitutive Epistemic Luck

The second type of constitutive epistemic luck deals with agents' luck in starting points. According to that dimension of assessment, the hybrid view will insist that agents are not fully responsible for their intellectual inheritance. There is no presumption of causal contribution or self-manufacture. There is no presumption that one's starting points are somehow responsive to one's reasons. We are, nevertheless, accountable for our intellectual resources. We can squander a rich inheritance. We can redeem a poor inheritance. The central insight that accompanies the recognition of epistemic luck is that the responsibilities meted out by the contingencies of individuals' histories are not equal. The fact that Sam is born into a familial and social environment that fosters closed-mindedness is not something for which Sam warrants full blame. Attributable though it is to causal factors that lie beyond the scope of Sam's control, the fact of his closed-

mindedness nevertheless warrants epistemic reproach. He is legitimately judged inapt for epistemic pursuits. Thus, while Sam does not warrant full responsibility for the resources that foster closed-mindedness, he is responsible for being the consequence of those resources.

It is an empirical fact that agents are allocated different developmental resources. Unfair though it might be, the epistemic virtues or vices we possess as a result of the different external causal contributors to our epistemic makeup are *our* epistemic virtues and vices. From the vantage point according to which responsibility is attuned to accurate predictions of future epistemic cohorts, one's epistemic heritage is invaluable to one's responsibility.

The hybrid standpoint enables attending to both the reasons-sensitive criterion for responsibility and the actual-possession criterion for responsibility. By implementing the reasons-sensitive requirement for full responsibility, the hybrid view alleviates some of the unfairness that attends the predictive warrant of holding persons responsible for the epistemic resources from which their agency begins. Agents are not fully responsible for their epistemic heritage until and unless the justificatory reasons for the resources transmitted inter-generationally are taken up and upheld by agents as their reasons for possessing particular traits.

Responsibility for one's character traits, temperament, and disposition vests in us simply in virtue of our epistemic interdependence. I can justifiably hold you responsible for your character traits just because you are a member of my epistemic community. Nevertheless, the affirmation of the reasons for the character trait one finds oneself with, or the rejection of the reasons one finds oneself with followed by the adoption of a

different character trait consistent with the reasons one does affirm, are vectors of epistemic responsibility.

Once they develop the capacity to believe according to their conception of the norms governing epistemic rectitude, agents become accountable to those upon whom they depend and who depend upon them in turn for their epistemically significant character traits, dispositions and temperament. Once agents-to-be become agents, they become responsible for their inherited character traits, temperaments and dispositions. This is so even though the character traits that agents develop are developed from initial states that the agents neither willed nor caused.

Agents become able to be fully responsible for their externally caused initial state once they have the capacity for self-assessment and self-control. It is not mere hazard that the capacity for self-assessment arises contemporaneously with the demand that others place on the agent to be answerable for her epistemic behaviour. Such answerability involves requirement that the agent be able to cite the epistemic norm that her behaviour is in conformity with. Once an epistemic character trait is sufficiently entrenched that the proximal causal explanation for the agent's behaviour is her possession of that character trait, it becomes reasonable and prudential to conform evaluations of the agent's aptness for epistemic cooperation to the presence or absence of the character trait. It is at the juncture of the agent's ability to solicit and affirm or deny the reasons underlying epistemic character traits, dispositions and temperament, and the concurrent entry into an epistemic community characterized by agents being answerable to one another for their aptness as epistemic cooperators that objective accountability for agents' initial epistemic states is vested in agents.

In so far as the traits out of which agents develop into epistemic agents are concerned, where those traits remain epistemically inaccessible, or are not authenticated but remain unaltered, or are authenticated but not reflected in the agent's manifestation of the character trait, the agent fails to be fully responsible. That is, even agents who are able to authenticate or repudiate possessed character traits may be unable to conform their characters to the epistemic norms they endorse. In such cases, agents suffer from constitutive epistemic luck. Though they cannot be held fully responsible for those character traits that they have not authenticated, they are nevertheless actual sequence accountable for possession of said traits.

Agents are responsible, in the backward-looking juridical mode, for taking responsibility for their epistemic traits. That responsibility is vested in the agent-in-training alongside the capacity for reflection and evaluation of those character traits, temperamental features and dispositions that constitute the initial state of the agent's epistemic makeup. One does not control the initial state of one's character, temperament and dispositions. But this initial state that one has neither willed nor caused is integral to our careers as epistemic agents. Acquiring full responsibility for one's initial condition is the cost of admittance for membership in a community characterized by interdependent epistemic agency. It also affords the means to evaluate the initial state of one's own character, temperament and dispositions relative to the end of having justified true beliefs. Once an agent gains cognitive access to her own epistemic initial state, and data sufficiently rich to afford her a moderately reliable appraisal of her own character traits, dispositions, and temperament relative to the end of having justified true beliefs, she is in a position to reform or revise or uphold and build upon her own epistemic initial condition. Thereafter, the project of self-manufacture numbers among those projects the

results of which the agent may be fully responsible for, or may be circumstantially or resultantly lucky for. The agent may be successful at inculcating open-mindedness in herself, in which case she may be counted fully responsible for the modifications she has wrought in herself. She may also be unsuccessful, in which case the agent warrants some responsibility for the reasons she upholds but is unable to conform to.

7.3.7. Constitutive Epistemic Luck and Societal Responsibility

There are two different ways in which constitutive luck affects attributions of responsibility. The first variety of constitutive epistemic luck involves agents' lack of control over their epistemic characters, temperaments or dispositions. An agent who realizes that she is intellectually lazy, berates herself for being so, and yet, despite every effort, is unable to make herself otherwise exhibits lack of control over her epistemic temperament. Though her laudable target warrants approval, her inability to stop herself from being intellectually lazy warrants disapproval. Though she is worthy of praise on grounds of her reasons, she is worthy of blame on ground of her possession of the epistemic vice. Where both her reasons for and possession of a character trait are in conformity, the agent does not exhibit lack of control. She therefore merits approval, or disapproval, for her epistemic character trait.

The second way in which lack of control exerts influence over agent's characters temperaments and disposition is by exerting influence over both the reasons an agent solicits in evaluating and approving a particular character trait of hers, and her possession of the character trait. This type of luck is ubiquitous. All agents are formed by factors beyond their control. There is, nevertheless, a sense in which some agents are lucky in their epistemic training while others are not. From this perspective on luck and agency,

responsibility for agents' consequent epistemic character traits is not attributable to the agents, but to those forces that caused the agents' character traits.

