



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e.g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

**Luck in Crime and Punishment:
Essays in Metaphysics and Legal Theory**

Di Yang

**PhD in Philosophy
University of Edinburgh
2018**

Abstract

This thesis examines some of the legal philosophical issues that are implicated in the problem of outcome luck. In the context of criminal law, the problem asks whether we should hold agents criminally liable for the consequences of their actions given that those consequences are never wholly within anyone's control. I conclude that outcomes should matter to an agent's liability and punishment, and I make this argument indirectly by examining some of the foundational questions in legal theory.

The thesis begins by considering a current trend in some areas of philosophy. This trend involves attempts to address philosophical problems surrounding luck by doing conceptual analyses on the nature of luck. **Chapter 1** critically examines modal theories, which conceptualize luck, as well as the related concept of risk, in terms of close possible worlds rather than probabilistic likelihood. I argue that not only are modal theories uninformative, but conceptual analyses on luck are unhelpful in addressing philosophical questions surrounding luck. **Chapter 2** then returns to the traditional notion of luck as lack of control, and focuses on the relationship between luck, risk, and culpability. Some theorists argue that culpability, for any offence, is in part a function of the degree of risk the agent imposes on others. In the context of criminal law, degrees of luck and risk can both be understood in terms of degrees of control, so the suggestion that culpability is a function of the level of risk imposed (and thus of the degree of control an agent exercises) is attractive for insulating culpability judgments from luck. However, I argue that this view is mistaken because culpability is only sensitive to risk in reckless actions, but not in purposeful actions.

The problem of outcome luck may raise different questions for reckless actions and purposeful actions. **Chapter 3** looks at the *mens rea* element of criminal attempts, which is crucial for understanding the problem of luck in the context of purposeful actions. I discuss a variation of what are sometimes referred to as impossible attempts, which have helped shape current English law. I argue that the current doctrine is largely correct, and that perhaps with the exception of few paradigm sexual offences, the *mens rea* element for attempts should require a direct intention as to the consequence element of an offence, and knowledge or belief as to the circumstance element of that offence.

Chapters 4 and 5 then look at normative justifications of criminal punishment. In order to understand whether outcomes should matter for punishment, we must first understand whether and why punishment is an appropriate response to criminal offending. Here, I defend a communicative theory, where punishment is a communicative process between the offender, the political community, and the victim. What punishment communicates is the appropriate degree of censure that is warranted in response to the

offender's wrongdoing. And in doing this, it publicly recognizes the wrong that has been committed by the offender. **Chapter 4** offers a detailed explanation of this account, and argues that the political community's recognition of wrongdoing is a valuable aim of communication. **Chapter 5** then takes up a crucial challenge against communicative theories of punishment, which is that such theories fail to take crime prevention seriously. Against this criticism, I will show that general prevention can in fact be an essential part of communicative punishment. And I will show that it is specifically the political community's recognition of wrongdoing that entails punishment's preventive aims.

Acknowledgements

I am deeply indebted to my brilliant supervisor, Andrew Cornford, for his unwavering encouragement, support, and patience, from the very confused beginning of this project to the slightly less confused end. I have learned more from him than I could have hoped for from any teacher, and I thank him for showing me what it means to approach difficult questions with intellectual integrity and rigor. I am also incredibly grateful to Elinor Mason for taking on this project at a critical time. This thesis has benefited greatly from her valuable insight and criticism, and I thank her for treating every bit of it with her utmost care and attention. Many thanks are also due to Duncan Pritchard, whose work on epistemic luck has taught me much about the role of luck in law and philosophy. This project may never have taken off without his enthusiasm and support, and I am thankful for everything I have learned from him.

Like many people who work in this area, I owe a great debt to Antony Duff and his dedication to our academic community. His work has influenced my thinking on many important issues, including every chapter of this thesis, and I cannot thank him enough for his generous criticism of my work.

This thesis has also benefited from the support I received from friends and colleagues. Special thanks are due to Martin Smith, Debbie Roberts, and Katie Keltie, for their encouragement throughout; to Lee Whittington, for entertaining all of our disagreements about luck and risk; to Rosa Hardt, who has an amazing ability to add to my confusion on these topics in a very good way. Thanks are also due to Daniel Calder, Clarin Chan, Ana-Maria Cretu, Giada Frantantonio, Patrick Kaczmarek, Cristina Lafuente, Richard Stöckle-Schobel, and Matthew Scott for their discussions.

