

Syracuse University

SURFACE

Philosophy - Dissertations

College of Arts and Sciences

6-2012

Moral Luck

Patrick A. Beach
Syracuse University

Follow this and additional works at: https://surface.syr.edu/phi_etd



Part of the [Philosophy Commons](#)

Recommended Citation

Beach, Patrick A., "Moral Luck" (2012). *Philosophy - Dissertations*. 68.
https://surface.syr.edu/phi_etd/68

This Dissertation is brought to you for free and open access by the College of Arts and Sciences at SURFACE. It has been accepted for inclusion in Philosophy - Dissertations by an authorized administrator of SURFACE. For more information, please contact surface@syr.edu.

ABSTRACT

The problem of moral luck arises because of an apparent conflict between the luck that humans face, the pervasiveness of moral considerations in our lives, and the idea that moral responsibility requires a level of control that luck robs them of. The claim is that we cannot be governed by luck and be morally responsible. The notions of luck and moral responsibility are developed, and several solutions to the problem of moral luck are discussed and rejected. Finally, a solution is offered that dissolves the problem. In particular, the luck that affects an agent's moral standing comes from the choices that are presented to her from given character traits. This variety of moral luck is found to be genuine and yet to not conflict with a principle that moral responsibility requires control. The agent has control over what choices she makes from her palette of live options, no matter if the range of options is given to her without her choice. So while she lacks complete control, she has the right kind of control for moral responsibility. Hence, there is moral luck, but no *problem* of moral luck.

MORAL LUCK

By

Patrick Beach
B.A. Boise State University, 1984
M.A. Miami University, 1987

DISSERTATION

Submitted in partial fulfillment of the requirements for the
degree of Doctor of Philosophy in Philosophy
in the Graduate School of Syracuse University

June 2012

Copyright 2012 Patrick Beach
All rights reserved.

Acknowledgments

I first came to Syracuse University in 1988, so the completion of my doctoral degree has been sidetracked for one reason or another for more than 15 years. That's a long time, and it's been a difficult road getting my philosophical legs under me again, and carving out time to actually write, especially given that I've worked full time and taught half time during most of the dissertation process. Nothing in my life has been done alone, and it's proper to give thanks to those who have helped me.

I want to thank, first and foremost, my core committee members, without whom this dissertation would not have been possible: Ben Bradley, Ken Baynes, and Mark Heller. Special thanks are due to Ben Bradley who, as my dissertation adviser, did the yeoman's work. Thanks also to John Robertson for serving on my defense committee. Special thanks to Michael Zimmerman for the time and effort to come to Syracuse as a committee member. His work on moral luck has been central to my thinking, and I'm fortunate to have him on my defense committee.

I would also like to thank the professors and fellow graduate students from my early days at Syracuse for their help and guidance. In particular, I would like to single out Jonathan Bennett as the most influential teacher I've ever had the pleasure of studying under. The seeds to this work were sown in his seminar on consequentialism.

Closer to home, and in the 21st century now, I'd like to thank the members of the Boise State University Philosophy Department for their support. Andrew Cortens has given me courses to teach to keep me active in the profession, and Stephen Crowley, in particular, offered much late-night encouragement.

More than anything, however, I was able to complete this with the support of many friends: Carol Dorsey, Della Hawkins, Jerry Luther, Chris Knight, Deborah Knight, Carrie Morgan, Denae Barowsky, John Keller, Rob Russell, Ken Reed, and Travis Moats. Hud Hudson and Randy Olson have remained true friends and

philosophical heroes. I'm lucky to have had Lisa Farnsworth helped me with some of the mechanics of getting through the dissertation process after so many years' absence.

Thanks goes to my family which has accepted that I disappear for months at a time.

Finally, I'd like to thank John Liebenthal for year after year of Sunday morning coffees.

I dedicate this to the memories of two: Gerry Frachiseur and Marc Smith.

Table of Contents

Chapter 1: Introduction:	8
1. Moral Luck: A Preliminary Characterization	8
2. The Plan	11
Chapter 2: The Nature of Luck: A Causal-Contrastive Account	13
1. Introduction	13
2. Preliminary Points	15
2.1. Dictionary Definitions	16
2.2. Experimental Philosophy	18
2.3. The Possibility of Varying Spheres of Luck	22
2.4. Metaphysical and Epistemic Considerations	23
3. The Lack-of-Control Account of Luck	26
3.1. Paradigmatic Cases	29
3.2. The Lack-of-Control Account Evaluated	37
3.3. The Big Problem for the Lack-of-Control Account	43
4. The Modal Account of Luck	45
4.1. The Scalar Objection	48
4.2. Lackey’s Buried Treasure Objection	49
4.3. The Whimsical Trip Objection	52
4.4. The Non-Causal Objection	53
5. The Probability-Value Account of Luck	54
5.1. Critique of Rescher’s Analysis	58
6. The Probability-Control-Value Account of Luck	60
7. The Sensitive-Causal-Chains Account of Luck	63
7.1. Formulating the Causal-Contrastive Account	68
8. Objections and Replies	73
Chapter 3: Moral Responsibility:	77
1. The Nature of Moral Responsibility	77
2. Hume’s Fork	79
2.1. Free Will and Causal Determinism	80
2.2. Free Will and Causal Indeterminism	82
3. Understanding Free Will	83
4. Challenges to Libertarian Free Will	87
4.1. P.F. Strawson’s Reactive Attitudes	88
4.2. Harry Frankfurt on Alternative Possibilities	91
5. Functions of Praise and Blame	94
6. Accountability and Excuses	95
Chapter 4: The Problem of Moral Luck:	99
1. Three Varieties of Moral Luck	99
2. Kant and Moral Evaluation	101
3. The Problem of Moral Luck	105
4. CP and the Causal-Contrastive Account of Luck	108
5. Williams on Moral Luck	109
6. Critique of Williams on Retrospective Justification	112
7. Williams on Agent Regret	118

8. Nagel on Moral Luck	124
Chapter 5: Solutions to the Problem of Moral Luck:	133
1. Kinds of Solutions	133
2. Richards' Epistemic Solution: Resultant Luck	134
3. Richards' Epistemic Solution: Constitutive and Circumstantial Luck	138
4. Zimmerman's Argument Against Moral Skepticism	144
5. Zimmerman and Genuine Moral Luck	147
6. Does PCP Neutralize Luck?	158
7. Understanding Constitutive Moral Luck and the Control Principle	162
8. The Problem of Moral Luck Revisited	166
Chapter 6: Concluding Thoughts:	170
1. What We've Covered	170
1.1. The Nature of Luck	171
1.2. Moral Responsibility	172
1.3. The Problem of Moral Luck	172
1.4. Solutions to the Problem of Moral Luck	174
2. Final Thoughts	178
Works Cited	181
Vita	185

Chapter 1

Introduction

1. Moral Luck: A Preliminary Characterization

Suppose that Fred murders George and that Ida unsuccessfully attempts to murder Kevin with the same zeal and reasonable expectation of success as Fred. Due to the intervention of luck, a bird flying in the path of the bullet perhaps, Ida is unsuccessful in her attempt. While not a universal moral reaction, many people judge Fred more critically than they do Ida.

Take another case. Suppose that Della and Helen have the same generous characters, but only Della faces circumstances in which her generosity is called upon. While they have the same admirable dispositions, Della has done something praiseworthy, and Helen has not. We can point to Della's generous act, and there is no corresponding act of Helen's that we can point to. It seems, then, that their moral records differ by the influence of luck.

As a third example, consider a son who has certain character traits that his father lacks--for example, a tolerance of people from other cultures. Suppose again that the difference between them can be attributed to luck. The son has virtues that his father lacks because the son was raised in a culture that values tolerance. Again, the son is morally better in some ways than his father (he is after all more tolerant), but the reason has nothing to do with a lack of moral genuineness on his father's part. Rather, the father was raised at an inopportune time in human history, and his intolerant character is directly attributable to that.

In short, our lives are governed by luck, and our moral records seem to differ because of it. The phenomenon of moral luck can be characterized as any situation in

which the moral status of the agent is affected by the intervention of luck.¹ If luck affects the moral status of her actions, intentions, or character traits, and thereby the praise or blame that can be correctly attributed to her, then she is subject to moral luck. On the other hand, these cases may be instances of only apparent moral luck. If luck is only apparently morally significant, we see the difference that luck sometimes makes but ought to ignore its influence in our moral evaluations.

It is reasonable to believe that our constitutions, the circumstances in which we act, and the effects we are able to bring about are riddled with luck. Yet when we consider ourselves merely as parts of the furniture of the universe, it poses no special philosophical problem. When we consider ourselves as having free will, as being part of the moral order, as having moral responsibility, however, the issue of luck raises significant philosophical problems.

On the Kantian understanding of morality in particular, it's never correct to judge an agent for what is not in her control. According to Kantianism, one's moral record can never be worsened or augmented by the intervention of luck. We can thus present the problem of moral luck as arising because we are committed to an inconsistent triad. Each of these propositions seems to be true, and yet they can't all be true: (1) very little of our lives are under our control because of the ubiquity of luck; (2) we are not morally responsible for what is not in our control; and (3) we are morally responsible to a significant degree.

¹ My focus is on whether the agent's moral responsibility--primarily her praiseworthiness or blameworthiness--is affected by luck. As Michael Zimmerman points out, we could also discuss how luck affects our aretaic judgments of the agent (whether she has certain virtues or vices) or our deontic judgments (our evaluations of the agent's moral obligations). Thus in Zimmerman's view, my discussion will be only partial (see (Zimmerman, "TLS" 554–55) and (Zimmerman, "Map" 585)).

If we cannot somehow resolve that inconsistent triad, we're stuck with options each of which is unpalatable. First, we could deny that our lives are largely subject to luck or find some understanding of luck in which the amount of luck in our lives is minimized. The problem with this solution is that it's very easy to motivate the thought that luck is ubiquitous in a way that is ripe for moral significance.² We're lucky to the extent that certain factors are out of our control: (a) our given *character traits* are good or bad, (b) the *circumstances* in which we find ourselves engage or fail to engage our moral characters, and (c) the chance *consequences* we bring about are good or bad. If we consider these three kinds of influence (the way we are, the circumstances in which we act, and the downstream effects of our choices), there's little--perhaps nothing--that is under our complete control and therefore immune from luck. Thus, as I'll show more fully in Chapter Two, luck is ubiquitous.

Second, we could deny that moral evaluation applies only to what is under the agent's control, yet the Kantian idea that we're morally evaluable only for what is under our control is well entrenched. Third, we could accept the ubiquity of luck and the control constraint on moral evaluation and deny that we are morally responsible, or at least deny that our range of morally evaluable actions, intentions, and character traits is nearly as extensive as we have supposed. At worst, then, we'd be forced into a moral skepticism. At best, we'd have some limited use for the notion of moral responsibility, and we'd have to extensively revise our notion of ourselves. In order that we might solve the problem of moral luck, it will be helpful to get clear on the nature of luck and on the nature of moral responsibility.

Let me pause to highlight one central virtue of my account of luck over others extant in the literature. As we'll see, some purported examples of moral luck deal with

² Some luck is not morally significant, of course. Being born into wealth, for example, is lucky but no moral good is inherently at stake.

the causally downstream results of the agent's choices and some pertain to an action's causal predecessors, either the circumstances in which he finds himself or the character traits he currently has. My account of the nature of luck is causally contrastive: it's a causal account that measures the amount of luck by how else things might have turned out. Thus, it's able to account for lucky events whether they are causal contributors to or causal effects of action. In addition, it invites reflection not only on what the agent did but also on what he did not do, but might have done instead. Furthermore, it invites us to compare the values produced by the agent with the values he might have produced instead. Thus, luck is not only causally contrastive; it is also value contrastive.

No other account of luck has such rich resources since they (a) point downstream only and (b) compare the value of the outcome with the value of the initial event, rather than comparing the values of the alternative ways in which the world might have unfolded.

2. The Plan

Chapters 2 and 3 are devoted to luck and moral responsibility, respectively. Chapter 4 introduces the problem of moral luck as it was presented in a pair of important papers by Bernard Williams and Thomas Nagel in 1976.³ I examine and reject Williams' presentation of the problem, but accept Nagel's conception of it. Nagel doesn't offer a solution, however, so the discussion of proposed solutions commences in Chapter 5.

While there are many avenues of discussion of the problem, two discussions particularly interest to me. First, I discuss Norvin Richards' epistemic solution according to which there is no genuine moral luck but only luck in our *epistemic access* to whether agents deserve praise or blame. His discussion is interesting and correct so far as it goes,

³ References to these papers will be to reprinted versions contained in a standard anthology of readings on moral luck (Statman, D.).

but he doesn't show that the moral luck to which we are subject is *only* epistemic. I then discuss a line of argument presented by Michael Zimmerman that endorses genuine moral luck. I object to his argument, but think that his central idea is salvageable. I expand on it and present what I take to be an adequate answer. Finally, Chapter 6 concludes this work and provides further reflections.

Chapter 2

The Nature of Luck:

A Causal-Contrastive Account

1. Introduction

The concept of luck invades several areas of philosophy. It holds a place in the metaphysics of free will; in epistemology, especially since Gettier's famous paper but also in reference to radical skepticism; and in moral philosophy going back to Ancient Greece, but revived primarily in the decades since Bernard Williams and Thomas Nagel wrote their now famous papers on moral luck. In many of the discussions of luck in metaphysics, epistemology, and ethics, one shortcoming is the lack of a thorough or accurate presentation on the nature of luck. It seems desirable, however, to have a firm grasp of the nature of luck and then to see what influence that notion has in those areas. My goal in this chapter is to capture a number of the prominent explications of luck, and to diagnose precisely what goes wrong with them. Then I present a new view that avoids the traps into which the other accounts fall and which provides an intuitively plausible grasp of the nature of luck. A proper understanding of luck should, in subsequent chapters, help us to come to proper terms with the literature on moral luck and will help us to see the true moral significance of luck in evaluating ethical behavior.

Many of the competing accounts, however, while not fully illuminating, are *illuminating enough* to be valuable in certain discussions. Let me give an example. One standard line in discussions of free will is that if luck influences an agent's action, the agent doesn't have free choice in that instance. In discussing that claim, one need only a *correct enough* account of luck to identify clear examples in which luck influences the agent's behavior. In order to understand, evaluate, and participate in those discussions perhaps one needs but a basic intuitive sense of the nature of luck.

A detailed and philosophically defensible account of luck may not take some discussions significantly beyond what can be achieved by using an in-the-ballpark account of luck, though it may make *some* aspects of the discussions at least a bit clearer. While a small advance is less exciting than a large one, small steps are still good steps, and even the possibility of a small step is worth striving for. Be that as it may, however, coming to grips with the nature of luck is in itself intrinsically valuable, and the worth of this chapter ought to be judged at least in part on its own terms. When it comes to moral luck, however, the correct account of luck represents more than a moderate advance. As I argue, there is no luck *simpliciter* but only luck relative to our concerns. If responsibility and luck are related--our responsibility can only begin where our luck leaves off--, then responsibility is also relativized. This would be no small departure from the standard view.

When Williams helped to launch the contemporary discussion of moral luck, he took the term as primitive: "I shall use the notion of 'luck' generously, undefinedly, but, I think, comprehensibly" (Williams 37).⁴ Many others have followed suit, but I think that it is better not to take the notion as primitive unless we cannot make progress in illuminating the concept--that is, unless it is truly a primitive notion. Many authors, however, don't take the notion as primitive, and their work contains a number of interesting attempts to come to grips with the nature of luck. In thinking about (and beyond) those accounts we can make enough progress for our efforts to be worthwhile.

Luck has two parts. One part is constant between almost all accounts: for an event to be lucky for an agent it must have positive or negative effect on his or her life.⁵

⁴ Williams was writing an article, not a book, and luck is *primitive enough* to be treated as primitive in his discussion.

⁵ The one exception is Michael Zimmerman's account which includes the possibility of neutral luck (Zimmerman, "Map" 585). I take the possibility of neutral luck

Four notable conceptions of luck in recent literature examine the second, non-affective component as either (1) a lack of control in given situations, (2) a property best explicated in terms of modality and possible worlds, (3) an element of unforeseeable chance or accident in life, or (4) a view that's a hybrid of (1) and (3).

I offer a new conception of the nature of luck which is explicated in terms of what David Lewis calls "insensitive casual chains" and which also owes a debt to Walter Sinnott-Armstrong's doctrine of contrastive comparisons. My complete account contrasts an actual event and an alternative, averted event along two axes. First, it contrasts the event's positive or negative effect on the agent to the value associated with the averted event--that is, it is *value contrastive*. Second, it contrasts the likelihood of the given event with the likelihood of the alternative event. Thus, it is also *causally contrastive*. It allows us to make comparative judgments such as the following: my winning a million dollars in a public lottery is more lucky than winning a hundred dollars but less lucky than winning ten million dollars, and my making a basketball shot from the free throw line is more lucky than making a lay up but less lucky than making a full-court shot.

2. Preliminary Points

Before I discuss the various specific accounts of luck, there are some more general points to be made which I have gathered in this section. I think that this organization is less clunky, on the whole, than interrupting the flow of discussion several times for side bar and footnote conversations, and that ought to excuse the clunkiness it exhibits.

to be a simplifying assumption on his part, and nothing of consequence hangs on his making it. If we're trying to capture the notion of luck in its common or philosophical usage, though, that simplifying assumption is a mistake.

2.1. Dictionary Definitions

One may react to philosophers' attempts at getting clear on luck by claiming that it is an imprecise folk notion, a piece of ordinary language, and that it is a mistake to put any philosophical weight on it. Such an opponent to my project may think that "luck" has a historical meaning and that philosophical analysis does more to manipulate language than to elucidate it. The thought may be that relying on a dictionary definition leads to bad philosophy because of its imprecision. To make the meaning more precise, however, is to co-opt the word, and our understanding of "luck" will bear only an incidental resemblance to its historical meaning. One further point against this project is that there is no such thing as *the* meaning of the term. Thus, to try to get a fix on the meaning of the term is at best to distort it.

One way to address these points is to see if the word has at least one central use that is precise enough and resilient enough to be useful in philosophical explanations. If so, then we can leave it well enough alone, and just use the ordinary concept in substantive philosophy. If the term is not robust enough for philosophical work, we have several options. First, we can work to *uncover its meaning* by more clearly understanding how it functions in ordinary language and in philosophical contexts. This is a process of discovery not invention, and in this case, we will in effect leave the word as it stands and amend or correct the word's "dictionary meaning." After all, if the philosopher's understanding of a word goes beyond the lexicographer's understanding, ought not the lexicographer's be expanded rather than the philosopher's being truncated? In making this point, I take the primary lexicographic task to be to describe the actual use of terms in both ordinary and technical contexts, and that function is therefore not prescriptive or only secondarily so. Second, philosophy can *capture the word as a term of art* and give it a more precise meaning. After the term comes to have a certain amount of life as a term of art, a sort of philosophical "street credit," then this situation reverts to the first option,

and the lexicographer should endeavor to incorporate the philosophical understanding of the term. Third, if its meaning is too much of a tangle or cannot be usefully applied to or extended for philosophical work, *we can abandon its use altogether in philosophical explanation* and rightly think that philosophy is better off without it. The third option ought to be a haven of last resort, and while an appropriate response to certain situations, not one to which we will need access in this work.

According to the *Oxford English Dictionary*, luck is “[f]ortune, good or ill Good fortune; success, prosperity or advantage coming by chance rather than as the consequence of merit or effort.”

As is evident from that passage, luck has two components, the element of chance and the success or prosperity or advantage (or failure or impoverishment or disadvantage) that results from the influence of chance in our lives. That accords very well with the central philosophical accounts of luck. There are, however, at least three issues with regard to the definition and its bearing on philosophical contexts. First, it does not tell us whether the chance involved is subjective or objective, that is, whether luck is primarily an epistemological category or whether it is primarily a metaphysical category. Second, packed into the *OED*’s entry is the idea that the fortune that comes from luck is in contrast to the fortune that results from merit or effort. The definition, therefore, seems to rule out the possibility that an agent can, in relation to the same event, be both lucky and meritorious; thus, luck and merit are mutually exclusive. It’s reasonable to ask if that claim is true. This *question* appears in discussions of moral luck: is the behavior of the drunk driver who harms another worse than the behavior of someone who similarly drives drunk but “gets away with it”? Both are blameworthy for endangering others, but the question is whether the unlucky driver is blameworthy for something more. If luck can augment or diminish the moral status of the agent, her status is due to both merit *and* chance, *pace* the *OED*. That’s a substantive moral issue about which there’s no wide-

ranging agreement, and philosophical argument ought to be prominent in deciding whether luck and merit are mutually exclusive.⁶ Third, there are a number of philosophical methods for evaluating chance, and, as is evidenced by the diverse accounts outlined below, different methods for reckoning how chance contributes to lucky events and in evaluating the degree of luck involved. Since the *OED*'s entry doesn't help us decide between various methods for telling lucky from nonlucky events or the more lucky from the less lucky, it needs to be expanded.

2.2. Experimental Philosophy

There's a related issue to the question of whether we ought to rely on a dictionary definition of "luck." There's a growing trend in philosophy, experimental philosophy, that notes that how "luck" ought to be used depends on how it is in fact used. A few moments ago, philosophical investigation was presented as a potentially useful supplement to dictionary glosses. Many contemporary philosophers would take that as an invitation to engage in some armchair activity. In contrast, according to experimental philosophy, a term's meaning is an empirical question for which experimental, rather than armchair, methods ought to be used.

Experimental philosophy would advise that the question of whether a given event is an instance of bad or good luck for an agent (or not luck at all) should not be answered by one person. Rather, the proper way to *discover* which events fit into which buckets is to appropriately poll people, and run statistical analyses of the results. From the statistical analyses we can abstract out what people mean by "luck." Since it is an

⁶ This chapter won't discuss the proposed answers to questions concerning moral luck or any other substantive philosophical question concerning luck but rather seeks to pave the way for such discussions by coming to a more precise understanding of luck.

empirical question what a non-stipulative term like “luck” means, we ought to approach the question experimentally.

The experimental philosophy movement is not inherently *anti-intuitionistic*.⁷ The least revisionistic of its practitioners think that the method of intuition as usually applied in conducting thought experiments is especially open to personal bias and that thought experiments fall short of the rigor of, and thus are not, real experiments. One central type of claim in experimental philosophy is that philosophers may be perfectly sincere when they claim that a given term is “obviously used this way by all competent speakers” and yet be wrong about that. Suppose, for example, that a moral philosopher considers two scenarios which differ in what she believes to be a small and relatively unimportant way. At least sometimes experiments seem to show that a surprising number of people when presented with the two cases see a much greater moral significance in the difference of small details than the philosopher would have predicted (Knobe 190). Thus, some conclude that we ought to distrust or at least question the propositions delivered by intuition.

Some philosophers fight this approach. An article for *The Chronicle of Higher Ed* quotes Timothy Williamson as saying, “If anything can be pursued in an armchair, philosophy can.” Judith Jarvis Thomson adds, “A philosophical problem is not an empirical problem” (Shea B9). The thought here is that it’s a good thing that philosophy is an armchair activity because that is its strength. On the other hand, Ernest Sosa is somewhat sympathetic to certain aspects of experimental philosophy but insists that the role of intuition in philosophy is too central and the results are too well grounded for it to be discarded as a valuable approach (Sosa 101). My sympathies, at least for this

⁷ Some of its practitioners are, however, but that seems to be the minority position. See (Bishop and Trout). In reply, see Sosa (Sosa 105–07).

investigation of the meaning of “luck,” lie in the more traditional camp.⁸

In the sections that follow I present a number of intuitions about specific examples of luck. Then I evaluate each particular account of luck according to whether it makes the correct call in categorizing each of the particular examples as lucky, unlucky, or nonlucky. An account, then, is rejected if it doesn’t classify cases in the way that I pretheoretically think they ought to be classified. (Or at least it’s rejected if it misses the mark on any *central cases*--on what *I* decide are central cases. If the best account available misses on a small number of peripheral cases, I might seek to save the account by revising my evaluation of those cases, by bracketing those intuitions, etc.) Virtue for an account is thus measured by how coherently and deeply rooted its explanation is for a wide variety of cases. It’s not that I cannot be wrong in classifying any of the cases; it’s that I’m more sure of my classification than I am (at least initially) of any particular explanation of the phenomena, so the test cases are taken as fundamental, but not inviolable, data.

Suppose that I were to augment the methodology outlined above with empirical data. Two outcomes are possible. First, the experimental data might confirm or very nearly confirm my pretheoretic classifications as reasonable by the lights of the subject population. In that case, the experiment would have no practical effect on my classification. Second, the experimental data could be at significant odds with my armchair introspections. Had I not given the matter due care, I might take this as an opportunity to revisit my thinking about some of the test cases. But by hypothesis, I’ve

⁸ Some fascinating results are coming from experimental philosophy on the nature of luck and what influences people’s judgments concerning luck, particularly how people’s judgments concerning the moral value of an action affect their judgments of the action’s intentionality (see, for example (Sousa and Holbrook)). What is true is that traditional armchair methods can be used to find a conception of luck, for example, that’s worth testing by experimental methods.

given the issue due diligence. Thus, as one reasonable response to that scenario, I might give myself the following speech.

Well, the experimental subjects obviously think of luck differently from how I do, but my way is philosophically correct on grounds that my classification of the test cases coheres with a deeply rooted explanatory account. Thus, what they are calling “luck” is different from what I call by that word. I’m not going to insist that my usage is the correct usage, so, officially, we may take my investigation to be of luck* and its relation to metaphysics, epistemology, and ethics. My main investigation is thus into the nature of moral luck* which is somewhat aligned with what the ordinary Jane on the street might call “moral luck.” She and I will appear to agree on some important cases and appear to disagree on others. This is to be expected since the difference between her notion and mine are real (so we will appear to disagree on occasion), but also relatively small (because we will appear to agree on most occasions). We won’t be disagreeing, however, or agreeing, since we will be in important ways talking at cross purposes.⁹

My method in what follows thus is to give each example careful evaluation and anytime I suspect that my classification might be idiosyncratic, label it as such and be explicit about what considerations speak in its favor, and what considerations speak against it. Also, I believe that I’m investigating luck, but if you think that my target is really luck* that’s fine with me.

⁹ A third possibility is that luck and luck* are the same concept but that interference from other concepts is causing experimental subjects’ answers to vary from our predictions. (See Knobe and Cushman.)

2.3. The Possibility of Varying Spheres of Luck

One further potential block to finding a *general* account of luck is that it's at least *epistemically possible* that there are different kinds of luck in each of the different spheres of inquiry. For example, it may be that what differentiates moral from epistemic luck is *more than* the fact that the former represents a gain or loss of some moral desideratum--moral praiseworthiness, perhaps--and the latter a loss of justification or some other epistemic good. It's possible that the species of luck in moral luck differs from the species of luck in epistemic luck, and they both differ from the species of luck in endeavors requiring the use of practical rationality (practical luck). It's possible thus that there's no common genus under which all species fall. There are two ways this might happen. First, there may be a two-part account of moral luck, a three-part account of epistemic luck, and so on. Second, if all accounts have two parts, for example, it may be that neither part of moral luck shares any more than a superficial resemblance to either part of epistemic or practical luck.

The idea that there are varying spheres of luck, however, gets very little traction with me. The primary consideration against there being *sui generis* spheres of luck is that it's clearly not the simplest explanation for the range of intuitive phenomena that we will consider. If that consideration is found wanting, I don't know how to further *argue* that there's one genus of luck and that what differentiates the species of luck are the kinds of positive or negative effects for the agent. I can report, however, that in examining a wide variety of purported instances of luck, nothing of that sort has ever seemed to me like a reasonable explanation for their differences.

Of particular interest for philosophy is luck as it presents itself in philosophy's three central evaluative contexts: moral luck, epistemic luck, and practical luck. A gambler's good luck, for example, is measured within a context of practical reasoning, betting against the odds, and coming out on top despite the odds not being in her favor.

On the other hand, the luck of finding a hundred dollar bill blowing down the street has nothing to do with how an agent performs in any evaluative context, so the reach of the genus of luck is wider than the disjunction of the three central philosophically interesting species.¹⁰ Though it is good luck to find a hundred dollars and ought to be seen as such by a proper account of luck, no particular philosophical insights or issues are raised by it. It's good to keep such examples in mind, however, as a check against the scope of our investigation inadvertently narrowing.

2.4. Metaphysical and Epistemic Considerations

One notable feature of lucky events is that they have an element of surprise to them, and so they must be in some sense rationally unpredictable.¹¹ Nicholas Rescher distinguishes two ways in which we may not be able to rationally predict future events: (1) an event is rationally *unpredictable* when *the event's occurrence* is blocked by the objective matters of the case and thus we are unable to predict that it will occur, and (2) an event is rationally *impredictable* if *we as agents* are blocked from being able to rationally predict its occurrence (Rescher 216n).

Unpredictable events are unpredictable for everyone, including God. Events which are impredictable but not unpredictable (which we may call *simply impredictable*), however, may be closed off for only some agents, though not for omniscient beings or

¹⁰ I know of no species term for luck of this type, but that is no matter.

¹¹ One may, of course, *irrationally* predict all manner of events, but these don't concern us. What makes a prediction rational is that it is or can be justified in the right ways by appeal to conceptual truths, the available evidence, laws of nature, etc. Suppose that Dave predicts that he will find a diamond larger than which has not been perceived, and suppose that the prediction is true. The truth of the claim is irrelevant to the rationality of the prediction, since he ought to be judged by the justification of his belief that he will make such a discovery. If in this case his belief is not justified, the prediction is fanciful and not rationally supported.

perhaps even some finite agents. Simple unpredictability is relative to the agent: an event may be simply unpredictable for Bob but predictable for Zelda. Furthermore, it's a necessary truth that unpredictable events are unpredictable. If Tom has a ticket in a fair lottery in which there is an element of objective chance (causal determinism thus being false), then whether he wins is unpredictable. Because it's unpredictable, it's also unpredictable. Whether he will win is not something that even God can rationally predict.

The converse, however, does not hold. Suppose, for example, that a lottery has been rigged in Ann's favor, but she is kept in the dark. The result is for her unpredictable: she cannot rationally predict that she will win because she does not have all the relevant evidence. This is true even though before the drawing there is good evidence--known at least to the people who have put in the fix--that she will win. It's not unpredictable because if she had all the relevant evidence, if she knew what her benefactors know, she could rationally predict that she will win. Furthermore, an event may be unpredictable to all human agents even if it's not unpredictable. The upshot is that unpredictable events (and hence some unpredictable events) are shut off from human knowledge in a stronger sense than are events which are simply unpredictable.

Rescher believes that simple unpredictability is a strong enough requirement for an event being lucky for an agent, and he counts Ann lucky if she wins. He thinks that what matters for luck is that the agent cannot predict the event. This can be because (1) she doesn't have enough information about the past or the current state of affairs, (2) she doesn't know what generalities reliably link past to future events, or (3) she knows the facts and the true conditionals that link past to future events, but she doesn't have the cerebral power to make the connections. To elaborate on clause (1), the agent may not have enough information because the data is too complex to understand or because there's no record of it or because the data is distributed behind people's inscrutable

intentions (think of the people rigging the lottery). Thus, it's a question of what data is available to the agent with her natural limitations, and it's not a matter of what data would be available to an agent who has a better epistemic situation. When any of conditions (1)-(3) is true, it's as if the world was causally open-ended, and that, Rescher thinks, is enough for luck to make its way into our lives (Rescher 41–42).

Other authors have the stronger requirement of unpredictability. Some, for example, think that luck is a matter of the agent lacking control over whether the beneficial or harmful event comes about, where that lack of control is not merely not knowing that the event is likely to happen, and not merely not knowing how to bring it about. They mean that the event is unpredictable in a metaphysical sense.

One further difference between the view of luck as epistemic and the view that it's metaphysical is what is required for the agent to reduce her dependence on luck. In the metaphysical sense, it would be to gain more skill or more power and hence more control of the outcome. In the merely epistemic case, it would be to gain more knowledge.

To put matters another way, is luck compatible with causal determinism? One such as Rescher may answer in the affirmative. After all, on his view it's sufficient that we can't predict the outcome and hence we are surprised at Ann's winning the lottery. On the metaphysical account, however, unpredictability is required, and there is no luck in a causally deterministic world.

Rather than trying to settle *this* dispute, let me say that moral responsibility requires a kind of free will that is incompatible with causal determinism (as I argue in the next chapter). Hence *moral luck*, if there is any, requires that the world contain unpredictability and not merely impredictability, that is, that it contain metaphysical and not merely epistemic luck.

3. The Lack-of-Control Account of Luck

The lack-of-control account is usually presented in the literature on moral luck as a rough and ready stand-in for a full analysis. Some philosophers who endorse this account give a nod toward what may be behind our notion of luck, but don't take the discussion very close to completion. For example, note how Nagel presents his understanding of luck in his seminal paper, "Moral Luck": "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" (Nagel 59). Luck, then, includes a lack of control, and bad moral luck is a lack of control that does not excuse the agent from receiving attributions of negative moral judgment. Good moral luck, on the other hand, does not bar the agent from receiving whatever positive attributions he may normally have incurred in situations which have no admixture of luck. What Nagel's discussion lacks, however, is any deep revelation as to the nature of luck.¹² We can take *lack of control* as having an intuitive sense, but we ought to come to a more detailed account of the phenomenon, if we can (and we can).

While Nagel's target is *moral* luck, we can take what he says more generally. That is, we can differentiate between the lack of control, which is presumably neutral between various practical or philosophical areas in which luck has an influence, and the kind of boost or set back the agent experiences because of her lack of control. Thus, if it is moral approbation or disapprobation that accrues to the agent, the particular instance is of *moral luck*. Other areas of endeavor will have their own correlates to moral favor and disfavor: epistemic luck may show itself as good luck when the agent's belief is unexpectedly true, and a gambler's luck may be felt as financial gains and losses in the betting parlor. It is easy to see the differences between luck in war and luck in love,

¹² Again, he was writing an article, not a book, and some abbreviation is understandable and unavoidable.

between good luck in buying a good car for a reasonable price and bad luck in coming down sick the first day of spring break. These differences show themselves in the various kinds of goods and ills they bestow upon the agent.