Though disparities in access to epistemic resources occur within socially constructed boundaries of class, race, ethnicity, gender and disability, they are most pervasive and most pernicious across such boundaries. Subordinated persons have less access to education and fewer opportunities and fewer resources for engaging in discursive justification. They are more often placed in epistemically unpropitious circumstances. And their development of epistemic virtues, such as intellectual bravery, is stunted both by their circumstances and by their preclusion from forums in which such virtues are exhibited. Social identities, to use Miranda Fricker's words, "constrain participation in epistemic practices—practices of asserting, denying, telling, asking, giving reasons etc." (2000, 147).

First Nations Canadians experience moral and epistemic harms on grounds of their epistemic marginalization. Western denigration of epistemic virtues that are responsive to oral traditions, teaching-myths, intellectual apprenticeships, the partitioning of knowledge claims among secret societies, as well as the interweaving of scientific with spiritual claims diminishes First Nations Canadians' discursive authority in both western and traditional knowledge contexts.⁹² The policy of patriarchal assimilationism outlined in the British North America Act of 1867 professed as its goal the destruction of traditional epistemic knowledge claims, practices, and character traits.

Epistemic character traits that are virtuous in traditional contexts are construed as vices in Western individualistic epistemic contexts. For example, among members of many First Nations Persons, responses to challenges to one's truthfulness are not met with indignation but with silent acceptance of the charge accompanied by trust that the truth

will out. Self-advocacy and insistence in the face of distrust are perceived as individual epistemic faults, leading to discord and distrust, in societies where social knowledge is paramount. In consequence, First Nations Canadians are perceived as ignorant and untrustworthy by their western interlocutors. In judicial contexts, such epistemic injustice has led to dramatic moral injustices. In *Canada's Courts*, Peter McCormick outlines in blood chilling detail the difference in incarceration rates between First Nations and non-First Nations Canadians. More worrisome still, from the perspective that refuses to acknowledge or sanction the existence of different epistemic practices, are his data supporting the claim that much of that injustice is a direct consequence of mis-application of western epistemic norms to First Nations Peoples.⁹³

Generally speaking, a young First Nations person, when asked a question by someone of equal or greater social power will remain silent, even when they know the answer. No elder would ask them, in their own traditional society, to speak with authority. The development of discursive authority in a young person proceeds by being invited to listen. The young achieve the discursive authority of elders if or when even younger persons acknowledge them as elders by seeking knowledge from them. Though generalizations across dramatically different First Nations Peoples are commonly to be mistrusted, this stratification of epistemic behaviour along lines of entitlement is widespread across North American First Nations Peoples. There are serious impediments to a Judge, in a courtroom operating in an adversarial judicial system, receiving an answer from a young First Nations person, much less an answer that inspires any degree of confidence. Moreover, social cues about honesty and trustworthiness in western circles are, in the juridical context, inverted among First Nations persons. To look an elder in the

eye is an affront. It is presumptuous and disrespectful. To respond to a charge with indignation would be unthinkable.

In “Epistemic Injustice and a Role for Virtue in the Politics of Knowing” Miranda Fricker distinguishes between culpable and non-culpable epistemic injustice of just the sort experienced by First Nations Canadians before Canadian courts.⁹⁴ The epistemic injustice at issue is a hearer’s mistrust in persons’ credibility for reasons of their social otherness. It is culpable when there exists a “sound deliberative route [from the] veridical motivation to the conclusion that he should doubt his lack of trust in...[the speaker]” (2003, 171). It is non-culpable when there does not exist such a route.

Fricker concludes that, where there are no reasons that one could cite in favour of an alternative epistemic habit, distrust in social others is not something for which agents can be deemed fully responsible. However, insofar as the trait is liable to lead to justified but false beliefs the trait warrants epistemic disapproval. In Fricker’s example of non-culpable epistemic injustice, Mr. Greenleaf is unable to distrust his distrust of his son’s fiancée, Marge, as a reliable source of information about his missing son. Fricker condemns the epistemic character trait as bad from the perspective of someone who genuinely seeks a true answer. It is, however, not culpable because “[a] setting in which there is little critical awareness of gender is a setting in which no-one is in a position to possess the virtue of Reflexive Critical Openness *vis-à-vis* gender prejudice” (170). The epistemic fault, a lack of Reflexive Critical Openness, is not grounds on which Mr. Greenleaf can legitimately be blamed.

Allow that Mr. Greenleaf has met both the subjective and objective requirements on full responsibility for an epistemic character trait outlined by the hybrid conception of responsibility. In order to allay any worry that Mr. Greenleaf lacks control over his

epistemic character trait, we may suppose he has critically evaluated, and approved, his trait of discounting womens' testimony. He wholeheartedly approves the reasons underlying his distrust of women as reliable epistemic agents. Thus, according to the hybrid view on the first kind of epistemic responsibility, Mr. Greenleaf is fully responsible for his bad epistemic character trait. There is, however, something important that underlies Fricker's defence of Greenleaf. There is something to the idea that he is not culpable in a way that someone would be "who lacked the virtue of reflective Critical Openness with regard to women speakers all along, but the relevant advance in collective consciousness will mean that this shortcoming in his epistemic conduct...will have become blameworthy" (171).

According to the hybrid view, Mr. Greenleaf is fully responsible for his epistemic trait. However, because his epistemic character and the reasons available to him to evaluate his character were formed under conditions that virtually preclude him being otherwise, he is unlucky to be so. He is unlucky in the second sense of constitutive epistemic luck: He is epistemically unlucky in his starting points.

There seems something discomfiting about exonerating persons from responsibility for their epistemic injustice simply on grounds that they do not have a 'sound deliberative route' to the fact that they are being unjust. A similar disquiet attends exonerating persons reared in racist regimes for their racism because there exists no sound deliberative route to recognizing the injustice. Internalists about epistemic responsibility would not feel such disquiet. According to the control condition, so long as there are no means accessible to the agent for assessing her belief acquisition mechanism as unreliable when confronted with testimony by social others, there are no reasons for casting aspersions on the agent. Externalists about epistemic responsibility will claim

that such agents would be responsible for her epistemic injustice. The hybrid view enjoins a moderate position between internalists and externalists. It insists that someone who has no reason available to her on which to ground her mistrust of her own mistrust is culpable for this failing, though not subjectively and objectively culpable for doing so, as someone who understood that mistrusting social others constitutes an epistemic injustice, but approved her own continued mistrust would be.