I would like to thank Joy and Christopher Sworn for their generous care during the final days of writing this thesis. And I am tremendously grateful to Ben Sworn, for his wealth of insightful comments, extensive discussions, and editorial support.

Finally, I dedicate this work to my parents as thanks for everything they have done for me.

Declaration

This thesis, including the published chapter, has been composed by me and is entirely my own work. It has not been submitted for any other degree or professional qualification.



30th August 2018
Preston, Rutland

Note on Publication

Chapter 1 has been published, in its entirety, as:

Yang, D. (forthcoming). 'What's Wrong with Modal Conceptions of Luck and Risk', *Erkenntnis*.

Contents

Abstract	1
Acknowledgements	3
Introduction	9
1. What's Wrong with Modal Conceptions of Luck and Risk	21
1. Probability within the Modal Account of Luck	21
2. Lottery Propositions and the Modal Account of Luck	29
3. What's Wrong with Modal Conceptions of Risk	33
2. Luck and Culpability	41
1. Outcome Luck and Control	42
2. Luck as Lack of (Degrees of) Control	47
3. Risk and Culpability	51
4. Outcome Luck, Risk, and Fair Impositions of Liability	59
3. Recklessness and Circumstances at the Attack- Endangerment Divide	63
1. Attacks and Endangerments	69
2. Law of Attempts as a Law of Attacks	72
3. Recklessness, Circumstances, and Attempts	77
4. From Attempts to Endangerments	82
5. Recklessness and Rape	92
6. Conclusion	94
4. Recognition and the Significance of Formal Apologies	97
1. Communicative Punishment and Formal Apologies	98
2. The Content of Recognition	107
3. Punishment and Recognition	111
4. The Significance of Formal Apologies	118
5. Censure and Prevention in Communicative Punishment	125
1. Censure and Hard Treatment	126
2. Censure and Prevention in the Justification of Punishment	133
3. Censure and Prevention in the Distribution of Punishment	146
Conclusion	153
References	157

Introduction

Somewhere in the world right now, a few people are about to leave a pub, restaurant, or party and drive themselves home after having had too much to drink. Some of these people will make it home without any incidents. At most they will be caught for impaired driving, an offence that is typically punishable with a fine and perhaps a temporary license suspension.¹ Others won't be so lucky. Unable to react in time, they will get into an accident and cause death or serious injury to other drivers or pedestrians. If convicted, they will typically receive a prison sentence.² Elsewhere in the world, a few people are plotting to partake in some kind of criminal activity. Some of these people will succeed and others will fail. If caught and convicted, the ones who succeed will most likely be given a harsher sentence than those who have failed.³

Many of us, upon first glance, may not object to discrepancies in the severity of punishment for merely reckless drivers and reckless drivers who cause harm, or between successful criminals and failed attempters. Perhaps we think that those who have caused harm—be it reckless agents or purposeful actors—have committed a more serious offence than those who have not caused harm, and are thus deserving of harsher punishment than their respective counterparts.⁴ But, irrespective of the occurrence or non-occurrence of harm, each pair of reckless agents or purposeful actors are equally culpable. Differences in the outcomes of their actions are, to some extent, beyond their control and do not add to or detract from the level of fault we can attribute to an agent. Had the unlucky drivers decided to take a different route or leave a few minutes later, they might not have found themselves at a busy intersection and might have been able to get home without an accident.⁵ And had the successful criminals done something even slightly different during the course of their criminal activities—e.g. aim the

¹ See Duff (2016).

² Ibid.

³ See Feinberg (1995).

⁴ See Cornford (2011).

⁵ See Williams and Nagel on Moral Luck (1976).

gun slightly to the left or burgle a different house—they might not have succeeded.

The problem of outcome luck is a problem about what we ought to do upon the realisation that, like most aspects of our lives, the consequences of our actions are, to some extent, always a matter of luck.⁶ This problem arises out of a conflict between our initial acceptance of two intuitive and independently plausible principles.⁷ On the one hand, we normally adhere to the Control Principle, which is roughly that agents should be held responsible only for what is within their control. But on the other hand, we seem to accept the fairly widespread practice of holding agents criminally liable for the consequences of their actions, which are never wholly within anyone's control. In light of this conflict, should we adhere strictly to the Control Principle and consider outcomes irrelevant for liability, or should we abide by existing legal tradition and admit some degree of luck in the criminal law?

Many foundational questions in law and philosophy are implicated in the problem of outcome luck. This thesis offers an examination of those issues, particularly in the context of criminal law. The discussions here should help shed light on some of the questions that must be answered in order for us to address the problem of outcome luck, but it also leaves many questions unanswered. As such, my discussion should not be treated as an attempt to offer any kind of “solution” to this problem, but rather as an attempt to answer some key legal philosophical questions, which would concern anyone who cares about the influence of luck in the criminal law.