Nagel is not alone in endorsing luck as a lack of control. Indeed, it is perhaps the most prominent account. Here is Daniel Statman:

Let us start by explaining what we usually mean by the term ‘luck.’ Good luck occurs when something good happens to an agent *P*, its occurrence being beyond *P*’s control. Similarly, bad luck occurs when something bad happens to an agent *P*, its occurrence being beyond his control (Statman, D. 146).

For Statman, the divide between lucky and nonlucky events is drawn by the agent’s lack of control in determining whether the events in question come about, and the distinction between good and bad luck is drawn, much as it is in Nagel, based on the effect those events have on the agent. Neither Nagel nor Statman indicate, however, what we are to make of those occurrences which are beyond the agent’s control but which are either neutral in their effect on the agent or which do not affect the agent at all. That is, they both leave open the possibility of ‘neutral luck’.

A third version of the lack-of-control account of luck comes from Michael Zimmerman: “[S]omething which occurs as a matter of luck with respect to someone *P* is something which occurs beyond *P*’s control” (Zimmerman, “Luck and MR” 231). Zimmerman’s words are consistent with their being luck *simpliciter* which divides between good, bad, and neutral luck, and in another work he explicitly claims as much (Zimmerman, “Map” 585).¹³ No philosophical issue hangs on whether we count three varieties of luck or only two.

¹³ Zimmerman does admit that neutral luck “is not very interesting” (Zimmerman, “Map” 585), but the very phrase ‘neutral luck’ sounds harsh to my ear.

The sensible (and I think charitable) way to fill out their individual accounts, though, is to take (1) the lack of control and (2) the positive or negative difference in the agent's life as individually necessary and jointly sufficient conditions for an event being lucky. Thus, there is on this nomenclature only good and bad luck, and we have neutralized neutral luck. I will pretend that Zimmerman's words pertain to only a proper subset of the necessary conditions for luck (and not to its whole nature), the other necessary condition being that lucky and unlucky events have a positive or negative influence on the agent's life, and I will pretend that Nagel and Statman would agree to amend their accounts along the same lines. I will take it, then, that Nagel, Statman, and Zimmerman are all of a piece in their understanding of the nature of luck, and I will take them as endorsing a unified view which I will call the *lack-of-control account of luck*.¹⁴

Let us note one point about the notion of control used here. Zimmerman distinguishes between *restricted control* and *unrestricted control*. Restricted control is control of at least some factor upon which the event is contingent, such that it's within the agent's power to bring about the event, and the agent can also prevent its occurrence. Unrestricted control is the most complete control possible, that is, restricted control of every event that contributes to the target event coming about (Zimmerman, "Luck and MR" 219). It's not plausible to think that any finite being has unrestricted control over any event: no one has any control over whether she was born, for example, and that is a factor pertaining to any candidate event over which the agent might have unrestricted control. We can say with Zimmerman, then, that an event is *strongly out the agent's control* if and only if the event is out of her restricted control, and it is *weakly out of the agent's control* if and only if the event is out of her unrestricted control (Zimmerman,

¹⁴ If they don't agree, that leaves the relatively unimportant question concerning the status of purportedly neutral luck, and nothing of great importance is at stake.

“Luck and MR” 219–20).¹⁵ Thus, for temporally and causally finite beings like ourselves, all events are out of our unrestricted control, that is, out of our control in the weaker sense, but some events are out of our restricted control, that is, out of our control in the stronger sense.

3.1. Paradigmatic Cases

One way to get a clearer picture of what these writers mean in accounting for luck in terms of a lack of control is to consider some detailed cases of luck. One paradigmatic case of luck on these accounts is winning a fair lottery where the odds of winning are such as to render that occurrence unlikely and where the payoff is non-negligible. One salient fact about lotteries is that they are not contests of skill, and the primary influence of whether the agent wins is the objective chance of winning where this is not within the agent’s restricted control. Whether she wins need not be absolutely out of her control, though, because in many state-run lotteries, for example, the agent usually has to purchase a lottery ticket, and in many lotteries she will have the option of choosing which numbers to play and typically has the option to purchase more than one ticket.¹⁶ Having multiple tickets does not increase the odds of winning for any one ticket, but does

¹⁵ In “Moral Luck: A Partial Map,” Zimmerman uses the terms *partial control* and *complete control* to mark the same territory as *restricted control* and *unrestricted control*, respectively (Zimmerman, “Map” 591).

¹⁶ Reflection on these facts further underlines the notion that luck requires that the event be out of the agent’s restricted control: that the agent controls some aspects relevant to winning does not mean that she controls whether she wins. Furthermore, while she can control whether she doesn’t lose (e.g., by not buying a ticket), she does not thereby have restricted control of winning. Remember, restricted control of an event requires that the agent controls some salient feature which determines whether the event occurs or does not occur. The lottery case highlights the fact that for the agent to control the situation, she may need to have restricted control of more than one salient feature-- buying the ticket and influencing which numbers are drawn--and that in a fair lottery no agent has control of both those features of the situation.

increase the overall odds of winning for any player with multiple tickets. So while the odds of winning may be *somewhat* in the agent's control, it's only minimally so, and not enough to correctly say that she has restricted control of whether she wins, since picking her own numbers does not increase her odds when compared to randomly generated numbers, and purchasing a relatively small number of tickets does not greatly increase her chance of winning.¹⁷

Having your lottery ticket drawn is not always a good thing: military drafts in times of war and scenarios such as described in Shirley Jackson's "The Lottery" are instances where "winning" is a case of bad luck. Consider too the case of the unlucky soldier who, having drawn the short straw, has to play David to the enemy's Goliath. Call such scenarios "negative lotteries" and the more typical scenarios "positive lotteries". For simplicity of locution in what follows, I shall use only examples of positive lotteries and will use the general term "lottery" and its cognates in that connection. Any of the examples could easily be supplanted by examples of negative lotteries. Thus, in the most usual case, winning the lottery is a good thing for the agent, but minor shifts would have the agent experience the bad luck of "winning" a negative lottery.

Even though winning the lottery is a prime example of good luck, perhaps we can get more clear about what features of the case make it lucky. Perhaps these can be brought into relief if we consider what conditions *defeat* the claim that an occurrence is a matter of luck. First, consider a lottery has been rigged, either by Ann, or by one of her benefactors--whether or not the rigging is known to Ann. When *we* are made aware that

¹⁷ I have a colleague who does not buy lottery tickets but says he keeps a watchful eye for them blowing down the street. His point is that *finding* a winning lottery ticket has about the same odds as *buying* one. True enough. But only in the latter case is one playing the lottery, and in this example the point is that in playing the lottery, the resources the agent can use to control the outcome are severely limited.

the lottery has been rigged we stop seeing the fact that Ann won as a matter of luck (*pace* Rescher). In the case of her having an unknown ally, *she* may still regard herself as a lucky winner, and in that she'll be partly right. She is lucky to have benefactors who would (and could) rig a lottery in her favor, but it doesn't seem to be a matter of luck for her to win. So the truth to her belief is that she is lucky, though she's mistaken about what she is lucky for: she regards herself as lucky for winning the lottery, but she's really lucky to have the benefactors she has. She's lucky that the fix is in, but given that the fix is in, how can the outcome be a matter of luck?

From an intuitive point of view, therefore, Ann is not lucky. According to the conception under examination, however, she lacks restricted control over winning the lottery, and the event of her winning is a good for her; thus it *is* a matter of luck for her. What's driving the intuition that she's not lucky is that if her allies rig the lottery in her favor, then it's not in the least bit chancy that she wins (assuming that they're competent cheats). The reader need not agree on this case, but what ought to be without dispute is that if *Ann herself* rigs the lottery in her own favor, then it's clearly not a matter of luck that she wins. The difference, then, between a fair lottery and a rigged one tracks whether the agent herself (or perhaps her benefactor) exerts significant control of the outcome.

In relation to this defeating condition, we also ought to consider non-rigged lotteries. These lie on a continuum from cases in which numbers are drawn absolutely randomly to ones in which there is no element of randomness. At one end of the continuum, of course, lies causal determinism. If determinism were true, there would be no non-epistemic randomness at all: armed with knowledge of how the world was, say, ten million years ago and knowledge of the actual laws of nature, one could know whether Ann wins the lottery on this drawing. No human is in such an epistemic position, however, and perhaps no human is even in the position to know whether the past and the laws of nature causally necessitate the future. If the outcome is causally

determined such that Ann wins the lottery, and the causal flow doesn't conspicuously travel through her will, the lack-of-control account would hold that she is lucky since it is out of her restricted control whether she wins the lottery.¹⁸

If causal determinism is false, one *may not* be able to deduce from the past as it was ten million years ago and the laws of nature that our agent would win. What if, however, causal determinism is not strictly true but very nearly true in terms of a particular state of affairs obtaining? For example, suppose that given the way the world was five minutes ago and the laws of nature as they are it *is* causally necessary that our agent wins the lottery. It may be *as if* causal determinism were true in such a case if either (a) there are no causally indeterminate events in the five minutes preceding the drawing or (b) none of the causally indeterminate events have any effect on the drawing perhaps because either they are on a sub-atomic scale and their effects do not "echo up" to effect the outcome of the lottery or they are not local to the drawing (perhaps even being out of its light cone). Such a case reverts to the same answer given if causal determinism is true: the agent is lucky because, considered from the point of view of the last five minutes,¹⁹ it is not in her restricted control whether the event happens.

On the other hand, suppose that the event is truly random, as the decay of a radioactive atom is thought to be. In that case, the event is clearly not under the agent's control, and if we build our lottery drawings around such an event, then if the agent wins it is out of her control, and hence a piece of good luck. The upshot is that she is lucky only if the fortunate event is not up to her in some meaningful sense--at the very least, it

¹⁸ I find this result counter-intuitive. How can a causally determined event be lucky? I will develop that critique in a few pages.

¹⁹ This clause foreshadows one aspect of my considered view, to be explained later.

is not determined by the past and the laws of nature are not random either.²⁰

A second defeater is seen when we shift examples and consider feats which are usually associated with some level of skill, such as making a difficult basketball shot.²¹ There are two ways in which we can develop an account of luck in reference to such cases. On one understanding of luck, making a difficult shot is classed as lucky if the agent lacks the skill to perform the feat within an acceptable range of regularity, though the range's extent is often left almost entirely vague. The luck here results from the agent's inability to channel the outcome to the desired end, again within some unspecified range of acceptable regularity. Magic Johnson need not hit every jump shot to have the ones he does hit be matters of skill and not luck; his shooting percentage, compared against the best players in the world, reveal that he is in some strong sense in control of whether the basket is made. If an agent can reliably hit only one percent of the same shots under the same conditions that Magic Johnson hits with 80% regularity, each of the agent's few baskets is rightly regarded as being made by luck and not by skill. Indeed, his skill may be so minimal as to make it a matter of luck whether he even hits the rim, let alone makes the shot. Thus, even in cases in which both the skilled and the unskilled player miss the shot, the difference in skill will often show itself in *how* each misses the shot. For example, if we give the skilled player and the unskilled player ten shots each of such difficulty that they both miss all of their shots, we will be able to tell which player is

²⁰ Readers should note that the reasoning here parallels the sort of reasoning that makes free will look mysterious--free will seen as incompatible with both determinism and indeterminism. See van Inwagen on the mystery of free will (van Inwagen, "Free Will Remains a Mystery").

²¹ Sousa and Holbrook, in their paper on intentional action and immoral and amoral luck, identify five features as pertinent to intentional action: desire, belief, intention, skill, and awareness (Sousa, et al. 351–52). It's plausible to think that the factor relevant to a lack of control is whether the agent has skill.

the more skillful by how closely he misses as compared to the unskilled player. In exhibiting or judging skill, a miss is really *not* as good as a mile.

Skill in this context is to be seen as having some measure of control of the outcome, where that level of control need not be complete (it's restricted control that we're interested in) and need not be precisely defined. If luck in such matters is a lack of control (plus the positive benefit of making the shot) and the cure for luck is to have skill appropriate to the task at hand, then there can be a large swath of vagueness between that which comes about through luck and that which is the product of skill. Magic Johnson presumably didn't start as a skilled shooter, but it seems overly precise to say that at a precisely given level of ability, luck was no longer a factor. For our purposes, however, it doesn't matter how large the separation between luck and skill. It's enough that there be paradigmatic cases at each end of the scale and that what explains the ordering of the scale is the presence or absence of the agent's skill and hence control of the outcome.²² On this account there are three replies to the question of whether the player made the shot by skill or by luck: (1) by skill, (2) by luck, or (3), this question lands between skill and luck, and so is properly neither, given our current usage.

One question that can be asked of that account is the following: If Magic Johnson makes a shot through skill that is admittedly *not* absolute, what other factors besides his skill contribute to him the making shot? One answer is that luck fills in where Magic Johnson's skill leaves off. Thus, the second view of the interplay between luck and skill is that for any agent and any event, whether that event obtains is to some extent

²² Luck, so understood, suffers from vagueness understood *de dicto*. I follow Lewis on vagueness: "The only intelligible account of vagueness locates it in our thought and language. The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word 'outback.' Vagueness is semantic indecision" (Lewis, On The Plurality of Worlds 213).

(including zero) up to the agent's influence, and luck is the contribution made by the world where the agent's contribution leaves off.²³

On this understanding we need a more fine-grained understanding of control. We can keep Zimmerman's division between restricted and unrestricted control. We then ought to say that even for events that are within the agent's control, there are gradations of control.²⁴ Magic Johnson has a high level of restricted control, but that level of control isn't maximal, since luck still contributes somewhat to his making the shot. Control and luck are therefore both scalar notions. This seems correct. If in attempting a particular shot, a player relies on both skill and luck, it would be exceedingly odd for one of those notions to be scalar and the other binary. Furthermore, in saying that Magic Johnson made a particular basket through skill, we aren't denying that luck played somewhat of a role; we're just labeling the shot according to the factor that contributed most prominently.

Armed with the contrast between skill and luck, we can develop a nuanced appreciation for another paradigmatic kind of luck: hunting for and finding buried or lost treasure. Even if the treasure hunter has maps and other clues and thus knows the general area in which to look, even if he is armed with the best equipment for the task, finding the treasure is more a matter of luck than skill. To be sure, there is some skill involved in treasure hunting. Deep sea salvage experts spend a great deal of time off the water learning what they can about where a bullion-laden vessel was lost at sea. But it's a big sea, and even a careful investigation may only narrow their search area down to scores if not hundreds of square miles. For land searches, it is not often that we can get detailed,

²³ If the event has a neutral or no effect on the agent, we can't describe it as luck, since luck is only positive or negative. Where skill leaves off, the agent has a lack of control, and the event is lucky if it positively or negatively impacts the agent.

²⁴ These remarks apply to restricted control for finite agents.

credible evidence of the treasure's location. "X" rarely marks the spot. Finding the treasure, then, is largely out of the agent's control and only minimally a matter of skill, and this view explains why we take it to be luck. If the agent finds the treasure because of useful and very specific information, the specific act of finding the treasure is not itself lucky, though finding and understanding the specific information may be. In this case, finding the treasure *inherits* the luck from finding the clues that give the specifics of the treasure's location.²⁵

We can also distinguish one who *hunts for and finds* treasure from another who *merely stumbles upon* it, the one being engaged in a more or less specific kind of activity and the other just making his way through the world, either engaged in no specific activity or in one that's unrelated to finding treasure. If the well-prepared treasure hunter is lucky, then the one who happens upon an unsuspected treasure is luckier still since by no one's lights is any skill involved.²⁶ Between treasure hunters, we can distinguish their differences in luck by how much their skill contributes to converting clues to found treasure, and we can differentiate their skill in finding treasure by how much luck they need beyond the clues given.

Indeed, we might distinguish *three* ways in which an event may be brought about: (a) through the agent achieving a goal by the use of skill, (b) through the intervention of

²⁵ Should we say that the person who finds a specific clue that leads definitively to the treasure has two instances of luck, or only one? We can say two, because finding the clue was lucky and then finding the treasure was lucky. Or we can say that there's one, finding the clue, and it's just that the payoff is temporally downstream from the act of finding the clue. Either works for me, and I don't see that a great deal hangs on it.

²⁶ In 2009, David Booth from Scotland bought a metal detector as a new hobby. *On his first outing and only seven steps from his vehicle*, he found golden torqs dating to the Iron Age that are worth \$1.65 million dollars. While his case may be in the border between hunting for and finding treasure and merely stumbling upon it, his luck is without dispute.

luck, and (c) through brute-force solutions. For example, the agent is neither lucky nor skillful in winning the lottery if she's bought one of every possible combination of lottery numbers. What *usually* separates those who need not rely on luck from those who do is the cultivation of skill, and brute-force solutions are not matters of skill, though they are matters of control. Indeed, what makes the brute-force method an initially attractive option in some cases is that there is no way to genuinely improve one's skill in the task at hand, for example, at picking randomly drawn numbers. In many lotteries, however, the brute-force method, while effective, is not a sound strategy, since the cost of covering all the possibilities is greater than the expected payoff. In lotteries with progressive pay outs, the money from one drawing with no winner is carried over to the next drawing and added to the prize for that drawing. Thus, after several consecutive drawings with no winner, the amount won on a single ticket can be greater than the cost of investing in one ticket for each possible combination of numbers. In those cases, what one is wagering against are not the odds of being *a winner*, but the odds of being *the sole winner*, since if two or more people match the winning numbers, the pot is usually split among the winners, and in that case the expected pay off is again less than the cost of covering all the possibilities.

3.2. The Lack-of-Control Account Evaluated

At the very least, any account of luck needs to give a plausible analysis of the lottery and buried treasure examples. Any serious contender for the correct account also needs to give some reasonable explanation of what marks these cases off as paradigmatic kinds of luck, as opposed to, for example, the case of a psychologically normal agent finding her keys exactly where she left them ten minutes ago. The lack-of-control account is plausible at first blush, because agents in our paradigmatic cases genuinely lack a level of control sufficient to significantly improve their odds for success, where this lack of control is especially contrasted with the use of skill in achieving a goal.

Even though the lack-of-control account appears to be initially plausible, there is good reason to think that any account in terms of lack of control has such serious defects as to warrant rejection. As Jennifer Lackey points out, there are clear examples which show that a lack of control is neither necessary nor sufficient for regarding such occurrences as instances of luck. The argument divides into two parts, denying the sufficiency claim and denying the claim that a lack of control is necessary for an event being lucky.

What is denied in the first part is that the lack of control in bringing about an event, taken with the event's positive or negative influence on the agent's life, is sufficient for the event being lucky. Jennifer Lackey gives this example: her husband, she tells us, has no control over whether she picks up their child from school as she has committed to do, and it is certainly significant to him, but it is clearly not a matter of good luck for him that she does as she has promised to do (Lackey 257–58). Lackey's husband is neither skilled in bringing that event about, nor has his 'solution' to the problem been one of brute force. Thus, on the lack-of-control account, he's lucky that she has picked up their daughter. Her very reliability, it seems from the armchair, rules out it being lucky for him that she picks up their daughter. On the lack-of-control account, however, he would be lucky. Thus, the result from the lack-of-control account is at odds with an intuitive sense (or at least my intuitive sense) that he's not lucky. Or if we take luck and skill to slide along the same scale, we ought to say that he demonstrated a minimal amount of skill in communicating to his wife that she ought to pick up their daughter. The skill here, though, is minimal, and the amount of luck involved would be great. This answer is still at odds with our intuitions.

(That is not to say that there are not cases of genuine luck in the vicinity.

Suppose, for example, that he's forgetful and tired, and when he wakes up, he's in a panic because he thinks he's forgotten to pick up their daughter. When he gets his bearings,

however, he realizes that it wasn't his turn, and he feels lucky. Considered from his point of view of a less-than-ideally reliable parent, this feeling makes sense, given that whether he picks up their daughter as scheduled is not always in his restricted control, even in cases where it ought to be.)

From our discussion of the lack-of-control account, however, our three options seemed not only mutually exclusive but exhaustive of the corner of logical space we were considering. That is, for every event which positively or negatively affects the agent, either it is lucky for her, or it is a matter of her exercising skill, or it is a matter of her exercising brute-force control. Lackey's case, however, introduced a new possibility: the beneficial event may be out of the agent's control and yet not be lucky if it is under the control of a benefactor whose actions are aligned with the agent's interests. Since the three original options are not truly exhaustive of logical space, the account cannot be correct. The missing category is the class of events which are out of the agent's restricted control but within someone else's control where that person will reliably bring about a token-event in the appropriate circumstances.

From the benefactor's point of view, what counts as "the appropriate circumstances" will be what will benefit or at least is likely to benefit the agent who is downstream of the intended action. Lackey's husband does not control her, but it seems accurate to say that her love for him and their daughter influences her behavior: she feels the influence of her affections for her family and thus exercises restricted control of her behavior. Her restricted control, however, is not his restricted control, though in some sense he does have some control--he's got the sort of relationship with her such that his interests become to some extent her interests.²⁷ There are, therefore, two senses of

²⁷ This example features their daughter, so Lackey's interests is in an important sense independent of her husband's since she'd presumably look after her daughter no matter how it affected others. If this is a hang up for any reader, the example could be

restricted control: agent directed causality--good old restricted control--, and causality directed through the will of a benefactor.²⁸ (We might call the restricted control a benefactor uses to influence our lives “restricted control once removed”, but perhaps no such special term is needed.)

If our agent performs an action that’s within his restricted control with no thought of benefiting another, but he benefits his friend inadvertently, then that is lucky for her. The important difference is that his desire to act is not influenced by the thought of the benefit she receives.

Proponents of the lack-of-control account may try to work around the issue of benefits coming from the exercise of restricted control by a benefactor. They can claim that the lack of control required for a genuine instance of good or bad luck should be understood to be complete in this sense: it’s neither in the agent’s restricted control nor in anyone’s restricted control as a conscious extension of the agent’s desires that the event comes about. That is, it’s not done for the agent’s benefit. Even in light of such a clarification, there are still clear instances of significant events that are under no one’s control (and hence under neither the agent’s nor her benefactor’s control) which seem counter-intuitive to call lucky or unlucky. Let us consider one such case.

The fact that the sun came up this morning was beneficial to Jim, under no finite being’s restricted control, and yet not lucky for him or for anyone else. Jim is fortunate

streamlined to something that only importantly involves Lackey and her husband, and for which the content of her interests is dependent on the content of his. A good example would be making reservations for an anniversary weekend.

²⁸ More accurately, a benefactor delivers goods upon the agent, and a nemesis delivers negative influence. I focus only on the case of the benefactor to achieve a stylistic fluency. These remarks apply both to agents who share our agendas and those who wish to block them.

that the sun came up this morning, since he has experienced many good things that he would not have experienced had the sun not come up.²⁹ The fact that Jim is fortunate, however, does not seem to require luck. The good fortune of the sun rising is seen in contrast to the bad fortune of the sun not rising, and all that is indicated by calling an event fortunate is a discernible difference in the relative value between it and an alternate event, not whether the event is within anyone's restricted control or whether the event was causally determined. In other words, fortune is just one component of good luck.³⁰

In the second part of the argument against the lack-of-control account, let me review Lackey's example which shows that a lack of control is not necessary for an event to be lucky. Lackey gives the example of a demolition worker who, when she presses a button in her office, blows up a building that her company has been setting charges in for a month. She is in control of the blast and the blast is significant for her and her company, and so would not seem to be a matter of luck. The crucial aspect of Lackey's example, though, is that while she's in control of the blast, it is only by luck that she's in control. This is because a rat has chewed through the wires connecting the button in her office to the charges, taking away her control of the blast, but just as she's pressing the button a coworker hangs his coat on a nail in the office which moves the tip of the nail penetrating the wall just enough to make the connection that the rat's chewing had inadvertently severed. Lackey's point is that the demolition worker is in control of the blast, but only through luck, and hence the blast itself is a lucky event (Lackey 258–60).

²⁹ The likely scenarios that would stop the sun from coming up would be something like a gigantic asteroid striking the earth and halting its rotation, or that for some surprising reason the sun would explode. In either case, were Jim even alive this morning, he probably would not be having a good day.

³⁰ The concepts of fortune and luck overlap, and this is one instance. What's more interesting are the cases where they don't overlap: the area in a Venn diagram of the lucky-but-not-merely-fortunate.

The upshot of her example is that a lack of control is not necessary for the event in question being a matter of luck since in this case the agent has control of whether the blast comes off and yet is lucky that it does. She has no control over being in control, and the re-established continuity is due to luck. The lucky occurrence of the blast gets or inherits its luck from the part of the situation that is more clearly lucky. Let me explain how that works.

Suppose two events, A and B, are such that the occurrence of B causally depends on the occurrence of A (B won't occur unless A occurs, but A may not be the only event upon which B depends). Furthermore, it's a matter of chance for the agent that A occurs, and once A occurs, it is a further matter of chance that B occurs. The chance involved in the A-B sequence is the chance involved in A occurring multiplied by the chance involved in B occurring. To see this, suppose that the agent has 50% control over whether A occurs, and, if A occurs, 50% control over whether B occurs. Given that, she has 25% control over whether A and B both occur; that is, there's a 75% chance that the A-B sequence won't occur. On the other hand, if she had 50% control over A occurring, and A necessitated B, then she'd have 50% control over B occurring.

Returning now to the demolition example, since the re-establishment of continuity is out of the demolition worker's restricted control, the continuity is lucky, and that which is causally downstream from the return of control inherits that luck, so the blast is lucky. The proponent of the lack-of-control account has a natural reply. In the normal case, the demolition worker has control of part of the situation: whether the explosives are wired correctly and whether she pushes the button. When the agent has control of those aspects of the demolition, she controls the blast. Hence, when it comes off it's not a matter of luck.

This case is different, however, because the rat chewed through the wires and made it necessary that the demolition worker be in control of more aspects of the blast

than in the normal case if the blast is not to be a lucky event. In particular, for the event to not be a case of luck, the demolition worker would have to be in control of wiring the building, pressing the button, *and* ensuring that the wires have the appropriate continuity. She does not have control of each of those factors, and so does not have control of the sum of those factors. In particular, she doesn't have control of the wires having continuity, so when continuity is restored, that's good luck for her, and hence, it's lucky when all the factors come together and the demolition comes off.

3.3. The Big Problem for the Lack-of-Control Account

This solution to the problem, however, brings to the fore a fatal problem for the lack-of-control account. According to the reply, the luck in Lackey's example is generated by the fact that she doesn't control a crucial part of the situation, namely whether the continuity of the wires becomes re-established. Now consider the case in which no rat has chewed through the wires, the demolition worker wires the blast correctly, and presses the button. Call this the *standard demolition case*. The blast happens, and it seems that there is no luck involved. But that's not correct, according to the lack-of-control account. In this case, the demolition worker did not have control over the fact that no rat chewed through the wires severing the continuity between the button and the explosives, or the fact that the button used as the trigger didn't have a factory defect, or the fact that the blasting caps weren't factory duds. The point is that in addition to those parts of the event over which the demolition worker has control, the situation has integral physical parts and physical processes that she relies on but over which she has no control. Thus, if we measure the luck of an event by whether it's significant to the agent and whether the conspicuous contributors to it are under the agent's control, nearly every significant event in the human sphere is a matter of luck. That seems to be the wrong result. That a significant event is lucky should be the exception rather than the rule

(standard cases need on any analysis to remain standard cases), and diligently applied, the lack-of-control account yields the wrong result. I'm not pushing the claim that the agent needs to have unrestricted control, because I think that Zimmerman is correct in thinking that this presents an unreasonable standard. There is, however, a reasonable understanding of what goes into the standard demolition case, and the agent is not in control of all significant parts of that.

The obvious reply is to accept that the standard demolition case is lucky for the demolition worker, just not lucky enough to take much note of. Indeed, the lack-of-control proponent can endorse that almost every significant event in life is lucky at least to some minimal extent. What, though, separates the standard demolition case from the non-standard case, if in both cases the demolition worker is lucky that the wires have continuity when she presses the button? In the standard demolition case, she's lucky that the continuity has not been interrupted, and in the other case, she's lucky that the continuity has been restored. On the lack-of-control account, she's lucky in both cases, which we can grant for the sake of argument. The rub, however, is that she's clearly much luckier in the case where the co-worker restores continuity by hanging his coat on a loose nail than in the standard demolition case. Both are out of the agent's restricted control, and so presumably both ought to be equally lucky.

The problem is that the lack-of-control account isn't properly scalar: more or less luck accrues to the agent depending on the degree to which he is benefited or harmed by the event, but given a fixed amount of benefit or harm, the agent is lucky if a significant aspect of the event is out of his control. Here is Statman again: "Good luck occurs when something good happens to an agent *P*, its occurrence being beyond *P*'s control. Similarly, bad luck occurs when something bad happens to an agent *P*, its occurrence being beyond his control" (Statman, D. 146). Whether an occurrence is beyond *P*'s control is a binary, not a scalar question. The lack-of-control account cannot distinguish

the little luck needed in the standard demolition case from Lackey's case in which the rat chewing created a need for a greater amount of luck. Thus, the lack-of-control account as presented is fatally flawed.

My account, developed below, is properly scalar and while different in important ways from the lack-of-control account, it is a close descendent of it. Before we can turn to it, we need to clear away a couple more rival accounts.

4. The Modal Account of Luck

One important attempt to nail down the nature of luck has been championed by one of the most prolific writers on moral and epistemic luck, Duncan Pritchard.

Pritchard examines the lack-of-control account of luck and concludes that while it hits the mark when evaluating lottery cases as instances of luck, it misses the mark in some cases. For example, he also cites the example of sun's rising discussed above. Pritchard thus contends that the lack-of-control account of luck is far too inclusive. His ultimate diagnosis is that it doesn't explain *at a deep level* what's lucky about lucky events, and that's why it performs poorly in categorizing such cases. The alleged metaphysical shallowness that Pritchard attributes to the lack-of-control account contrasts, of course, with the account of luck which Pritchard takes to have the proper metaphysical depth.

Pritchard's analysis falls into two parts, a modal clause and a clause linking the target event with a sense of significance for the agent. Here is the version he presents in "Epistemic Luck":

(L1) If an event is lucky, then it is an event that occurs in the actual world but which does not occur in most of the nearest possible worlds to the actual world (worlds which most resemble the actual world) (Pritchard 197).

(L2) If an event is lucky, then it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts) (Pritchard 199).

If we consider the case of the agent winning the fair lottery and the case of the sun rising, we can see how application of clauses L1 and L2 helps us to identify the lottery case but not the sunrise case as an instance of genuine luck. According to Pritchard, winning the lottery is an unlikely event which makes it a clear candidate of a case of good luck (plus the fact that it's significant for the lottery winner). He understands the unlikelihood of the event in terms of how the world under consideration compares to the nearest possible worlds. Pritchard also credits his account with explaining why the rising of the sun is not lucky: while the sun's rising is indeed out of the agent's control and its rising is a significant event in the life of the agent, it is not a matter of luck because the sun rises every morning not only at this world but also at most of the nearby, and hence relevantly similar, possible worlds. Similarly, suppose that Jennifer Lackey promises her husband that she will pick up their daughter from day care, and does (Lackey 257–58). This is another instance where the lack-of-control account gets the wrong result: since what Lackey does benefits her husband but is out of his control, that analysis holds that he's lucky that she does so. On Pritchard's account, her husband is not lucky because at most of the relevant possible worlds in which she (or her counterpart³¹) commits to picking up their daughter, she makes good on her promise. According to Pritchard, therefore, whether an event is out of anyone's control is *only incidental* to the event being a matter of luck--the lack of control is an occasional symptom, not the true cause.

³¹ From here on, locutions such as "or her counterpart" will be omitted. This is for ease of explication and not to indicate any particular stand for or against counterpart theory. It's not that I don't have a stand, but a wide array of answers is compatible with our present investigation.

We can also meaningfully contrast luck with skill, and a jump shot made by a skilled shooter differs importantly from one made by a lucky shooter. On this account, skill (and the lack of skill) can be understood counterfactually: an agent performs an action, X, as a matter of skill if and only if the agent can perform X reliably, that is, at most of the nearest worlds. Thus, the level of skill exemplified by Magic Johnson is explained in terms of his ability to control where the ball goes, and *that* is explained in terms of his reliability, and *that* is explained counterfactually. When an unskilled player takes the same shot and makes it as a matter of luck, however, the claim that it was a lucky shot is, on this account, a claim about his reliability in making such shots, which is to be understood in terms of the proportion of nearby possible worlds in which he makes the shot.³²

Finally, Pritchard's account does not count the standard demolition case as lucky, even though there are central parts of the case that are beyond the agent's control. Since blasting caps, wires, triggers, etc. are reliable to a high degree, the demolition comes off as planned in most nearby worlds. Should any component be constructed in a way that lends a high likelihood to its being defective, the blast will be lucky since in most of the nearest worlds the blast will not come off. If the demolition worker does not recognize that she's gotten away with using an unreliable component, she may not realize her luck, even though she has been lucky. And if she mistakenly thinks that one of the components is of poor quality and it's actually reliably constructed, then she'll mistakenly think that

³² If there are infinitely many possible worlds (and there are), then figuring proportions may seem problematic, but it is not. Suppose I draw a circle and color one third of it yellow and two thirds of it blue. The number of yellow points and the number of blue points is the same: aleph-one of each. What's crucial is to not look at the problem of getting a proportion between two sets of points but of getting a proportion between two regions. Think of logical space, then, in terms of regions or areas and not of points, and the apparent problem goes away.

she's been lucky.³³

While Pritchard claims success at presenting a deep explanation for the line between luck and a mere lack of control, there are at least four issues with the modal account. The following sections detail those defects.

4.1. The Scalar Objection

The first problem is that the modal account, like the lack-of-control account, is not properly scalar. That is, a beneficial event is lucky according to Pritchard if it occurs at the actual world but does not occur at most of the nearest possible worlds. But of two events with the same beneficial effect on the agent, the one which happens at fewer of the nearest possible worlds ought to be counted as luckier, but Pritchard's analysis does not take that into account--it tells us *that* an agent is lucky, not *how lucky* the agent is. It seems, however, that a proper account of luck ought to not only enable its adherents to *spot* luck but also to appreciate the *magnitude* of the luck.