An additional benefit of the hybrid view is that the apparatus of the hybrid account for attributions of epistemic responsibility can apply not only to the agent who inherits reliable or unreliable epistemic resources, but also to the agents who have made substantial causal contributions to the agent's epistemic initial condition, both deliberately and inadvertently. Thus, the causal contribution of one's physical and social environment can be credited or discredited. Those traits that are imbued in agents-in-training as a result of external causal factors lie within the scope of the imbuer's responsibility. When the normatively significant state of affairs about which responsibility ascription is at issue is the initial state of the agent-to-be's epistemic character, those to whose reasons the initial state is responsive can and do bear responsibility.

There is little question that the more information accessible about an agent's epistemic inheritance, the better able we are to make sense of an agent's epistemic character traits, dispositions and temperament. Nevertheless, the hybrid approach upholds a centrally important role for voluntariness in character formation, and a correspondingly substantial role for responsibility with respect to the external determinants of one's character.

The central emphasis of the naturalistic approach to responsibility upholds the view of agents as embedded in physical and social environments that are not of our making or of our choosing. Even those successful purposive actions for which we warrant full responsibility are caused by at least some factors that we do not control, and result in consequences that are contingent on at least some factors that we do not control. The naturalistic hybrid approach divorces responsible agency from the ideal, if it is an ideal, that human agents are self-manufactured, that they control the problems and circumstances that they face, and that they bring about all and only states of affairs and events that they will. In consequence, human agents are responsible for being who they are, even when they neither made themselves nor willed themselves to be the way that they are.

7.4. Hybrid Responsibility and Causal Epistemic Luck

Causal epistemic luck is luck in the causal antecedents that determine one's epistemic character, temperament and circumstances, luck in the causal antecedents that determine the circumstances and problems one faces, and luck in the causal antecedents that determine the particular beliefs one holds. By highlighting the apparent conflict between a person's belief being caused and it being under her control, causal epistemic luck jeopardizes those features of responsible believing, such as judicious inquiring and justifying, that agents are commonly thought to control. Causal epistemic luck imperils the very possibility of legitimately holding persons responsible by subsuming even those facets of beliefs over which agents wield control under the rubric of epistemic luck. The driving motivation for this threat is the supposition that if agents' character-responsive justified true beliefs are caused, then agents are not in control of their justifications, what they believe in response to various circumstances, or whether or not they actively partake

in the formation of their epistemic characters. If they are not in control of these things, then it seems irrational to hold them fully responsible for anything!

7.4.1. Causal Epistemic Luck and the Problem of Free belief and Determinism

The problem of causal epistemic luck applies the issues from the freewill and determinism debate to the question of responsibility for one's epistemic triumphs and failures. Causal luck concerns the compatibility of universal causal determinism with persons' responsibility for their epistemic characters, their intellectual conduct under epistemically fraught evidentiary circumstances, and the truth or falsity of their beliefs.

There are only two possible positions regarding free belief and causal determinism. One can be a compatibilist who sees no impediment to responsibility attributions in beliefs being caused, so long as they are appropriately caused, or one can be an incompatibilist who views the causation of beliefs as defeating responsibility. Incompatibilists come in two varieties. Hard determinists declare everything, including agents' justifications of their beliefs, to be caused and, hence, immune from legitimate responsibility attributions. Libertarians insist that at least some things are not caused, namely agents' purposive undertakings, and that agents may be legitimately held responsible for those things.

7.4.2. Compatibilism and Causal Epistemic Luck

If one is a compatibilist about causal determination and epistemic responsibility, causal epistemic luck presents no problem additional to those presented by resultant, circumstantial, and constitutive epistemic luck. Causal epistemic luck is luck in how one is determined by causal antecedents that are beyond the agent's control but which compel agent's belief. Causal *luck*, accordingly, is nothing over and above the following varieties of luck: Resultant luck, Circumstantial Luck and Constitutive Luck.

If one is a compatibilist, the remaining causal antecedents of agents, for which we treat them as objects of epistemic judgment, do not constitute luck. The agent's reasons for believing a proposition, which cause her to believe that proposition, are not lucky causal antecedents. The agent's character traits, temperament and dispositions, which, together with the agent's inquiring and justifying mechanisms, generate the agent's reasons for believing a proposition, and hence generate her belief in that proposition. These are not lucky causal antecedents. And the agent's identification with the reasons, causal, emotive and cognitive, that underlie and comprise her epistemic character traits, dispositions, and temperament, and generate the agent's ordered sets of reasons for believing, are not lucky.

The remaining causal antecedents for beliefs, such as illusion, hypnosis, and brain washing either fall into the category of non-moral luck because persons so caused fail to be appropriate objects of epistemic judgment, or they fall into one of the categories of epistemic luck to which the hybrid solution to problems of epistemic luck applies. If one is a compatibilist, causal luck does not present any problem in addition to what has been presented, diagnosed, and solved by the naturalist hybrid approach to responsibility.

7.4.3. Incompatibilism and Causal Epistemic Luck

If one is an incompatibilist, causal epistemic luck presents no problem additional to those presented by resultant, circumstantial and constitutive epistemic luck. If one is a hard determinist, then one cannot be held responsible for something that is at least partially beyond one's control, for one cannot be held responsible at all. If one is a libertarian, then one will insist that those justifying and inquiring procedures that are under the purview of the agent's will are not caused by antecedent states of events or affairs, but are caused instead by the agent.

The question that vexes the judicious libertarian is whether or not those facets of agents' epistemic endeavours for which responsibility seems attributable are caused. If, however, the process whereby agents justify their beliefs, employ their epistemic virtues or vices to particular circumstances, and evaluate and authenticate their own character traits is uncaused, the incompatibilist will happily attribute responsibility to those agents.

Whether one is a compatibilist or a libertarian, the problem posed by causal epistemic luck adds nothing over and above the individual problems of resultant epistemic luck, circumstantial epistemic luck and constitutive epistemic luck. For compatibilists, one's beliefs should be antecedently caused, albeit by appropriate reasons, one's belief-forming processes under epistemically fraught circumstances should be caused, albeit in response to one's epistemic character, and one's identification with one's epistemic character traits should be antecedently caused, albeit by reasons in favour of that trait. Causal *luck*, for compatibilists, would consist in the absence of the appropriate reason-sensitivity of the causal relation between one's beliefs, one's belief forming mechanisms, and one's epistemic character and their antecedents. For libertarians, the absence of causal *luck* involves not only the absence of the conditions emphasized by compatibilists, but also non-deterministic relations between reasons and belief, belief processing mechanism and epistemic character, and epistemic character and reasons. Therefore, whether one is a compatibilist or a libertarian, causal luck adds nothing to the problems of resultant, circumstantial and constitutive epistemic luck, and it detracts nothing from the hybrid solution offered here.