▪

The problem of luck is not specific to legal philosophy. In ethics, the problem of moral (outcome) luck is almost identical, except questions about responsibility concern moral rather than legal responsibility. In epistemology, the truth or falsity of our beliefs may be down to luck, but luckily true beliefs are thought to be incompatible with knowledge.⁸ Legal theorists have gained valuable insights by looking at the philosophical debate on moral luck, but

⁶ See Duff (1996), chapter 12.

⁷ See Greco (1995); Nagel (1979).

⁸ See Pritchard (2005).

significantly less attention has been paid to discussions in epistemology, which have had a lot to say about the notion of luck in recent years. This thesis begins by looking at epistemologists' attempts to deal with the problem of epistemic luck in order to see whether legal theory can gain anything from discussions in epistemology.

▪

The problem of (veritic) epistemic luck goes like this. Suppose my clock reads midnight and, based on this evidence, I form the belief that it is midnight. It is in fact midnight, but unbeknownst to me, the clock has stopped the previous day and it was by luck that I formed a true belief. Had I looked at my clock an hour earlier, my belief that it is midnight would have been false. Even though my belief is true, I do not *know* that it is midnight. This is known as the anti-luck platitude in knowledge; that in order for a subject *S* to know a proposition *p*, it must not be the case that *S*'s belief in *p* is true as a matter of luck. In other words, knowledge is thought to exclude luck.⁹

Contemporary epistemology has produced a great deal of work on luck, much of which is concerned with giving us a descriptively accurate definition of the concept. The motivation behind this work is the belief that in order for us to address various philosophical problems that arise out of this concept, it would be helpful first to have a precise definition of luck.¹⁰ For example, in what are referred to as anti-luck theories of knowledge, we see a two-step process for defining knowledge. The first step involves conceptual analysis on luck; it involves an attempt to understand the nature of the phenomenon and to specify its defining features. The second step then goes on to use our understanding of the concept to specify the sense in which knowledge must exclude luck.¹¹

This kind of approach is contrary to traditional philosophical discussions of luck in a couple of ways. First, rather than using the concept of luck in a loosely defined manner, anti-luck epistemology has given us a

⁹ Ibid., chapters 5 & 6.

¹⁰ Ibid., chapter 5.

¹¹ Pritchard (2014).

plethora of nuanced and sophisticated definitions of luck.¹² Second, rather than taking the concept to be synonymous with lack of control, many epistemologists have convincingly argued that lack of control, by itself, is not sufficient for luck.¹³ For instance, if luck is merely the absence of control, then an event such as the sun rising every morning would, counterintuitively, be a matter of luck. The fact that we do not typically consider sunrises to be matters of luck suggests that something else—perhaps in addition to lack of control—is necessary for luck.¹⁴

Many have argued that this something else is the notion of chanciness, or the notion that matters of luck are in some sense unlikely to occur. For instance, what makes a paradigmatic event, such as winning the lottery, a matter of luck is that it is typically an unlikely event. The specific sense in which matters of luck are unlikely is the subject of debate between those who favour a probabilistic conception of luck and those who favour a modal conception of luck.¹⁵ A probabilistic account holds roughly that an event *E* is a matter of luck if and only if *E* is probabilistically unlikely. And a modal account holds roughly that an event *E* is a matter of luck if and only if *E* occurs in this world but does not occur in a specified range of close possible worlds.¹⁶

Recently, this disagreement has also extended to the related concept of risk. Some have argued that luck and risk are essentially the same concept, except our judgements about luck are made *ex-post* while our judgements about risk are made *ex-ante*.¹⁷ Quite often, the degree to which an event is a matter of luck is determined by the degree of risk that surrounds the event. For instance, if subject *S* is at high risk of developing cancer, then *S* would be very lucky if he does not develop cancer. But if *S* is at low risk of developing cancer, then we would not typically use the language of luck to describe *S* if he does not develop cancer.

¹² See Chapter 1.

¹³ See e.g. Latus (2000); Pritchard & Smith (2004); Pritchard (2005, 128).

¹⁴ See Latus (2000).

¹⁵ See Chapter 1.

¹⁶ The precise formulation of the modal account of luck will be articulated in Chapter 1.

¹⁷ See Pritchard (2014).