If we allow that luck is scalar along two axes--the amount of good accruing to the agent and the proportion of the nearest possible worlds at which the event occurs--then we have two options for understanding which events to count as lucky and for assessing the magnitude of luck involved. First, we can hold the line with Pritchard and insist that a lucky event must not occur at most of the nearest possible worlds. In that case, we will not count the sun rising, for example, as lucky. We can still insist, however, that the greater proportion of the nearest worlds at which the event does not occur--so long as it's

³³ It may be the only component she has, so she has to take a chance, and it may be the only member in an otherwise defective lot that was assembled after the worker figured out what he was supposed to do. Thus, it may be reliable, and yet she may be justified in thinking that it is not. So she may be justified in thinking that the resulting blast is lucky, when it is not. This is more evidence that luck is a metaphysical, not an epistemological, notion.

over half the worlds--, the greater the agent's luck. Call this the *semi-scalar version* of the modal account of luck. Second, we can say that a beneficial event is lucky if it does not occur at all of the nearest possible worlds. On this account, the sun rising will be counted as lucky, *pace* Pritchard, if at any of the nearest possible worlds it does not rise. The amount of luck involved will vary in proportion to the worlds at which it does not rise. The luck in this case is presumably minimal. On this version of the account, Magic Johnson making a particular shot out of skill is at least to a small extent lucky, given that his skill is not absolute. Call this the *fully scalar version* of the modal account of luck. Either expansion of Pritchard's account could work as a friendly amendment of his official account. Thus, Pritchard can resist the scalar objection.

4.2. Lackey's Buried Treasure Objection

The second criticism is that satisfying the nearest-possible-worlds criterion expressed in the modal account of luck (clause L1) is not necessary for an event to be lucky. Lackey offers the following extended example as a case in which the agent is lucky to find a buried treasure, and yet also finds the buried treasure at a large proportion of nearby possible worlds (Lackey 261–64). Suppose that Sophie lives on an island and desires to bury her personal treasure before she dies, which she fears is soon because she is in failing health. She has a deep and long-abiding love of roses and wishes to bury her treasure and plant rose seeds above the treasure to memorialize the spot. As it turns out, there is one small patch of ground in an isolated part of the island with soil suitable for growing a rose bush from a seed. She buries her treasure, plants the seeds, and never returns.

Vincent, who has no knowledge of this, comes along a few weeks later and discovers the patch of ground. It is meaningful for him because he's been scouring the island looking for such a spot to plant a fully grown rose bush, and he immediately

recognizes the unique properties of the soil. In his extensive search of the small island he has come to believe that this is likely the only suitable spot in which to plant a rose bush of the variety he possesses. His belief is in fact true. It is several weeks' remove from Sophie's visit, and he cannot tell from the look of the ground that it has been recently disturbed. When he digs the hole to plant his rose bush he discovers Sophie's buried treasure. Lackey's claim is that it's clearly intuitively plausible to describe Vincent's find as a matter of luck, and yet, given Sophie's reasons for burying her treasure there and Vincent's reason's for digging his hole there, at most of the nearest possible worlds, Vincent finds the buried treasure.

Vincent finds the treasure at most of the nearest worlds because the agents have separate but very similar passions for roses, and these passions lead each of them to dig in the only suitable spot. When we hold Sophie's and Victor's interests as fixed, the case has a strong intuitive appeal as an instance of luck, and yet there is not a lot of divergence in outcome across the relevant possible worlds.³⁴

Lackey presents a formula for generating such examples: (1) consider a paradigmatic case of luck, such as finding a buried treasure, (2) conceive of an example in which each end of the chain (the burying and the finding) are, in Lackey's phrase, *counterfactually robust* but such that the connection between them is not, and (3) add whatever further features of the case are necessary such that most of the relevant worlds fall within that band of worlds in which the connection is made between the two ends of the lucky chain. The result will be a lucky case that violates Pritchard's L1 (Lackey 263).

Pritchard can, of course, dig in his heels and claim that in this case finding the buried treasure is not a matter of luck and that what carries the intuitive weight is our

³⁴ I recognize that reasonable people will doubt that finding the treasure is Vincent's good luck. Were this the only thing I have to say against Pritchard's account, I would take it as still a plausible candidate.

expectation that all cases of finding buried treasure are lucky. He could claim that what this case shows is that our expectation, while a good rule of thumb, is indeed false as a universal rule. He can do better than dig in his heels, though. On the standard analysis of possible worlds, which sphere of worlds counts as nearby depends on what similarity relation we are using.³⁵ So Pritchard can say that if we use one similarity relation under which Vincent fails to find the treasure at most of the relevant worlds, he is lucky, but if we use a different similarity relation such that at most of worlds he finds the treasure, Vincent is not lucky.³⁶ There is nothing problematic here, as far as Pritchard's use of possible worlds is concerned. The upshot is that the luck of a given event is relative to what similarity relation we are using to assess the situation. Pritchard can then go on to insist that examples following Lackey's formula for generating purported counter examples aren't really counter examples because *ex hypothesia* the target event happens at most of the nearby possible worlds. Given that there are other similarity relations under which the target event does not happen at most of the nearest worlds and thus is lucky, it's natural for readers to be confused about whether the example under consideration is an instance of genuine luck.³⁷

³⁵ David Lewis endorses the idea that what counterfactuals are true is relative to what respects of similarity we hold fixed (Lewis, On The Plurality of Worlds 20–27).

³⁶ It is understood, of course, that similarity relations divide logical space into regions and not into points.

³⁷ I'm not sure that I'd endorse an account, however, that accused anyone who presented counter examples against it of committing such a basic confusion. Still, this is one line of argument open to Pritchard that is at least somewhat adequate as a reply, and discussions about possible worlds are notorious for sending people into confusion.

4.3. The Whimsical Trip Objection

Third, let us examine the claim that criterion L1 is not sufficient for an event being lucky. Suppose that on a whim Sue takes an unplanned trip which is significant for her since it is to a place that she has always wanted to visit. In such a case, Sue's whimsical trip fits both criteria of the modal account of luck: it's an event that does not occur in most of the nearest possible worlds because it occurs to her only on a whim, and it's significant to her as an agent. But there need not be anything lucky about the event (Lackey 264–66). That's not to say that a whimsical trip cannot be lucky. Sue may, for example, meet someone on the trip whom she would not otherwise have met and that person becomes one of her best friends. That would be a stroke of good luck. On the other hand, she may catch dysentery on her trip, and that would be a clear instance of bad luck. To put it another way, *what happens while she's on the trip* may be lucky or unlucky (or a mixture or neither), but *the fact that she's on the trip* is neither, even though it's a significant event and doesn't occur at most of the relevant possible worlds.

Furthermore, whimsical events may be counterfactually robust. Suppose, for example, that Craig gives into spontaneity quite regularly, and it's been awhile since he's had an attack of spontaneity, so it's true in most nearby possible worlds that he gives in soon and calls in sick to work as a "mental health day." That is, he is about at the limit he can endure in his current cycle of non-spontaneous living. Also, suppose that he's easily susceptible to suggestions of adventure and his coworkers have been talking the last couple of days about the fun they've had kayaking. In such a scenario, it's plausible for him to call in sick from work and go kayaking, and it's also plausible that he does so in most of the nearest possible worlds (Lackey 265).

Craig's kayaking is whimsical, counterfactually robust, and not lucky; Sue's trip is whimsical, not counterfactually robust, and also not lucky. Spontaneity is in character for Craig, out of character for Sue; Craig is predictably spontaneous, but Sue is

spontaneously spontaneous; his spontaneity is counterfactually robust, but hers is not. Neither case of spontaneity is a matter of luck. For Pritchard's thesis to hold, however, whimsical events which are not counterfactually robust would be lucky, but some, like Sue's are not. Craig's and Sue's cases match in terms of whether they are lucky, but don't match in terms of counterfactual robustness. The proper conclusion is that the notions of luck and whimsy are orthogonal.³⁸

Pritchard can, of course, reply that on the semi-scalar version of the modal account of luck Craig's predictable spontaneity is not a matter of luck, but that Sue's is. On the fully scalar version, both are matters of luck, though Sue's spontaneity is luckier than Craig's. The semi-scalar version doesn't strike me as plausible in these cases, but the fully scalar version does a better job of matching my intuitions in the case of whimsical decisions--at least having them both be to some extent lucky is not as jarring to my ear as having Sue be lucky and Craig not. Thus, while his official account is not plausible, he may have an adequate reply to Lackey in terms of the fully scalar version.

4.4. The Non-Causal Objection

Finally, Pritchard's account is open to a different kind of objection. What the lack-of-control account had in its favor was that it explained luck in terms of causality. That is, the agent was lucky if and only if (a) the event was favorable or unfavorable to him and (b) he didn't have restricted control over whether it occurred. As noted above, it's useful to contrast luck with skill, but that contrast only makes sense if the different

³⁸ It is a little surprising to me that Lackey's point has not been featured more prominently (she's the only one I've seen who takes whimsical events seriously in general discussions of luck). The reason for my surprise is that lucky events and whimsical events strike me as very close cousins, and one test for an adequate theory of luck is that it at least get approximately right the contours of overlap between the whimsical and the lucky.

levels of control between luck and skill are understood causally. Put another way, control is a causal notion, and it makes sense to understand luck in terms of not being in causal control of certain important aspects of the world and their being to some extent causally chancy.

Pritchard's account loses that intuitive appeal, however, since it's not clearly a thesis about agents having or lacking causal control of the world. The reason is that understood counterfactually, causal connections are relations between pairs of events. In terms of deterministic causation, striking the match causes it to light if and only if at all of the nearest worlds in which the match is struck, it lights. In terms of probabilistic causation (aka "chancy causation"), striking the match causes it to light if and only if at most of the worlds in which the match is struck, it lights. Pritchard's account, however, is not about a relation between events but about a single event; thus it's not a causal account, and he's picked the wrong counterfactual features of the world to focus on.

In the case of Vincent and the buried treasure, we saw that according to one similarity relation, Vincent was lucky, but according to another, he was not. Perhaps by an appropriate unpacking of the similarity relations, we can discover the causal stories which underlay our determination of whether Vincent is lucky. Perhaps, then, we can map Pritchard's official account onto an account in terms of causal connections. If we have to make that translation to properly understand luck, however, his account loses the claim that it offers a deep explanation of luck. Better to have an account of luck that is explicitly about the causal connections between events.

5. The Probability-Value Account of Luck

Nicholas Rescher offers the following account of luck: an event is lucky for an agent if and only if (1) "that as far as the affected person is concerned, the outcome came about 'by accident'" and (2) "that the outcome at issue has a significantly evaluative

status in representing a good or bad result, a benefit or loss” (Rescher 32). Rescher captures this notion in a formula for measuring luck: Where E is an event, $pr(E)$ is the probability of the event occurring, $\delta(E)$ is the difference the event makes to the agent, for better or worse, and $\lambda(E)$ is the degree to which E is a lucky event, then

$$\lambda(E) = \delta(E) * [1 - pr(E)].$$

Intuitively, the idea is that the degree of luck for an event is a product of the difference the event would make to the agent and the chance of the event coming to pass. Note that even if the overall effect of Sam losing money from a hole in her pocket is even as far as the overall distribution of wealth is concerned, it may still be unlucky for her but lucky for whoever finds the money. If two people find Sam’s money and split it, then each is at least approximately half as lucky as she is unlucky. Finally, if the money lost is a great deal to Sam because of her economic status, but a trifling amount for those who find it, then the overall balance of luck may be negative, since her bad luck would outweigh their good luck. Rescher’s analysis fares nicely when compared to these intuitively plausible judgments.

Let’s measure Rescher’s analysis against our two paradigm cases: winning a fair lottery and the standard case of finding buried treasure. First, let’s see how Rescher’s account fits with our intuitions in a variety of lottery scenarios. In public lotteries, we count one winner as luckier than another if the amount of money won by the former is significantly greater than the amount won by the other. A person winning \$10 million has about half the luck of the person winning \$20 million. One might think that the difference in value between \$10 million and \$20 million is not that great because one cannot fathom how a lottery winner would have a hard time adjusting to a lifestyle such that \$10 million would not be enough but \$20 million would suffice. While this may be true in some instances--\$10 million is enough, so having the extra \$10 million doesn’t represent an added good--in fact many people blow through all their winnings and the

extra \$10 million would make a considerable difference at least in terms of extending the time from getting their first check to the time of being penniless.

The second factor in the luck of winning a lottery concerns the objective odds of winning. Winning a lottery in which the odds are one in a million is much luckier than winning a lottery in which the odds are one in a thousand; in fact, it is exactly one thousand times luckier, with the caveat that the lottery prizes are of the same value.

One advantage of Rescher's account, then, is that it shows that these judgments are substantial and why they are accurate. His formula can allow us to show the luck of various lottery scenarios in which certain parts of the story are held constant (the odds of winning or the payoff of winning), and it can help us to make sense of ordering certain lottery events according to their degree of luck in cases in which no part of the story is held constant, that is, where the odds differ and where the expected benefits differ.³⁹ Finally, Rescher's account helps to explain how one agent can be very lucky in winning the lottery and yet it not be terribly unlucky in not winning. While not winning the lottery may represent a very significant event in my life, the odds of *not winning* are very high which makes $[1 - \text{pr}(E)]$ very low. Therefore, the negative luck attached to the event will be correspondingly low. This means that not winning a public lottery is not a neutral event--as far as luck goes--but is only very slightly a matter of bad luck.

Rescher's formula relating luck, chance, and value also helps us to make sense of cases of finding buried treasure. If we keep the odds of finding two given treasures constant, the greater luck is to find the treasure that has the greater value. Similarly, if we keep the value of the treasures constant, the greater luck goes to the one which is harder

³⁹ Rescher's scheme orders lottery cases only if the differences it makes to the individual agents-- $\delta(E)$ --are commensurable or at least comparable across the examples. This complication need not detain us here, though it deserves to be kept in mind since it means that if there are incomparable goods then Rescher's answer cannot be the perfect answer.

to find, that is, whose probability of being found is less. Rescher's formula also helps us to compare cases of finding treasures where neither the odds nor the values are equal. Furthermore, Rescher's formula helps us to compare the luck associated with disparate events, such as the luck of finding a given buried treasure as compared to the luck of winning a certain public lottery or as compared to the luck of drawing to an inside flush in heads up poker when the pot has \$2500 in it.

Rescher's formula also helps us to distinguish between lucky events and whimsical events (or, more accurately, merely whimsical events from the whimsical events that are also lucky). The merely whimsical events do have a relatively low probability and a certain value, so when Rescher's formula is applied, they will count as at least minimally lucky. While that may not seem ideal to one for whom whimsical events don't seem lucky at all, it may be that their level of luck is so slight that our intuitions are untrustworthy in appreciating whether in a given case there is no luck involved or only an insignificant amount of luck. So, if Sue's whimsical a trip holds some significance for her, we could without lapsing into nonsense regard it as at least a slightly lucky turn of events. On the other hand, should a similarly whimsical decision lead to meeting someone by chance who becomes a lifelong trusted friend, that's a clear case of good luck on her part. Rescher's formula will explain the difference between the two cases as differences in both value and chance as follows.

First, the whimsical-and-yet-lucky event is one which is less probable than the merely whimsical event. If going on the trip has a low probability, then going on a trip and meeting a future trusted friend has a much lower probability since the conjunction of two unlikely events has a probability that's equal to the product of the likelihood of the two unlikely events. So if taking a trip on a whim has a probability of 0.05, and meeting a new friend for the first time has a probability of 0.07, then the probability of going on the trip on a whim and meeting a friend for the first time has a probability of 0.0035.

Second, whatever value making the trip has, that value is presumably enhanced by making the trip and making the friendship. It's not a simple calculation, the way the figuring out the probability is. Consider, for example, the trip in which no new friendship is made. That trip will have certain experiences that, taken together, have a certain value. A trip making the friendship though would have at least a slightly different set of experiences, and it is impossible to see in the abstract how that might play out. Even though there's no neat metric by which we can figure the various counterfactual differences in value between the merely whimsical trip and the whimsical trip which resulted in making the new friend--the lucky trip--, it's clear that such examples are to be found, and it's clear that should Rescher's account of luck hold up, that it explains the difference between them.

As a final example, Rescher's account will explain the story of Vincent and Sophie and the buried treasure as a genuine instance of good luck for Vincent. The value of the treasure is a boon to Vincent. His good luck may or may not be bad luck to anyone else. The other component of the case that's relevant is that the likelihood of both Sophie and Vincent having such an intense interest in roses is fairly low. Indeed, as we can see from above, whatever likelihood there is for Sophie to pick that spot in which to dig and whatever likelihood there is for Vincent to pick the same spot in which to dig, the combined likelihood is far less, since it's the product of two events that are not independent of each other.

5.1. Critique of Rescher's Analysis

One strength for Rescher's account is that it treats luck as a fully scalar notion. It tells us of every situation that has a positive or negative effect on the agent and a non-zero probability of occurring how lucky it is for the agent. On the down side, however, Rescher's account does have problems. Consider the case in which Joe can do something

such that the outcome is of great value and the ability to bring the outcome about is rare. For example, suppose that Joe can perform a medical procedure on Mary that will save her life and in the situation in which he's called upon to do it he's the only person who can perform the procedure. As Andrew Latus points out, it's counter-intuitive to call his exercising that ability lucky (Latus 467). Latus further notes that while it is not lucky for someone like Joe to exercise his ability, it is lucky for Mary that Joe is there when she needs him. The same event may be lucky for Mary but not a matter of luck for Joe. Rescher's formula can handle this case, since it is implicitly relativized for the agent. There is a problem, however. Suppose that Mary's condition has this special complication: she will die if someone performs the procedure that Joe has knowledge of and will live otherwise. Joe does perform the procedure on her and she dies. From the point of view of her having the condition, it's very bad luck that Joe performs the procedure on her and she dies because there's a very low likelihood that someone will act such that she dies (remember, only very few people know how to perform the procedure). From the point of view of her having the special complication and Joe performing the procedure on her, however, it is very highly probable that she dies, and so not a case of bad luck.⁴⁰

The problem is that while Rescher's formula nicely relativizes the luck of an event to different agents, it doesn't recognize that assessing the probability of an event is sensitive to what background information we hold as relevant. Thus, there are different vantage points from which a given event for a given agent is lucky and other vantage points from which the same event for the same agent is not lucky. Rescher's formula has no mechanism for relativizing the luck of an event to different vantage points and so is

⁴⁰ This may be a great surprise to Joe who has absolute confidence in the effectiveness of the procedure. More evidence that luck is a metaphysical and not an epistemological concept.

improperly insensitive to the various values that can be reasonably assigned to the probability of an event.

Clearly, the issue is that we need to develop some sensitivity in saying precisely what in a situation is lucky (and for whom the situation is lucky). I don't see that Rescher's formula can be easily adjusted to relativize it to various vantage points. But that's what's needed. In the next section, we will see a hybrid view which combines the strengths of the lack-of-control account and Rescher's explanation of the degree of luck in terms of an interplay between probability and value.

6. The Probability-Control-Value Account of Luck

The hybrid view which I will now examine comes from Andrew Latus's "Constitutive Luck." He sees both the lack-of-control view and Rescher's view which focuses on probability and value as both wanting, but thinks that each is wanting because each lacks the strengths of the other view. Thus, he develops an account of luck that leverages Rescher's formula but which explains degrees of luck in terms of an interplay between probability, control, and value. The formula he proposes is this (Latus 468n): where E is an event, $\text{pr}(E)$ is the probability of the event occurring, $\text{con}(E)$ is the amount of control the agent has in the situation, $\delta(E)$ is the difference the event makes to the agent, for better or worse, and $\lambda(E)$ is the degree to which E is a lucky event, then

$$\lambda(E) = \delta(E) * [1 - \text{pr}(E)] * [1 - \text{con}(E)]$$

The advantages of this account of luck are the combined advantages of the lack-of-control account and the account in terms of chance and value taken separately, but with the added advantage that this account is able to avoid the problem cases that were faced by each of the component accounts. Let us revisit those problems in turn.

The first problem for the lack-of-control account was pressed by Lackey's example of her picking up her daughter after school and whether that counted as good

luck for her husband. From her husband's perspective, the event has value and is out of his control, but it is not a matter of luck that the daughter gets picked up and brought safely home. Given Lackey's conscientiousness as a mother, the probability of her keeping her word and picking up her daughter as planned is very high, so that brings the level her husband's luck with regard to this event to a very low level.

Now consider again Lackey's example of the demolition worker who blows up a building, but only does so because of a stroke of luck that returns to her the control that she presumed herself to have. What makes the nail under the coat returning the control to her a matter of luck is that the event is not something that she's got control of, and it's very unlikely that she should have gone from a lack of control of the blast to being able to control it in just the nick of time. The luck, then, comes in her having regained control of the explosives, and there is, given that the button now properly controls the explosives, no further luck involved in whether her pressing the button will cause the demolition.

A more serious problem for the lack-of-control account was in the standard demolition case since it is out of the demolition worker's control that the button was correctly constructed at the factory, that the blasting caps weren't factory duds, etc. On that view, the standard case turns out to be lucky, and that is counter-intuitive. On Latus's view, however, the lack of control *by itself* is insufficient for an event being lucky; it is also required that the event be of low probability. This breaks our example into two cases. In the first case, suppose that either the button or the blasting caps or both were unreliably constructed at the factory. If the reliability is low enough, say buttons of that type fail 50% of the time, then the demolition coming off as planned is lucky. It would be luckier if the reliability of the buttons were lower. In the second case, suppose that both button and blasting caps are reliably constructed. Then the probability of all the parts working as planned is very high, and the luck involved is very low. That is a significant improvement over the lack-of-control account taken in isolation.

Turning now to the issues facing Rescher's account, we see that one worry is the case of Joe having medical training that Mary needs on the spur of the moment and for whom Joe is the only hope. It's lucky for Mary that she's in the immediate vicinity of someone who's able to give the needed care. This is because it's out of her control that Joe is there (she hasn't hired Joe as her personal physician because she knows what particular skills Joe brings to the situation), *and* it's improbable that someone with the needed training should be in close proximity. From Joe's perspective, the situation is lucky in the sense that it's unlikely that he'll ever need to apply his training, and it's out of his control that on this particular occasion he needs to utilize his training (he hasn't poisoned her for example). Also from Joe's perspective, it's not a matter of luck that he's able to correctly apply his training since he received the training and has kept his skills up-to-date, and he thus has a high probability of success in performing the procedure. All of these determinations seem reasonable.

Conjoining these two insufficient explanations works in this case to *correctly categorize* a wide variety of cases, but this kludge cannot *correctly explain* them. That is, what makes the explanation work is that the conjuncts work together to save each problem case from being an apparent counterexample, and that there's no counterexample that applies to both parent analyses. The conjunctive explanation, however, doesn't illuminate why each particular conjunct is part of the proposed definition. Why is it that a lucky event is one that's significant to the agent and both a low probability event and out of the agent's control? What is it about the event being of low probability and out of the agent's control that befits it for being an instance of luck? To put the question another way: what is it that *integrates* the notions of low probability and lack of control for an event being lucky?

Clearly, then, giving a classification of events that's merely coextensive with the concept of "luck" isn't ideal. Were it the only alternative, we could be satisfied with

those results, but in the next section I present an explanation that not only covers all the proposed cases but which also explains why each case counts as lucky, unlucky, or not a matter of luck.

7. The Sensitive-Causal-Chains Account of Luck

Let us consider four related scenarios. I'll then give an account that *both correctly classifies and correctly explains* them as relevantly lucky, unlucky, or nonlucky. After that, I will explain what the common failure is among the four accounts already considered, and why my account doesn't fall prey to the same defect.

Scenario 1: John is walking through the business district at 10:02, on his way to meet a friend for coffee. A worker on the top of the eight-story building he is passing accidentally drops a hammer to the sidewalk below. At two seconds before the hammer reaches ground-level, John stops to count the coins in his pocket, and had he not stopped the hammer would have hit him in the head and he would have died. What prevents him from being killed by the hammer is the intervention of a whim. Judging from the perspective of two seconds before the hammer reaches ground level, John's sudden whim is one of the very few events that could have prevented him from being at the point of impact at the time of impact. Thus, at two seconds before impact, the causal chain between the hammer falling and John getting hit by it is relatively insensitive to interruption. And yet it does get interrupted, John lives, and John is lucky.

Scenario 2: John is walking through the business district at 10:02, on his way to meet a friend for coffee. A worker on the top of the eight-story building he is passing accidentally drops a hammer to the sidewalk below. At ten seconds before the hammer reaches ground-level, John is slowed in crossing the street by a driver who is making an illegal right turn. Judging from the perspective of ten seconds before the hammer reaches ground level, the fact that John slowed for the illegally turning driver is one of the very

few events that could have prevented him from being at the point of impact at the time of impact. (There are other events. Judging from that perspective, he might also have stopped at two seconds before the impact to count the coins in his pocket.) Thus, at ten seconds before impact, the causal chain between the hammer falling and John getting hit by it is insensitive to interruption, though not as insensitive to interruption as Scenario 1. And yet it does get interrupted, John lives, and John is lucky, but not as lucky as in Scenario 1.

Scenario 3: John decides to call a friend to meet her at a coffee shop several blocks from his work. It is something they do several times a month, and it's been two weeks since they've last met. He thinks it would be good for him to work till about 10:00, and then set out walking so that he can meet her at 10:15. He calls her at 9:05 to arrange the meeting, gets her voice mail, and on a whim doesn't leave a message. So instead of going to the coffee shop, he stays at his desk until 11:30 and goes to lunch. At 10:02, just as he would have been on his way to meet her, the hammer crashes to the sidewalk, but he's still at his desk. Even though John lives in this scenario, he's not lucky since the causal chain between his calling to arrange the meeting with his friend and his location at the time that the hammer crashes to the sidewalk is so variable and so sensitive to deflection that it is very unlikely that he would be at the wrong place at the wrong time. (On the fully scalar account of luck, his level of luck is only very minimal, almost to the point of being nil. It is certainly beyond anyone's ken to know that he's lucky in this small way.)

Scenario 4: John decides to call a friend to meet her at a coffee shop several blocks from his work. It is something they do several times a month, and it's been two weeks since they've last met. He thinks it would be good for him to work till about 10:00, and then set out walking so that he can meet her 10:15. He calls her at 9:05 to arrange the meeting, she answers her phone, and they agree to meet. At 9:57 he saves the

document at the end of the section he has just finished editing, then gathers his coat to leave. He gets ten feet from his desk and goes back to get his umbrella. He gets to street level, looks at what parts of the sky he can see, and decides it's just misting and doesn't look like it will rain hard. The cool mist will feel good, he thinks, so he doesn't open his umbrella. Along the way, he stops to look at a poster advertising a local band, to double check the return address on the bill he mails, and to pet a woman's dog. At the light he waits not only till the traffic is clear but also till he gets the 'walk' light. As he is crossing the street a driver attempts an illegal right turn, but stops because John is in his path. He considers stopping to count the coins in his pocket but remembers that it's his friend's turn to buy, so he doesn't. At 10:02, the hammer that very easily could have crashed to the sidewalk instead catches John on the crown of his head, and he dies. This string of events is very unlucky for John.

Now for the explanation of my classification of scenarios. In his postscript to "Causation" David Lewis discusses "insensitive causal chains" (Lewis, "Causation" 184–88). His target is different from mine, since he's working to distinguish between the case of an agent killing someone from the case of a death in which the agent is a part of the causal history but not enough to say that the agent killed him. Lewis points out that if the causal chain between the agent's actions and someone's death is insensitive to interruption, then it's true to say that he's killed him. For example, if the agent puts a loaded firearm against someone's sternum and pulls the trigger, and the person dies, then he has killed him. On the other hand, if the causal chain between the agent's behavior and the person's death is highly sensitive to interruption, then the agent did not kill him. On that view, the question of whether one person killed another may not have a neat solution, except at the endpoints where the agent's behavior directly leads to a death or has no causal connection to it. To say that a causal chain is insensitive to deflection is not to say that it cannot be deflected. Rather, it's to say that the connection between two

events along the causal chain can be broken and yet each way in which the causal chain could be broken is consistent with the past at the start of the causal chain and with the laws of nature. Causal instability, on Lewis's view, seems to be a rather mundane sort of connection between events.

Considering the line between causing a death and killing a person, from an ethical perspective as opposed to Lewis's metaphysical point, Jonathan Bennett says that killing someone requires "that the causal chain run through a stable and durable structure rather than depending on intervening coincidental events" (Bennett, "Morality and Consequences" 71; Bennett, The Act Itself 32). For example, suppose that Tricia lets her dogs into the back yard, they escape to the street, run in front of a car which swerves to avoid them, and the car kills a child. Tricia is *a* cause of the child's death (but not *the* cause) and the causal chain between her behavior and the death is sufficiently unstable so that it's false to say that she killed the child. (Though she may bear some legal responsibility since her dogs contributed to the child's death, that is a legal distinction and comes from neither metaphysics nor ethics). Let's see how this "stability requirement" can be brought to bear in the case of John's date with the hammer in each of the four previous scenarios.

Scenario 1: at two seconds before the hammer would have killed John had he kept walking in the direction and at the pace he was walking, he stops, and does not die. This is lucky for John because only a very small number of occurrences can interrupt the causal chain and spare his life. One of those is that he stops to count the coins in his pocket. There are others. He might have changed course to walk in the sliver of sunlight that was now reaching the sidewalk and avoided being hit. Or someone might have called to him because he dropped something or The point is that while the causal chain from at the two-second mark is relatively insensitive to interruption, it's not so stable as to be not interruptible.

Scenario 2: at ten seconds before the hammer would have killed John had he kept walking in the direction and at the pace he was walking, he changes his pace to avoid being hit by a driver making an illegal right turn, and does not die. While the causal chain from at the ten-second mark is insensitive to interruption, it does get interrupted, and John is lucky that it did. But the causal chain at the ten-second mark is more sensitive to interruption and less stable than the causal chain from the two-second mark. It includes all of the interruptions that could have happened at the two-second mark and many more. Thus, John is lucky, but not as lucky as he would have been had he slowed or stopped or changed direction at the two-second mark. Since the negative payoff in each scenario is the same--John's death--, the differences in luck between the two-second scenario and the ten-second scenario is explained by the relative sensitivities to interruption of the two segments of causal chain.

Scenario 3: at one hour before the hammer would have killed John had he started at a particular time and kept walking in a particular direction and at a particular pace, he decides to stay at his desk and does not even walk along that sidewalk. On a semi-scalar account of luck, he's not lucky because, given the instability in the many links of causality that might have existed between his decision to call his friend and when the hammer hit, it is very easy to interrupt the causal flow. The hour-long causal chain is so sensitive to interruption and so unstable that when it does get interrupted and he lives, it's not a matter of luck that he lives. (On a fully scalar version, he is lucky, but only minimally so.)

Scenario 4: at one hour before the hammer falls, John starts a train of events such that at 10:02 he dies. He was unlucky because the causal chain that existed between his decision to call his friend and when the hammer killed him was very easy to interrupt, and yet at every fork in the path that could have saved his life, he took the fork that led to his death. Had he not stopped to check so many details on the document he was editing (or

stopped to check more), he would have lived. Had he not gone back to get his umbrella, he would have lived. Had he not stopped to check the skies, read the poster, mail the letter, or pet the dog, he would have lived. Had he been a person who rushes across intersections against the warnings of the 'do not walk' sign, he would have lived. Had the driver making the illegal right turn forced the situation, he might have been hit and got a broken leg or gotten into a fist fight, but would have lived. The hour-long causal chain is sensitive to interruption and unstable, so that when it does not get interrupted and he dies, it's a matter of bad luck for John.

7.1. Formulating the Causal-Contrastive Account

There are two factors that determine how lucky an event is for an agent. First, as is almost universally recognized, for an event to be lucky it has to make some positive or negative difference in the agent's life. I need to make two clarifications about that point: (1) there is no such thing as neutral luck, but all luck makes at least a minimal difference to the agent, and (2) since difference is a comparative notion, the positive or negative difference is in contrast to how the agent would have been affected if the event hadn't come about and some other event had come to pass. The events could be both positive, both negative, or each of opposite polarity. What matters is that the various events represent a difference in value for the agent.

Second, as we have seen, there must be some sort of causal gap between a precursor event and the lucky or unlucky event. If there were no causal instabilities in Scenario 4 that could have averted his death, his death would have been unfortunate for him, but not unlucky. To say that the event is unlucky is to recognize not only the event that causes his death but also to implicitly recognize that there were other events that might have come about such that he would have lived. The difference in value for the agent is the difference in value between how the world unfolds causally and how it might

have, but did not, unfold. Causation on this view is a relationship between events with an open texture.

Now for a more careful formulation of my thesis. At an initial event, IE, the world can fork toward either the resulting event, RE, or an averted event, AE. The luck of an event for an agent, A, is the relation between IE, a resulting value for A, RV, which occurs concomitantly with RE, and an averted value for A, AV, which would have occurred if IE had led to AE, instead of RE. The amount of luck is determined by (1) the magnitude of the positive or negative difference between RV and AV, and (2) in contrast to AV coming about, the likelihood that RE (and hence RV) will come about given IE, where the likelihood is measured by the resiliency of the causal chain leading from IE to RE. The resiliency of a causal chain can be explained in terms of the stability of the causal chain and its insensitivity to deflection. For example, the more likely that the causal chain links IE to RE, the less luck is involved in A receiving the positive or negative difference between RV and AV. On the other hand, if the causal chain that links the IE to RE is unstable and open to easy deflection, then RV which comes about as part of RE is lucky or unlucky to the extent that RV (in contrast with AV) is positive or negative for A. If the difference for A between RV and AV is zero, then no luck is involved.

The formula I propose is as follows. Let IE, RE, RV, AE, and AV be as before. Let $pr(IE \rightarrow RE)$ be the conditional probability of RE occurring given that IE occurred, $\delta(RV - AV)$ be the difference the RE makes to the agent, for better or worse, as compared to what AE would have produced, and $\lambda \langle IE, RE, AE \rangle$ be an ordered triple meant to capture the luck involved in RE coming to pass as opposed to AE coming to pass. Thus,

$$\lambda \langle IE, RE, AE \rangle = \delta(RV - AV) * [1 - pr(IE \rightarrow RE)]$$

On my view, luck is contrastive along two axes. It makes no sense to ask how lucky the agent is in a given situation unless we determine (1) what initial event we're starting our accounting from (which determines the probability of getting to RE) and (2) what might have happened had the resulting event not come to pass (which gives the contrast between how things have turned out and how they might have turned out).