I think that there are sound reasons to reject the spectre of incompatibilism about responsibility and causal determinism. As should be clear by my insistence that responsibility outstrips even reasons-responsive causal control, the demand that

responsibility be limited to those beliefs over which the agent wields uncaused sole authorship is far too stringent to accommodate the types and degrees of responsibility that ordinary agents bear for the beliefs that they hold. Moreover, the evolutionary function of responsibility attributions and acceptances as means of determining agents' aptness for participation in cooperative and interdependent living would applaud, rather than deplore, causal relations among agents' reasons-sensitivity and the beliefs they actually hold. Nevertheless, there is no logical impediment to a libertarian construal of the naturalist hybrid solution to the problems of resultant, circumstantial and constitutive epistemic luck.

7.4.4. Free Belief and Determinism is a Different Problem from Epistemic Luck

There are important intersections between the problem of epistemic luck and the problem of free belief and determinism. However, they are different problems. The problem of epistemic luck concerns the coherence of norms of epistemic appraisal given the ineluctable penchant to hold persons responsible for their accidental beliefs, and the equally ineluctable commitment to the view that what matters for epistemic responsibility are non-accidental beliefs. The problem of free-belief and determinism concerns whether one can be responsible for beliefs that are caused. Not only are they different problems, the presupposition that responsibility requires freedom upon which incompatibilists and compatibilists concur is called into question by the problem of epistemic luck. Moreover, positions in the two debates do not align neatly. One may be an internalist about epistemic luck and a compatibilist about freedom and determinism. One may be an externalist about epistemic luck and an incompatibilist about freedom and determinism. Although being a hybridist about epistemic luck strongly indicates being a compatibilist about freedom and determinism, it is possible to be a libertarian hybridist.

Though the naturalistic hybrid view offered here has implications for the freedom and determinism debate, those implications are ancillary to the central objective to solve the problem of epistemic luck. The burden of this dissertation is to offer a well-reasoned, consistent diagnosis of the practical problem of the apparent incoherence of epistemic norms while doing justice to the intuitions underlying the combatant positions. The diagnosis offered meets that burden and more. It also provides a resolution to the practical problem of where and how to assign moral and epistemic responsibility when conflicts between these epistemic norms arise.

End Notes

Chapter 1

¹ Nicholas Rescher (1993) pays particular attention to both positive and negative luck.

² See Sulloway (1996).

³ As will be elaborated in Chapter 6, there are two varieties of constitutive luck. The first deals with the difference between two agents where that difference is attributable to control or lack of control over the target character trait. The second variant deals with the difference between two agents where that difference is not attributable to a difference in control wielded by the agents, but deals with a difference in the agents' starting points such that one wields control over the target character trait while the other does not.

⁴ See J. Earman and C. Glymour (1980).

⁵ See Giere (1997).

⁶ As will be further elaborated in chapter 2, justification involves both a dialectical component and an evidentiary component.

⁷ I thank Dr. Susan Campbell for this example.

⁸ See Westfall (1980).

⁹ For examples of how dimensions of moral and epistemic luck might play out in the real world, see Goodman, Allegra (2007).

¹⁰ The combination of all the varieties of moral luck and/or epistemic construed as an ongoing state would be problematic. However, because something that exhibited that state would not constitute an agent, it would not legitimately invite any sort of responsibility.

Chapter 2

¹¹ These authors are among many others who have written on this theme. These authors are particularly forthright in their analysis of the problem of moral luck as an epistemic problem, rather than a moral problem. Other authors such as Judith Jarvis Thomson (1993), Steven Sverdlik (1993) and Henning Jensen (1993) offer an important variation on the epistemic theme.

¹² Walker does not "ignore the existence of some lively, 'principled' intuitions" (245); she believes those intuitions to be wrong-minded.

¹³ See Walker (1998).

¹⁴ Norvin Richards (1993) expresses the problem with the epistemic approach as follows: "this new way of thinking is seriously incomplete: [because] some of our deepest feelings about deeds in which harm is done have no place in it" (178).

Chapter 3

¹⁵ Although the position did not receive its name until Alvin Goldman introduced the term in "The Internalist Conception of Justification" (1980), its currency dates at least as far back as Descartes.

¹⁶ See Armstrong (1973), Ramsey (1931), Unger (1968), Skyrms, (1967), and Dretske (1971)

¹⁷ Goldman also includes "remembering, good reasoning, and introspection" in his list of reliable processes.

¹⁸ This is the Müller-Lyer Illusion.

¹⁹ See Sosa (1991).

²⁰ See Greco (1999) and (2002), the later Goldman, Alvin, (1986) and (1988).

²¹ Williams employs his conception of justification in a context of combating 'Agrippan' and 'Cartesian' skepticism.

²² In addition, so long as a person's justification can, itself, be at least in part due to luck, there remains the sort of mismatch between a person's being justified and her being responsible for her justification that gives rise to conflicting intuitions about responsibility. A parallel case of epistemic luck involving justification can be constructed. Particularly when a healthy amount of skepticism about the truth-value of a proposition is appropriate, one might strive to be justified, on the subjective side, and actually be justified, on the objective side. An approach to justification that distinguishes between personal justification and grounding, as Michael Williams' does, seems particularly appropriate for dealing with such variants on resultant epistemic luck. So long as the subjective and objective conditions on full responsibility correspond, one may be fully responsible for being justified.

²³ Note that these conclusions follow from epistemic internalism, and can be undercut with externalism.

²⁴ Engel here follows Keith Lehrer and Stewart Cohen (1983).

²⁵ Kornblith (1999).

Chapter 4

²⁶ Notice that I blur Hans Reichenbach's distinction between the context of discovery and the context of justification here.

²⁷ Hybrid it is a bit of a misnomer because it implies the joining of two separate things, whereas the view of responsibility I endorse views paradigm instances of responsibility as one thing comprising two facets.

²⁸ Richmond Campbell, personal communication.

²⁹ The claim that cooperative inquiry is an evolutionarily significant group trait is defended by Boyd and Richerson (1985) and Goldman (1999).

³⁰ See Smith (1988).

³¹ This is a position I argue for in my "Socializing Evolutionary Epistemology" (2000).

³² Nietzsche, Friedrich. *The Genealogy of Morals*, essay II; see also Freud's *Civilization and Its Discontents*, Chapter VII.

³³ Presented at the symposium on the Evolution of Norms at the CPA. Winnipeg (2004).

³⁴ Richmond Campbell deserves all the credit for this example.

³⁵ See Campbell (2005).