Some proponents of a modal conception of luck have also articulated a modal conception of risk.¹⁸ They argue that, instead of conceptualising risk in terms of probability, we should measure risk in terms of counterfactual closeness. An unwanted event *E* is a high risk event if the possible world in which *E* could occur is very similar to our actual world. Conversely, *E* is a low risk event if it could only occur in a possible world that is very dissimilar to our actual world.

Chapter 1 takes a detailed look at the disagreement between modal theorists and probabilistic theorists about the concepts of luck and risk. Both of these concepts are relevant to the problem of outcome luck and the legal philosophical issues that it implicates. The scope of Chapter 1 is limited to conceptual analyses of these concepts, rather than a commentary on the significance of different conceptions of luck and risk for epistemology or legal theory. In my view, before we can appreciate the extent to which conceptual analyses on luck and risk can inform deeper philosophical questions, we must engage in conceptual analysis first. It may be the case that conceptual analysis will not help shed any light on the problem of outcome luck. But in order to know whether that is the case and *why* that may be the case, it would be helpful for us to engage with the debate, to understand why such disagreements matter to epistemologists, and to consider the differences between the problems that luck raises in law and in epistemology.

▪

As we will see in Chapter 1, I offer a detailed criticism of modal conceptions of both luck and risk. With respect to luck, I argue that modal accounts of luck are unnecessary because they are reducible to probabilistic accounts of luck, and that when it comes to risk, modal conceptions of risk are either unnecessary or trivial, especially if we already know the probability of a risk event. Chapter 1 concludes with a negative account of luck and risk; that neither of these concepts are modal concepts. I hint that probabilistic conceptions of luck and risk are superior, but I do not offer positive accounts of either concept. The reason for this is that, at least in the legal context, it is

¹⁸ See Pritchard (2015).

not clear what we can gain by having a precise definition of luck. I am skeptical that conceptual analysis on luck can inform deeper philosophical issues, especially in relation to the problem of outcome luck in law.¹⁹

Recall that proponents of chanciness conceptions of luck argue that the notion of chanciness, rather than lack of control, is better suited to capture our intuitions about luck. In this way, chanciness accounts of luck are descriptively accurate while control accounts are not. This is what the sunrise example is meant to show. Admittedly, the nature of luck extends beyond mere lack of control. But legal theorists are not particularly interested in the nature of luck or genuine matters of luck, such as winning the lottery. The problem of outcome luck is specifically about the alleged incompatibility between legal responsibility and lack of control, and legal theorists are only interested in luck to the extent that luck connotes lack of control.²⁰ Within legal philosophical discussions about outcome luck, we can replace talk of “luck” with talks of “accident” or “chance”, as long as we know that these terms are used to convey the absence of control.²¹

Proponents of chanciness conceptions of luck also argue that the notion of chanciness can easily capture degrees of luck, which is often neglected when we equate luck with lack of control.²² For instance, if we equate luck with lack of control, then it is a matter of luck when a skilled marksman hits his target *and* it is a matter of luck when he misses his target. The marksman lacks complete control over the outcome of his shot. But, as chanciness theorists about luck will point out, surely the two outcomes are not matters of luck to the same extent. It is *more* of a matter of luck when he misses. Given that he is highly skilled, the chances of him missing the shot are reasonably low.

This point about degrees of luck is relevant for the legal context and, if true, would give legal theorists good reason to at least take note of

¹⁹ But see Peels (2015) where a modal conception of luck has been applied to the problem of moral luck.

²⁰ See Enoch & Marmor (2007).

²¹ For instance, as Hart (2008, 131) puts the question of outcome luck: “why should the accidental fact that an intended harmful outcome has not occurred be a ground for punishing less a criminal who may be equally dangerous and equally wicked?”

²² See Pritchard (2014).

chanciness conceptions of luck. However, as we will see in Chapter 2, this point is false. In the legal context, where luck is stipulated to mean the absence of control, degrees of luck can be accounted for in terms of the degree of control that is exercised by an agent.²³ Skilled shooters have greater control over their shots and, compared to novice shooters, impose a greater risk of harm on their intended targets or victims. When skilled shooters miss their target, their failed attempt is more of a matter of luck compared to the failed attempt of a novice shooter.

▪

Traditionally, the problem of outcome luck concerns the significance of outcomes for liability; that is, whether an agent can be liable for the outcomes of his actions. If not, then successful criminals and failed attempters should receive the same level of punishment, as should reckless drivers who cause harm and those who do not.²⁴ But once we take note of degrees of luck, we realize that the same outcome can be more or less of a matter of luck for different agents depending on the level of risk an agent imposes. The same harmful outcome is less of a matter of luck for the agent who unleashes a high risk of harm and more of a matter of luck for the agent who unleashes a low risk of harm.