Let me shed light on the notion of contrastivism. Suppose it is raining and a child asks his mother why it is raining. The child's question is, on the face of it, ambiguous. Is he asking why it is raining as opposed to not raining? If so, then the answer is one that concerns relative humidity. Is he asking why it is raining as opposed to snowing? If so, then the answer is one that concerns temperature. Explanatory answers presuppose that what is desired to be explained is one feature of the situation as opposed to a given alternative feature of the situation (Sinnott-Armstrong 258). What counts as an answer is relative to the child's intention in asking.

In a similar way, asking how lucky a particular person is in winning the lottery is an ambiguous question. Winning the lottery is lucky in contrast with not winning it, if we judge from the background of having already drawn five of six numbers, but it's exceptionally lucky from background of merely having bought a ticket. Winning a given prize is lucky compared to winning a lesser prize and is unlucky compared to winning a significantly greater prize.

In saying that John is lucky to have missed being killed by the falling hammer, context helps us to pick out our initial events and our alternative events. Did he just miss being hit because he stopped to count the coins in his pocket or because his path was interrupted by a driver making an illegal turn? That distinguishes the two-second from the three-second scenarios.

What, though, of a case which context doesn't pick out a salient event that is the one which saved his life? Suppose a case in which he walks from his desk much as in

scenario 4, but in which he lives, and there was no conspicuous counting of the coins, for example, that we can conveniently point to as the event which averted John's death.

What I'm committed to is this: if John is lucky to have not died, there is an IE that leads to an RE of him living. As opposed to RE is a contrast class of AEs such that if any of these AEs had come to pass John would have died. For any AE in the contrast class, John's luck is figured from the difference in value (life versus death) and the likelihood that that particular AE would occur as opposed to RE. If there are nine members of the contrast class, for example, then there are nine different answers we can give to the question of how lucky John is.⁴¹ In an actual case, we won't know what members are contained in the contrast class (nor will we likely be adept at figuring the odds), so we'll have to live with the idea that our estimation of John's luck is more guess than estimation. Epistemically, it's the best we can do.

Thus, there's no such thing as luck *simpliciter*, since luck only makes sense in contrasting alternative events and their values from the perspective of an initial starting point. Change the starting point, and our estimation of the luck involved can change. Change the alternative event the resulting event is compared to, and our estimation of the luck involved can change.

Bill meets the woman of his dreams at the store. Lucky, we think. But in comparison to what alternative event? My claim is that if we compare the value of that trip to the store with what he had before (not having met the woman of his dreams) we are making the wrong comparison. We should be contrasting the value he actually realized with the value he might have received. It's not that he was single and wound up in a solid relationship that makes him lucky. It is that he was single before he went to the store and had he followed a different path through the aisles he most likely would remain

⁴¹ The details in this explanation of contrast classes is inspired by Bradley (Bradley 287).

single. The trip to the store changed his life in a welcome but unexpected way that makes him lucky. The luck, however, is not the backwards-looking contrast between his new life and the life he had, but a forwards-looking contrast between how his life goes as a result of meeting his new person and how it would have continued had he taken a different path.

Let me make a number of general points. First, on my view, if the world is causally determined, then there is no luck. Luck is a causal, and so an ontological, category. It is not an epistemic category. Second, events *simpliciter* are never lucky on my view and the magnitude of luck is relative to what point in the causal chain from which the luck is being judged. Winning the lottery is minimally lucky when judged from the point of view of someone who has matched five balls and is waiting for the sixth to drop, but it's very lucky considered from the time that the ticket is bought three days before the drawing.

Third, in many cases our ability to compare the probabilities involved is limited because we may not really know how stable the causal chains are. We may also not know what value the averted event held for the agent. Finally, as we've seen already, we may not know what events are in the contrast class. It's enough for my purposes if we have paradigmatic cases that we can work from in theory building, even if we cannot always deftly apply the notion to actual cases.

Fourth, as Thomas Jefferson said, "I'm a great believer in luck, and I find the harder I work, the more luck I seem to have." This quote only makes sense if hard work can lead to success by stabilizing causal chains. Rescher is wrong when he says that the enemy of luck is knowledge. The enemies of luck are skill and preparation (a kind of insurance against luck) since they make causal chains more stable. I'm not lucky to pass a test that I've studied hard for because I'm prepared to answer a wider variety of questions, so there's less that can upset my path to getting a good grade. I'm not lucky to

be financially stable if I follow sound financial principles. Or if someone insists that I'm at least minimally lucky, I can agree and attribute the residual luck to the residual instability: I'm lucky to pass the test given that the professor could have asked exam questions that were off topic from the material she said we would be tested over or were of a complexity beyond what I had prepared for.

8. Objections and Replies

Objection: Analyzing luck in terms of a multi-part relation doesn't accord well with common usage. When we say that Jones is lucky to have survived a fall, we don't have in mind all this *analytic machinery*. It is, in other words, a lot of technical explanation for something that's almost as common as water.

Reply: It's okay to speak with the vulgar, but philosophers also ought to think with the learned, and I think that this machinery is what's behind the true parts of common usage. The five-part relation seems to me to be not only true but to also provide us with a sufficiently deep understanding of the phenomena. Thus, even if it should not perfectly accord with common usage, I'd be happy to set aside examples from common usage if these seem like corner cases. On the other hand, I know of no case from the luck literature where the causal-contrastive account gives an answer at variance with (my) intuition. It does seem at variance with the details of explaining luck compared to accounts that seek to explain luck in terms of folk notions, like a lack of control. (Even though Rescher's analysis is as a formula not so different from mine, his explanation is firmly grounded and easily expressible in folk terms.) In addition, while luck may be as common as water, understanding water as the chemist does shows reality to be much more technically complicated than would occur to the unschooled. I'm not worried that it is technically more advanced than our folk notions would have us to understand. In fact, as a philosopher trained in the analytic tradition, I'd be more surprised if the explanation

were as straight forward as our folk judgments. That last point, however, is a bit of intellectual biography and not an argument. The argument, should one be wanting, is practical: how clearly does it allow us to think about judgments of luck? The causal-contrastive account accords well by that standard.

Objection: Luck on the causal-contrastive account is relativistic--that is, there's no sense in which John is lucky *simpliciter* in having the hammer miss him. Our common usage would say that John is lucky not to be killed by the hammer which lands at his feet full stop. A relativistic account, therefore, cannot capture our common usage.

Reply: Guilty as charged, but I don't take any part of that objection to be telling against the causal-contrastive account. There is no sense in which John is lucky *simpliciter*. He is lucky only insofar as RE is compared to an AE and the value of how things turned out for him (RV) is compared to how they might have turned out (AV). That answers--or at least stands up to--the first part of the objection. The second part amounts to claiming that ordinary usage makes absolutist pronouncements about luck. Is it the case that John is lucky to have been missed by the hammer when that event is considered in isolation from other events and other alternative outcomes?

Let's consider these two notions: space-time taken as absolute (Newtonian space-time) and space-time taken as relativistic (Einsteinian space-time). Newtonian space-time underpins Newtonian physics which gives a picture of the world that matches very well with common sense. The problem with that notion of space-time is that it doesn't accord with how the world is, so Newtonian physics gives us at best an approximate understanding of reality. On the Newtonian account, if a train is going away from you at x miles per hour, and a person is walking toward the front of the train at y miles per hour, he is moving away from you at a rate of $x + y$ miles per hour. The problem is that if the train is going at the speed of light, so the train's speed is c , he's not moving away from you at $c + y$ miles per hour. That makes no sense from a folk understanding of physics.

And yet, it's true. Thus, what gets rejected is the underlying conception of space-time and in the move to relativistic physics, the Newtonian account is seen as a very useful fiction for working on a very wide variety of real world engineering problems.

In the same way, a relativistic understanding of luck does much better than an absolutist conception. Suppose we grant that John just is lucky in having the hammer miss him by two feet, and we grant that he's lucky in having it miss him by ten feet. On the absolutist account, what explains the fact that in the first case he's luckier than in the second case? I'm at a loss to see how absolutist conceptions can explain those differences in luck. For example, on the lack of control account, John has no more conscious control in the second case than in the first. The other accounts are at least implicitly relativistic: one gets different spheres of worlds depending on what series of preceding events one takes as relevant, and the probability of an event occurring is relative to what events one takes as the relevant precursor events.

Objection: The details of the nature of luck don't greatly matter in discussions of moral luck, for example, because as long as each of the theories agree on the central cases of what events are lucky, each theory makes about the same contribution to discussions of moral (or epistemic or practical) luck. It's nice, perhaps, to have such a discussion to give a more filled out explanation of moral luck, the central topic of the present work, but it's not necessary. Whatever the final pronouncement on whether luck ever makes a moral difference, any account of luck will do as long as it allows us to pick out our central cases for discussion. Williams, for example, made a significant contribution to the discussion by taking the notion of luck as undefined but intuitively understood. Put another way, the discussion of moral luck will focus more on the nature of morality than on the nature of luck, and any theory of luck that maps the territory at all accurately will be sufficient for discussing the various moral questions and the answers to them. For example, one significant discussion of moral luck is whether there is anything that might

properly be called moral luck. The focus is on always on whether a correct moral theory can allow that the moral evaluation of an agent can vary depending on how lucky she is. It's never on whether there is any luck.

Reply: The lack-of-control account is causal but not relativistic. The other competitor theories are not explicitly causal but are at least implicitly relativistic. Since my account is both causal and relativistic (along two axes), it is markedly different from the competitor accounts. In addition, since one aspect of the overall discussion of moral luck is how luck can affect moral responsibility, it's more than a merely cosmetic difference. That is, if luck can effect moral responsibility and luck is a dually relative notion, then moral responsibility will also have to be understood relativistically. That is, our understanding of luck will affect our understanding of morality.

That's no small contrast with the competitor theories.

Chapter 3

Moral Responsibility

1. The Nature of Moral Responsibility

The focus of this chapter is the sense in which an agent can be morally responsible. On the view presented here, an agent, A, is morally responsible for some action, intention, or character trait, X, if and only if A is aptly subject to a range of participant reactive attitudes (such as praise, blame, anger, regret, guilt, etc.) for X. This is moral responsibility as accountability. The notion of being an apt subject of reactive attitudes can be fleshed out as it being fair to the agent to react to her in these ways. The appropriate standard of fairness to employ is that the action, intention, or character trait that we are reacting to in some fundamental way belongs to the agent. Thus, moral responsibility presupposes an attribution requirement.

One common way of discussing this attribution requirement is on analogy with the agent having a moral ledger in which X can be entered as an asset or a debit. The idea is to give ownership to A of X. Derek Pereboom describes the situation as follows:

for an agent to be *morally responsible for an action* is for this action to belong to the agent in such a way that she would deserve blame for it if the action were morally wrong, and she would deserve credit or perhaps praise if it were morally exemplary. . . . [F]or an agent to be morally responsible for an action is for it to be imputable to her (Pereboom xx).

This ledger or attribution requirement for moral responsibility can be further clarified if we address what defeating conditions prevent an action, intention, or character trait from properly being attributed to the agent. The standard answer goes back to Aristotle's *Nicomachean Ethics* (Aristotle, NE Bk. III, 1–5). Aristotle holds that two conditions are required for moral responsibility: (1) that the action, intention, or character trait be under

the agent's voluntary control, and (2) that the agent not be acting from a place of non-culpable ignorance. Thus, should we find that the agent has a microchip in her brain by which scientists are controlling her behavior, we would withdraw the claim that she is responsible for her behavior. In addition, behavior on her part which would normally cause us to resent her, for example, will no longer engender that reaction, at least toward her. We might, of course, resent the scientists who are controlling her behavior.

In addition, if we find that the agent has acted out of certain kinds of ignorance, we will drop the claim of moral responsibility and the attending moral reaction. As Brian Rosebury points out, however, there is a moral responsibility not to be ignorant of the relevant facts surrounding one's actions or the consequences of one's actions. Speaking of Hitler's and Stalin's accomplices he says,

It is not that they were simply less intelligent than the others who acted differently, if it can be shown that they failed to pursue or deploy relevant knowledge with the energy that (given their capacities and circumstances) would have been appropriate in view of the gravity of their actions or the actions with which they associated themselves, a gravity that increases the stringency of the requirement of deployed knowledge, then they are morally culpable (Rosebury 504).

Note that on Rosebury's explanation, it is not enough that the agent possess the morally relevant knowledge but that she also deploy it. Put another way, morally responsibility requires that one be actively informed. Rosebury also notes, interestingly enough, that certain philosophical problems dissolve if the knowledge requirement isn't met: Jim has no obligation to shoot one of the Natives, in Williams' famous example, because he doesn't know that the Captain will keep his word to let the other nineteen go.

The voluntary control requirement can be cashed out in terms of a libertarian notion of free will. An agent can only be morally responsible if her actions are under her

control in the right way. As we have seen with the example of scientists controlling an agent's behavior, it is not enough for the behavior to go through the agent's will. A further requirement is that the agent herself, through the exercise of her will, controls her behavior. In addition, for a behavior to have the ultimate sort of freedom that libertarians require for moral responsibility, the agent has to have a number of alternative options available to her and from which she chooses. For example, Helen cannot be morally culpable for slighting me if that was her only option. This follows from the dictum that “ought’ implies ‘can’”, and if she could not have refrained from slighting me, then it's false to say that she ought to have refrained from doing so. Thus, moral responsibility requires that at the time of choice the agent face a causally open future and that it be up to her which path the world follows. This is the traditional view of free will and moral responsibility. As we shall see, there have been a number of attacks on it, against which I shall briefly defend.

Finally, we can say that a person is a *morally responsible agent* just in case it is appropriate to hold her morally responsible for her behavior. Of course, that an agent is morally responsible does not require that we praise or blame her for every action. Sometimes, her behavior will be morally neutral, such as when she puts on her socks before her watch. In other cases, we might have good moral reason not to praise her for a praiseworthy act, if, for example, it will make her vain.

2. Hume’s Fork

The notion I have sketched above requires a particularly robust sense of free will, a *libertarian* notion of free will. There are, traditionally, a pair of arguments against libertarian free will that go by the name “Hume's Fork.” The idea is this. One can argue that causal determinism--the thesis that given the past and the laws of nature there is but one possible future--implies that no agent is ever free in the libertarian sense, since the

future does not represent to the agent a “garden of forking paths.” One can also argue that if causal determinism is false, which path the world takes is sometimes open, and there is a garden of forking paths. On one understanding of this view, the agent has no free will in the libertarian sense because it is never up to her which of those paths is taken. Either causal determinism is true or it is false, so the agent is never free in the libertarian sense. Because my claim is that the libertarian sense of freedom is required for moral responsibility, one can extend the conclusion of Hume's Fork to say that no agent is ever morally responsible. Let's pause to look at both sides of Hume's Fork.

2.1. Free Will and Causal Determinism

One fork of the argument that takes seriously the idea that the world is causally determined. This idea is often captured by the Consequence Argument, and my exploration of it follows that of Peter van Inwagen (van Inwagen, *An Essay* 56–82).

Let us consider a case as innocuous as whether I shaved this morning (I did). On the supposition that causal determinism is true, we can consider the state of the world three billion years ago and the immutable laws of nature that have governed the world's events from the beginning of time. Given those facts about the world, it follows that I shaved this morning.

Let me introduce the following monadic sentential operator, Np , which says that p is true and no human being has ever had any control over whether p is true. We can then construct the following argument in which “P” abbreviates “The past is as it was three billion years before January 4, 2012”, “L” abbreviates “The laws of nature are as they are”, and “S” abbreviates “Pat Beach shaved on January 4, 2012”:

NP

NL

$N(P \rightarrow (L \rightarrow S))$

Therefore, NS

The first premise is true, because it is true that the past was a certain way three billion years ago and that no human ever has had any control over that. The second premise is also true since the laws of nature, being immutable, are never such that their content has ever been controlled by any human. The third premise expresses the idea that my shaving this morning follows from P and L in a way that no human has ever had any control over. This branch of Hume's Fork assumes causal determinism, so we will take the third premise as provisionally true.

Given that the premises are true, all that remains to be seen is that there is an inference rule that justifies our moving from the premises to the conclusion. van Inwagen identifies such a rule, which he calls Rule Beta:

Np

$N(p \rightarrow q)$

Therefore, Nq .

He does not argue for Rule Beta; rather, he claims that Rule Beta strikes him as obviously true. Certainly it strikes me as true that if I'm powerless over a conditional and its antecedent, I'm powerless over its consequent. What Rule Beta does, in effect, is to legitimize that powerlessness is transferable in just this way.

The conclusion to the argument, NS, follows from two applications of Rule Beta. Hence, the argument is valid with true premises. If all of this is accurate, then it looks impossible for me to have done anything other than shave this morning. If that is the case, then, in the libertarian sense of free will, I wasn't free to not shave and hence my shaving this morning wasn't free either. For sake of illustration I have picked a particular behavior of mine (shaving this morning), but any of my behaviors would have done as well in illustration and indeed the behavior of any human being would have also served. Thus, we can generalize the argument to show that no human has ever been free in the libertarian sense, given that causal determinism is true.

2.2. Free Will and Causal Indeterminism

The other side of Hume's Fork employs what is sometimes called the Rollback Argument. We can consider the same event as before, that of my shaving this morning. If causal indeterminism is true and whether I shaved this morning was an undetermined event, then consider the following scenario. Suppose that God watches me shave but then rolls back the history of the world to some time prior to my shaving this morning. He could roll it back to the Big Bang and then let the history of the world unfold again. In such a situation, events which had previously led to my birth might not have happened, and hence it would be false on the unfolding that I shaved this morning. For clarity's sake, however, let's suppose He rolls back world to the ten seconds before I shave and lets it unfold. Let's also suppose that he does this some large number of times, say a thousand.

While there may be patterns that emerge in the unfolding (perhaps I shave 60 percent of the time), on any particular unfolding it is a random event whether I shave on that particular unfolding. If it is random, however, on any particular unfolding, it's not the case that I ever control how the world unfolds. Libertarian free will requires that I have options from which to choose, that is, various paths that my life can take, but the randomness inherent in causal indeterminism ensures that it is not *me* that guides the path from the present to the future from among the causally open paths. Since it is not me that guides the future, then if causal indeterminism is true, I don't have free will. Again, we could generalize this for any causally indeterminate event that any human being has ever been involved with.

We can conclude that, given both sides of Hume's Fork, we don't have free will in the sense required by moral responsibility, and hence we don't have moral responsibility.

3. Understanding Free Will

The libertarian, of course, has an answer to the above argument, which is to find an explanation for how the agent can exercise control even though the world is causally indeterminate. (The libertarian can endorse the Consequence Argument for the incompatibility of free will with causal determinism. What the libertarian cannot do is to endorse the existence of the libertarian account of free will and both the Consequence and Rollback Arguments.)

There are two families of reply that one who endorses libertarian free will can take. The first is endorsed by agent causation theorists who count among their members philosophers such as Aristotle, Kant, Reid, Taylor, Chisolm, and O'Connor. They hold that in addition to event causation there is an additional kind of causation, agent or substance causation, where it is the agent herself and not any event within the agent that begins a given causal chain. As Aristotle says, "Thus, a staff moves a stone, and is moved by a hand, which is moved by a man" (Aristotle, *Physics* Book VII, 5, 256a, 6–8). Chisolm calls such causation "immanent" and distinguishes it from event or "transuent" causation.

The standard objection to agent causation is that it's left as a mystery what this kind of causation is like, it being of a radically different kind from the event causation we experience in the day-to-day world. When we think and speak carefully about causation, we say that the event of the hammer striking the glass caused the event of the glass breaking, and we attribute to the hammer itself no intrinsic causal powers. If event causation is the correct model for physical causation of the hammer-and-smashed-glass variety, then introducing agent causation as a different kind of causal power seems hard to fit with physicalism about the human mind.⁴² Furthermore, as van Inwagen notes, all

⁴² This is not to beg the question against non-physicalist explanations of the human mind, but any explanation of free will ought to be amenable to a physicalist

philosophers have to accept a mystery somewhere--all questions must come to an end--and if agent causal theorists are comfortable accepting this mystery, then that's their line in the sand.

For those who take the phenomenology of free will as a given and who reject agent causation as too mysterious, however, there is another family of theories which Robert Kane (borrowing from Gary Watson) calls "teleological intelligibility theories," which seems to offer a better solution (Kane 221).

Kane regards preeminent examples of free will to be any occasion where the agent has strong reasons (in terms of psychological force) to act in incompatible ways. In such a case, if the reasons aren't sufficient to move the agent, and no internal or external causal mechanism moves the agent, then she must move herself by granting to one set of reasons a motivating force.⁴³

An important question is whether the relationship between an agent's reasons and actions is a causal one. This can perhaps best be seen through an example. Owen desires to take a day off work as a "mental holiday" and he desires to be a good and honest employee. He knows that the only way to get the day off, however, is to lie to his boss about being ill. In this situation, Owen feels a great deal of psychological conflict between his desire to be morally good and his self-interested desires. It is not just that he has the phenomenology of conflicting desires that is important for Kane, but also that he have a conflict of reasons. Owen has good reason to be a good and honest employee. He

philosophy of mind.

⁴³ Kane also connects this view with Kant's notion that the person who resists the impulse to rape is a morally better exemplar than one who doesn't have the impulse to resist. The person who doesn't have the impulse exercises no free will in not raping but the person who does resist the impulse has not raped as a genuinely free choice.

also has good reason to take care of himself, and he has good reason to think that taking a mental holiday is a good way of achieving that. Owen thinks through both sides carefully: he considers the carefree day and its restorative powers; he considers the fear he feels at the thought of getting caught lying to his boss; he considers the uneasiness he will feel at questions from co-workers about his absence. He puts all these into the mix and still the reasons going each way have the same amount of psychological force. It comes time to either get ready for work or call in. Owen terminates the debate and decides to call in sick. That, according to Kane, is a prime example of free will since Owen decides which set of reasons to give the most weight to. Owen intentionally terminates the debate. It is that act which is the instance of free will. Notice that Owen's character traits, motivations, and preferences were fully engaged. It's not that Owen's free action has to arise independently of any traits he may have been born with or acquired along the way; it's rather that he has to choose in spite of those things. They merely set the battle ground, but Owen settles the day by putting his will behind one or the other set of reasons. He acts for reasons that are of his own choosing.

Free action, therefore, isn't about choosing which actions to take but about choosing which reasons to make dominant over others; free action follows from reasons freely chosen. We can raise the question of why Owen chose this alternative rather than the other. Kane's answer is that in such cases "the agent came to believe at the time that these were all things considered the weightier reasons" (Kane 245). We can, Kane suggests, answer the question of why Owen chose as he did with either the cognitive event of coming to believe or with the volitional event of making a decision (Kane 246). He takes these explanations to be informative descriptions of the same event. In the volitional description, what kind of causation is Kane proposing? He is not accepting agent causation. Rather, he says,

I would suggest that the best model of the transition from indeterminate efforts of the will to choices as described in this paper--that is, the best model for exercise of free will--is the collapse of the wave packet in quantum theory. This incredibly puzzling phenomenon involves a like transition from indeterminateness to determinateness and is tied up with deeper ontological issues about determinism and indeterminism (Kane 249).

In other words, he eschews accepting the mystery of agent causation, and instead embraces a different mystery, but, he thinks, one more acceptable if only because, in an Occam-like turn, it doesn't multiply mysteries beyond necessity. The reasoning is this: we've got good theoretical reason to accept quantum mechanics--thus there's no further risk of violating Occam's Razor--it's rather a matter of using that theoretical commitment to explain as much as possible.

By the time we reach adulthood, our characters are so solidified in their motivational structures that instances of free action are probably relatively rare.⁴⁴ If free will is rare, however, then isn't moral responsibility correspondingly rare? No. Suppose that Owen has such a good time on his mental holiday that he acquires the well ingrained habit of taking off whenever he sees fit (though he's careful about the spacing and the timing of these holidays so as to not arouse suspicion). After the behavior has become a habit, the reasons for taking the occasional day off so dominate the reasons not to that it requires no mental effort on Owen's part in deciding whether to take a particular day off. If this reasoning decides the case for him, he is not acting freely. Does that mean that he's not morally responsible for his behavior? No, he's responsible because in the months previous he gave weight to the reasons which now guide his behavior, and his moral responsibility stems from there. The same considerations apply if Owen resists the

⁴⁴ van Inwagen holds a similar view about the rarity of exercising free will. See (van Inwagen, "When Is the Will Free?").

urge to play hooky and becomes a model employee. The choices that settle his dispositions are freely chosen, but after that his actions do not result from *new* free choices but from an established motivational structure that he (at least in part) freely chose. These later actions, though not free, are fit for attributions of moral responsibility because they derive from free actions.

Since Owen acts freely on his first spree, he's morally responsible for his action. If his boss catches him and fires him, he will have deserved it since he chose to act that way in spite of the fact that he had strong reason not to. On further holidays which Owen takes not as freely chosen but as following from his settled motivational structure, he's still morally responsible since he chose to embrace the motivational structure.

Galen Strawson has an intuitively simple test we can use. Would sending the agent to heaven or hell for her behavior be just or fitting? Would the agent be getting her *just deserts*? Of course, one need not believe in heaven or hell to feel the intuitive pull of this question (Strawson, G., "Impossibility" 9). On Kane's view of how moral responsibility tracks past free choices, Owen is responsible in a heaven-or-hell sense, and that's the sense that drives the question of moral luck, to which we turn in the next chapter. First, however, let's look at some fundamental challenges to this understanding of freedom and responsibility. The final two sections of this chapter will briefly touch on how praise and blame function and on how to understand moral excuses.

4. Two Challenges to Libertarian Free Will

Two seminal papers in contemporary moral philosophy are "Freedom and Resentment" by P.F. Strawson (Strawson, P.), and "Alternative Possibilities and Moral Responsibility" by Harry Frankfurt (Frankfurt). Their accounts are deflationary against libertarian accounts of free will since each says that moral responsibility is meaningful independently of any particular conception of free will, and so independent of the

libertarian view. I don't hold either deflationary view because I think the reasons for each are not compelling.

4.1. P.F. Strawson's Reactive Attitudes

Strawson's approach is to deny the importance of one central debate in moral philosophy and in metaphysics. Compatibilists argue that free will--and hence moral responsibility--is compatible with causal determinism. Incompatibilists claim that free will and causal determinism are incompatible. They hold that free will requires that the future be open in the sense that at some points in an agent's history there be junctures where more than one future is physically possible. If causal determinism is true, then at every juncture, given the past and the laws of nature, there is but one future that is physically possible--or perhaps it is best to say that there are no junctures, though it may seem that there are to cognitively finite beings such as ourselves.

Strawson seeks to avoid the conflict between compatibilists and incompatibilists by showing that the practical aspects of our reactive attitudes do not presuppose that we are free in a metaphysically rich sense. That is, he believes that our practice of holding people responsible is consistent with both the libertarian's rich free will thesis and the ersatz freedom of the compatibilist. He thinks that no matter what we discover about the metaphysics of free will, our practices surrounding praise and blame would stand. We hold people morally responsible on his view when we interpret their behavior as exhibiting toward others a lack of appropriate concern, as when they treat people merely instrumentally. Holding others morally responsible is independent of the metaphysical underpinnings of their behavior, but is rather an expression of our social natures. In praising or blaming someone, I am treating her as a member of my moral community. In other cases, where I don't hold the agent responsible because he is not a full-fledged member on my moral community, I treat him objectively, Strawson says, as someone to be treated or managed, but not as a moral equal.

What's good about Strawson's paper is the richness with which he deals with the complexities of our social interactions. For example, he shows the central position of our emotional natures in ascriptions of praise and blame. Strawson focuses on how we interpret and react to the behavior of others, and his focus is not really about the metaphysical or moral status of their behavior. If I take Linda to be responsible because I interpret her actions in terms of whether she treats people with the appropriate level of concern, it's still an open question whether she is really free and really morally responsible.

Let us suppose that he is correct in this assessment of our practices as *social* practices. Even though our practices are neutral between libertarianism and compatibilism, agents are either morally responsible in a deeply metaphysical sense or they are not. If they are our practices are correct, and if not then our practices are flawed--even though, if Strawson is correct, we have no hope of abandoning them. That is, I grant that he has accurately judged our behaviors as social practices, but I deny that he has correctly judged them as moral practices.

Consider this analogy. Humans may be wired such that they cannot help but think that dogs are conscious beings. Their behavior, while not as complex as ours, shows them, we think, to be conscious beings. They play, pout, and panic. It would be very hard for us to take their behavior as only seeming to play, as only seeming to pout, as only seeming to panic. Thus, it's an open question whether we even could think of them as non-conscious beings. That is, it may be psychologically impossible for any individual from any culture to hear a dog yelp and not think of its behavior as pain behavior.⁴⁵ Even so, it is an interesting and open philosophical and biological question whether dogs are sentient and whether their yelps really are indications of pain. The fact that we will

⁴⁵ If such is the case, stories about vivisectionists who regarded dogs as purely mechanical things would be apocryphal.

continue to treat them as sentient no matter what we should find out about their biology does not mean that when we do so we are correct.

Whether dogs are sentient is an interesting biological and philosophical question, and the fact that humans are compelled to think that dogs are sentient is quite beside the point. One upshot of the naturalness of the belief is that philosophers and biologists who seek to explain why dogs are sentient will be looked at as if they are trying to explain what is so common-sensical as to need no further explanation (and they'll be seen by non-academics as having found a way of milking out more than their fair share of taxpayer money). A common attitude will be that since it's *obvious* that dogs are sentient and that their yelps are pain behavior, that fact hardly needs any explanation. Another upshot is that researchers who believe that dogs are not sentient are going to be in the minority; they are going to be looked at askance by many of their colleagues and as absolute crackpots by the majority of lay people (as even bigger crackpots than the more middle of the road philosopher). But for all that, it doesn't mean that they are wrong or misguided. And it matters whether dogs are sentient because that fact is what determines whether we are correct in our beliefs.

The upshot of my critique of Strawson is this. I'll grant to Strawson that our *practice* of attributing responsibility to others is so firmly ingrained as to resist revision in light of any philosophical or biological theory to the contrary. His position, however, is, morally speaking, not to the point. Let me explain.

I hold that there are a cluster of intrinsic properties the possession of which is sufficient for moral agency. Membership in the moral community, on my view, requires that the agent have (1) certain cognitive capacities, (2) a range of emotional capacities, (3) the capacity to be motivated by other-regarding considerations, (4) a structure of the will that includes first-order and at least second-order desires (that is to have certain desires and to be able to desire to be the sort of person who has certain desires), and (5) a robust capacity to choose between various courses of action.

In the context of the five capacities necessary for moral responsibility, what I'm granting is that only the first four are required for our social practices of praise and blame. We will treat objectively, in Strawson's sense, beings who lack one of those requirements. The fifth capacity--the capacity to make genuinely free choices--is what determines whether our social practices count as correct moral practices. Strawson treats that result as inconsequential, and I think he goes astray there.

4.2. Harry Frankfurt on Alternative Possibilities

Frankfurt takes a different approach. At bottom his claim is that my fifth requirement for moral responsibility--that the agent have free will in the sense that the future is open to her and that she controls which future comes to be--is not irrelevant but false. In particular, his claim is that the free will which underlies moral responsibility does not require alternative possibilities, his term for what I have called an open future. He tries to show this by attacking the principle of alternative possibilities:

(PAP) An agent is morally responsible only if he could have done otherwise
(Frankfurt 829).

He gives an example in which it seems clear that the agent is responsible for what she has done and yet from the scenario she could not have acted otherwise than she did.

Frankfurt examples are common in the literature; here is one of my own devising.

Nancy sits on the city council and faces a vote on a zoning ordinance which would put a mental hospital across from an elementary school. Black is unsure how Nancy will vote, but it matters to him that she vote for it. Indeed, as the vote approaches, even Nancy isn't sure how she will vote. Though the balance of her reasons incline her to vote for it, she still has reasons that incline her against it. Black has implanted an instrument into her brain that can both read her intentions and guide her behavior. He has programmed it such that if Nancy decides to vote against the ordinance it will force her to vote for it, but

if she decides to vote for the ordinance, it will have no effect on her decision or her behavior. As it turns out, Nancy decides without any influence from the instrument to vote for the ordinance and does so.

Frankfurt's point is that it seems intuitively clear that Nancy is responsible for voting as she did, and yet, while she may have thought that she had an open future in this regard, she did not. Hence, concludes Frankfurt, the requirements on moral responsibility that the incompatibilist supposes are not correct. For Nancy to be morally responsible the well spring of her behavior must come from her--that she not be coerced, insane, physically forced, etc. There is no requirement, Frankfurt holds, that Nancy's future be a garden of forking paths.

The literature on Frankfurt's treatment of moral responsibility is enormous, and I cannot do justice to it here. A couple of comments will have to suffice to indicate how a libertarian might reply to Frankfurt's purported counterexamples.

First, let us follow Michael Zimmerman's distinction between two senses of "control" that are pertinent, regulative control and guidance control. He says that "one has regulative control over something just in case one can (in the relevant personal sense of 'can') bring it about and can also refrain from bringing it about. One has guidance control over something just in case one can bring it about, even if one cannot also refrain from bringing it about" (Zimmerman, "Map" 593).

While Zimmerman holds that moral obligation requires regulative control, he thinks that Frankfurt cases show that moral responsibility only requires guidance, but not regulative, control (Zimmerman, "Map" 602-03). Thus, he takes Frankfurt as having shown that moral responsibility doesn't require a libertarian sense of free will since moral responsibility doesn't require regulative control.

We can grant that Frankfurt has shown that PAP is false (that his purported counterexamples are the real deal), but then argue that the kind of moral responsibility

that the libertarian endorses isn't based on PAP. Since PAP at first blush seems required for moral responsibility, Frankfurt's opponent then needs to find one or more PAP-like principles which can underpin our intuitions about moral responsibility and then to argue that the additional set of principles is immune to Frankfurt-style examples.⁴⁶ Peter van Inwagen, for one, takes this approach, and offers three PAP-like principles:

PPA A person is morally responsible for failing to perform a given act only if he could have performed that act (van Inwagen, "Ability and Responsibility" 156).

PPP1 A person is morally responsible for a certain event (particular) only if he could have prevented it (van Inwagen, "Ability and Responsibility" 157).