³⁶ For a nice resolution to the problem the Libet drag poses for the problem of freewill and determinism, see Daniel Dennett's compatibilist adaptation of Robert Kane's view in Dennett (2003).

³⁷ See Edward (1990).

³⁸ This is a position that Richmond Campbell ably defends in expounding the inadequacy of primitive normative control to account for particularly human normative practices. Campbell (2005) and Campbell and Woodrow (2003) and (2004).

³⁹ This is on the assumption of iterated prisoners dilemmas. See Axelrod (1997).

⁴⁰ See Campbell and Woodrow (2003).

⁴¹ See Fricker (1998).

⁴² Fricker's primary concern is with this mismatch, and how social boundaries influence often erroneously influence credibility judgments.

⁴³ Woodrow, Jennifer. 2000. *Socializing Evolutionary Epistemology*. Unpublished.

⁴⁴ Quine (1969).

Chapter 5

⁴⁵ Similarly, examples of epistemic luck bring to light two categories of human epistemic behaviour for which praise and blame is only problematically meted out. I return to epistemic luck in the latter part of Chapter 5.

⁴⁶ Libertarians also emphasize the requirement that free actions are responsive to reasons. For example, in propounding his agency theory, Timothy O'Connor (2000) insists that reasons explain free action. O'Connor claims that "an agent causally initiates his action and that agent causation is conceptually tied to the agent's having a reason for acting" (88). Some indeterminists, Carl Ginet (2000), for example, believe that there can be free responsible action without there being any explanation that involves factors antecedent to the action in question.

⁴⁷ A first (and traditional) pass at what it means to say someone is in control of her act, A, is that she could have done other than A. However, as Harry Frankfurt pointed out, being able to do otherwise does not appear to be a necessary condition for responsibility. Frankfurt's examples have the following form: Suppose P decides to do A for reasons a, b, and c. Suppose, further, that Q also wishes for P to do A. Q, then, implants a device in P's brain which, should P's resolve to perform A waver, Q can activate and thereby cause P to perform A. As things turn out, P's resolve does not waver; P does A, just as intended, and Q does not activate the device. Now, it seems clear that P is responsible for A.

⁴⁸ Sufficient reason, here, is to be interpreted as justificatorily sufficient reason, rather than motivationally sufficient reason. An agent, thus, may concede that she has sufficient reason to do A without being sufficiently motivated to do A. This accounts for cases where an agent is too weak-willed to choose to perform what she believes she ought to choose to perform, as well as those cases where the agent makes an extra effort of the will to choose to do something she has sufficient reason not to do.

⁴⁹ Michael Hymers, personal communication.

⁵⁰ Here a state of affairs or event is norm-governed when it can be rated in accordance with norms rather than, merely, a state of affairs or event that is guided by norms.

⁵¹ Though I do not bring it up in this context, the issue of responsibility for intentions raises the issue of whether the intention A forms to bring about S/E are character-responsive. I try to flag that character-responsiveness in terms of the 'understandable pattern of reasons' that motivate A and to which her deliberative-acting mechanism is responsive. At the first level of moral assessment, A is morally responsible for intending to bring S/E about. She is a suitable object of moral assessment on grounds that she

intended to cause S/E. The additional question about whether A's reasons-responsive and reasons-reactive attempt to cause S/E is in character or out of character for her also influences the responsibility A bears for intending to bring S/E about.

⁵² Inductive inferences are legitimated by similarities between the things about which such inferences are made. Bruce Hunter, personal communication, pointed out the importance of this point for attributions of responsibility for accidental acts and beliefs.

⁵³ The example is Richard Feldman's (1998).

⁵⁴ The issue of agents being manipulated to perform particular actions exacerbates this problem. We require an account capable of noting the similarities between cases of agent manipulation and agent accidents.

⁵⁵ The view of action I espouse is meant to be ordinary and not 'theoretical'. It is not meant to offer a reductionist theory of action. I do not think that one can capture, without remainder, everything important about action by describing action in, e.g., neuro-chemical terms.

⁵⁶ The child may also be partially assessable as causally responsible, although the appropriateness of doing so will depend upon the child's position qua member of her moral community.

⁵⁷ A may also be a prospective participating member of a community for whom S/E is norm-governed. We can justify holding children consequence-responsible even though they are not full-fledged participants on pragmatic grounds that one is teaching the normative structure. Participation in moral communities comes in different types and with differing degrees.

⁵⁸ Quoted in Watson, (2003, 79).

⁵⁹ Bernard Williams articulates this objection as follows: "if in full consciousness I could will to acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e., as something purporting to represent reality" (1973, 148)

⁶⁰ Alston is not alone in claiming that belief is not under voluntary control. The view is also espoused by Jonathan Bennett (1990), and Richard Feldman (1988).

⁶¹ One could also embark on a regimen of information designed to produce false beliefs, such as soliciting information from disreputable and skewed sources. This would be epistemically irresponsible, but it would nevertheless be voluntary.

⁶² Lorraine Code claims that “freedom, both in action and in belief, is constrained to some degree, but that this does not render it nonexistent or trivial” (1987, 85).

Michael Stocker claims that “if we can be responsible and active for various sorts of interesting and important physical acts, we can be responsible for mental goings-on, including beliefs” (1982, 417)

⁶³ If, so far as I can tell, there are facts that strongly support the supposition that *p*, then surely it is *all right* for me to give my assent to *p*. What more could be *demand*ed of me? I have done all I can. What the actual facts are over and above what I am most justified in believing is something I cannot be held *responsible* for. Once I have marshalled all the cognitive resources available to me to determine the matter I have, in my body of justified beliefs, the closest approximation I can make to the actual facts. That is the best I have to go on, and it would be quite unreasonable to suggest that I ought to be going on something else instead (Alston, 2001, 80).

⁶⁴ It is not vanishingly improbable that if there were a sufficient reason not to believe as *A* believes, *A* would take up that reason, and so choose not to believe as *A* believes.

⁶⁵ *A* may also be a prospective participating member of a community for whom *S/E* is norm-governed. We can justify holding children consequence-responsible even though they are not full-fledged participants on pragmatic grounds that one is teaching the normative structure. Participation in moral communities comes in different types and with differing degrees.

Chapter 6

⁶⁶ Agent regret, reproach, indignation and other emotional responses particularly substantiate this.

⁶⁷ 1 and 2 tend to collapse for Kant, but not entirely since one can get it wrong about what act is right according to the categorical imperative, even though one intends to do what is right.

⁶⁸ Williams, Walker, Card, Nussbaum, and Aristotle, to mention a few.

⁶⁹ See Scanlon (2003).