In Chapter 2, I will explain the relationship between luck, risk, and control, and how we can account for degrees of luck and risk by accounting for degrees of control. An agent may not have *complete* control over anything, but he does have *some* control over the risks that his actions unleash. The level of risk an agent unleashes speaks to his culpability, and, in Chapter 2, I will consider the extent to which culpability can be sensitive to the degree of risk an agent imposes.

Some legal theorists—most notably Alexander and Ferzan—have argued that culpability is in part a function of the level of risk an agent (knowingly) imposes on others.²⁵ And for Alexander and Ferzan, rather than giving weight to outcomes, liability is grounded solely in an agent's

²³ See Duff (1996), chapter 12.

²⁴ See Feinberg (1995).

²⁵ See Alexander & Ferzan (2009).

culpability. If we are concerned about the influence of luck in criminal law, then given what I have said thus far, their theory of culpability is attractive for insulating culpability and liability judgements from degrees of luck. The reckless driver who imposes a greater risk of harm on others is more culpable than another who poses a lower risk of harm. While their account of culpability is attractive for being sensitive to degrees of luck, especially in the context of reckless actions, I argue that it fails to give us a proper account of culpability for purposeful actions. I will show that while degrees of risk are relevant for the culpability of reckless actions, the same is not true for purposeful actions.

▪

Some theorists have noted that the problem of outcome luck may be different for reckless actions and purposeful actions.²⁶ Chapter 3 looks at the influence of luck in the context of purposeful actions. A crucial part of the debate here concerns the *mens rea* element for criminal attempts and what that should require. In attempts, the intention is the principal ingredient of the offence; an attempt requires an intention to commit an offence even if intention is not required for the completed offence.²⁷ Chapter 3 offers a detailed discussion on what the *mens rea* element should require. The point of controversy has to do with whether recklessness as to a circumstance element of the offence is sufficient for intention. I argue that, typically, it is not.

I will discuss this issue against the backdrop of a problem which is sometimes associated with the problem of outcome luck for purposeful actions—i.e. the problem of impossible attempts. In so-called impossible attempts, an agent sets out to commit an offence and does everything he believes is necessary for the successful completion of the offence.²⁸ But given the objective state of affairs, it is impossible for the agent to commit the intended offence. For instance, suppose an agent intends to appropriate what he believes is another's property. The offence of theft requires that the

²⁶ See Duff (2016).

²⁷ See Simester & Sullivan (2016), chapter 9.

²⁸ See Duff (1996), chapter 3.

property in question belongs to another person. But unbeknownst to the agent, the property that he tries to appropriate is actually his own. And given that the property actually belongs to him, it is impossible for him to commit theft by taking his own property.

Allegations of impossibility have never been enough to exempt an agent from criminal liability.²⁹ Despite the fact that external factors are such that it was impossible for the agent to succeed, he can still be blamed for trying to commit an offence. This has to do with the fact that, in attempts, the focus is on whether the agent acted with intention to commit an offence. If the agent had acted with an intention [to take *another's* property] then he did act with an intention to commit theft. Not only did this agent intend merely [to take property] or [to take that thing], he also acted with intention to the circumstance element—i.e. that it is *another's* property.

The cases that concern us in Chapter 3 are ones in which the agent acted with an intention [to take that thing] and was reckless about whether the thing belonged to another person. When the thing in question turns out to belong to the agent himself, it is sometimes thought that the agent has committed an impossible attempt. The key rationale behind this thought is that the agent did act with an intention to commit theft when he acted with recklessness as to ownership of the property. If this is true, then we can very well think of such cases as cases of impossibility.

However, I argue that recklessness as to circumstance is not sufficient for intention, and that the *mens rea* element for attempts should require at least knowledge or belief as to the circumstance elements of an offence. I will proceed on the assumption that criminal attempts are a species of *attacks*. That is, a failed attempt to commit theft is a failed attack on another person's property interests. As I will explain, attacks are in contrast to acts of endangerment.³⁰ In attacks, the agent *aims* to harm some protect interest, whereas in endangerments, the agent merely *risks* harm to that interest. In the case where an agent acted with an intention [to take that thing] and is reckless about ownership of the thing, I argue that he has not attacked, but

²⁹ See Chapter 3 for possible exceptions to this rule.