PPP2 A person is morally responsible for a certain state of affairs only if (that state of affairs obtains and) he could have prevented it from obtaining (van Inwagen, "Ability and Responsibility" 161).

van Inwagen then argues that these three principles together support our intuitions about moral responsibility and each is immune to Frankfurt-style counterexample. We hold people responsible (1) for their moral failings (things they ought to have done but didn't) and (2) for the consequences of what they have done. PPA is intended to capture (1) and PPP1 and PPP2 are intended to cover (2).⁴⁷

As Susan Wolf notes, there is an asymmetry between praise and blame such that blameworthy acts and omissions require that the agent could have done otherwise, but praiseworthy acts don't require that the agent could have done otherwise (Wolf 156). In Zimmerman's terms, blameworthiness require regulative control and praiseworthiness

⁴⁶ A Frankfurt example is a counter example to PAP. A Frankfurt-style example is an example that purports to be a counter example to some minor variant of PAP.

⁴⁷ We have two principles so that we can remain neutral to count events or states of affairs as the consequences of an action. Bill shoots Tom and Tom dies. Is what Bill is morally responsible for the Tom's death (an event) or Tom being dead (a state of affairs)? PPP1 covers Tom's death; PPP2 covers Tom's being dead.

require merely guidance control. Be that as it may, van Inwagen's three principles require regulative control and hence Frankfurt's conclusion is to be rejected.

There are other at least initially plausible replies to Frankfurt's challenge. One reply is to doubt the cogency of Frankfurt's examples by asking *when* Black's device would activate. If it would activate after her decision to vote against Black's wishes, his device would prevent her from voting that way, but only by thwarting her free will. In that instance, her freely choosing how to vote demonstrates that she has alternative possibilities. On the other hand, if he springs his trap too soon, not allowing her to decide to vote against the ordinance, she has no alternative possibilities (as Frankfurt imagines), but neither does she plausibly have free will. There are many replies to this objection, and I'll rest content to leave off here and throw my weight behind van Inwagen's answer.

Let me observe, however, that both Frankfurt and I agree that agents are morally responsible when their actions flow from their intentions. That agreement is enough to carry my project forward, at least provisionally. A full treatment, however, would involve taking us even further afield, so cannot be justified. If my final results are only provisional, that's not such a bad thing since all philosophical claims are provisional.

5. The Functions of Praise and Blame

There is an ambiguity in speaking about moral responsibility, however, owing to what we take to be the proper functions of praise and blame. One option is a *consequentialist function of moral responsibility* in which we properly mete out praise and blame not on the merit of past actions but as a way of influencing future actions. The legal function of punishing trespassing, for example, may be to discourage future instances of trespassing and not primarily as punishment for the offender (Ripstein 121). This function of moral responsibility is essentially forward looking. The problem with it is that it can become divorced from the notion of desert: neither an actual horse theft nor

an actual horse thief is required for the hanging to serve its social purpose. The consequentialist function of moral responsibility, however, is not what I have in mind in this discussion. Responsibility can serve those reformatory and deterrent functions even if there is no free will as a libertarian might conceive it.⁴⁸

The second account of culpability we might call the *merit-based sense of moral responsibility*. That sense of responsibility is essentially backward looking. Its function is to bring focus to the agent, and to put the weight of what she has done back on her shoulders. The function of praise and blame that is important in this discussion is to *correctly* ascribe desert. If there is no free will, the merit-based sense of moral responsibility is a complete sham. It is the sense of moral responsibility that gives the question of moral luck its interesting bite. Of course, a particular instance of praise or blame may be *both* forward *and* backward looking. If we praise Kim for teaching the illiterate to read, we can give her the good attention she deserves and encourage her and others to be generous and helpful in the future.

6. Accountability and Excuses

Important to a discussion of accountability is the notion of an excuse. We can acknowledge two broad kinds of legitimate moral excuse: (a) appeals to ignorance--the agent didn't have the factual knowledge not to make the moral mistake (e.g., Paul owed John money but paid what was promised to John's evil twin, Don, because Paul didn't recognize that he wasn't fulfilling his obligation to pay John); and (b) appeals to inability--the agent could not have avoided the mistake (for example, Raymond promised

⁴⁸ Remember the tongue-in-cheek buck-passing line: I didn't say it was your fault; I said I was going to blame you. The core to that one-liner is the forward-looking sense of accountability.

his daughter, Amy, that he would take her to the ball game but his plane was delayed and he could not get back in time).

If these appeals refer to genuine inability (where we take ignorance as perhaps a special case of inability), we can think of these as *exculpating excuses*, that is, as excuses that block the claim that the agent did anything wrong, thus making the question of accountability a moot point.⁴⁹ The underlying appeal in exculpating excuses is the moral dictum that *ought* implies *can*, so if Raymond could not have arrived in time to take his daughter to the ball game then it's not the case that he ought to have. What distinguishes exculpating excuses from other justifications would be, I think, the appeal to *ought* implies *can* as the reason given for what the agent did or failed to do.⁵⁰

Amy might reject her father's excuse if he *could have* arrived in time for the game, for example, by buying a seat on a different airline. Raymond's reply might be that while he *could have* in the sense of it being physically possible, given that the cost was prohibitive, he *could not have* as a fiscally responsible person. Similarly, John may complain that Paul could have checked Don's driver's license and found the mistaken identity before he gave his evil twin the money, and Paul might counter that expecting those sorts of precautions are unreasonable.

It seems that the offended party in each example is appealing to one standard for judging whether the agent could have fulfilled his obligation and the agent is appealing to a different standard in his own defense. Legal contexts may help us out of this mire. The

⁴⁹ Michael Zimmerman denies that this is moot and has sketched a theory of *accuses*, that is, cases in which the agent has done nothing morally wrong but still deserves blame (Zimmerman, "Accuses"). While an interesting philosophical exploration of the logical space of justification, excuse, and responsibility, Zimmerman's argument is not one we can take the leisure to follow.

⁵⁰ Since an exculpating excuse implies what the agent did was morally permissible, exculpating excuses are a species of justifications, and calling them excuses at all is a misnomer but one well entrenched in common usage.

self-defense justification in a killing doesn't require that the defendant prove that his life was actually in danger but merely that a reasonable person would have believed it to be so. So, if Susan shoots and kills an intruder who has a pistol in his hand, her defense is complete even if the intruder's pistol is a reasonably genuine-looking toy.

Thus, if an agent claims that she couldn't have done X and so is not obligated to have done X, perhaps the standard for *what the agent could have done* ought not be what was physically possible, but what a reasonable agent would do, given the circumstances. So it's reasonable for Paul to pay Don since it was the best choice available given Paul's evidence. Of course, Paul still owes John the money (which he may or may not recover from Don). Also, Raymond may or may not be morally culpable for missing the opportunity to take his daughter to the game. Given what Raymond knows about his finances, bills that are coming due, etc., he can decide whether he can get a different flight while being responsible to his other moral obligations. It's those sorts of details that will decide whether Raymond has acted reasonably in not taking a different flight.

This is to put the matter in terms of what we subjectively judge to be reasonable, based on our current evidence. We can put the matter in the more objective terms of what sanctions would deter the behavior for an agent with a reasonable regard for self-interest. As Ripstein puts it, "objective tests for excuses suppose that in circumstances in which a person of ordinary fortitude could not have been deterred by the prospect of punishment, the act is properly attributed to the circumstances rather than to the agent" (Ripstein 168). In such a case, he says following Justice Blackstone, we treat the agent's decision as if it were "morally involuntary, so we treat it as though the agent wasn't involved at all" (Ripstein 166). If, for example, Wendy robs a bank because her child is being held hostage, then by this standard we ought not to hold her accountable because it would take heroic and not ordinary fortitude to resist the demands of her child's abductor. Similarly, if Amy's disappointment at missing the game wouldn't move a person of ordinary

fortitude to purchase a second ticket, then Raymond ought not to be blamed for breaking his promise if the only way to keep it would be to buy the second ticket.

Thus we can have a theory of excuses that ranges from exculpating excuses (ones which absolutely absolve the agent from moral stain on his record) through mitigating excuses (ones which leave some but not full moral stain on his record) through lame excuses (ones which leave the full moral stain on the agent's record).

Chapter 4

The Problem of Moral Luck

1. Three Varieties of Moral Luck

Let us consider three pairs of siblings. Mike and Jim are raised by parents who instill in them a reverence for and an absolute obedience to authority. Their upbringing is such that they would feel at home in the tightly structured life the army offers and they both enlist. Mike is stationed in a remote area and sees no combat. Jim, however, finds himself in a war zone. His captain orders him to kill the family of a captured militant in order to break the man's resolve and get him to reveal sensitive information. Jim shoots the man's wife and two of his five young children before the combatant breaks and gives up the information. Jim is later convicted at court martial and sentenced to life in prison for the three murders. Had Mike been given the same orders in the same circumstances he would have exhibited the same brutality that Jim possessed. By many people's lights⁵¹, Jim's moral record is worse than Mike's. The primary moral difference between them is that, unlike Mike, Jim had the bad luck of being put in circumstances in which his bad character produced disastrous results. Though he is no worse a person, his record is worse than his brother's by bad luck.

Our second siblings, Sarah and Jane, live on opposite sides of town and meet downtown Wednesday afternoons for lunch, which usually includes the consumption by each of several high-octane cocktails. Let me stipulate that, given the routes they take home, each poses the same risk to public safety by driving home while at least mildly

⁵¹ Interesting work is coming from experimental philosophy such that philosophers' data can be harvested from what people really think rather than what philosophers think that people think. (See, for example, work by Joshua Knobe (Knobe).)

intoxicated. One such Wednesday, Sarah arrives home safely, but Jane hits and kills a small child who has flitted from between cars and into her path. Had she not been drinking, her reactions would have been crisper, and the child would have been unharmed. Clearly, both Sarah and Jane have acted badly in driving drunk, but Jane's alleged moral record is worse than her sister's because her bad luck resulted in a child's death.

Finally, Carol and Don were separated from each other at a very young age during their parents' bitter divorce. Carol goes with her father who is a cruel man, and he raises her to be cruel. Don, on the other hand, goes with his mother, who is patient and loving, and as a consequence of his mother's values, Don grows into a patient and loving adult. Their moral records are reasonably predictable given their deeply ingrained moral traits--Carol's moral record reflects her heinous dispositions and Don's moral record, while not perfect, shows him to be a decent and caring human being. Let me stipulate that had the parent-child pairings gone the other way, Carol would have been patient and loving and Don cruel. What matters in this case is that neither Carol nor Don had any control over which parent they lived with and thus no real control over their constitutions--Don is lucky to get his virtuous character, and Carol unlucky with her vicious character. What's more, Don's moral record is better than his sister's, since he had the good fortune to be raised by the parent who provided better moral guidance.

These siblings illustrate three ways in which luck can purportedly make a moral difference in our lives. In the case of Mike and Jim, their moral differences result from the circumstances in which they find themselves; they are differentiated by *circumstantial moral luck*. Sarah and Jane differ morally because of the downstream results of their

behavior--this is a case of *resultant moral luck*. Finally, Carol and Don fall prey to *constitutive moral luck*, or luck in how their characters are constituted.⁵²

This chapter focuses on the difference luck has made in each sibling's life. Jim, for example, has been asked to kill, and has killed, and Mike has neither been asked to nor killed another human being. In addition to the differences that luck has made to the details of their lives, many people also judge them to differ from a moral standpoint. The question is whether their moral records really differ--whether people are *correct* in judging them as having different moral records. That is the *problem of moral luck*.⁵³

2. Kant and Moral Evaluation

Some ethical theories do not judge agents by their causal contributions to the world. Kantians, for example, measure the moral rightness or wrongness of an action by the agent's intention. They don't give weight to the effects of the agent's actions, nor to the circumstances in which she finds herself, nor to her natural gifts or developed character traits. That is, they don't give moral weight to any of the kinds of luck outlined in the previous section. The pairs of siblings differ by the influence of luck, but it is not a moral difference.

Thus, as is widely recognized, Kantians hold that moral luck is oxymoronic.

Consider Kant's famous passage:

A good will is not good because of what it effects or accomplishes--because of its fitness to attain some proposed end: it is good through its willing alone--that is,

⁵² These categories come from Nagel (Nagel 60). A fourth category is sometimes also discussed, *causal moral luck*, which is a combination of circumstantial moral luck and constitutive moral luck. As such, it is redundant and will get no further explicit attention.

⁵³ I will present the problem of moral luck more formally in Section 3.

good in itself. . . . Even if, by some special disfavor of destiny⁵⁴ or by the niggardly endowment of a stepmotherly nature, this will is entirely lacking in power to carry out its intentions; if by its utmost effort it still accomplishes nothing, and only the good will is left (not, admittedly, as a mere wish, but as the straining of every means so far as they are in our control); even then it would still shine like a jewel for its own sake as something which has its full value in itself. Its usefulness or fruitlessness can neither add to, nor subtract from, this value⁵⁵ (Kant Ak. 394).

The problem of moral luck is often framed against this facet of Kant's moral theory. Since Kant holds that the only thing that is good without qualification is the good will, the moral status of an agent depends only upon his will. Thus, it depends only on the maxims of action which drive his behavior, the maxims being the fundamental expressions of his will. No disposition of character determines the agent's moral worth. It is irrelevant how the world frames the circumstances in which he acts. It does not matter if he is unable to achieve the results he aims toward, so long as he strains toward realizing his intention with every morally legitimate means under his control. Excluded from figuring into the moral evaluation of the agent, therefore, are the causal predecessors of and the causal consequences of his action.

That is not to say that Kantians cannot recognize that some consequences are better than others, since that would clearly fly in the face of the facts. But on the standard

⁵⁴ McGregor's translation uses the term "fortune" here, which perhaps brings out the idea that Kant is opposing moral value to any value that can be gained through luck. James W. Ellington uses "some especially unfortunate fate," again connecting this passage at least intuitively with bad luck.

⁵⁵ Notice that Kant covers Nagel's categories of luck: "some special disfavor of destiny" is circumstantial and constitutive luck; "endowment of a stepmotherly nature" is constitutive luck; and "by its utmost effort it still accomplishes nothing" signals resultant luck.

Kantian line, while some consequences are better than others, that evaluation ought not to affect our *moral evaluation* of the agent's actions. This is because while the consequences can be evaluated in terms of usefulness or fruitlessness or pleasure or pain or . . . , none of these evaluations are *moral evaluations*. Moral value and hedonic or instrumental value are from the Kantian perspective incommensurable and incomparable. It does not make sense to add moral values and non-moral values to get an overall moral evaluation. Thus, no action and no agent is made *morally* better by producing good consequences or *morally* worse by producing bad consequences, though in a particular case an agent may be morally better for aiming at good consequences than if she had aimed at bad ones.⁵⁶ No wrong action--wrong in terms of the will involved--can have that deficit made up by engendering desirable consequences; and no right action can be tarnished by having undesirable results. Since Kantianism is not a moral theory that takes an action's upstream precursors or its causal consequences into account, if luck is at bottom a causal notion, it's no surprise that luck makes no moral difference within a Kantian framework.

The thought that moral evaluation is immune to the world's accidents holds an intuitive place in the thinking of many people. The thought seems to be something like the following. I may be lucky in love or lucky in the car I buy or lucky in the talents I have or in being able to enjoy the talents of those who are close to me (and I may be

⁵⁶ For Kant, an action's moral worth is measured by whether the intention that drives the action can be universalized under the categorical imperative. Sometimes, at least, those intentions can be to bring about a particular good consequence, such as preventing a murder. An intention-based consequentialism, on the other hand, may dictate that the agent ought to strive to maximize the goodness of the world (however that is measured) but doesn't require that a morally good agent be successful in realizing that intention. It is clear, however, that Kant himself is not an intention-based consequentialist (or indeed any sort of consequentialist, traditionally understood). My point is rather that this Kantian idea--focusing on the intention rather than the results--need not be barred from use by consequentialists.

unlucky in all those things). While good luck makes life enjoyable, it doesn't make one person better than another. One sphere in which I am individually responsible for my own evaluative status, however, is the moral sphere. I am responsible for just what I can control, and luck plays no part there.

Given this Kantian intuition, the idea that morality *could* collide with the contingencies of the world is *prima facie* problematic. Modern discussions of moral luck, beginning with Bernard Williams and Thomas Nagel, emphasize that point (Williams; Nagel). For Williams, two factors are ingrained in the standard conception of ethics which help to explain the seemingly oxymoronic status of moral luck. The first factor is that moral evaluation does not follow causal lines either in terms of upstream influences or downstream effects, as we have seen. The second factor is that, of all the standards of evaluation that we bring to a situation (moral, aesthetic, rational, etc.), moral evaluation reigns supreme.⁵⁷

Morality, therefore, is immune to luck because of how those factors work together. Morality is *immune from within*, since moral evaluation is limited to what the agent can control--what is up to *her*--, and it is *immune from without*, since no other sphere of evaluation in which luck has an influence can trump moral evaluation and thus allow luck to influence moral value. To put the issue another way, the first factor insulates morality from the *direct* influence of luck by ignoring causal factors in making moral evaluations. The second factor insulates morality from the *indirect* influence of luck by never subjugating morality to any sphere in which luck has a direct influence.

Suppose, for example, that morality could be legitimately subjugated on occasion to self-interest. It might be that the agent would have to suffer a great personal loss to persevere in a minor moral matter. If self-interest could legitimately trump morality in

⁵⁷ Nussbaum makes a similar point about Kant's supremacy of the moral (Nussbaum 4–5).

such cases, and it's a matter of luck that the agent finds himself in such a situation, then it's a matter of luck that his moral record is worse than it might have been. (This assumes, which seems plausible, that if self-interest could trump morality, it would not therefore cancel the moral obligation.) If self interest in this case were overriding, the agent could justifiably claim that his otherwise morally bad behavior deserved excuse because some more important value had forced his hand. The mere fact that asking and reasonably expecting to be morally excused shows us, I think, that his moral record is *prima facie* worse in this scenario than it might have been. On the other hand, suppose morality is the supreme sphere of value. If the agent chooses to take the personally expedient road at the expense of doing the right thing, he cannot legitimately claim that his hand was forced by the dictates of an overriding sphere of self-interest. His moral record would differ, as before. In this scenario, however, the agent's own selfishness is to blame and not the combination of an overriding self-interest and bad luck.

3. The Problem of Moral Luck

Let us start with a preliminary explanation of the phrase "the problem of moral luck." It is often explicated in terms of a moral fact, an empirical fact about agents, and a definition of the term "moral luck." From these three parts, we can present an argument that captures *the problem of moral luck*.

Roughly stated, the moral principle involved in this discussion holds that the limits of moral responsibility are defined by what is within the agent's control. Here is the Control Principle, as formulated by Dana Nelkin:

(CP) We are morally assessable only to the extent that what we are assessed for depends on factors under our control.⁵⁸

⁵⁸ In Chapter 2, I denied the cogency of lack-of-control accounts of luck. In the next section, I'll discuss how the causal-contrastive account of luck relates to CP.

CP is intended to be a statement of the Kantian principle discussed in the last section.

Remember, however, that the Kantian intuition is consistent both with Kantianism *per se* and with certain intention-based forms of consequentialism.

Even if we reject Kant's position that we are *only* morally responsible for what we directly control--the intentions that drive our actions--we still have good reason to accept CP. After all, the central idea of moral responsibility is that we can reasonably be expected to own only our contributions to the world. The less credible it seems that a particular fact about the world was under my control, the less credible is the claim that I'm morally responsible for it.

The empirical fact concerns the ubiquity of luck in human endeavors. We act as agents in the causal stream and, as such, have less than full control over our constitutions, over what causal lines we find ourselves on, and over the downstream effects of our actions.

Additionally, an explication of the problem of moral luck needs a definition of the phenomenon under question, that of moral luck. Here is Nelkin's formulation of what it is for an agent to be subject to moral luck:

(ML) moral luck occurs when an agent can be *correctly* treated as an object of moral judgment, despite the fact that a significant aspect of what he is assessed for depends on factors beyond his control (Nelkin).

Now consider the valid argument whose premises are *prima facie* plausible and whose conclusion conflicts with the intuitively plausible CP.

The Basic Problem of Moral Luck:

1. Agents are sometimes morally responsible for their actions, intentions, and character traits.
2. No action, intention, or character trait is ever within any agent's control.
3. Thus, agents are sometimes morally responsible for actions, intentions and character traits that are *not* under their control.

Premise (1) expresses the idea that we are at least sometimes morally responsible. This is a substantive philosophical thesis, and while at least some philosophers will deny it,⁵⁹ I take (1) as expressing a truth. While I will later link premise (2) to the kinds of luck that Nagel describes in his work on moral luck, it should be clear from Chapter 2 that no fact about us is entirely clear from the influence of luck. The conclusion (3), however, conflicts with CP, which we have good reason to endorse.

The *problem of moral luck* is that deliberation leaves us in the position of endorsing each member of an inconsistent triad of propositions: (1), (2), and CP. Thus, we are compelled to choose among three individually unpalatable options: deny CP, deny that luck is ubiquitous, or embrace moral skepticism. CP, however, gets to the heart of many people's intuitions about moral responsibility, and if moral responsibility loses its moorings, moral theory would be a rudderless ship; our lives are shot through with luck; and moral skepticism is generally taken to be repugnant, leaving many philosophers thinking that moral skepticism ought to be the refuge of last resort.⁶⁰ None of these options are pretty, but that's why the problem of moral luck is a problem.

The goal of this chapter is to come to grips with the problem of moral luck. Toward this end, let me first discuss the relation between CP and the causal-contrastive account of luck. Then we will benefit from examining the contributions of the authors who jump started modern discussions of moral luck: Bernard Williams and Thomas Nagel. Finally, in the next chapter we will consider other proposed solutions to the problem.

⁵⁹ Galen Strawson is skeptical about free will and hence about moral responsibility (Strawson, G., "Impossibility"; Strawson, G., Freedom and Belief).

⁶⁰ Some, like van Inwagen, have argued that given that we deliberate we cannot even really endorse moral skepticism. We may verbally endorse moral skepticism, of course, but van Inwagen's point--well taken--is that our actions, particularly our moral reactions, reveal that our true beliefs are otherwise.

4. CP and the Causal-Contrastive Account of Luck

In Chapter 2, I argued that the lack-of-control account of luck is wrong. What CP expresses is the Kantian idea that moral responsibility requires control. That explains, I think, why the lack-of-control account of luck seems so natural for those thinking about moral luck. The world's contingencies and moral responsibility clash because the former contains what the latter excludes: a lack of control on the part of the agent. I have no problem with this way of talking so long as the lack of control is understood relativistically.

Let me explain. Recall the formulaic expression of the causal-contrastive account of luck:

$$\lambda_{\langle IE, RE, AE \rangle} = \delta(RV - AV) * [1 - \text{pr}(IE \rightarrow RE)]$$

The “lack of control” maps onto the part of the formula that expresses the susceptibility of the causal chain to deflection in going from the initial event to the resulting event:

$$[1 - \text{pr}(IE \rightarrow RE)]$$

The proponent and the opponent of moral luck need not differ about that value. Their disagreement is over whether the resultant value differs from the averted value. The proponent says that an account of luck can accommodate a moral difference and the opponent says that it cannot. Let us consider those claims in some detail.

First, let's consider the case of Jane who kills a child while driving drunk. If we consider the case from two seconds before she strikes the child with her car, there's very little that could deflect what eventually happens, so $[1 - \text{pr}(IE \rightarrow RE)]$ will be a small value. Given that the difference between AV and RV is high, the luck of the situation will be correspondingly low, but not zero, owing to the low value of $[1 - \text{pr}(IE \rightarrow RE)]$. Even a proponent of moral luck, however, might deny that Jane has moral luck in this case because she has no control over whether RV or AV is realized. Moral responsibility

requires some positive level of control on the part of the agent. What's at dispute between the proponent and the opponent of moral luck is whether moral value can be assigned to the agent if she has less than full control.

Now consider the situation from ten minutes before Jane kills the child. Here [1 - $\text{pr}(\text{IE} \rightarrow \text{RE})$] a large value because there are many ways that the causal path might have been deflected from the event of the child being killed, some of which are under Jane's control: she might have had one more cigarette or one more drink, she might have taken a different route home, etc. Or, quite importantly, she may have decided to consciously lowered her chances of killing a child by not driving drunk.

There are two relevant questions in this scenario. (1) Can downstream result affect an agent's moral value? If the answer is Yes, then (2) does Jane have in this case the right kind and amount of control for her moral value to be affected? The latter question connects relevantly with CP, and it gets a fair hearing two chapters hence.

5. Williams on Moral Luck

I want to highlight two aspects of Williams' position as he expresses it in his paper "Moral Luck": his fictionalized example of the painter Gauguin and his conception of agent regret. He concludes his paper by taking an anti-theoretical stance toward moral theory, but not against what he calls "the ethical." Since I deny the soundness of his reasoning as regards the first two points, I see no need to extend the discussion to his ultimate, anti-theoretical conclusion.

Gauguin was a competent but uninspired painter living in Paris with his wife and small children. He decided that he could reach new levels with his art if he were to leave his family, move to Tahiti, and paint the natives and the native sights. The question Williams wants us to consider is whether Gauguin was rational to embark on that adventure. According to Williams, if Gauguin gets to Tahiti and succeeds as a painter,

then we can know that his choice was rational (though his deserted family still may not endorse his decision). To succeed as a painter, Gauguin need not be commercially successful or even revered by other accomplished painters. He has to at minimum live up to his potential, either judged objectively or subjectively. That is, he has to become the painter that he can be or at least the painter he thinks he can be. There are two ways for Gauguin to fail, however, each of which ought to be examined.

First, suppose that Gauguin makes the move but has miscalculated his artistic ability or the inspiration that the South Sea Islands will afford. In such a case, Williams would say that Gauguin has failed as a painter. Gauguin has misunderstood something about his talent or his propensity to be inspired by a given *milieu*. That is, his failure reveals a mistake in how Gauguin thinks either about his raw artistic talent or about his ability to leverage immersion in the culture into works of artistic excellence. The latter can happen in several ways. He might be too much of an artistic dullard to be inspired no matter what the setting (even if he has a great deal of technical talent). On the other hand, he may lack the discipline to apply himself in such an exciting and interesting setting. In any of these situations, he has as genuine an opportunity to thrive as a painter as he could hope for, yet he fails to excel. In these sorts of scenarios, Williams would hold that *Gauguin himself* has failed and that the reasons for his failing are *internal* to him.

Second, suppose that he makes the move, and on the trip he is injured, and his painting career ends before it takes root in Tahiti's soil. Gauguin may or may not have had what it takes to succeed: the raw talent and the ability to turn a personally inspiring subject matter into a painting that meets or exceeds a certain objective or subjective standard. It doesn't matter. Gauguin neither succeeds nor fails *as a painter* in this case. Since his project was to blossom into a great painter, however, *the project* has failed. Gauguin, however, cannot have failed as a painter because he was never tested as a

painter.⁶¹ He can only fail at a project if he fails the test that the project presents. The project can fail, however, either by his failing the test or by his being prevented from taking the test. Williams regards the former kind of failure as *internal* to Gauguin and the latter as *external*. So while the project has failed in this instance, *Gauguin himself* has not failed, and it is only his failure due to reasons internal to the project can show that his plan was irrational.

Gauguin's failure in the case of injury is external to the rationality of his undertaking, so while that outcome would not justify his decision, it does not show that his decision was a mistake. External misfortune does not *unjustify* the decision. Again, Williams' view is that whether Gauguin succeeds (i.e., whether his decision is rational) or fails (i.e., he has made an irrational decision) is given by how well he does in the test he has set for himself. If injury prevents the test, then whether his endeavor is rational or irrational remains an unanswerable question. That is, Gauguin's endeavor can be shown to be mistaken or irrational only if he fails *as a painter*.

Williams concludes that whether Gauguin is justified in going down this adventurous path is an *essentially retrospective* judgment. That is, Gauguin cannot know whether his move was a mistake until the dust settles and he either fails or passes the test he has given himself. (The case of injury in effect blocks his ability to make this determination and invalidates the internal evaluation.) Thus, Williams says,

Gauguin could not do something which is thought to be essential to rationality, and to the notion of justification itself, which is that one should be in a position to

⁶¹ True, he was a painter in Paris and failed there. The historical painter Gauguin did not fail in Tahiti, so we know that failure in one set of circumstances does not necessitate failure in another. Thus, returning to the fictionalized Gauguin, let us take the expression "failed as a painter" to mean failed as a painter in Tahiti.

apply the consideration at the time of the choice and in advance of knowing whether one was right (in the sense of coming out right) (Williams 38–39).

Thus, whether Gauguin is justified is, from a perspective prior to the adventure, a matter of luck. In the next section, I evaluate Williams' stance on prospective versus retrospective judgment.

6. Critique of Williams on Retrospective Justification

I don't find Williams' position convincing regarding retrospective judgment. I do not deny that people jump into projects and hope for the best. Neither do I deny that some people claim that a fortunate outcome retrospectively vindicates their behavior. They are mistaken, however, for it does not. A good or desirable outcome can never retrospectively justify an impulsive decision.

My plan of action is as follows. I want to first consider a case whose analysis is formally similar to Gauguin's case of trying to maximize expected value in a situation of uncertainty. The reply to my case will be that I've stacked the deck too much in my favor. That is, I've picked a case in which too much is known before the decision point to be relevantly similar to Gauguin's experience. Then I will expand my case to one that is closer to Williams' Gauguin example and show that prospective judgment in that case is available to my agent. The upshot is that if we can reasonably hold my agent to the standards of prospective judgment, we ought to be able to make the same judgments about the rationality of Gauguin's decisions.

Perhaps the most ubiquitous examples that highlight the distinction between prospective and retrospective rationality (as judged from late-night cable television) are betting decisions in the poker game Texas Hold'em. The set-up is simple: each player receives his own two cards (the value of which is private to him), and five communal cards are dealt face up. The play is punctuated by several betting rounds. During betting

rounds, each active player can raise the current bet, call it, or fold the hand and forfeit his or her interest in the pot. The winner of each hand is the currently active player who displays the best hand as judged by the rules of the game.

We can distinguish the *expected value* of an option from the *actual value* of an option. The expected value of an option is a function of how many chips are being risked and the mathematical odds of the risk paying off by having the best hand active in the round of play.⁶² In contrast, the actual value is the amount of chips actually won or lost by that play.

Poker can, of course, be played for a wide variety of reasons such as to spend time with friends or as an innocuous hobby. Competitive poker, however,--the kind one sees on television--is an exercise in practical rationality. Since players have three options open to them (raise, call, or fold), the goal is to pick the option that has the highest expected value.⁶³

The difficulties in maximizing expected value are that (a) a player has incomplete information about the relative strength of his hand as compared to the other hands still in play and (b) players use behavioral cues to deceive each other about the strengths of their respective hands. Thus, it is very often impossible for a player to be certain prospectively if his decision is the correct one. Still, we can have these general guidelines: (1) if a choice has a high expected value (that is, it is likely to maximize

⁶² I intentionally ignore under- or over-playing a hand as part of a larger strategy meant to increase one's overall take in a particular game by confusing one's opponents. Thus, the expected value of an option may stretch across hands.

⁶³ This is a slight fiction in two ways. First, for the sake of simplicity, I've ignored the fact that one can raise by varying amounts, thus making for more than just three options. If it helps, consider a case where the only raise option will put all of the player's chips in play, so there are really only three options there. Second, a player may use a satisficing strategy, rather than trying to maximize gains. Satisficing may lead to less drastic peaks and valleys, and thus may be suited to those who are more risk averse.

value), it's rational to follow that path, (2) if a decision has a low expected value (it's likely to not maximize value), it's rational to avoid that path, (3) if two or more options are approximately equal in their likelihood of maximizing value, then it's as rational to take one option as the other, and (4) if the waters are still too murky to tell, then it's rational to shore up one's decision procedure beyond just employing the theory of expected value.

In cases (3) and (4) in which no course of action is recommended by the theory of expected value, our player, Jerry, is still required to take some action. The theory of expected value has, in such cases, delivered all the benefit it can. Jerry has two options. First, he can flip a coin, so to speak, over which action to take. If the stakes are low for the hand relative to Jerry's comfort level, he may not go beyond a coin flip because he judges protracted effort to be unwarranted by the potential payout. Second, if he sees this as a significant decision, he can employ a secondary decision procedure, one that says "in the case of tied or uncertain expected values, act such as to minimize potential losses." Or he can follow a decision maxim that enjoins him to "... act such as to maximize potential gains." Which secondary maxim he adopts may be a matter of temperament.

The point is that even if there are ties or uncertainties in the expected value, that doesn't mean that Jerry has exhausted all his rational resources. He may not be able to convince one who is not risk averse that the minimizing losses maxim is the one to use, for example, but he *will* be able to say that given that his conception of rationality includes minimizing losses he has made a rational decision in folding the hand.

If the theory of expected value and the complementary maxim yield ties or the situation remains unclear, then Jerry can either introduce another maxim by which to further supplement his decision procedure or say that all his decision theoretic resources have been exhausted. When all is said and done, either he'll have a single recommended course of action or multiple equally well recommended courses of action (justifying a

coin toss) or no recommended course of action. In that final case, since he is forced to act, it would be unfair to criticize his decision, given his due diligence. Neither, however, is praise justified if he wins, given that his guess was merely fortunate.⁶⁴

Notice that on this conception whether a particular choice is indicated is knowable, if at all, before any action is taken.⁶⁵ Thus, there seems little at this point to recommend Williams' analysis.

Sometimes a rational choice (one that is likely to maximize value) results in a loss of value. On this conception of practical rationality, if the play had the highest expected value, the player acted rationally, even if he loses the hand. The fact that his decision has had a negative effect on him (in terms of losing chips) does *absolutely nothing* toward showing that his decision was irrational. The question of whether the player has acted rationally is merely a function of how well he used the information available to him. The actual value of the decision has absolutely no determination in whether the decision was rational.

Suppose Jerry chooses to bet a large number of chips on a hand that he believes he is extremely unlikely to win; he has made a bad decision. If Jerry gets lucky and wins the pot, he still made a bad bet. Jerry may be very happy that he bet as he did since money won from bad decisions spends as well as money won from good decisions (though usually there's a lot less of money from bad decisions to be happy about). Jerry is mistaken, however, if he congratulates himself on the rationality of the decision.

⁶⁴ It would be reasonable to praise him, however, for his due diligence.