⁷⁰ This is the point of the virtues of grace and lucidity that Walker outlines.

⁷¹ It is important to underline that the relationships outlined here are not merely evidentiary.

⁷² Although some see opportunities where others don't, and people can create situations.

⁷³ See Latus (2003).

⁷⁴ See Adams (1985).

⁷⁵ They also sometimes deliberately and successfully worsen themselves.

⁷⁶ The dearth of the concept of sexual harassment is another example in which there fail to exist reasons for the unease and debasement suffered by women who underwent unwanted advances.

⁷⁷ It invites equivocation on at least two fronts. The first is attribution. The second is responsibility in terms of origination and voluntarism.

⁷⁸ For example, Anita Allen writes "we are responsible people who deserve self-esteem and the esteem of others, not because we have blindly adhered to communal norms but, rather, because we have chosen to regulate our conduct in accordance with reasonable social expectations" (1998, 107).

⁷⁹ That the society into which one is born influences what sort of moral agent one will become is, of course, an instance of moral luck.

⁸⁰ The same goes for non-agents, though the latter are not the appropriate subjects of moral judgment. As Margaret Walker highlights, "Even if the burdens morality assigns are not equal...nonetheless we are all *equally* judgeable in the *same* deep way in light of such burdens as have fallen to us" (1985, 323).

⁸¹ Not to mention physiology, psychology, biology...etc.

⁸² Lugones endorses 'ontological pluralism' claiming that multi-personhood, though engendered by the dual demands of one's membership in both oppressed and oppressing society, is also a means to combat oppression.

⁸³ The character trait in question might be antithetical to the oppressed person's character-trait or inimical to the subordinate group of which she is a member. Moreover, the character trait, disposition, or temperament might not be a virtue after-all, but merely a trait commonly possessed by oppressors. These are important variations on the intersection between oppression and constitutive moral luck. Unfortunately, they must await a different forum.

⁸⁴ That enjoined by Robert Kane as well.

⁸⁵ Recall that the central premise of incompatibilism is that one cannot be both free and caused. One could be a libertarian who accepts that one can be responsible for some of the things one does not do freely.

⁸⁶ The hybrid view of moral responsibility in its ordinary functioning serves to remedy some compatibilist dialectical infelicities. The major infelicity is the regress argument. By demonstrating that responsibility is not transitive, and cannot be conjoined

cumulatively, the hybrid view of moral responsibility in its ordinary functioning solves the major problem that compatibilists have in holding sway over the debate. It is not the devil in the details; it's the solution.

Chapter 7

⁸⁷ There are other determinants of worthiness for epistemic collaboration. Particularly reflective reasoners, for example, might look at their history of past beliefs and find many to have been false. This legitimates a certain amount of chariness about affirming the truth of current beliefs.

⁸⁸ Harman (1996).

⁸⁹ See Sir Alexander Fleming. *Encyclopaedia Britannica Inc* (1958).

⁹⁰ See Chamberlain. *Encyclopaedia Britannica Inc* (1958).

⁹¹ Zagzebski (1999), Code (1987), Fricker (2003), and Dalmiya (2002).

⁹² Louis Soop, Elder, Blackfoot Nation, personal communication, and Kevin Healy, Horn Society Elder, Blackfoot Nation, personal communication, Lori Healy, Motokiiks Society Elder, personal communication.

⁹³ See McCormick (1994).

⁹⁴ Fricker (2003).

Reference List

- Adams, Robert M. 1985. "Involuntary Sins." *The Philosophical Review*, XCIV No.1: 3-31.
- Allen, Anita. 1998. *Debating Democracy's Discontent: Essays on American Politics, Law and Public Philosophy*. New York: Oxford University Press.
- Alston, William. 2001. "Internalism and Externalism in Epistemology." In *Epistemology: Internalism and Externalism*, ed. Kornblith, Hilary. Malden, Mass.: Blackwell Publishers Inc.
- Anscombe, G.E.M. 1957. "Direction of Fit." In *Intention*. Oxford: Blackwell.
- Armstrong, David. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Axelrod, Robert. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton: Princeton University Press.
- Babbitt, Susan. 1996. *Impossible Dreams: Rationality, Integrity, and Moral Imagination*. Boulder, Co: Westview Press.
- Bennett, Jonathan. 1990. "Why is Belief Involuntary?" *Analysis* 50: 87-107.
- Benson, Paul. 2000. "Feeling Crazy: Self-Worth and the Social Character of Responsibility." In *Relational Autonomy: Feminist Perspective on Autonomy, Agency and the Social Self*, MacKenzie, Catriona and Stoliar, Natalie, eds. New York: Oxford University Press.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, MA: Cambridge University Press.
- . 1987. "Nozick, 'Externalism and Skepticism.'" In *The Possibility of Knowledge*, ed. Luper-Foy, Steven. Totawa, New Jersey: Rowman & Littlefield.
- Boyd, R. and Richerson, P. 1985. *Culture and Evolutionary Process*. Chicago: The University of Chicago Press.
- . 1992. "Punishment Allows the Evolution of Cooperation (or Anything Else) in Sizeable Groups." *Ethology and Sociobiology* 13: 171-195.
- Broad, William and Wade, Nicholas. 1982. *Betrayers of the Truth: Fraud and Deceit in Science*. Oxford: Oxford University Press.
- Campbell, Richmond. 2005. "Two Faces of Moral Judgment." Unpublished.
- Campbell, Richmond and Woodrow, Jennifer. 2003. "Why Moore's Open Question is Open: The Evolution of Moral Supervenience." *The Journal of Value Inquiry* 37: 353-372.
- . 2003. "How to be both Docile and Autonomous: The Drive to Make Reasons Explicit." Unpublished.