³⁰ See Duff (2007), chapter 7.

rather endangered, another's property interests. The discussion in this chapter should help shed light both on debates surrounding how we should deal with instances of impossibility, as well as the *mens rea* requirement for criminal attempts more generally.

▪

Finally, in Chapters 4 and 5, I turn my attention to the issue of punishment. The problem of outcome luck is, to a large extent, a problem about how we should punish two similarly culpable agents, one of whom has caused harm while the other has not.³¹ In order to make progress on this issue, we must address the question of *why* punishment is an appropriate response to criminal wrongdoing, a question that takes us to the debate on the justification of punishment.

In these last two chapters, I will focus on what I consider to be the most plausible accounts of punishment: instrumentalist justifications of punishment and communicative theories of punishment. Instrumentalists have typically argued that punishment is primarily justified in terms of its preventive functions.³² If our penal institutions proved to be an ineffective deterrent against criminal offending, then we would have good reason to abolish it. Communicative theories, on the other hand, are a species of expressivist theories of punishment, which takes the expression or communication of censure to be the central justifying aim of punishment. For these theorists, we punish offenders in order to express or communicate to them the moral criticism they deserve for their offence.

The aim of Chapters 4 and 5 is to defend a communicative account of punishment against an important criticism from instrumentalist critics, which is that communicative punishment fails to take crime prevention seriously. In order to address this criticism, I will begin in Chapter 4 by giving a detailed overview of Antony Duff's communicative account of punishment. On Duff's view, punishment is a communicative process between the offender, the victim, and the political community to which they belong. The message of communication is one of censure, where the offender receives the level of

³¹ See Feinberg (1995).

³² See Tadros (2011).

moral criticism that is appropriate in response to his offence. In censuring the offender, the political community recognizes the wrong that he has committed against the victim. I argue that recognition can be valuable for the victim regardless of any instrumental effects it may have, but I do not argue that the value of recognition is enough to justify our penal institutions.

Throughout Chapters 4 and 5, I proceed on the assumption that instrumentalists are right about the normative significance of crime prevention. In Chapter 5, I argue that communicative punishment does in fact have an essential preventive dimension. However, communicative theorists like Duff aim for a different type of prevention than instrumentalists. As I will show, communicative punishment entails general prevention, where the mode of prevention is one of moral dissuasion. This can be differentiated from the mode of prevention that is prevalent in the literature on punishment, which concerns general deterrence through prudential disincentives. Communicative theorists can take crime prevention seriously, or so I argue, without aiming for the latter kind of prevention.

By the end of Chapter 5, we will see that, if we care about crime prevention, there are unresolved questions about the limits of pursuing a communicative account of punishment. These questions point to further work that needs to be done on this topic. Nonetheless, Chapters 4 and 5 constitute a serious attempt to address an important challenge that has been raised against communicative theories. It may not be enough to convince instrumentalist critics, but I hope it will shed some light on how we might advance the debate surrounding punishment.

▪

In the Conclusion to this thesis I will make some brief suggestions about how we should punish similarly culpable agents who cause different degrees of harm. My suggestions will be tentative since, as I mentioned earlier, there are still important questions that need to be addressed in relation to the problem of outcome luck.³³ For now, let us begin by turning to the debate on the

³³ For example, this thesis has not touched on the debate between subjectivists and objectivists about the proper foundations of criminal liability.

nature of luck and risk, and what epistemologists have had to say about the concept of luck and risk.

Chapter 1

What's Wrong with Modal Conceptions of Luck and Risk

The modal account of luck has become very popular and influential in the past decade. More recently, some of its proponents have also put forth a modal account of risk and argued that we ought to apply it to problems both in and out of philosophy.¹ Here, I try to show that such conceptions are mistaken.² I start, in section 1, by arguing that modal conceptions of luck are parasitic on the probability of the event in question. I clarify some key components of modal accounts of luck and show how such accounts ultimately reduce to probability. Then in section 2, I consider some existing discussions surrounding so-called “lucky” true beliefs in lottery style scenarios. I argue that contrary to some modal theorists’ suggestion, lottery scenarios do not show that luck is a modal notion. Finally, in section 3, I turn to the modal account of risk. Here I argue that modal theorists have offered us no reason to be concerned about modal risk. Depending on the level of probabilistic risk in question, modal theorists’ approach to risk assessment and reduction is either unnecessary or trivial.