⁶⁵ A professional player and an unskilled player (a fish) will have the same information available to them. The professional's skill allows him to make better use of the information, and to convert what is mere noise to the fish into useful strategy. Thus, what is rationally recommended is relative to one's skill or one's maturity as a player. One may be a very good beginning player, for example, while not being a good player, all things considered.

Next, suppose that Jerry faces a tough decision on whether raising or calling will maximize expected value. Then it's a mistake if he folds, and he's rationally permitted to raise or to call. In such a situation there is nothing that deserves to be called "the rational thing to do" because the reasonable options are multiple. If nothing before the decision indicates whether it's more rational to raise than to call, then nothing after the decision will indicate it either.

These thoughts have direct application to Gauguin's situation. Suppose that he has the two options of staying with his family or following his artist's heart. If the likely value of him going is high (ignoring as Williams does his obligations to his family), it's rational for him to go whether he ends up succeeding or failing. If the expected value is low, then deciding to go is irrational, even if he should turn out against expectation to have a tremendous career. If, however, there is not enough information to recommend one course over the other before the fact, there will be nothing after the fact, and again there will be nothing that deserves to be called "the rational choice." This is not to say that the theory of expected value is useless in this situation, as Williams seems to conclude. Indeed, it is a virtue of that theory to show to Gauguin (and to us) that he can't determine that one expected value overrides the other and thus that he faces a rational dilemma. It's not as though the theory of expected value has *failed* Gauguin. It's telling him that given his paucity of information, no action is recommended as rational. That seems to me to be an excellent result. Also, as we've seen where the theory of expected value leaves off, Gauguin can employ secondary strategies, such as a strategy that recommends limiting potential losses. Alternatively, he could follow a maxim of maximizing potential gains. My purpose here isn't to adjudicate between these or any other secondary strategies. It's to explain where I think Williams has mishandled the case.

One reply to my analysis is that the case of Jerry considering his options in a single poker hand is too dissimilar to Gauguin's situation to be of use. A decision such as Gauguin faces has a paucity of information, so Williams scoffs at Gauguin's chances for prospective justification. The response to Williams' position, of which he is aware, is to try to find some rule or principle or datum point that Gauguin had access to prior to setting off such that he could have known *prospectively* whether the decision was rational. The most telling principle which Williams discusses is the idea that one is justified in making such a trip "if one is reasonably convinced that one is a great creative artist." Williams sees no hope with that kind of rule and asks,

What is reasonable conviction supposed to be in such a case? Should Gauguin consult professors of art? The absurdity of such riders surely expresses an absurdity of the whole enterprise of trying to find a place for such cases within the rules (Williams 39).

I fail, however, to see what is wrong with Gauguin doing precisely what Williams derides, namely, getting professionals to evaluate his chances of success.

Consider the case of Jerry desiring to become a professional poker player. He can do no better than to find experts to assess his play and to advise him on whether he has the psychological and constitutional capacities for success. They may play against him, or give him certain tasks to perform such as completing on-line tutorials or reading certain books, or they may take him through a particular course of skills challenges. They may have other methods, but the point is that a teaching professional won't be without resources for determining Jerry's chances of success. If they advise him that he should not pursue that end, then he will not be rationally justified in doing so, even if he does so out of stubbornness and it turns out well for him. So, against Williams, my claim is that Gauguin could do no better than to get expert advice on his talents and his capacity to capitalize on them.

By my lights, therefore, at least this fictionalized Gauguin behaved irrationally if he set off on a whim. Even if the project should succeed, his behavior is irrational since he does not have the prospective justification for taking such a costly risk, where the cost is measured financially, professionally, and personally. Similarly, it would be irrational to commission Gauguin to complete an important and costly project with absolutely no vetting of his talent beyond his own dead reckoning of his likelihood of success.⁶⁶

If Gauguin succeeds, we might take a revisionist view of his life and attribute to him more awareness of his protean talent than we had given him credit for, attribute to him the tenacity to meet the challenge, etc. Such a charitable revision attributes to him information that would have prospectively justified his behavior. Retrospective justification plays no role here. Rather we are granting retrospectively what we denied him before he left: prospective justification of his plan of action.

I deny, therefore, that luck can retrospectively justify Gauguin's or anyone's decision. My point is that if there is no justificatory luck--the ends will *never* retrospectively justify the decision--then there is no moral luck, at least of a kind that shows itself retrospectively. I cannot, of course, claim at this point to have shown that there is no moral luck, having denied the cogency of this one argument.

7. Williams on Agent Regret

Let us now turn to a second consideration that Williams brings against the standard conception of deliberative rationality. My criticism of Williams' position regarding the retrospective justification of Gauguin's behavior was based on a given

⁶⁶ Don Levi holds that the historical Gauguin's project was different from how Williams and I have been conceiving it. He holds, rather, that Gauguin's project is to live a certain lifestyle and to have certain freedoms--a bohemian lifestyle (Levi 116–18). Of course, Williams is free to fictionalize the case in any way that meets his needs.

understanding of practical rationality as being only prospectively action guiding.

Williams' denial of this conception of practical reason has a motivation that needs to be addressed. He is concerned with how Gauguin would evaluate himself from a first-person perspective rather than the third person perspective from which we evaluate Gauguin.

The thought is something like this. If Gauguin goes to Tahiti and fails as a painter, his regret about his decision will differ in a special way from how his family may regret the decision. Williams calls the first-person evaluation Gauguin may feel toward himself *agent regret*.

In explicating the notion, Williams turns toward a different motivating example to illustrate agent regret. He has us imagine a lorry driver who, even though he is driving responsibly, hits and kills a child. While the spectators certainly regret that it happened and wish that something else had happened, the lorry driver has a more intimate, less detached, less external reaction to the child's death. The lorry driver's reaction includes the wish that *he* had acted differently, not merely that things had turned out differently (Williams 43).

As Williams notes, in trying to console the driver others will tell him that there was nothing that he could have done, and Williams takes these consoling remarks as the others trying to move him from the thought that he could have acted differently to a more external and detached perspective. It is from the internalized perspective that the agent feels the kind of *remorse* which Williams characterizes as *agent regret*. It is a response to the situation that is available to the lorry driver, but not to all. Once the lorry driver moves from agent regret to the regret appropriate to a mere bystander, remorse is no longer possible for him.

The child's mother might also experience agent regret, if, for example, she had sent him on the errand that led to his death. Her friends might also intervene and seek to

move her to the point of view of not thinking of herself as responsible for her child's death. The lorry driver and the child's mother need to be and ought to be moved from the perspective in which they believe that their causal involvement is sufficient to establish moral responsibility to a perspective in which that inference is severed.⁶⁷

Williams emphasizes that regret attaches not to whether the agent has acted to maximize expected value but to whether the decision procedure he has employed has worked out for him. In speaking of an agent such as Gauguin, he says,

If he succeeds, it cannot be that while welcoming the outcome he more basically regrets the decision. If he fails, his standpoint will be one of one for whom the ground project of the decision has proved worthless, and this (under the simplifying assumption that other adequate projects are not generated in the process) must leave him with the most basic regrets. So if he fails, his most basic regrets will attach to his decision, and if he succeeds, they cannot. That is the sense in which his decision can be justified, for him, by success (Williams 50).

The reason that Williams regards agent regret as so important is that it distances success or failure from the notion of rationality that I endorse. The case of the lorry driver brings to our attention, Williams argues, a mistake in the analysis of action given by the standard conception of deliberative rationality.

On that conception, one may feel regret over any unfortunate state of affairs, a regret that the world has turned out in a certain undesirable way, but one ought only to feel remorse if one is responsible for the unfortunate state of affairs. A corollary to this

⁶⁷ Even though mere bystanders cannot experience agent regret, they *could* make the same illicit inference from the lorry driver being causally involved in the child's death to him being morally responsible for it. That is, a kind of moral outrage runs parallel to agent regret. If the bystanders are confusing upstream causal contribution with moral responsibility, and are experiencing misplaced moral outrage, they ought to be talked out of that, too.

idea is that if one deliberates well and acts in the way that rationality recommends, it's at least conceptually possible to live a life with no deserved remorse. Williams quotes Rawls much to this effect: "the guiding principle [is] that a rational individual is always to act so that he need never blame himself no matter how things finally transpire" ((Rawls 422) cited at (Williams 48)). Williams thinks that even though the lorry driver has not acted badly, recklessly, or irrationally, it's still appropriate for him to feel *remorse*--agent regret--for what had happened, which is contrary to the standard conception of deliberative rationality.

On my view, agents have acted badly when the expected outcome is bad, and they have acted well when the expected outcome is good, irrespective of the actual outcome. Agents' assessments of their own actions ought to track whether they have acted well or badly and not whether their endeavors have proved successful or not.

Agent regret is reasonable if and only if the agent has acted irrationally or in a way that is morally reprehensible. (The lorry driver, for example, ought to feel agent regret only if he was driving in some derelict manner.) Similarly, if an agent reasonably believes that his project is likely to succeed, then even if it fails (for whatever reason) he ought not to regret what he has done, though he can reasonably regret the outcome. Thus, on my view, if Gauguin succeeds, he ought to rejoice in his success, but he also ought to recognize that his setting off as he did was irrational and irresponsible. If Gauguin responds to success without more basically regretting the decision, then he has a blind spot in his self-assessment.

Likewise, if Jerry makes an inadvisable bet and lucks out, he can rejoice in his winnings without rejoicing in the decision. Or if he rejoices in his behavior, he ought to be rejoicing in the luck of the outcome, not in the rationality of his choice, retrospectively justified. If his response is, "I knew I would win," then he's a lucky fool.

Sometimes agent regret is an *admirable irrationality*, one that agents ought to be talked out of, but not too quickly. The lorry driver ought to feel regret at having caused pain, but it is irrational to blame himself for the accident, if he has done nothing morally wrong or acted irrationally. Still, it's good that agents confuse feeling badly about having caused pain with moral regret since that's good evidence that they take morality seriously and are willing to subject their own behavior to moral scrutiny. The idea to be emphasized here is that we can distinguish between regretting the outcome of a decision and regretting the decision itself.

Once all the facts are in and it's clear that the lorry driver's contribution to the scene was involuntary, he ought to drop his stance of agent regret and shift to a stance of regret appropriate to a bystander. If he cannot, that is a good sign that therapy would be appropriate. In other words, a true understanding of his contribution as involuntary ought to preclude his feelings of remorse (though it is both rational and morally commendatory for him to regret the outcome).

Finally, what are we to make of Williams' contention that it's psychologically impossible for the agent to feel regret if his decision yields pleasing results? When Williams says that "if he fails, his most basic regrets will attach to his decision, and if he succeeds, they cannot" (Williams 50), I take him to mean that the agent cannot rationally attach regret to successful decisions. It seems to me on the contrary that the more rational and level-headed and introspectively honest an agent is the more likely it is for this phenomenon to occur.

Suppose, for example, that Derek has a tyrannical manager who is on friendly terms with the CEO. His manager fires a good worker basically because he, the manager, is having a bad day. Derek sees the injustice of this and weighs the risk of going to the CEO to stand up for his co-worker against the safety of keeping his head down. Justice is an important value to Derek, however, and he takes the risk, knowing that if taking a

stand fails, he will likely be fired and his family will be put in a very hard position. As it turns out, his co-worker is reinstated and their tyrannical manager is reassigned to a position that has no managerial duties. All has turned out better than Derek could have hoped for. While he may be very happy at the outcome, he may regret the potential hardship to which he exposed his family, which, all things considered, wasn't worth the potential payoff. He may come to recognize after the fact that his behavior was brash, and that if his co-worker had wanted to keep her job (which may not have been the case), she ought to have fought for it. Thus, he may regret his decision as not being driven by anything as noble as justice but by his irrational need to rescue the underdog (and he may realize this only after the fact when his emotions have leveled off). Even while he regrets his decision on a number of different points--as a rational decision measured in terms of expected payoff, as letting an emotional response override what should have been a more reflective decision--he can still be ecstatic about the success of going to the CEO. Since I find Derek's reaction to his behavior to be natural, or at least not to be insane, then I find good reason to deny Williams' claim that the success of a project can retrospectively justify what is on balance a bad decision. I also reject his claim that a bad decision cannot be regretted if it turns out well for its author.

Again, this is not to deny that agent regret is sometimes rationally justified, as it would be for someone who drives drunk and kills. Furthermore, the homicidal drunk driver ought not to be talked out of agent regret since the behavior really is blameworthy.

The last several sections lay out Williams' critique of the Control Principle. Gauguin's case urges us to laud as rationally justified decisions made "in the dark," so to speak, and the consideration of the lorry driver pushes us to commend agent regret of the involuntary. I found nothing to prospectively recommend Gauguin's behavior, and retrospective justification is absurd. Furthermore, I find that the lorry driver's agent regret shows an admirable display of moral sentiment, but not an admirable rational

response. Indeed, I hold that a reasonable estimation of the situation ought to move the lorry driver from a stance of agent regret to the kind of regret appropriate to bystanders. The lorry driver's response is thus an admirable irrationality. My conclusion is that Williams' assault on CP is unsuccessful, and the problem of moral luck remains a problem.

8. Nagel on Moral Luck

As noted previously, Nagel distinguishes four specific types of luck:

Resultant Luck: Luck in how things turn out;

Circumstantial Luck: Luck in the circumstances in which the agent finds herself;

Constitutive Luck: Luck in the sort of person the agent is and in her natural and developed talents, her dispositions to behave, her character; and

Causal Luck: Luck in how the agent fits into the causal order.

Nagel's categories divide nicely between luck affecting the downstream causal consequences of an agent's action (resultant luck in Nagel's terms) and causal luck which affects the upstream influences of an agent's actions. The upstream causal influences further divide into (1) which causal paths upon which the agent finds herself (circumstantial luck) and (2) what given dispositions or character traits she brings to those situations (constitutive luck).

The causal-contrastive account of luck presented in Chapter 2 helps to make sense of these categories. Resultant luck is the easiest category to understand: Agent A brings about initial event, IE, which contributes to resultant event, RE. In contrast to this, IE may have but didn't lead to averted event, AE. If IE is causally unstable in bringing about RE instead of AE, it may be a matter of luck for A. It's good luck if the resultant value, RV, from RE compares well to the averted value AV from AE, and bad luck if it compares poorly. We see this pattern if we compare our siblings from the start of the

chapter, Sarah and Jane, who drive drunk. Jane's bad luck is in killing child (rather than killing nobody) because, as illustrated by comparison with her more fortunate sister, she might have arrived home without causing any harm.⁶⁸ That is, even though Jane and Sarah have relevantly similar initial conditions, luck intervenes and their resulting situations are importantly and tragically different.

Next, let's deal with circumstantial luck. As with resultant luck, circumstantial luck is figured from the connection between IE and RE, where RE has a certain resulting value, RV. RE compares well or poorly in contrast to an adverted event, AE, which has an associated averted value, AV. In circumstantial luck, the initial event is the agent carrying with him a given set of character traits, and the resultant event is the agent finding himself in a particular situation in which those character traits lead to a morally significant result (in contrast to being in a particular situation that leads to a morally better or worse result).

Mike finds himself in a circumstance in which his willingness to torture an enemy combatant presents itself and his brother, Jim, (who has the same disposition to do whatever he thinks will be required) does not. The brothers differ in the value they bring to the world, and what drives that difference in value comes not from the choices they make (or would make) but in the situations life lays at their feet. In that case, the torturing brother has bad luck because he has been placed in a situation in which he finds torturing others the natural solution. On the other hand, he may also have good luck if he might have been called on to do something even more heinous.

In constitutive luck the idea is much the same except what distinguishes one agent from another are the character traits that each possesses. The perspective we are using in declaring a person to be lucky to have a particular character trait, for example, is from

⁶⁸ She might have also killed two people on this drive home. If so, then she has the *good luck* of only having killed one person.

some time before the character trait becomes prominent. We can either judge immediately before it becomes operative in the situation, from a perspective prior to it becoming ingrained as a durable trait, or indeed from any arbitrarily chosen prior point.

Let me walk through an example to illustrate what I mean. Suppose that Adam finds himself having the virtue of courage in a battle. Given that being courageous has more value than not being courageous, it's lucky for Adam to have the character trait of being courageous if there is an IE which contributes to his courage at RE, such that IE does but need not lead to RE. Also, he has good luck relative to having a worse set of character traits, and he has bad luck if he might have acquired a better character, given the time frame from IE to RE.

So, if Adam is courageous but might not have been, then he is lucky when a situation which requires courage visits him. The amount of luck in having the virtue of courage in that situation is seen (1) by contrasting how else he might have turned out from IE and (2) by seeing how different the outcome is for him because of the courage compared to not having it in that situation, that is, how RV compares to AV.

It's not lucky for Adam to act courageously on the battlefield if we consider the narrow slice of his life in which he is in the battle and has the character trait of courage from which to act. Or at least he's not lucky in having the virtue, but perhaps he's lucky in living through enough of the battle for his courage to exhibit itself. If, on the other hand, we consider him as a young boy, taking a wider perspective on his life, he is lucky to have received the training that was so useful in the battle. If we take the narrower or the wider perspective of Adam, we still see his behavior as courageous. His courage is not at question. The question is whether we see his courage as lucky.

If we take the narrower view of his life, we are likely to attribute to Adam responsibility for his courage. This is perhaps laziness on our part. After all, if we don't know the source of his courage, we don't know if he is in any sense ultimately

responsible for having it. On the wider view, however, we attribute less responsibility to him and more to his training and to the people who trained him, that is, to influences and forces outside of him. On the wide view of Adam's life, for example, we may see how his natural tenacity contributed to his becoming courageous. If the tenacity is innate, rather than developed, we may see courage as at least partly due to his innate constitution rather than to any *choices* which we can attribute to him. What we will be looking for, from a biographical point of view, is evidence that Adam embraced his courage in a way that indicates choice. Even this biographical datum, however, would be insufficient to establish moral responsibility, since the visage of him embracing courage is consistent with his being causally determined to embrace courage.

If the attribution of luck from the wider view can diminish or preclude an attribution of moral responsibility, the question is, therefore, whether Adam is *really* morally responsible for his battlefield courage. Ought we really to praise him, not merely as an encouragement to others to emulate his behavior?

From a pragmatic point of view, we can praise Adam for his courage viewed within the context of the battle, and not worry about whether he was responsible in any ultimate sense. That is, from a principle of charity, we may simply presume that if his behavior seems to us to be meritorious, we ought to praise him. This may be merely speaking with the vulgar, as Berkeley would have it. Or if we judge his behavior more critically (thinking with the learned) or from a wider point of view (from a biographer's view, for example), while we may emphasize his battlefield courage, we may note that he was just performing as trained and that since he wasn't responsible for that he wasn't responsible for the resulting courage. That is, if we judge his behavior under a more critical eye, we may be hard pressed to see where his training leaves off and where his choices begin. To take another example, Donald Trump's children exhibit a fine business sense, but are they praiseworthy for it? After all, given their circumstances, what other

choices did they have? It's doubtful, for example, that forming an indie band was ever given breath as a live option.

The change of perspective does not alter our (or at least my) estimation of Adam's courage, but it does alter whether he is seen as being responsible for the cluster of character traits that drove his behavior on the battlefield. The more outside influences that we find to be relevant to what he does, and here think of the contributions to his durable character traits, the less responsible he seems for them. Since moral responsibility varies from what we count as our initial event, we can say that he's morally responsible from one perspective and not morally responsible from another. The dog that fetches slippers on command impresses, but it seems less like an act of the dog's will than the master's when we consider the training put into getting that response on cue.

This is to take what Nagel calls the *external view* of Adam, Trump's children, and the well-trained dog. If we figure in constitutive, circumstantial, and resultant luck and take a wide view of people, we see little that we can unambiguously attribute to them as agents. Nagel says,

But as the external determinants of what someone has done are gradually exposed, in their effect on consequences, character, and choice itself, it becomes gradually clear that actions are events and people things. Eventually nothing remains which can be ascribed to the responsible self, and we are left with nothing but a portion of the larger sequence of events, which can be deplored or celebrated, but not blamed or praised (Nagel 68).

In contrast to a view of Adam which emphasizes the external determinants of his courage there is a view which focuses on the internal determinants of his choices. On a naturalistic view of the human person, the internal/external distinction won't amount to much. If there are naturalistic forces that impel Adam one way or the other, then he is not in any relevant sense ultimately responsible. Thus, we ought to take Nagel's distinction

to be between naturalistic and non-naturalistic forces.⁶⁹ The issue is that once we discount as praiseworthy or blameworthy any behavior or character trait which has an explanation in terms of external or naturalistic forces, we find nothing about Adam which is a fit object of moral evaluation. Thus, Adam has ceased being seen as a moral agent, about which Nagel says,

The area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point. Everything seems to result from the combined influence of factors, antecedent and posterior to action, that are not within the agent's control. Since he cannot be responsible for them, he cannot be responsible for their results (Nagel 66).

One can, of course, in spite of these considerations insist that Adam's courageous behavior is rightly *his*. Nagel says that to take that position is "to brazen it out and refuse to accept the results" (Nagel 66). Many will do so, and we can consider a couple of ways in which such a view might be justified.

First, one may cozen a traditional compatibilist view of moral responsibility in which one is responsible for the contribution of one's will to one's character traits or behavior if that contribution comes off in the right way. That is, Adam's behavior will be meritorious given that it is driven by his wishes, desires, affective impulses, etc., without consideration for where these volitional artifacts might have come from. He is responsible for his battlefield heroics provided that he wasn't coerced, wasn't sleepwalking or deceived, didn't pull the trigger due to an involuntary spasm, and the like. I have rejected such a view of moral responsibility in the previous chapter, however,

⁶⁹ By naturalistic forces, I mean forces that are part of event-causal explanations. Naturalistic forces can be causally deterministic or causally indeterministic. By non-naturalistic forces, I include agent causation (in contrast to event causation), Kane's teleological intelligibility theory, noumenal selves, and explanations of free will or moral responsibility that rely on mind-body dualism.

and have said that if that's *all* there is to moral responsibility, it's not a responsibility worth having.

Second, one may claim that even after we subtract out the external influences on Adam's behavior--including dispositions or character traits that are innate or have a genesis that points beyond Adam's choices--there has to be something that is not accounted for--namely, the contribution of Adam's choices to his traits or his use of his will. Much more would have to be said, however, for this to not just stand as a bare assertion, but there are a number of philosophers who try to defend an incompatibilist view in which Adam himself is the cause of his choices, not the influences which impinge upon him from within or without. Reid, Taylor, Chisolm, O'Connor, and others, for example, defend agent causation, that is, the claim that not only do events cause other events to obtain, but also that *agents themselves* have causal powers that are not reducible to the model of event-causal explanations. Similarly, Kant posited a noumenal self, outside the naturalistic order, as the source of free will and moral responsibility. The drawback to such views is in understanding *how* it is that agents have this power without accepting it as a mystery.

In addition to the external view of human behavior, Nagel also recognizes an *internal view* of human nature, according to which we extend to ourselves attributions of responsibility for our actions. It is the view that we also extend to others in considering them as moral agents. In this view, we do not merely look on ourselves as things and our behaviors as events. Rather, we are agents who act in the world, and when we morally judge ourselves and others, we are judging ourselves and others as agents. We are not merely judging the outcome of the person's behavior. We are judging *him* as an agent and as a person responsible for his actions and their effects. As Nagel says,

About ourselves we feel pride, shame, guilt, remorse--and agent-regret. We do not regard our actions and our characters merely as fortunate or unfortunate episodes--

though they may also be that. We cannot *simply* take an external evaluative view of ourselves--of what we most essentially are and what we do. And this remains true even when we have seen that we are not responsible for our own existence, or our nature, or the choices we have to make, or the circumstances that give our acts the consequences they have. Those acts remain ours and we remain ourselves, despite the persuasiveness of the reasons that seem to argue us out of existence (Nagel 68).

Nagel thus sees the problem of moral luck as growing out of the conflict between these two views of ourselves, neither of which can be easily abandoned. If we take seriously the thought that we are part of the natural order, it's hard to see how we're more than that, that we're special in that we have genuine free will in spite of our similarity to those creatures that don't. On the other hand, it's very hard to think of ourselves as beings without free will, to think of our apparent free will as only illusory. Even though it may seem obvious to us that we have free will, however, the fact that we seem free doesn't guarantee that we are free.

On Nagel's view, the solution to the problem will not come from showing that one of these views is chimerical, that the external view is only apparently true or that the internal view is illusory. Rather, Nagel thinks that the solution will be to show that the truths embodied in these views are only apparently in conflict. Nagel's preferred solution is to accept the premises of the basic argument, while denying the conclusion of moral skepticism, but Nagel himself does not offer a solution. (It is sometimes service enough to frame an interesting problem without offering a solution.)

In the next chapter, I want to examine several proposed solutions to the problem of moral luck, which I will evaluate not only on their own merits but also in terms of how well they meet the challenge on the terms which Nagel has dictated: do they respect both the internal and external views of human nature without casting off one or the other as

illusory. (It is interesting to take Nagel's problem as expressing a conflict between two views of human nature, but only to the extent that one preserves the cogency of each of the contrasting views.)

Chapter 5

Solutions to the Problem of Moral Luck

1. Kinds of Solutions

As we saw in the last chapter, Thomas Nagel presents an interesting problem for moral philosophy: our moral responsibility is limited to our sphere of genuine control, that sphere is empty since our lives are riddled with luck, and thus our moral responsibility has vanished. A more moderate version, still robust enough to engage us, is as follows:

The Moderate Problem of Moral Luck

1. Agents are often morally responsible for their actions, intentions, and character traits.
2. Most actions, intentions, and character traits are not within any agent's control.
3. Thus, agents are sometimes morally responsible for actions, intentions, and character traits that are not under their control.

The upshot of the moderate problem is to emphasize that to generate *a* problem of moral luck we need merely for the range of moral responsibility to be called into question in some important way. If we are rarely morally responsible, for example, or responsible only for relatively trivial actions, then we may wonder what of importance remains of moral responsibility. That is, even if we are not rationally required to table our notion of moral responsibility, we may be required to scale it back beyond a point at which there remains a robust moral theory.

I aim to solve both the Basic Problem of Moral Luck which was presented in the last chapter and the Moderate Problem presented here. In Sections 2 and 3, I first consider, and reject, an epistemic approach from Norvin Richards. Sections 4 and 5 address important ideas from Michael Zimmerman with regard to constitutive and

circumstantial moral luck. The last several sections present and defend a principle of evaluating actions that is consistent with the existence of genuine moral luck with regard to constitutive luck. There is, I claim, no genuine moral luck with regard to resultant luck or circumstantial luck. Furthermore, I argue that while there is genuine moral luck there is no *problem* of moral luck since the existence of genuine moral luck neither prompts us to deny the Control Principle nor forces us to embrace moral skepticism. Neither the Basic Problem nor the Moderate Problem need worry us. Finally, in Section 8, I sketch a picture of moral development in which genuine moral luck can fit comfortably.

2. Richards' Epistemic Solution: Resultant Luck

In his paper "Luck and Desert," Norvin Richards argues against moral luck as a genuine phenomenon within morality. Instead, he claims that the apparent moral difference between two agents which we attribute to luck really reflects an epistemic difference in what we know about them. In doing so, he denies the second premise of both the Basic and Moderate arguments--agents do have the right kind of control over their interiors to be morally responsible. His central claim, however, is more modest, namely that Nagel is wrong in his assessment of the examples used to motivate his position. We often lack evidence concerning people's deserts, but this does not imply that they are morally responsible for what is out of their control.

One of the agents is unlucky to have been shown to be a scoundrel, and the other, though equally a scoundrel, has escaped our detection. Less negatively, one has been fortunate enough to have been seen in a good light, and another, perhaps an equally good person, has escaped our moral notice.

Let's take, as Richards does, an example of resultant moral luck.⁷⁰ Recall our

⁷⁰ I discuss Richards' stance on constitutive and circumstantial luck in the next section.

siblings Sarah and Jane from previous chapters. Every Wednesday, they meet for lunch, and over the course of a couple of hours get at least mildly drunk before driving home. One Wednesday, Sarah arrives home safely, and Jane kills a child on her drive home. *Prima facie*, there are two possibilities in how luck might make a moral difference between Jane and Sarah. First, luck can create a difference in the wrongness of each sibling's behavior. If one counts the resultant value of an action, it is reasonable to believe that Jane's behavior is worse than Sarah's. This follows from how badly things turn out in one scenario as compared to the other. Second, luck might make a difference in the blame that each sister deserves.⁷¹ If Jane's behavior is worse than her sister's because of the worse outcome, Jane deserves a level of blame that Sarah does not.

Let's pause to reflect briefly on the nature of praise and blame. There are typically three factors to consider relative to a single episode for which an agent may deserve praise or blame. First, there is whether the agent has done something morally admirable or morally offensive for which she is morally responsible. If we were aware of all the relevant facts, would we be justified in administering praise or blame? In Chapter 3, I said the notion that the agent may correctly deserve praise or blame is the idea that each agent has a moral ledger. It is an objective notion.

Second, there is our judgment that the agent is blameworthy or praiseworthy. A judgment that someone is praiseworthy or blameworthy ought to be based on evidence, rather than, say, superstition or prejudice. As with all empirical judgments, a moral judgment is liable to error since an agent may be praiseworthy, for example, even though our evidence suggests she is blameworthy. In addition, agents may have all sorts of credits and debits on their moral ledgers about which no one knows. As a judgment, this

⁷¹ I deny this possibility, but I also deny resultant moral luck. For one who endorses resultant moral luck--Richards' target--this claim will seem plausible.

is an epistemic affair, and making the judgment is correct or incorrect according to epistemic norms.

Third, there is our overt act of administering (or withholding) praise or blame. As with any act or omission, it is liable to judgment by moral norms. My distinction between the second and third features of moral responsibility echoes a similar distinction made by Michael Zimmerman. He calls the epistemic judgment that an agent is praiseworthy or blameworthy *inactive* and says that it is *internal* to the judging agent. Zimmerman calls giving praise or blame *active* and says that it is *external* to the agent who might confer praise or blame, in the sense of his making or suppressing an external action (Zimmerman, “Luck and MR” 218–19). Inactive praise and blame are subject to epistemic norms, and active praise and blame, being actions, are subject to moral norms.

Even if the agent is morally accountable (her behavior is reflected in her objective moral account) and even if we are epistemically justified in believing that she’s blameworthy, we may not be morally justified in administering blame. For example, there may be good consequentialist grounds for withholding it, even if we’ve got good reason to believe that it’s deserved. Or we may have good consequentialist reasons to overtly blame or praise even if we know that the agent is undeserving. For ease of exposition, however, I will assume that if the person objectively deserves to be held accountable, we will have good epistemic reasons to hold her accountable, and we will have good moral reasons to administer praise or blame in exactly the proportion that she deserves it. Armed as we are now with that map of moral accountability, we can return to our discussion of Richards’ position.

In claiming that there is no resultant moral luck, Richards denies that luck has made a difference either in the wrongness of Jane’s and Sarah’s respective behaviors or in what each sister deserves, *pace* one who holds that there is genuine moral luck. All moral luck is, on his account, of the epistemic variety. His task, therefore, is to explain (1) the

intuitive appeal of thinking that they differ at least in the wrongness of their behavior, if not also in their blameworthiness, and (2) how a defensible moral theory can accommodate the intuition without accepting that there is genuine moral luck.

To these ends, Richards identifies a difference in our epistemic positions with regard to the sisters, rather than finding a moral difference between them. In making up the story of Sarah and Jane, I availed myself of an omniscient third-person point of view. Each was equally wrong and deserved equal blame because each had endangered others to the same extent.⁷² As author of the scenario it is reasonable to insert myself into that perspectively superior position.

However, that “pretended omniscience,” as Richards calls it (Richards 169), is not true to life. We are not (generally) in a position to know whether real people differ in the risks they pose to public safety, even though it is tempting in this case to think that we are. Richards’ idea is that we should not let our lack of epistemic access into the reality of who’s blameworthy lull us into thinking that we have good evidence to think that there is genuine moral luck. We aren’t omniscient, though sometimes we pretend to be even in real life, and when we make judgments from this point of view, we make logical errors. We lapse into thinking that we have sufficient epistemic access to responsibly make differential moral judgments, and we don’t.

We can capture my position as follows. On one hand, knowing only the difference in drunkenness between A and B (assuming the difference to be significant in terms of their ability to drive safely), I can easily assess the moral difference between them. If A is significantly drunker than B, it’s morally worse for A to drive. If there is no difference in threat that each poses, then A and B are on an equal moral footing. On the

⁷² I have maintained throughout that the wrongness of behavior follows the intention that it embodies, not its causal consequences. This view is consistent with both Kantianism and intention-based consequentialism, to say the least.

other hand, knowing only the difference in outcome between C and D (C drives without incident and D kills a pedestrian); I can easily say which outcome is worse, but not which driver deserves more moral blame. Perhaps D, like Williams' lorry driver, deserves no blame. Thus, considerations of outcome are neither necessary nor sufficient for assessing blame. It's not necessary because I can assess blame with knowing only the difference of risk posed. It's not sufficient because I can know that an agent caused a particular bad outcome without being able to assess whether he or she is morally responsible for the outcome.

Richards' point is well taken. We rightly blame the drunk driver who seriously injures or kills others. The threat to others they posed is undeniable. For "successful" drunk drivers, we have less evidence of their threat to others, so less justification for blaming them. Thus, it may *appear* that the bad luck of injuring or killing others has made a moral difference, but that is a false appearance. It's still cogent to say that they deserve exactly the same blame and thus that luck has made no difference in their moral ledgers.

With regard to resultant luck, Richards can claim that *if* our drunk drivers have the same level of control over their actions, intentions, and character traits, their level of moral responsibility does not differ. On Richards' view, it is agents' durable character traits, primarily, for which they are morally assessable and which punishment and reward serves to support or to check. If agents have some control over their character traits, therefore, the problem of resultant moral luck is, on Richards' view, resolved.

3. Richards' Epistemic Solution: Constitutive and Circumstantial Luck

It's important that we see how Richards responds to the other two categories of alleged moral luck: constitutive luck and circumstantial luck. Not surprisingly, Richards treats these much as he has treated resultant luck. Generally, he looks at the agent who

evades detection as lucky in our epistemic access to her character rather than as lucky in the rightness or wrongness of her behavior or in the amount of praise or blame that she deserves.