- Card, Claudia. 1996. *The Unnatural Lottery: Character and Moral Luck*. Philadelphia: Temple University Press.
- Chamberlain, Neville. 1958. *Encyclopaedia Britannica Inc.* Vol. 6. 170.
- Code, Lorraine. 1987. *Epistemic Responsibility*. Hanover: University Press of New England. 85.
- Collins, Harry and Pinch, Trevor. 1993. *The Golem: What Everyone Should Know about Science*. Cambridge: Cambridge University Press.
- Coyne (Walker), Margaret Urban. 1985. "Moral Luck?" *The Journal of Value Inquiry* 19: 319-325.
- Craig, Edward. 1990. *Knowledge in the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Davion, Victoria. 1991. "Integrity and Radical Change." In *Feminist Ethics*, ed. Card, Claudia. Lawrence: University Press of Kansas. 180-194.
- Dennett, Daniel. 2003. *Freedom Evolves*. New York: Penguin Group Inc.
- Descartes, Rene. 1955. "Meditation 4." In *Vol. I of Philosophical Works of Descartes*, Haldane and Ross eds. Dover.
- Dretske, Fred. 1971. "Conclusive Reasons." *Australasian Journal of Philosophy* 49:1-22.
- Dretske, Fred and Berent Eng. 1984. "Causal Theories of Knowledge." *Midwest Studies in Philosophy*, vol. 9: 517-528.
- Earman, J. and Glymour, C. 1980. "Relativity and Eclipses: The British Eclipse Expeditions of 1919 and their Predecessors." *Historical Studies in the Physical Sciences*, 11 (I): 49-85.
- Ekstrom, Laura Waddell. 2001. *Agency and Responsibility: Essays on the Metaphysics of Freedom*, ed. Ekstrom, Laura Wadell. Boulder Colorado: Westview Press.
- Engel Jr., Mylan. 1992. "Is Luck Compatible With Knowledge?" *The Southern Journal of Philosophy* Vol.XXX, No.2: 59-75.
- Feldman, Richard. 1988. "Epistemic Obligation." In *Philosophical Perspectives* 2, ed. Tomberlin, James E. Atascadero, CA: Ridgeview.
- Fischer, John M. and Ravizza, Mark. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fleming, Sir Alexander. 1958. *Encyclopedia Brittanica Inc.* Volume 9. 370
- Frankfurt, Harry. 1969. "Alternate Possibilities and Moral Responsibility." *Journal of Philosophy*, 66/23: 829-839.

- . 1988. *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press.
- Freud, Sigmund. 1989. *Civilization and Its Discontents*, Chapter VII. W.W. Norton & Company.
- Fricker, Miranda. 1998. "Rational Authority and Social Power: Towards a Truly Social Epistemology." Meeting of the Aristotelian Society.
- Fricker, Miranda. 2000. "Feminism in Epistemology: Pluralism without Postmodernism." In *The Cambridge Companion to Feminism in Philosophy*. Fricker, Miranda and Hornsby, Jennifer, eds. Cambridge: Cambridge University Press. 146-166.
- . 2003. "Epistemic Injustice and a Role for Virtue in the Politics of Knowing." *Metaphilosophy* Vol. 34, Nos. 1/2.
- . 1985. "Moral Integrity and the Deferential Wife." *Philosophical Studies* 47: 141-150.
- Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" In *The Theory of Knowledge: Classical and Contemporary Readings*, 2d ed. Belmont, Ca: Wadsworth Publishing Company. 142-143.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Giere, Ronald. 1997. *Understanding Scientific Reasoning*. 4d. ed. Belmont: Wadsworth Group.
- Ginet, Carl. 2000. "Epistemic Requirements for Moral Responsibility." *Nous-Supplement: Philosophical Perspectives*, vol. 14: Action and Freedom: 267-277.
- Goldman, Alvin. 1999. *Knowledge in a Social World*. Oxford: Clarendon Press.
- . 1998. "Reliabilism: What is Justified Belief?" In *The Theory of Knowledge: Classic and Contemporary Readings*. 2d ed. Ed. Poiman, Louis P. Belmont, CA: Wadsworth Publishing Company.
- . 1988. "Strong and Weak Justification." *Philosophical Perspectives* 2: 51-69.
- . 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- . 1980. "The Internalist Conception of Justification." In *Studies in Midwest Philosophy V: Studies in Epistemology*, French, P et.al. eds. Minneapolis: University of Minnesota Press.
- Goodman, Allegra. 2006. *Intuition*. Random House.
- Greco, John. 2002. "Knowledge as Credit for True Belief." In *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Michael DePaul and Linda Zagzebski, eds. Oxford: Oxford University Press.

- . 1999. "Agent Reliabilism." *Philosophical Perspectives* 13: 273-96.
- Greenough, W.T., and F. R. Volkman. 1972. "Rearing Complexity Affects Branching of Dendrites in the Visual Cortex of the Rat." *Science* 176: 1445-1447.
- Haman, Gilbert. 2000. *Explaining Value: And Other Essays in Moral Philosophy*. Oxford: Clarendon Press.
- Hart, H.L.A. and Honore, A.M. 1959. *Causation in the Law*. Oxford: Clarendon Press.
- Hill Jr., Thomas E. 1979. "Servility and self-respect." In *Today's Moral Problems*. 2d ed. Ed. Wasserstrom, Richard A. Macmillan, New York: 133-147.
- Hurley, Susan. 2006. "Justice, Luck, and Knowledge." *Ethics: An International Journal of Social, Political, and Legal Philosophy*, vol. 116, no. 2: 428-432.
- . 1993. "Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 57-71.
- Jacobs, Jonathan. 2001. *Choosing Character: Responsibility for Virtue and Vice*. Ithaca: Cornell University Press.
- Jenson, Henning. 1993. "Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press.
- Kane, Robert. 2002. *The Oxford Handbook of Free Will*, ed. Robert Kane. New York: Oxford University Press.
- Kant, Immanuel. 1979. *Mortal Questions*. New York: Cambridge University Press.
- Kornblith, Hilary. 2001. "Introduction." In *Epistemology: Internalism and Externalism*, ed. Kornblith, Hilary. Malden, Mass.: Blackwell.
- . 1999. "How Internal Can You Get?" In *The Theory of Knowledge: Classical and Contemporary Readings*, 2nd Edition. ed. Louis P. Pojman. Belmont, Ca: Wadsworth.
- Korsgaard, Christine. 1996. *The Sources of Normativity*. New York: Cambridge University Press.
- Latus, Andrew. 2003. "Constitutive Luck" *Metaphilosophy* Vol. 34, No. 4.
- Lehrer, Keith and Cohen, Stewart. 1983. "Justification, Truth, and Coherence." *Synthese* 55: 191.
- Lehrer, Keith. 1999. "A Critique of Externalism." In *The Theory of Knowledge: Classical and Contemporary Readings*, 2nd Edition. ed. Louis P. Pojman. Belmont, Ca: Wadsworth.
- Lugones, Maria C. 1990. "Structure/Antistructure and Agency under Oppression." *Journal of Philosophy* 87: 500-507.