1 Probability within the Modal Account of Luck

Winning a fair lottery with long odds is a paradigmatic instance of luck. It is a matter of good luck if the win confers significant benefits upon the winner and it is a matter of bad luck if the “winner” has been selected to meet some dreaded fate. Most accounts of luck contain a significance condition such that

¹ The most prominent proponent of modal accounts of luck and risk has been Duncan Pritchard, and I focus primarily on his articulation and defence of such accounts here. See also Whittington’s manuscript on risk.

² There has been a plethora of conceptual analyses on luck and risk within the past decade, partly due to the thought that such investigations will shed light on important questions in epistemology and other areas of philosophy. This kind of approach calls for a two-part solution to the problems of “luck” in philosophy. First, we need to understand the nature of luck. Second, equipped with such an understanding, we then specify the sense in which various philosophical issues (e.g. knowledge, responsibility, free will, etc.) are incompatible with luck. This paper is concerned with the first part of this project—i.e. the nature of luck. But see fn. 23 on why this approach may not be as useful as its proponents seem to believe.

an event is a matter of luck only if it is of some value for some specified being.³ In addition to a significance condition, different theories will account for luck with either a chanciness condition or a lack of control condition. Proponents of a lack of control account argue that matters of luck are significant events that are also beyond an agent's control.⁴ And those who endorse a chanciness account argue that matters of luck are significant events that are also in some specified sense unlikely to occur. This paper is not concerned with the significance condition or lack of control accounts of luck. The focus is on competing notions of chanciness—i.e. on the differences between modal and probabilistic accounts of luck.

According to probabilistic accounts, lottery wins are matters of luck roughly because they are significant events that are probabilistically unlikely.⁵ The chances of winning a national lottery is so probabilistically farfetched and stacked against any given ticket that it is far more likely for one to lose than it is for one to win. Given different odds (e.g. 0.9), winning the lottery may not be considered a matter of luck—or it would be less of a matter of luck—because the event is more likely or almost guaranteed to occur.

When we talk about the probability of winning the lottery, we typically talk about objective chance rather than epistemic probability or subjective credence. Objective chance refers to features of the world and is independent of our beliefs about the event in question.⁶ In contrast, epistemic probability refers to the evidence we have in support of a particular proposition, whereas credence measures our degree of belief in a proposition irrespective of our evidence for that proposition.⁷ Our beliefs can track objective chance, but the objective probability of winning a national lottery is probabilistically farfetched regardless of our beliefs about its likelihood.

Modal accounts of luck *appear* to offer a different explanation for why some events are matters of luck, but the explanation do not differ, in

³ But see Pritchard (2014) who argues against the necessity of a significance condition.

⁴ Riggs (2009); Broncano-Berrocal (2015).

⁵ McKinnon (2013); Rescher (2014); Steglich-Peterson (2010).

⁶ In talking about objective probability, I make no commitments to any of the different interpretations of objective probability (e.g. frequency or propensity), though I will limit my discussion to the frequentist account.

⁷ See Hacking (2001, chapter 11); Mellor (2004, chapter 1).

essence, from the explanation offered by probabilistic accounts. Consider Pritchard's (2005) initial formulation of the modal condition on luck:

"If an event is lucky, then it is an event that occurs in the actual world but which does not occur in a *wide class of nearest possible worlds* where the relevant initial conditions for that event are the same as in the actual world."⁸

To see why the modal account does not differ meaningfully from the probabilistic account of luck, we need to begin with a clear understanding of the modal condition on luck.

Crucially, two sub-conditions make up the modal condition on luck. The first sub-condition has to do with the notion of modal distance, which concerns degrees of similarity between possible worlds and the actual world. Call this sub-condition_D. The second sub-condition has to do with the notion of modal frequency, which refers to the prevalence of an event across a specified set of close possible worlds. Call this sub-condition_F. An event is modally robust if it occurs across a wide class of specified close possible worlds, and it is modally fragile if it occurs across a limited number of specified close possible worlds. Together, sub-condition_D and sub-condition_F make up the modal condition on luck. Sub-condition_D specifies the relevant set of possible worlds that we are interested in (e.g. *nearest possible worlds*), and sub-condition_F specifies the proportion of those worlds in which an event must not occur (e.g. *a wide class* of those nearest possible worlds). On the modal account of luck, lottery wins are matters of luck because they are modally fragile across the set of possible worlds where we participate in the lottery.

Notice that the notion of modal distance (i.e. sub-condition_D), by itself, tells us nothing about luck. It simply measures the extent of the similarities between possible worlds and the actual world.⁹ For instance, a possible world in which the glass of water on my desk is five inches closer to me than it actually is at this moment is closer to the actual world than another possible world in which I am a barrister. And the possible world in which I am a

⁸ Pritchard (2005, p. 128), my emphasis.