In the case of *constitutive luck*, Richards argues that it is fair to treat people with identical characters in the same way only if we know that each has made the same contribution to his own character traits.⁷³ If Tom and Bill act on a given occasion equally bravely or equally cowardly, for example, we ought to praise or blame them equally only if we are sure that each had an equal hand in choosing his durable character traits. If Tom had a merely given constitution, with no part of it attributable to his free choices, then we ought to neither praise nor blame him for actions which arise from those durable traits. On the other hand, Bill deserves praise if, for example, he consciously developed his courage. That is, we can distinguish between an agent's character as given (his durable dispositions to behave that he came with, so to speak) and his moral character (his character traits which significantly affect others and which the agent has chosen or consciously contributed to).⁷⁴

Since for all practical purposes we don't know that a particular action flows from a given character trait or a developed character trait, Richards' suggestion is that we ought to administer praise and blame as best we see fit, but realize that often we judge matters incorrectly. These judgments are fraught with error, of course, but the luck which accrues to the agent is purely with regard to our epistemic positions and our ability to judge him correctly. That being so, that one is seen as morally courageous and another is

⁷³ Of course, given our lack of epistemic access, there is also a practical barrier to our knowing *that* two agents have identical characters.

⁷⁴ One oddity of my taxonomy is that courage will count in some people as a moral virtue and in others will be a mere character trait. Zimmerman makes a similar distinction between our given characters and our developed characters (Zimmerman, "Luck and MR" 223–24).

not is not a matter of genuine moral luck. The agent's moral condition or the response he deserves is not a matter of luck. Rather, the agent's luck comes in our knowing his true nature--both in the nature of his character traits and in their origin--and so in our ability to give him what he truly deserves.

Note that there are two ways for us to err in our assessment of the agent. First, our evidence may conceal whether he has a particular character trait. An agent may be durably brave and yet his disposition never be revealed in his behavior. Second, it may be clear that he is brave but unclear *how* he came to be brave. His bravery may have been inborn or forced upon him, or he may be ultimately responsible for its development, given the choices he has previously made. The question remains as to whether the agent is *ever* in control of his constitution. Richards says,

This argument succeeds, I think if one's character is to *no extent* one's own artefact. But if the individual makes any contribution whatever to the sort of person he is, that contribution can be the basis for his deserving praise or blame for what he does. . . . It could be that one's character is shaped *entirely* by forces beyond one's control. If so, the practice of attributing responsibility is undermined, with no need for us to refer to the other sorts of luck Nagel brings to our attention (Richards 172).

Thus, Richards rejects the second premise of the Basic Problem--that no action, intention, or character trait is ever within any agent's control. We can put his view a bit more modestly: either agents can control their actions, intentions, and character traits, and there is no genuine moral luck, or they don't have that kind of internal control, and there's no moral responsibility and therefore no genuine moral luck.

Finally, Richards gives a predictable gloss on *circumstantial luck*, given his earlier commitments. One way in which an agent may be lucky is whether we have evidence of his true character because of how circumstances give him or fail to give him the

opportunity to act on his character. Consider, as Richards does, Nagel's example of the person who avoids the moral condemnation due a Nazi concentration camp guard because his circumstances led him to Argentina in 1930 (Nagel 58–59). Richards is skeptical that a person with that damning callousness of character, that utter disregard for human suffering, and that predilection to blindly follow authority will avoid our notice. He says, "My central contention will be that if the potential agent is as much like the actual one as we are imagining, then there will be something else in his behavior which will call for the same response" (Richards 174).

I doubt that Richards is correct here, but it may be a minor quibble, despite the fact that he describes it as his "central contention." Part of what would allow the character of a Nazi concentration camp guard its full expression is that the guard would be encouraged in his behavior both within the ranks of the other guards and from his superiors. Thus, except for the thought that the Germans might lose the war and he may end up being held accountable for his actions (perhaps a remote thought, especially before the tide turned in 1944), a concentration camp guard had very little reason to mute his characteristic responses and was even rewarded for giving his cruelty full reign. On the other hand, a businessman in Argentina may have the same cruel streak but may radically suppress his natural reactions so as to hide his true nature and to evade detection and the ensuing censure, or at least suppress it to the extent that observers would have little knowledge of its depth. So, given that the Argentinian businessman has the same cruel dispositions as the Nazi guard, nothing in his behavior need call for the same response since his behavior may not stand as evidence for the same depth of cruel disposition.

Thus, I think we can allow that the Nazi and the businessman in Argentina may have the same characters and yet act vastly differently. The behavior of both may be signs of their durable character traits, but quite different and perhaps misleading signs, reflecting the vast imperfection in our epistemic access to what drives their various behaviors.

Be that as it may, I think Richards is safe with the more modest claim that the Argentinian businessman may find it difficult to conceal all of his true nature. While he may not give full reign to his cruelty, it will be very difficult for him to keep it completely concealed, so the luck of the circumstances will again effect not what he deserves but our ability to know what he deserves. If luck is on his side, we will see only the tip of the iceberg and have little reason to guess at its depth. This again is not luck in his level of blameworthiness, but luck in the epistemic features that allow us to discern his blameworthiness.

The causal-contrastive account of luck sorts these cases nicely. It figures the amount of luck by the likelihood of the alternative event and the goodness or badness of the actual results as compared to the alternative results. To see this, let us consider the case of the Argentinian businessman. We see someone going about his business in Argentina, someone who perhaps has a bit of a stern demeanor, one who is perhaps a little over-fawning to superiors and a little too ready to accept orders on authority. Let's suppose that it's by chance that he is in Argentina rather than in his native Germany. Thus, the result is that he spends the war away from Germany and the value to him of that outcome is that he commits some fairly small moral errors because of his circumstances.

On the other hand, he may have chosen not to go to Argentina or have chosen to return to Germany before the war or to help with the war. In these circumstances, his behavior, let's say, would have been atrocious, and his moral record would have suffered greatly. In cases like this, the agent's character is held constant as part of the initial event. His circumstances and his behavior within various circumstances are allowed to alter according to the various counterfactual suppositions, giving us the resulting and averted events.

With regard to a real case, or a case more realistically imagined (where our perspective is from within the scene), it will be opaque to us what the agent would do

were he to return to Germany. We just wouldn't have the epistemic access to the details of the agent's character traits to responsibly make judgments about what he may be responsible for.

The amount of luck on this account is objective and yet epistemically very difficult to know. The opacity of luck is, I think, a welcomed result for the theory, since it urges us to be humble and cautious in our non-textbook, no-pretension-of-omnipotence moral evaluations.

Overall, Richards' account seems reasonable to me. We ought to restrain our moral judgments to the limits of our epistemic reach. Thus, our epistemic access to Sarah and Jane's true desert is meager; we don't know how much blame is owed to each one. In addition, circumstances can reveal or fail to reveal moral responsibility, and it is difficult to discern the character traits agents have. It is even more difficult to tell whether they are morally responsible for having chosen those character traits.

Richards describes the luck of the ascription, however, and not the luck of the desert. That is, we can grant luck in the appearances, but is there luck *only* in the appearances? Richards thinks so: as we hold someone responsible our responses to him "only reflect our epistemic shortcomings, and the agent's good or bad fortune in those" (Richards 169). To put the issue another way, even if we grant that there is epistemic moral luck--luck in how our epistemic access allows us to or prevents us from making accurate moral evaluations--we can ask whether luck ever makes a difference in our actual moral worth, that is, whether there is objective moral luck. It's not clear that Richards has directly addressed this issue.

So, while I generally agree with the core of what Richards says--luck may affect the level of praise and blame we are epistemically justified in doling out--it is important to see that this does not directly address the problem of moral luck as I presented it in the last chapter. In particular, it does not address whether we have the appropriate level of

control needed for moral responsibility or whether we are ultimately responsible for our actions, intentions, and character traits.

This is not to criticize Richards. His goal is not to show whether there is or is not genuine moral luck (though it's clear from his writing that he thinks there is not). His target is more modest: refute Nagel's *paradox* of moral luck. Nagel presents certain examples (like the example of the would-be-Nazi living the war in Argentina) as examples of our lack of control in our moral records. He then argues that if we lack control of our moral records, we are subject to genuine moral luck. Richards holds, on the contrary, that the examples don't constitute a lack of control over our moral records, but a lack of control over the *evidence* of our moral records. That, he holds, is perfectly consistent with there being no genuine moral luck.

Richards' point, therefore, is modest. He's not out to *refute* Nagel's paradox but to *accommodate* it or *escape* it (Richards 177, 180). He does this by showing that in Nagel's motivating cases the lack of control is not over our moral desert but over people's perceptions of our moral desert. That is, Nagel's allegedly motivating cases ought not to move us toward the claim that there is genuine moral luck.

To put the issue another way, Richards hopes to explain how the appearance of moral luck can occur in the absence of genuine moral luck. If there are arguments which take us beyond the appearances, however, we ought to address those and try to get to the core of whether there is genuine moral luck. The remainder of the chapter is devoted to that task.

4. Zimmerman's Argument Against Moral Skepticism

Michael Zimmerman has interesting ideas on the problem of moral luck, and it would be instructive to look at his views at this juncture. In particular, I want to look at a related pair of his arguments, since it is through discussion of that line of reasoning that I can complete my account.

In “Luck and Moral Responsibility,” Zimmerman’s preliminary move is to distinguish between what he calls *restricted control* and what he calls *unrestricted control*. An agent has *restricted control* over an event if and only if he is able to bring about its occurrence or prevent its occurrence. An agent has *unrestricted control* over an event if and only if he has restricted control over both it and every causal factor on which its occurrence depends (Zimmerman, “Luck and MR” 219). I may have restricted control over whether I have a cup of coffee at this time of night, but I do not have unrestricted control over it, since I do not control whether I was born, whether coffee evolved as an ingredient suitable for potables, etc. Indeed, there are no events over which any finite being has unrestricted control.⁷⁵

With this distinction in mind, Zimmerman cleaves a traditional argument for moral skepticism into two related arguments. The first argument, which I will call Argument A, runs as follows. Where P is any person and *e* is any event:

- 1a. P is morally responsible for *e*’s occurring only if P was in restricted control of *e*.
- 2a. No event is such that anyone is ever in restricted control of it.
3. Therefore, no event is such that P is morally responsible for its occurring.

The second argument, Argument B, goes as follows:

- 1b. P is morally responsible for *e*’s occurring only if P was in unrestricted control of *e*.
- 2b. No event is such that anyone is ever in unrestricted control of it.

⁷⁵ Note that Zimmerman’s distinction between restricted and unrestricted control is not relativized to past and present (and perhaps future) times. Such a distinction might give meaning to finite beings such as ourselves having unrestricted control: an agent has unrestricted (time relative) control over an event at time *t* just in case at *t* he has restricted control of both it and every causal factor at or after *t* on which its occurrence depends. If I thought that an argument cast in these terms would fare better than the arguments I’m about to consider, I would pursue this option, but I don’t.

3. Therefore, no event is such that P is morally responsible for its occurring (Zimmerman, “Luck and MR” 220).

Once those arguments are distinguished, Zimmerman claims, it is easy to reject them as unsound and hence to not accept their common conclusion, which is statement 3.

Argument A has a false second premise because all sorts of events are under our restricted control in normal circumstances. If any event is under my control after any fashion, I have restricted control over it.

On the other hand, Argument B has a true second premise, since whether I was born is not under my restricted control, and any event over which I might have unrestricted control would require that I have restricted control over whether I was born. Zimmerman, however, rejects the first premise. He denies that moral responsibility for an event requires that it and all causal factors contributing to it be under the agent’s restricted control.⁷⁶

Finally, we can split the Basic Problem (and the Moderate Problem) along similar lines. That is, we can give an “A” version in terms of restricted control and a “B” version in terms of unrestricted control. In the “A” version, the second premise will be false, since some actions, intentions, and character traits are within the agent’s restricted control. The “B” version, on the other hand, will fail because it commits us to the claim that unrestricted control is required for moral responsibility, and it’s not.

Thus, it doesn’t look like we can generate a general argument from moral luck to moral skepticism. The question remains, however, whether there is any moral luck.

⁷⁶ This, by the way, seems to be where Galen Strawson parts company with Zimmerman and me. In Strawson’s Basic Argument for moral skepticism, he requires that for an agent to be ultimately responsible for being brave, for example, he must have restricted control over every aspect of the world that contributes to his bravery. That is, he must have unrestricted control over his bravery, and no one has that kind of control (Strawson, G., “Impossibility” 5–6).

After all, just because there's not sufficient moral luck to warrant moral skepticism doesn't mean that there's no moral luck.

5. Zimmerman and Genuine Moral Luck

Zimmerman rejects, as I do, genuine moral luck of the resultant stripe, and he does so for similar reasons. Instead of rehashing these reasons, we can move forward and discuss his account of what he calls *situational luck*, which encompasses what Nagel refers to as *constitutive* and *circumstantial luck* (Zimmerman, "Luck and MR" 228). In explaining situational luck, Zimmerman distinguishes between external influences on behavior and internal influences. He elaborates on these influences by saying,

It is against them as a background that one makes the decisions that one does; indeed, without such a background, no decision could be made. Nevertheless, as long as the decision, for example, to collaborate [with the Nazis] is made freely, then one is surely, *ceteris paribus*, to blame for such collaboration (Zimmerman, "Luck and MR" 228).

Thus, our evolving characters and the various situations in which we find ourselves serve as the background for our actions, and moral responsibility requires free action within that causal scaffolding. The question, therefore, is whether free action within those causal confines can produce genuine moral luck.

Zimmerman's goal is to *neutralize*, as he calls it, the role of luck in moral responsibility. This goal drives his search for a principle which will serve that purpose. That is, in regards to situational luck, he is looking for a principle by which the Nazi collaborator and the would-be collaborator who spends the war in Argentina are equally morally responsible and hence equally to blame.⁷⁷ To this end, he proposes for our

⁷⁷ He agrees, however, with Richards' point of view according to which it may be beyond our epistemic ken to know whether they are equally morally responsible and

consideration the following principle, which I will call the Counterfactual Principle.

Where P and P^* are any people, d is any decision, and s is any situation:

If (i) P made d in what he believed to be s ,

(ii) P^* would have made d if he had be in a situation that he believed to be s ,

and

(iii) P^* 's being in a situation that he believed to be s was not in his restricted control,

then whatever moral credit or discredit accrues to P for making d also accrues to P^* (Zimmerman, "Luck and MR" 225).

In trying to understand this principle in action, let's first consider the case of the Nazi conspirator and his Argentinian counterpart who have the same character traits which drive their behavior. This will be a case of what Nagel would call circumstantial luck since we are supposing that they have the same constitution and yet differ in their circumstances and opportunities (and hence in whether and to what extent their dispositions are activated). The collaborator makes a decision to act atrociously in the situation in which he finds himself, e.g., within his role in Germany in 1944. Suppose that the would-be collaborator would have made the same decision had he believed he was in the same circumstances, and it's merely a matter of luck that he's not in the same circumstances. In that case, the luck of not being in the same circumstances does not shield him from whatever moral credit or discredit the Nazi collaborator has earned.

Let us now consider a case of what Nagel would call constitutional luck. Suppose that Amy has decided to develop a certain disposition for which counts to her moral discredit, for example, a propensity to steal. Brenda has not made a decision to acquire

equally to blame. Since Zimmerman is worried about the phenomenon of moral luck and not our epistemic access to it, we can set that consideration aside.

the disposition because, by her good luck, she has never been in circumstances in which trying to acquire that disposition was a live option.⁷⁸ That is, she was never in a situation in which she believed that she could and should develop that character trait. What's more, had Brenda been in a situation which she believed would be conducive to acquiring that character trait and believed that she should acquire it, then she would have decided to acquire it. Finally, whether Brenda ever found herself in such a situation is a matter of luck for her. In such a situation, Zimmerman's position is that whatever moral discredit accrues to Amy for her acquired propensity to steal ought to accrue to Brenda for her propensity to acquire a similar character trait.

We have to be careful here. Suppose that Brenda is in a situation similar to what Amy believes herself to be in when she acquired the character trait that drives her stealing. If that situation stirs a compulsion to steal within Brenda, then she would steal just as Amy did. If Brenda's desire to become a thief is a compulsion, however, then she's not free with regard to it, and hence has no free will in the situation. Even if the desire reaches to level of an all-things-considered desire (but short of a compulsion), then by the belief-desire-action triangle, there's little room for free will, and hence for moral responsibility, to wedge itself in. Rather, it seems, for Brenda to be morally responsible for her decision to acquire a thieving disposition, it must be desirable enough to consider, but not so desirable as to wedge out the option of not becoming a thief. Thus, acquiring the disposition and not acquiring it must both be live options.

⁷⁸ For x to be a live option for agent A, doing x must be both desirable and seen as possible. Hence, while it's physically possible for me to step in front of a speeding bus, it's not a live option for me since I have no desire to do so. How likely A sees himself as being able to pull off doing x should he try and how desirable doing x is to A are both scalar notions. We may speak, therefore, of a particular option being more or less alive (or more or less vivid) for the agent.

In each case, it's not the would-be collaborator or the would-be thief's decision which earns the moral discredit, because, by the example, the would-be collaborator *never decided to collaborate* and the would-be thief *never decided to acquire a thieving character*. Rather, it is the counterfactual decision that each is primed to make which is his or her moral downfall. Zimmerman is aware that this is a *prima facie* odd if not problematic result. He says,

But if the noncollaborator is just as much to blame, what is he to blame *for*? Not collaboration, clearly; and in this case there is not even the decision to collaborate.... Perhaps we should simply say that the noncollaborator is to blame but just not to blame *for* anything; or perhaps we should say that he is to blame for being such that he would have made the decision to collaborate had he been in a situation that he believed to be *s* (where *s* is the situation that the collaborator believed himself to be in) (Zimmerman, "Luck and MR" 228).

Zimmerman realizes that both of his options are problematic. In the first case, the idea that the noncollaborator is to blame but he is not to blame *for* anything rubs against the ordinary grammar of blame. If P is to blame then there ought to be something that he's done or left undone, some intention realized, or some character trait acquired that we can point to and say *that's* what P is to blame for.

We can approach this objection from a different angle. While it makes sense to say that the collaborator and the noncollaborator have the same black mark on their souls (since they have the same characters, they are responsible to the same degree), it seems counterintuitive to say that the would-be collaborator is morally responsible but that there is nothing for which he is responsible, that his responsibility has zero scope. That is, the concept of moral responsibility includes being responsible to a positive degree but denies the possibility of being responsible with zero scope. After all, we already have a distinct concept that covers that territory, the evaluation of one's character as good or bad. We

can admit that the correct evaluation of the collaborator's and the noncollaborator's characters shows them to be to be equally bad, irrespective of the circumstances in which they find themselves. Evaluating them in terms of moral responsibility is a different matter, however, and requires that there be something for which they are responsible. Moral responsibility is always responsibility for something, and never just responsibility *simpliciter*.

In "Taking Luck Seriously," Zimmerman makes the same distinction between the degree and scope of moral responsibility. There he explicitly ties the notion of control to this distinction: we can distinguish the degree of control (how much control the agent has) from the scope of the control (how many things the agent controls) (Zimmerman, "TLS" 562). Again, this is a perfectly fine distinction, but it breaks down in the same limiting case in which the distinction between the degree of moral responsibility and the scope of moral responsibility breaks down: when the scope is zero. If A controls three events and B controls two, it's meaningful to say that A has a wider scope of control than B and yet they have the same degree of control. But if C controls no events, it's a mistake to say that A has a wider scope of control than C, and yet that they have the same degree of control. Thus, I don't see how it helps to discuss the noncollaborator as having a positive degree of moral responsibility (or control) but zero scope. If the agent's scope of moral responsibility or control is zero, then so must be his degree.

In sum, I disagree with Zimmerman when he says, "Degree of responsibility counts for everything and scope counts for nothing, when it comes to such moral evaluation of agents" (Zimmerman, "TLS" 568). Again, I disagree when he says in a different work that "How one is to be judged, from the moral point of view, is thus entirely a function of the degree, and not of the scope, of one's culpability" (Zimmerman, "Map" 599). Or, rather, I agree with such attributions when both scope and degree have non-zero values, such as when we compare the responsibility of an assassin from a would

be assassin who sneezes just as she pulls the trigger and thus misses the mark. I disagree with Zimmerman's wider point, however, since when the scope of responsibility or control is zero, the degree of responsibility or control cannot be anything but zero.

Zimmerman's other suggestion is to blame the noncollaborator for his character being such that he would have become a collaborator if placed in the same position in which the collaborator found himself, namely believing that it is a situation of kind *s*. There are three issues with this proposal, as Zimmerman well knows. First, there are persistent problems with determining the truth values of counterfactuals such as the Counterfactual Principle proposes. This is a problem with our accurately assessing whether an agent is blameworthy, not in whether she is blameworthy. That is, it is merely an epistemic problem.

Second, the counterfactuals about what the non-collaborator would have done in given situation may be without truth value; indeed, Zimmerman recognizes this reply as having some merit (Zimmerman, "TLS" 572–73). If the agent is disposed to cruelty, then it seems that if he were in what he believed to be *s*, for *some* value of *s*, he'd be more likely than not to act cruelly. That is, in at least some cases the following would be true: if P believes he is in situation *s*, he would bring it about that *e*.

A third worry about the counterfactual analysis is that all the relevant counterfactuals may turn out to be false. On the Lewis-Stalnaker *possible worlds* account of counterfactual semantics, the statement "if P believes he is in situation *s*, he would bring it about that *e*" is true if and only if at all the closest worlds that P believes he is in situation *s*, he brings it about that *e*.⁷⁹ As I claimed in Chapter 3, free will requires that

⁷⁹ There are a broad spectrum of view that fall under that heading and each individual interpretation has its own detractors. See Jonathan Bennett's *A Philosophical Guide to Conditionals* (Bennett, Conditionals) for a detailed discussion of this family of views.

the agent have a genuinely open future, a garden of forking paths, and that it be up to the agent which path is followed. One way to state this is to say that freedom requires that the agent have multiple live options and that which live option is acted upon is, in a fundamental sense, up to him. Suppose that at all the closest possible worlds, however, the agent brings it about that *e*. If that is so, then it doesn't seem that the agent has multiple live options. So the truth of the libertarian thesis requires that such counterfactuals be false.

Zimmerman's worry can thus be recast in two parts: either (a) the Counterfactual Principle violates the principle that free will requires that the agent have genuine options or (b) the Counterfactual Principle is false since there is no one decision that P would have made at all closest believed-to-be-*s* worlds⁸⁰.

Zimmerman recognizes the worries about losing the connection between free will and moral responsibility, if the noncollaborator is blameworthy because he is such that he would collaborate if placed in the right situation. He says (Zimmerman, "Luck and MR" 228),

Many seem to have supposed that if P is morally responsible for *e*'s occurring, then *e* was either a free action or a consequence of a free action of P's. But the noncollaborator's being such that he would have made the decision to collaborate under the conditions specified was clearly not an action of his and might very well not have been a consequence of an action of his; and, certainly, [the Counterfactual Principle] does not require that this characteristic of the noncollaborator have been either of these things in order for him to be as much to

⁸⁰ Nothing crucial hangs on the Lewis-Stalnaker account being correct. I use it simply because belief in it is something I appear to share with Zimmerman. The point is that on any semantics in which "if P believes he is in situation *s*, then he will bring it about that *e*" is true, it's false if free will requires the sorts of live options I consider it to require.

blame as the collaborator. Is there no essential link between freedom and moral responsibility?

That is, the Counterfactual Principle could be true and the noncollaborator morally responsible for his character traits even if he had never had any voluntary control over anything.

Zimmerman hopes to remedy this problem by proposing a modification to the Counterfactual Principle in which the term *made* is replaced with the phrase *freely made*. Thus, we have Counterfactual Principle*:

- If (i) P freely made *d* in what he believed to be *s*,
 (ii) P* would have freely made *d* if he had been in a situation that he believed to be *s*, and
 (iii) P*'s being in a situation that he believed to be *s* was not in his restricted control,
 then whatever moral credit or discredit accrues to P for making *d* also accrues to P*

The addition of the adverb *freely* doesn't help here. It doesn't meet the counterfactual worries. At best, Control Principle* has no truth value or has no truth value that we can discern. At worst, it is false. That is, given a fairly standard interpretation of counterfactual semantics, it's false to say that if P believes himself to be in an *s*-world, he would freely bring it about that *e*. The issue is that for P to freely bring it about that *e* he has to have the genuine options to both bring it about that *e* and to not bring it about that *e*, and that which option is realized is fundamentally up to him.

If bringing it about that *e* is a live option for the agent and so is not bringing it about that *e*, then we can imagine the following which shows Counterfactual Principle* to be false. We place the agent in a situation just *before* she chooses whether to bring it about that *e*. We let that situation unfold, and she chooses, say, to bring it about that *e*. If

her choice is free, then we could rewind the scenario again and again and sometimes she'd choose to bring it about that *e*, and sometimes she'd choose otherwise.

Why do I insist that the agent *couldn't* always *freely* choose to bring it about that *e*? Suppose that the agent is *very* hungry, and she's presented with a plate of live bugs to eat. If it's been days since she's eaten, with no other meal that she can anticipate, eating the bugs will be a live option. Now, suppose that she's presented instead with a plate of waffles (which she loves) and the plate of live bugs. I predict that she would eat the waffles no matter how many "rewinds" we would care to test. That result, I think, shows that while the plate of bugs was a live option when the alternative was starvation, it's not a live option when a plate of waffles is also available. If one option so dominates another that the one is *always* chosen at all the nearest possible worlds, I claim that we have one live option rather than two. Hence, it's false that the one live option is freely chosen.

Indeed, over a large number of "rewinds," the proportion of times she chooses to bring it about that *e* as compared to the times she chooses to not bring it about that *e* would mirror the strength of each option. Even so, on no particular rewind would we be able to predict which way she would choose. That is, given the initial scenario, there's no option that she'd *freely and consistently* choose. Thus, the Counterfactual Principle* is unable to neutralize luck, the objection goes, since it is false.

The worry can be staved off, of course, by adopting a chancy or probabilistic account of causation. In presenting the causal-contrastive account of luck in Chapter 2, I've already adapted such a view, and we can briefly revisit it here.

On a deterministic understanding of causation, if event X causes event Y, then the occurrence of X is sufficient for the occurrence of Y. In other words, given the laws of nature at the actual world and a sequence of past causes that include X, Y must occur.

According to a probabilistic account of causation, on the other hand, causal connections are not absolute but vary in strength. Thus, if event X causes event Y, then

the occurrence of X is probabilistically but not absolutely correlated with the occurrence of Y . That is, given the laws of nature at the actual world and a sequence of past causes that include X , Y is likely to occur. The more likely it is that Y follows X , the more strong the causal connection between X and Y .

Given this understanding of probabilistic causation, we can propose an amendment to Zimmerman's principle, which we can call Proportional Counterfactual Principle (PCP):

- If (i) P freely made d in what he believed to be s ,
- (ii) P^* would likely have freely made d if he had be in a situation that he believed to be s , and
- (iii) P^* 's being in a situation that he believed to be s was not in his restricted control,
- then whatever moral credit or discredit accrues to P for making d accrues to P^* in proportion to the likelihood that P^* freely makes d .

The immediate advantage of PCP is that we don't need to hold that there is some e such that at all of the closest worlds if P believes he is in situation s , he would bring it about that e .

A second important advantage of PCP is that it allows us to more sensitively judge the moral responsibility due to agents. Zimmerman's previous suggestion has been that the correct luck-neutralizing principle puts both the collaborator and the noncollaborator, for example, on exactly the same footing. They are each equally responsible. On PCP, however, while the collaborator is morally responsible to degree X , for example, the non-collaborator is responsible to perhaps some greater or lesser degree.

Additionally, not all noncollaborators are created equal, and if two collaborators differ in the likelihood with which they would have become collaborators, then their relative degrees of blameworthiness ought to track those differences.

It's not clear that Zimmerman would agree with me on PCP. He briefly considers the case in which Would Be Assassin fails to kill where Successful Assassin does not, but merely through the intervention of luck. He says, "Suppose that there is a probability of .99 that [Would Be Assassin] would have freely killed [Victim], had he not sneezed. Then one of two things follows: either [Would Be Assassin] is 99% as responsible as [Successful Assassin], or there is a 99% chance that [Would Be Assassin] is as responsible as [Successful Assassin]" (Zimmerman, "TLS" 573). While Zimmerman admits that the case isn't decisive, he opts for the latter on the following analogy. Suppose two mounds of white powder are before me, one being sugar and the other being sugar or salt. If there's a 0.99 probability of the latter being sugar, it's more cogent to say that it's 99% as likely to be sweet as the other, rather than it's 99% as sweet as the other (Zimmerman, "TLS" 573n).

Suppose that we have two agents, A and B, where both are cruel, but A is twice as cruel as B. There are several ways for B to be significantly less cruel than A, and for none of which is the sugar/salt analogy appropriate.⁸¹

First, B may be, on his worst behavior, willing to do things that are less cruel than A is willing to do; A will drown kittens but B won't. Perhaps B will walk by their empty water bowl on a hot day and hope that their thirst is particularly painful. Here is an apt analogy. Suppose that I've got two kinds of jeans, where the pairs of one kind are twice as durable as the pairs of the other kind. If I wear and wash them all regularly, after a year, say, I don't expect twice as many of the non-durable jeans to be worn out as compared to the durable jeans. Rather, I expect them to be twice as worn, where perhaps none of either kind are worn out.

⁸¹ The case of Would Be Assassin and Successful Assassin is a case of resultant luck, and the sugar/salt analogy appropriately applies there, but it cannot be extended, as I will show, to constitutive and circumstantial luck.

Second, B may sometimes be as cruel as A is, but he's more conflicted about his cruel sprees, and may not have the nonchalance about causing kitten deaths that A has. The analogy here isn't one pile of sugar in comparison to another pile of white granules that is either purely sugar or purely salt. Rather, it's a pile of sugar in comparison to a pile of granules that's some mix of salt and sugar. As such, we're comparing agents who have similar but not identical character traits, and the difference isn't that A and B differ in how cruel they are, but in which other character traits are mixed with the cruelty and affect how the cruelty manifests. We can thus set this case aside since it doesn't fall directly under PCP.

Third, B may be less frequently willing to be as cruel as A often is willing to be, though on his worst behavior B can be as cruel as A. He'll drown fewer kittens, as it were. A and B have relevantly similar character traits, and the difference is in the relative robustness of their respective cruel streaks. If a particular set of circumstances has A acting cruelly, and B, through luck, avoids a similar set of circumstances, how ought we to judge B? According to PCP, we ought to judge him proportionally to A in so far as his cruel streak compares in proportion to A's cruel streak. An appropriate analogy is the following. Suppose that we have two glasses of water in front of us, one of which has a teaspoon of salt in it. The other is more or less salty than the first, depending on whether it has more than or less than a teaspoon of salt in it. This seems correct and the analogy gives some legitimacy to PCP.

6. Does PCP Neutralize Luck?

The question at hand is whether PCP neutralizes luck. In answering this question, I want to exploit a distinction already made by Zimmerman, and I want to add a parallel distinction of my own.

Zimmerman distinguishes between restricted control and unrestricted control. The concept of unrestricted control is an absolute notion. That is, for an agent to have unrestricted control over an event he has restricted control over whether it occurs and over every event upon which the event is contingent. On the other hand, the agent has restricted control over an event if and only if he has control over whether it occurs or not. An agent, however, may have restricted control in the very minimal sense that he has one avenue of access in which to control the event. Suppose, for example, that Rich has control over whether a certain button is pushed, but circumstances are such that he can push the button (a) only with his right index finger and (b) only between 12:00 and 12:01 on a given day. Rich has, in this case, about as little control as one could have, without narrowing his window of opportunity to some small slice of a minute.

In contrast to Rich's circumstances, let's suppose that Jeremy also has restricted control over pushing a given button, but he has many more options. In Jeremy's case, for example, he may have access to the button for many years, and he may have many ways in which he could push the button. In addition to pushing it with many parts of his body (fingers, nose, elbows, etc.), he has various implements by which he could achieve the button pushing: sticks, pencils, icicles, and a number of entertaining and inventive Rube Goldberg machines. Furthermore, let's suppose that Jeremy has several minions who will push or not push the button at his request. Finally, let's suppose that Jeremy is aware of all the ways in which he controls whether the button is pushed.

Jeremy has much more control over whether his button gets pushed than Rich has over his. Both Jeremy's control and Rich's control, however, are instances of restricted control on Zimmerman's account. Relative to every agent and every event, Zimmerman's account has just four buckets into which the event can fall for the agent: it is (a) under his unrestricted control, (b) not under his unrestricted control, (c) under his restricted control, or (d) not under his restricted control. On my view however, for a given agent

and a given event, it may be under his unrestricted control (he has control of level 1.0), not under his restricted control (he has control of level 0.0), or under his restricted control on a scalar level (he has control greater than level 0.0 and less than level 1.0). Thus, whether an event is controlled by a given agent is a scalar notion.

I want to develop a similarly scalar notion of free will. It is based on the notion that everything that can be freely chosen by an agent must be a live option for her. Thus, a being would have unrestricted freedom within a situation if and only if every logical possible continuation of the situation was a live option for her. A being would have restricted freedom within a situation if and only if at least two possible continuations were live options for her.⁸² On my account, the more live options one has--that is, the more continuations one has which are to some extent desirable to her--the more freedom one has in the situation.

Clearly, no finite being has unrestricted freedom within any situation. Presumably, no omnibenevolent being has unrestricted freedom either, since some possible continuations will be too horrific to be desirable and hence won't count as live options. On the other hand, an omniscient and omnipotent being who was morally apathetic could have unrestricted freedom. For finite beings, however, we'll have in any situation a degree of freedom from level 0.0 to some level below 1.0. Some finite beings will count as having more freedom than others in a given situation, given differences in what they think is possible and desirable, that is, given the extent of their live options.

Note, however, that if an agent is free to level x given the live options that she has, she need not have freely chosen which of those options are live options. Thus, two agents will count as equally free within a given situation if they have the same

⁸² Even the most minimal restricted control has two options: the option of bringing the event about and the option of refraining to bring it about.

continuation sequences as live options, no matter whether those live options were given at birth, inculcated through brainwashing, the results of prior free choices, etc.