- Mackie, J.L. 1974. *The Cement of the Universe*. Oxford: Oxford University Press.
- McCormick, Peter. 1994. *Canada's Courts*. Toronto: James Lorimer & Company.
- Mele, Alfred. 2000. "Goal Directed Action: Teleological Explanations, Causal Theories, and Deviance." In *Philosophical Perspectives 14, Action and Freedom*. Tomberlin, James, ed. Atascadero: Ridgeview Publishing Co.
- Milikan, Ruth. 1995. "Pushmi-Pullyu Representations." In *Philosophical Perspectives 9: AI, Connectionism, and Philosophical Psychology*. Tomberlin, James, ed. Atascadero: Ridgeview Publishing Co.
- Nagel, Thomas. 1993. "Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 57-71.
- Nietzsche, Friedrich. 1996. *The Genealogy of Morals, essay II*. trans. Douglas Smith New York: Oxford University Press.
- Nussbaum, Martha. 1993. "Luck and Ethics." In *Moral Luck*, ed. Statman, Daniel. State University of New York Press. 73-108.
- O'Connor, Timothy. 2000. *Persons and Causes: The Metaphysics of Free Will*. New York: Oxford University Press.
- Plantinga, Alvin. 2003. "Epistemic Luck." *Journal of Philosophical Research* 28. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- . 1999. "A Critique of Internalism." In *The Theory of Knowledge: Classical and Contemporary Readings* 2d ed. Ed. Pojman, Louis P. Belmont, CA: Wadsworth Publishing Company.
- Pritchard, Duncan. 2006. "Epistemic Luck." *Metaphilosophy*, vol. 37, no. 5: 728-736.
- Quine, W.V.O. 1969a. "Epistemology Naturalized." In *Ontological Relativity and Other Essays*. New York: Columbia University Press. 68-90.
- Ramsey, Frank. 1931. *The Foundations of Mathematics, and Other Logical Essays*. Routledge.
- Rescher, Nicholas. 1995a. "Prediction." In *The Oxford Companion to Philosophy*, ed. Ted Honderich. Oxford: Oxford University Press. 713.
- . 1995b. *Luck: The Brilliant Randomness of Everyday Life*. New York: Farrar, Straus, and Giroux.
- . 1993. "Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 141-166.
- Richards, Norvin. 1993. "Luck and Desert." In *Moral Luck*, ed. Statman, Daniel. State University of New York Press. 167-180.

- Rorty, Amelie O., and Wong, David. 1990. "Aspects of Agency and Identity." In *Identity Character and Morality*. Flanagan, Owen and Rorty, Amelie O., eds. Cambridge: MIT. 19-36.
- Rosebury, B. 1999. "Moral Responsibility and Moral Luck". *The Philosophical Review* 104: 499-524.
- Scanlon, T. M. 2003. "The Significance of Choice." In *Free Will*, ed. Watson, Gary. Oxford University Press.
- Skyrms, Brian. 1967. "The Explication of 'X knows that p'". *Journal of Philosophy* 64. 373-389.
- Smith, John Maynard. 1988. "Models of Cultural and Genetic Change." *Games Sex and Evolution*, Hemel Hempstead, UK: Harvester.
- Sorensen, Roy. 1998. "Logical Luck." *The Philosophical Quarterly*, Vol. 48, No. 192: 319-334.
- Sosa, Ernest. 1991. *Knowledge in Perspective*, Cambridge: Cambridge University Press.
- Statman, Daniel. 1993. "Introduction." In *Moral Luck*, ed. Daniel Statman. State University of New York Press.
- . 1993. *Moral Luck*, ed. Daniel Statman. State University of New York Press.
- . 1991. "Moral and Epistemic Luck." *Ratio (New Series)* IV 2.
- Stocker, Michael. 1982. "Responsibility Especially for Beliefs." *Mind* 91, No. 363.
- Sulloway, F. 1996. *Born to rebel: Birth Order, family dynamics, and Creative Lives*. New York: Pantheon.
- Sverdlik, Steven. 1993. "Crime and Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press.
- Taylor, Richard. 1999. "Freedom and Determinism." In *Introduction to Philosophy: Classical and Contemporary Readings*. 3d ed. Perry, J. and Bratman, M. eds. New York: Oxford University Press. 437-449.
- Thomson, Judith Jarvis. 1991. "Morality and Bad Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 35-55.
- Unger, Peter. 1968. "An Analysis of Factual Knowledge." *The Journal of Philosophy* 65: 157-170.
- Vahid, Hamid. 1993. "Are Old Wives' Tales' Justified?" In *Feminist Epistemologies*, Alcoff, Linda (ed). New York: Routledge.
- . 1990. "Coherence, Truth, and the Omniscient Interpreter". *Philosophical Quarterly*, vol. 40, no. 158, pp. 86-94, January 1990

- Walker, Margaret Urban. 1991. "Moral Luck and the Virtues of Impure Agency." In *Moral Luck*, ed. Daniel Statman. State University of New York Press.
- . 1998. *Moral Understandings: A Feminist Study in Ethics*. New York: Routledge.
- . 1985. "Moral Luck?" *The Journal of Value Inquiry*. 19: 319-325.
- Watson, Gary. 2003. *Free Will*, 2d ed. Oxford: Clarendon Press.
- Westfall, R.S. 1980. *Never at Rest: A Biography of Isaac Newton*. Cambridge: Cambridge University Press.
- Williams, Bernard. 1993. "Moral Luck." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 35-57.
- . 1985. *Ethics and the Limits of Philosophy*. Glasgow: William Collins Sons & Co.
- . 1973. *Problems of the Self*. Cambridge: Cambridge University Press.
- Williams, Michael. 2001. "Contextualism, Externalism and Epistemic Standards." *Philosophical Studies* 103: 1-23.
- . 1999. *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford University Press.
- Wilson, George. 1999. "Reasons as Causes for Action." In *Contemporary Action Theory, Vol. 1: Individual Action*, In *Philosophy in Review, Vol. 19*. Holmstrom-Hintikka G. and Tuomela, R. eds. Dordrecht: Kluwer. 65-82.
- . 1989. *The Intentionality of Human Action*. Stanford: Stanford University Press.
- Woodrow, Jennifer. 2000. *Socializing Evolutionary Epistemology*. Unpublished.
- Zagzebski, Linda. 1999. "What is Knowledge?" In *Epistemology*, Greco J. and Sosa, E. eds. Oxford: Basil Blackwell. 92-116.
- . 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.
- . 1994. "The Inescapability of Gettier Problems." *Philosophical Quarterly* 44: 65-73.
- Zimmerman, David. 2002. "Reasons-Responsiveness and Ownership-of-Agency: Fischer and Ravizza's Historicist theory of Responsibility." *The Journal of Ethics* 6: 199.
- Zimmerman, Michael J. "Luck and Moral Responsibility." In *Moral Luck*, ed. Daniel Statman. State University of New York Press. 217-234.