This is an important consequence of my theory of scalar freedom, however, because it's what ultimately allows us to block Galen Strawson's Basic Argument against moral responsibility. The thrust of his argument is that if I do not initially have moral responsibility, I cannot somehow "bootstrap" myself into being morally responsible. Thus, if a finite agent begins his existence (as a fetus, for example) without moral responsibility, there is no way on Strawson's view that he can become morally responsible. On my view, one does not need to be initially morally responsible in order to make decisions for which one is morally responsible. One need only face a palette of live options and choose between them.

If we allow that freedom can vary according to the live options available to the agent within a situation, however, we ought to allow that moral responsibility can also vary according to the freedom available to the agent within the situation. Let us consider a range of cases. Suppose agents A, B, and C face option arrays from among options i, ii, and iii, where i is clearly morally better than ii, which is moderately morally better than iii. Thus, i is even more clearly morally preferable to iii. A has all three options on his palette, B has options i and iii, and C has ii and iii. A deserves praise if he chooses i and blame if he chooses ii or iii. B will deserve praise for choosing ii and blame for choosing iii, but not as much as A would for choosing iii since the loss of moral value by B choosing iii is comparatively small. C, of course, ought to choose i, and would deserve blame for choosing iii. Since A has an intermediate value, ii, to consider, it may blunt his perception of the gulf between i and iii, which for C may appear rather stark. If that is the case, then it would appear that A would be *less* blameworthy than C for choosing iii. A fourth agent, D, might have the same live options as A with the same ordering, but with different spacing. That is, i may be very vivid for D and ii and iii considerably less so.

Thus, if both A and D choose ii, it will reflect more negatively on D's moral record than it does on A's moral record.⁸³

On my understanding, therefore, what's required for freedom and hence moral responsibility are live options. Since luck can affect what options are live to an agent in a given situation, not to mention how vividly each option appears to her, luck can affect an individual's moral responsibility.

If we consider a case of constitutive luck, such as two people placed in a situation in which courage is the proper response, they may differ in the amount of praise or censure due them, if through luck taking the brave course of action appears more live to one agent rather than the other. Of course, two people may differ in those respects with no amount of the difference due to luck, for example, if they've consciously developed the constitutions they have. Thus, genuine moral luck occurs when the differences in the live options they see before them (or the vividness of the options) is due to factors beyond their control. If they've had opportunities to develop better character traits, but did not, that would diminish the amount of moral responsibility that is due to luck. As people go through their lives, therefore, the amount of genuine moral luck would tend to diminish, rather than to increase. PCP, therefore, does not neutralize luck.

Two questions remain. First, what are the influences on whether luck is morally significant? Second, how well does the existence of genuine moral luck sit with the Control Principle? The answers to those questions are addressed in the following section.

7. Understanding Constitutive Moral Luck and the Control Principle

Important, therefore, to the question of whether a situation is morally lucky for an agent is whether the agent had a palette of live options and the vividness of the options on

⁸³ Such scenarios would be excellent fodder for experimental philosophy.

that palette. If the agent sees only one live option in a given situation or has multiple live options such that the vividness of one of them dominates the rest (in which case the other options may not truly be live), then there is not much room for moral responsibility to wedge itself in, unless it did so at an earlier time. He may have, for example, consciously chosen to promote (or stifle) certain kinds of options within himself. A person who has fully developed the character trait of honesty will not even consider being dishonest (or it will be the smallest flicker) in a situation in which dishonesty is possible. He will nonetheless be praiseworthy since his prior choices have added his honest character trait to his moral ledger.

Similar remarks apply also to the person who is compulsively honest, through no choice of her own. Each will act honestly in a given situation, and yet only the person who developed the character trait will count as morally honest, and only the one who developed the trait of honesty will be praiseworthy. Thus, if Claire is honest to a fault (as the saying goes), giving us ample evidence that her honesty is a well ingrained character trait, we cannot know from that fact whether she is morally commendable.

A person's dispositions are not irrelevant to how we ought to judge him. It seems plausible that someone is worse for being inclined to rape (even if he successfully resists his inclinations) than if that behavior is never a live option for him. When someone has dispositions that actively inclines him toward rape, perhaps what we are judging is whether he knows he has the psychological disposition and is not working vigorously to rid himself of it. Someone who sits with the disposition risks acting on it. The same reason exists, therefore, to judge him harshly as it is to judge the drunk driver: both behaviors put others at grave risk.

On one standard reading of Kant, he famously judges the matters differently. Someone who is not inclined to rape is less morally worthy than one who is strongly inclined to but chooses not to. Someone who is abhorred by the thought of rape will find

it easy to act *in accordance with* the moral law. On Kant's view, the would-be rapist who struggles against his compulsion is acting *from* the moral law, which is what gives his action moral significance. If an agent has only morally correct dispositions, or at least lacks morally abominable dispositions like the disposition to rape, then acting in accordance with the moral law will be the default mode for the individual and will not require that he or she *embrace* the moral law, which is what Kant counts as morally significant. To embrace the moral law is, on Kant's view, to actively throw one's will behind it, rather than to have one's will be so formed as to leave no choices for the agent.

Thus, on the Kantian view--and this seems correct to me--the most significant moral choices we can make are the most genuine choices. Consider a case, for example, where we see the right thing to do, and we see a perhaps personally advantageous option in conflict with the right thing to do such that the balance of considerations does not tip one way or the other.⁸⁴ In such a situation, we must choose whether to put our wills behind the morally correct live option or the personally advantageous one. That choice presents us with the ultimate moral worth about which Galen Strawson worries.

Consider an agent that has a certain constitution which is out of his control. A different agent, with a different constitution, would have a different array of live options present in similar circumstances. Hence, their moral records can differ for reasons beyond their control, and we have an instance of genuine moral luck.⁸⁵

Finally, there is the question of how well the existence of genuine moral luck sits with the Control Principle. As we saw from Chapter Four, the Control Principle can be stated as follows:

⁸⁴ The inclination need not be exactly balanced in such a case, but if so, makes for the most vivid example of genuine moral luck.

⁸⁵ I do not see how we could have genuine resultant moral luck or genuine circumstantial moral luck. Thus, the only kind of genuine moral luck is constitutive.

(CP) We are morally assessable only to the extent that what we are assessed for depends on factors under our control.

The existence of genuine moral luck does not violate CP. Suppose that Geneva has as live options both telling the truth in an important situation and telling a lie to save face. Let's also suppose that her inclinations toward each are well balanced and that the vividness of each option is out of her control--suppose one parent tried to raise her to be scrupulously honest and the other to be connivingly self-serving. In such a case, it falls beyond her control what her palette of live options contains and how vivid each option is. Suppose that she decides to tell the truth. Does CP therefore bar us from morally assessing her as praiseworthy? No, it does not. We are not assessing her for the aspects of the situation that are not under her control. That is, we are not praising her for having the impulse to tell the truth, nor are we blaming her for wanting to tell a self-serving lie. Those are, as the story was told, beyond her control. Rather, what we can legitimately praise or blame her for is which of those live options she embraced. That is, this is a case of genuine moral luck, and yet there has been no violation of the Control Principle.

Sometimes we assess a character trait as good or bad *per se* without regard to whether the agent had any contribution to or control over whether she had that character trait (Enoch and Marmor 426). Since these judgments strike us as not misguided, they seem to present a different kind of genuine moral luck than I've been endorsing since we praise the brave and blame the cowardly for how they are, irrespective of whether the trait was consciously acquired or merely inborn.

I claim that this is not a species of genuine moral luck since the judgment merely assumes that the agent's will was engaged in acquiring the character trait. That is, if I see someone who exhibits a brave character I would be inclined to praise him, and yet, if I should learn that his having the character trait had nothing to do with his striving to be brave, I might still point his behavior out for emulation, but I would want to withdraw my

praise of the agent (or I might praise him for consequentialist reasons). That is, while I would still admire and encourage the agent, I would cease to see the quality as the agent's in any morally significant way. It would be for the agent as much a moral assessment of him as would be his genetically acquired good looks--that is, none at all.

Let me address one final worry. If Geneva's strongly inclined to tell the truth and weakly inclined to lie, then she may well tell the truth only because of the strength of her inclinations, and hence deserve no praise. On the other hand, she may tell the truth with its moral importance fully in view. She tells the truth not merely because it's the stronger impulse, but with such conviction that even if it were not the strongest impulse, she'd still opt for the truth. This way of choosing is praiseworthy.

Is Geneva able to choose among her impulses or does she merely follow the current of her inclinations? If the former then she's got a level of control and can be morally responsible. In these cases, there can be genuine moral luck. If she's merely swept along by her inclinations, and cannot *not* be, then she lacks the control necessary for moral responsibility. In such cases, she experiences no genuine moral luck.

8. The Problem of Moral Luck Revisited

Let us consider the two arguments presented thus far which I've called the Basic Problem of Moral Luck and the Moderate Problem of Moral Luck. Let us examine them in turn.

The Basic Problem of Moral Luck:

1. Agents are sometimes morally responsible for their actions, intentions, and character traits.
2. No action, intention, or character trait is ever within any agent's control.
3. Thus, agents are sometimes morally responsible for actions, intentions and character traits that are *not* under their control.

The Basic Problem loses traction once we reflect on the nature of free choice as I've presented it. In order for a choice to be free, the agent needs to have a palette of live options available to her. In choosing one of the options, she forms an intention to act on it. Thus, the second premise of the Basic Problem is false. It may not be under her control what array of live options she has, but given that she has multiple live options, it can be under her control which cause she promotes. Let's now turn to the Moderate Problem.

The Moderate Problem of Moral Luck

1. Agents are often morally responsible for their actions, intentions, and character traits.
2. Most actions, intentions, and character traits are not within any agent's control.
3. Thus, agents are sometimes morally responsible for actions, intentions, and character traits that are not under their control.

This problem can remain, even if the Basic Argument is deemed not to be sound. It may be that the intentional structure that supports free will and moral responsibility is too thin to scaffold a robust moral theory. Let me paint a picture of moral development for you.

If we consider the development of an agent, we see a being who begins with an array of characteristics, none of which are attributable to her choices. For a time, her actions flow either from those character traits, from merely aping those around her, or from her testing the waters of intentional action. If we consider her from that point in her life, before she begins acting as a moral agent, the first premise is false. She's making some inroads to intentional action, but her motivations and considerations will be amoral at this early stage of development.

Gradually, however, she begins to embrace certain character traits and to act in certain ways from moral motivations. She is told that certain things are right or wrong, and she begins to internalize those precepts. This gives stronger support to suppose that

the first premise is true, but as it does, we have reason to lose confidence in the second premise, especially since we are concerned with restricted rather than unrestricted control. If we judge agents from their current palette of live options, it's false that most of their lives are out of their control. Thus, depending on the developmental stage of the agent, we have different reasons to doubt the soundness of the argument.

The view I adopt of the human agent (I have no theory of divine or angelic agents) respects both aspects of our nature that Nagel emphasizes. We begin life deeply controlled by the influences that are not under our control. Eventually, however, we come to have an array of live options that we face as we go through our day-to-day situations. Sometimes the considerations will be balanced, and we'll have to choose to actively endorse a certain option, or we can leave it to a coin toss, or to the vagaries of our whimsical natures. At other times, however, one option or the other will have more vividness, more life, to it. In that case, we *can* just let nature take its course and allow it to dominate our other options. If in a given situation I am both inclined to tell the truth and inclined to lie, I can just let the stronger of the impulse carry me away. I'll have character traits, just not one's that I've developed or chosen.

To rest in this moral sloth, however, isn't to *not* have a choice between the live options. Rather, it's to exercise the choice to be slothful. Even though the stronger impulse drives my behavior, it still, as Hume would have it, inclines without necessitating. That is, the only reason the stronger inclination is the one acted upon is that I've chosen not to take a more commanding control of what inclines me. In just accepting and following my inclinations as they are given to me, it's not that I can't control my actions, intentions, and character traits. Rather, it's that I choose not to.

On the other hand, I may reject my moral slothfulness and reflect on which of my impulses *should* drive my behavior. In that case, however, I am not taking my current array of inclinations as given, but choosing to promote some inclination or character trait

over the others. Still, there is no moderate problem of moral luck since the second premise is false.

Chapter 6

Concluding Thoughts

1. What We've Covered

According to the nomenclature I've followed, an agent is subject to moral luck when she can correctly be subjected to moral judgment even though luck has contributed to that for which we are judging her. Contrast this with a Control Principle (CP) which says that we are only subject to moral assessment for what is under our control. If we combine those ideas, we can generate a significant problem for moral philosophy.

The Basic Problem of Moral Luck:

1. Agents are sometimes morally responsible for their actions, intentions, and character traits.
2. No action, intention, or character trait is ever within any agent's control.
3. Thus, agents are sometimes morally responsible for actions, intentions and character traits that are *not* under their control.

The *problem of moral luck* is that deliberation leaves us in the position of enjoining each member of an inconsistent triad of propositions: (1), (2), and CP. Thus, we are compelled to choose among three individually unpalatable options: deny CP, deny that luck is ubiquitous, or embrace moral skepticism. CP, however, gets to the heart of many people's intuitions about moral responsibility, and if moral responsibility loses its moorings, moral theory would be a rudderless ship; our lives are shot through with luck; and moral skepticism is generally taken to be repugnant, and many philosophers think that moral skepticism ought to be the refuge of last resort.

In the previous chapters, I've offered my understanding of the nature of luck and of moral responsibility, and I have staked my claim to an answer to the question of whether there is genuine moral luck. Given these findings, I've been able to chart a

course according to which moral luck is compatible with a robust moral theory, something which seems impossible given the Basic Argument. I deny the second premise, at least as regards to our morally significant character traits, and I affirm that there is genuine moral luck. Let me briefly highlight my central findings.

1.1. The Nature of Luck

I argue that luck is a causal-contrastive notion, one which relies on indeterminacy in causal chains that connect events such that it is possible for an event to turn out better or worse than it actually does. Thus, luck is contrastive along two axes. First, luck is determined in part by how likely various outcomes are, given an initial starting point. Luck will figure differently, depending on the end points for the causal chain we are considering. For example, the magnitude of luck required to draw a second ace in poker is greater if there is only one more card to be drawn, compared to if there are three. Second, luck is determined in part by how the agent is affected by one outcome as compared to how he might have been affected by an alternative outcome. If an agent breaks his leg, that may be lucky if the situation could have killed him and unlucky if he might have gotten off without injury. Thus, I maintain that there is no luck *simpliciter*, and all luck is comparative.

Since luck is a contrastive or comparative notion, moral luck is luck in how an agent's actual moral status compares with what it might have been, were it not for factors out of her control.

It is important to use a contrastive notion of luck in discussing moral luck for the following reason. I claim that the only genuine moral luck comes from how the agent freely chooses from the palette of live options present from her given character traits. Her luck is seen by comparing her actually given constitution to the given constitution she might have had. The causal-contrastive analysis is able to clarify such contrasts.

1.2. Moral Responsibility

In Chapter 3, I argue that moral responsibility follows from a libertarian understanding of free will such that an agent is free if and only if she has a range of live options from which to choose--a garden of forking paths-- and the ability to choose among them.

Libertarian free will follows on the possession of certain intrinsic properties: (1) certain cognitive capacities, (2) a range of emotional capacities, (3) the capacity to be motivated by other-regarding considerations, (4) a structure of the will that includes first-order and at least second-order desires (that is, to have certain desires and to be able to desire to be the sort of person who has certain desires), and (5) a robust capacity to choose between various courses of action.

For my purposes, the fifth capacity is the most important. My account of moral luck requires that in addition to having a palette of given character traits, the agent has the ability to choose which of those character traits to endorse. That is, an agent can be born with a given range of live options that influence her behavior, and yet can choose which to adopt not merely as her given constitution but also as her chosen constitution. This is important because it allows for the agent to be born without moral responsibility and yet to come to be morally responsible.

1.3. The Problem of Moral Luck

Chapter 4 discusses the problem of moral luck as it was first presented by Bernard Williams and Thomas Nagel in a pair of papers presented to the Aristotelian Society (Williams; Nagel). In that chapter, I confine my discussion to those two papers.

Williams' paper presents two interesting ideas which he discusses in relation to the example of a fictionalized Gauguin. Gauguin abandons his family to go to Tahiti to paint the native scenes. Williams says that there are two ways in which Gauguin's project

of becoming a significant artist can fail. First, his project can fail because of some *external* misfortune, such as being injured on the voyage. In that scenario, while his *project* has failed, *he* has not failed. Second, his project can fail because of some *internal* reason, for example, because he doesn't have the artistic talent to produce good art. In that case, not only has the project failed, but *he* has failed.

Williams makes two interesting points from this example. First, he says that there is no possibility of Gauguin justifying his decision before he departs. That is, Williams rejects a standard conception of rational justification according to which if one's decision can be justified at all, it can be justified before the fact, *prospectively* rather than *retrospectively*.

The second point that Williams makes has to do with whether it is rational for Gauguin to *regret* his decision if it turns out well for him. Williams says that it is not and that the question of whether Gauguin can more basically regret the decision if it turns out well cannot even arise.

I argue that Williams is mistaken on all points. First, I attack him on his notion of retrospective justification. To this end, I consider a number of variations on a single theme, betting under uncertainty in the poker game of Texas Hold'em.

Second, I deny that the player who makes an irrational decision cannot more basically regret that decision. That is, while he may have his winnings to celebrate, he ought to take his winning as undeserved, since his action was prospectively irrational. In other words, success will never make an irrational decision into a rational one, and there is no retrospective *rational* justification.

Let me now turn to a discussion of Thomas Nagel's contribution to the debate. Nagel distinguishes four specific types of luck:

Resultant Luck: Luck in how things turn out;

Circumstantial Luck: Luck in the circumstances in which the agent finds herself;

Constitutive Luck: Luck in the sort of person the agent is and in her natural and developed talents, her dispositions to behave, her character; and

Causal Luck: Luck in how the agent fits into the causal order.

Nagel's categories divide nicely between luck affecting the downstream causal consequences of an agent's action (resultant luck in Nagel's terms) and causal luck which affects the upstream influences of an agent's actions. The upstream causal influences further divide into (1) which causal paths upon which the agent finds herself (circumstantial luck) and (2) what given dispositions or character traits she brings to those situations (constitutive luck).

The power of Nagel's contribution is that he frames the discussion in terms of two conceptions of human nature. First, Nagel recognizes that we are parts of the natural order. As such, upstream contributors to our actions and the downstream effects are subject to luck. On the other hand, as moral agents, our moral worth ought to be insulated from the influences of luck. Thus, Nagel says,

But as the external determinants of what someone has done are gradually exposed, in their effect on consequences, character, and choice itself, it becomes gradually clear that actions are events and people things. Eventually nothing remains which can be ascribed to the responsible self, and we are left with nothing but a portion of the larger sequence of events, which can be deplored or celebrated, but not blamed or praised (Nagel 68).

Thus, the problem of moral luck is that as we come to appreciate the ubiquitous influence of luck in our lives, we find little or nothing about ourselves for which we can be truly responsible. Thus, the problem of moral luck raises for Nagel the specter of moral skepticism. To this end, Nagel says,

The area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point. Everything seems to result

from the combined influence of factors, antecedent and posterior to action, that are not within the agent's control. Since he cannot be responsible for them, he cannot be responsible for their results (Nagel 66).

I do not argue against Nagel in any substantive way, but I take his conception of the problem seriously.

1.4. Solutions to the Problem of Moral Luck

I consider two solutions to the problem of moral luck, one by Norvin Richards and another by Michael Zimmerman (Richards; Zimmerman, "Luck and MR"; Zimmerman, "TLS"; Zimmerman, "Map"). I use Zimmerman's solution as a springboard for offering my own solution.

In his paper "Luck and Desert," Norvin Richards claims that the apparent moral difference between two agents which we attribute to luck really reflects an epistemic difference in what we know about them. One of the agents is unlucky to have been shown to be a scoundrel, and the other has escaped our detection.

It's easy to imagine how this might be the case. One drunk driver kills a child, and her risk to others is exposed. Another who perhaps posed the same risk gets home safely, and we are none the wiser. I also explain how Richards considers instances of constitutive luck and circumstantial luck and how they owe their initial plausibility as instances of moral luck to our lack of epistemic access to the reality of things.

I think that Richards is correct in his diagnosis of the ways in which our limited epistemic access can make it *appear* that luck has made a moral difference without that being the case. What is lacking from his account, however, is the claim that there is *only* luck of the epistemic variety.

I then discuss a line of thought in Zimmerman. Zimmerman's initial argument is against moral skepticism. He distinguishes between what he calls *restricted control* and

what he calls *unrestricted control*. An agent has *restricted control* over an event if and only if he is able to bring about its occurrence or prevent its occurrence. An agent has *unrestricted control* over an event if and only if he has restricted control over both it and every causal factor on which its occurrence depends (Zimmerman, "Luck and MR" 219).

Given that distinction, he's able to cleave a popular argument for moral skepticism into two component arguments, one in terms of restricted control and the other in terms of unrestricted control. Once those arguments are distinguished, Zimmerman claims, it is easy to reject them as unsound and hence to not accept their common conclusion, that no agent is ever morally responsible.

In examining whether to reject genuine moral luck, Zimmerman considers various principles which might, as he says, *neutralize* luck. He finds good reason, as I do, for rejecting resultant luck as being morally significant. The reason is that there's good reason to evaluate the intention behind the act rather than the downstream consequences of the act. Thus, Zimmerman's real goal is to find a principle which will neutralize constitutive and circumstantial luck, which he calls by the general term *situational luck*. Here is a Counterfactual Principle that he considers:

Where P and P* are any people, *d* is any decision, and *s* is any situation:

If (i) P made *d* in what he believed to be *s*,

(ii) P* would have made *d* if he had be in a situation that he believed to be *s*,

and

(iii) P*'s being in a situation that he believed to be *s* was not in his restricted control,

then whatever moral credit or discredit accrues to P for making *d* also accrues to P* (Zimmerman, "Luck and MR" 225).

There are a number of problems with this principle. The most serious objection is that the principle is just false. Consider the standard counterfactual analysis of the following

statement: if P* had been in a situation he believed to be *s*, he would have made *d*. What that means is that at all believed-by-P*-to-be-*s* worlds, P* makes *d*. If that is the case, however, P* lacks robust free will at such worlds.

Zimmerman considers a second version, Counterfactual Principle*, which focuses not on the decisions that agents make but on the ones they *freely make*. This version is also problematic.

In response to these problems, I present a variation on Zimmerman's idea that I believe is defensible. First, I consider what can be called a chancy or probabilistic account of causation, in which it can be true that event X causes event Y even if not all X worlds are Y worlds, but that the proportion of X and Y worlds is large compared to the X and not-Y worlds. (This is the notion of causation I appeal to in my explanation of the nature of luck.) Given this understanding of probabilistic causation, we can propose an amendment to Zimmerman's principle, which I call the Proportional Counterfactual Principle (PCP):

- If (i) P freely made *d* in what he believed to be *s*,
- (ii) P* would likely have freely made *d* if he had been in a situation that he believed to be *s*, and
- (iii) P*'s being in a situation that he believed to be *s* was not in his restricted control,
- then whatever moral credit or discredit accrues to P for making *d* accrues to P* in proportion to the likelihood that P* freely makes *d*.

While PCP is a true principle, it does not neutralize luck. Let us consider an agent who has a palette of live options on which he can act, but the number, variety, and relative strengths of those options are merely given to him and involve no choice of his own. It's out of the agent's control that he has that palette of choices and not another. Yet, no matter what his palette of choices or their origin, it's up to him to choose. It's

plausible that the fair way in which to dole out moral responsibility is in proportion to the obstacles that might bar him from doing the right thing. If his given constitution is vicious (either because of genetics or his upbringing), then he'll have a hard time choosing to do the right thing. On the other hand, if his given character is virtuous, then it will be easy for him to do the right thing (though still possible for him to act on an immoral option). Thus, luck has made a difference to his moral record.

Finally, in that chapter, I sketch an account of moral development. I focus on how agents start without moral responsibility and can later come to acquire it through choosing to embrace or endorse some of their given character traits and to repudiate others. They are lucky in what palette of live options they are given, and morally responsible for their choices within that range of options.

2. Final Thoughts

The question that's loomed largest as I've written is this: is what I've done sufficient to stave off moral skepticism? That's an interesting question, because the way in which I pitched the problem of moral luck from the beginning was that unless one found a way of combating the problem of moral luck, being forced toward moral skepticism was a possibility. In particular, I view Galen Strawson's *Basic Argument* as very powerful. It's powerful, in part, because it's so stark. We're born without being morally responsible. If we're born without moral responsibility, there's no way for us to acquire it. Hence, we can never *become* morally responsible, and no one is ever morally responsible. My response has been that free will merely requires that we've got the power to choose between live options, no matter where those live options come from.

As an illustration, let's consider a standard case of teleportation. Suppose that instead of taking a plane to defend my dissertation, I opt to teleport. On the Physical Continuity Criterion for personal identity, the person, Pat*, who steps out of the teleporter

isn't me, but rather a mere replica of me, one who has all my memories⁸⁶, and more importantly for this discussion, all the same dispositions to behave. For the sake of this discussion, let's assume that this criterion for personal identity is correct. Pat* inherits all those dispositions to behave, and he's morally responsible for none of them since he's had no choice in any of them. The moral status that Pat* has when he steps out of the teleporter is the same as what Galen Strawson thinks is Pat's moral status or what we would think of as the moral status of an infant.

Suppose Pat* comes to a situation in which an impulse to steal is present, but not overriding--that is, it is but one of a palette of live options for Pat*. Let's say that he chooses to embrace honesty and so chooses to shore up one set of impulses over the other. On my account, he's morally responsible in proportion to the strengths of the various live options he overcame in order to do the morally right thing.

On my account, he's got a level of freedom that's not absolute. He could have had more freedom than he had in the situation either by having more evenly balanced live options or more live options that are no more out of balance than the live options he currently has. Thus, while I reject Strawson's argument, I pay a certain price. In particular, I have a non-standard understanding of the level of freedom that we have. On the standard view, every being is either an agent or not, and if the being is an agent, then he or she has all the freedom that he or she could have. Put another way, the standard conception of freedom--the kind that underwrites moral responsibility--is not a scalar notion, but on my conception, it is.

The final question that I'd like to raise is whether the existence of genuine moral luck is a bad thing. We can divide this question into two parts: is it bad for a moral

⁸⁶ Of course, Pat*'s doesn't really remember writing the dissertation, because he didn't, but he has mental states--quasi-memories--that are internally indistinguishable from genuine memories.

theory that it grants the existence of genuine moral luck? and is it bad for a person if she experiences an instance of moral luck?

Some think that any moral theory which accedes to moral luck is to that extent undesirable. Others think that any theory that *does not* include the possibility of genuine moral luck misses the point somehow that moral theories are meant for people like us, upon whom a plurality of moral values impinge. The very rare person might go through life with no occasions for moral luck to make a difference, but most would not.

To those who lament moral luck as undesirable, if their worry is that it conflicts with the Control Principle, my answer is that moral luck and the Control Principle are compatible. The others, the ones who think that moral luck is an obvious, and perhaps even central, fact of human existence may be less than satisfied with this work. I chose to address a more narrowly defined version of the problem, so I've missed including interesting work from primarily the Aristotelian tradition.

Let me now turn to my second question: is moral luck bad for the individual? My argument has been that the source of moral luck is in our given constitutions, how we came pre-wired or how we were conditioned before we could really choose what sort of people we desired to be. That is a bad condition to be in, not so much because of the moral luck involved, but that it's an unreflective life. We're told that the unexamined life is not worth living. To the extent that that's true, we ought to consciously develop, and not merely accept, our moral characters. Robert Louden argues, persuasively, that that is the most important task facing us (Louden). Thus, it seems to me that moral luck is bad, perhaps not in itself, but, particularly for adults, as a symptom of a deeper malady, an insufficiently reflective or morally inactive life.

Works Cited

- Aristotle. Nicomachean Ethics. The Complete Works of Aristotle, Volume Two. Ed. Jonathan Barnes. Trans. W.D. Ross, Rev. J.O. Urmson. Princeton: Princeton UP, 1984.
- . Physics. The Complete Works of Aristotle, Volume One. Ed. Jonathan Barnes. Trans. R.P. Hardie and R.K. Gaye. Princeton: Princeton UP, 1984.
- Bennett, Jonathan. The Act Itself. New York: Oxford UP, 1995.
- . "Morality and Consequences." The Tanner Lectures on Human Values, Volume II. Ed. Sterling McMurrin. Salt Lake: University of Utah Press, 1981.
- . A Philosophical Guide to Conditionals. New York: Oxford UP, 2003.
- Bishop, Michael A. and J.D. Trout. Epistemology and the Psychology of Human Judgment. New York: Oxford UP, 2004.
- Bradley, Ben. "Virtue Consequentialism." Utilitas 17.3 (Nov 2007): 282–98.
- Enoch, David and Andrei Marmor. "The Case Against Moral Luck." Law and Philosophy 26 (2007): 405–36.
- Frankfurt, Harry G. "Alternative Possibilities and Moral Responsibility." The Journal of Philosophy (1969).
- Kane, Robert. A Contemporary Introduction to Free Will. New York: Oxford UP, 2005.
- Kant, Immanuel. Groundwork of the Metaphysics of Morals. Trans. H.J. Paton. New York: Harper and Row, 1964. 1784.
- Knobe, Joshua. "Intentional Action and Side Effects in Ordinary Language." Analysis 63.279 (July 2003): 190–94.
- Lackey, Jennifer. "What Luck Is Not." Australasian Journal of Philosophy 86.No. 2 (June 2008): 255–67.
- Latus, Andrew. "Constitutive Luck." Metaphilosophy 34.4 (July 2003): 460–75.

- Levi, Don S. "What's Luck Got to Do with It?" Moral Luck. Ed. Daniel Statman. Albany, New York: State U of New York P, 1993. 109–21.
- Lewis, David. "Causation." Philosophical Papers, Volume II. New York: Oxford UP, 1986. 159–213.
- . On the Plurality of Worlds. Oxford: Blackwell, 1986.
- Louden, Robert B. Morality and Moral Theory: A Reappraisal and Reaffirmation. New York: Oxford UP, 1992.
- Nagel, Thomas. "Moral Luck." Moral Luck. Ed. Daniel Statman. Albany: State Univeristy of New York Press, 1993. 51–71.
- Nelkin, Dana K. "Moral Luck." Ed. Edward N. Zalta, 2004. Stanford Encyclopedia of Philosophy URL=<<http://plato.stanford.edu/entries/moral-luck>>.
- Nussbaum, Martha C. The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy (Updated Edition). New York: Cambridge UP, 2001. 1986.
- Pereboom, Derek. Living Without Free Will. Cambridge: Cambridge UP, 2001.
- Pritchard, Duncan. "Epistemic Luck." Journal of Philosophical Research 29 (2004): 193–222.
- Rawls, John. A Theory of Justice. Cambridge, MA: The Belknap Press of Harvard University Press, 1971.
- Rescher, Nicholas. Luck: The Brilliant Randomness of Everyday Life. New York: Farrar, Straus, and Giroux, 1995.
- Richards, Norvin. "Luck and Desert." Moral Luck. Ed. Daniel Statman. Albany: State U of New York P, 1993. 167–80.
- Ripstein, Arthur. Equality, Responsibility, and the Law. Cambridge: Cambridge UP, 1999.
- Rosebury, Brian. "Moral Responsibility and 'Moral Luck'." The Philosophical Review 1995.4 (Oct 1995): 499–524.

- Shea, Christopher. "Against Intuition." Chronicle of Higher Ed 55.6 (2008, Oct. 3 2008): B9.
- Sinnott-Armstrong, Walter. "A Contrastivist Manifesto." Social Epistemology 22.3 (July-Sept 2008): 257–70.
- Sosa, Ernest. "Experimental Philosophy and Philosophical Intuition." Philosophical Studies 132.1 (Jan 2003): 99–107.
- Sousa, Paulo and Colin Holbrook. "Folk Concepts of Intentional Action in Context of Amoral and Immoral Luck." Review of Philosophical Psychology 1 (Mar 2010): 351–71.
- Statman, Daniel, ed. Moral Luck. Albany: State U of New York P, 1993.
- . "Moral and Epistemic Luck." Ratio 4 (1991): 146–56.
- Strawson, Galen. Freedom and Belief. New York: Oxford UP, 1986.
- . "The Impossibility of Moral Responsibility." Philosophical Studies 75 (1994): 5–24.
- Strawson, P.F. "Freedom and Resentment." Free Will. Ed. Gary Watson. Oxford: Oxford UP, 1962. 59–80.
- van Inwagen, Peter. "Ability and Responsibility." Moral Responsibility. Ed. John Martin Fischer. Ithaca: Cornell UP, 1986. 153–73.
- . An Essay on Free Will. New York: Oxford UP, 1983.
- . "Free Will Remains a Mystery." Philosophical Perspectives 14 (2000): 1–19.
- . "When is the Will Free?" Philosophical Perspectives 3 (1989): 399–422.
- Williams, Bernard. "Moral Luck." Moral Luck. Ed. Daniel Statman. Albany: State U of New York P, 1993. 35–55.
- Wolf, Susan. "Asymmetrical Freedom." The Journal of Philosophy 77.3 (Mar 1980): 155–66.
- Zimmerman, Michael J. "Luck and Moral Responsibility." Moral Luck. Ed. Daniel Statman. Albany, New York: State U of New York P, 1993. 217–33.

- . "Moral Luck: A Partial Map." Canadian Journal of Philosophy 36.4 (Dec 2006): 585–608.
- . "A Plea for Accuses." American Philosophical Quarterly 34.2 (Apr 1997): 229–43.
- . "Taking Luck Seriously." Journal of Philosophy XCIX.11 (Nov 2002): 553–76.

VITA

NAME OF AUTHOR: Patrick Beach

PLACE OF BIRTH: Twin Falls, Idaho

DATE OF BIRTH: December 22, 1959

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

Boise State University, Boise, Idaho

Miami University, Oxford, Ohio

DEGREES AWARDED:

Bachelor of Arts in English with a Writing Emphasis, 1984, Boise State
University

Master of Arts in Philosophy, 1988, Miami University

PROFESSIONAL EXPERIENCE:

Graduate Assistant, Philosophy Department, Miami University, 1985-1987

Teaching Assistant, Philosophy Department, Syracuse University, 1988-1993

Adjunct Professor, LeMoyne College, 1992

Adjunct Professor, Philosophy Department, Boise State University, 1987-1988,
1994-1998, 2000, 2007-present