⁹ See Lewis (1979) for a discussion on the ordering of possible worlds.

barrister is closer than another possible world in which Edinburgh has switched places with Montreal.¹⁰ Both winning and losing the lottery are modally close events since not much has to change about the actual world in order to bring about the respective counterfactual events.

With respect to assessments of luck, the notion of modal distance is only relevant within sub-condition_D, where we pick out a set of close possible worlds that are relevantly similar to the actual world. For instance, consider Pritchard's (2007b) later formulation of the modal condition, where "an event is lucky *iff* it obtains in the actual world but does not obtain in a wide class of *near-by* possible worlds".¹¹ The only difference between the 2005 and 2007 formulation is the relevant set of possible worlds that we need to specify—i.e. *nearest* vs. *nearby* possible worlds. Once we have picked a relevant set of possible worlds, we still need to apply sub-condition_F and assess whether an event is sufficiently infrequent across those possible worlds so as to count as a matter of luck.

Sub-condition_F is a key part of the modal condition on luck. Without it, a modal account would be over-inclusive. For instance, suppose an employer is hiring new employees on the basis of a lottery and John has been selected. The possible event of John not being selected is modally close. But the fact that the counterfactual event is modally close does not tell us anything about whether the actual event is a matter of luck. On a modal account of luck, it would be a matter of luck for John to be selected *iff* there is a sufficient number of close possible worlds where he is *not* selected (i.e. the event is modally fragile or infrequent). If five people are selected from a pool of one-hundred applicants, then presumably John has not been selected in a wide class of relevantly close possible worlds. However, if ninety-five people are selected from a pool of one-hundred applicants, then presumably John *has* been selected in a wide class of close possible worlds. This is exactly how a modal account of luck is able to explain lottery wins. What makes an instance of winning the lottery a matter of luck is not that the counterfactual event is modally close, but that there is a sufficient number of close possible

¹⁰ Pritchard (2015, p. 443).

¹¹ Pritchard (2007b, p. 279), my emphasis.

worlds where the counterfactual event obtains instead. This is the concern of sub-condition_F; it specifies the frequency of the counterfactual event within a set of possible worlds. Without sub-condition_F, lottery losses would, counter-intuitively, be counted as matters of luck since the counterfactual event of winning the lottery is modally close.

There can, of course, be disagreements amongst modal theorists over the correct formulation of sub-condition_F—i.e. the requisite modal frequency or fragility of matters of luck.¹² However, some formulations of sub-condition_F (and sub-condition_D) are clearly too vague and too wide to offer an accurate account. For instance, in Pritchard's latest (2014) articulation of the modal condition, he writes that “what makes an event lucky is that while it obtains in the actual world there are—keeping the initial conditions for that event fixed—close possible worlds in which this event does not obtain.”¹³ This formulation differs from earlier formulations in that the relevant set of close possible worlds (sub-condition_D) has widened from *nearest* possible worlds (2005) to *nearby* possible worlds (2007b) to *close* possible worlds (2014). Additionally, the proportion of these possible worlds in which the actual event must *not* occur in order for the event to count as a matter of luck (sub-condition_F) has narrowed from a wide class of nearest or nearby possible worlds (2005, 2007b) to presumably at least one or two close possible worlds (2014). The (2014) formulation of sub-condition_F is clearly too wide since on this formulation, not only are lottery wins a matter of luck, but lottery losses are also matters of luck, assuming there are one or two close possible worlds in which the losing ticket is the winner.¹⁴

¹² See, e.g., Levy (2011) who argues that the requisite modal infrequency for matters of luck may differ depending on the significance of the event in question.

¹³ Pritchard (2014, p. 599).

¹⁴ Pritchard has suggested that all lottery events are a matter of luck, including lottery losses (personal communication, February 16, 2016). I think we should reject this suggestion. While we would not hesitate to use the language of luck to describe lottery wins, it would be odd to say to someone who has lost the lottery that his loss was a matter of (bad) luck. Of course, there can be negative lotteries, where a winner is selected to meet some dreaded fate. We would consider the “winner” of this lottery to be extremely unlucky, and we would also consider the “losers” of this lottery to be lucky. However, what explains this judgment is not that all lottery outcomes are a matter of luck, but that whether an outcome is a matter of luck also depends on the significance of the outcome in question.

Notwithstanding any vague formulations of the modal condition on luck, the problem I want to raise is more fundamental and concerns sub-condition_F in particular. Suppose modal theorists agree on the formulation of sub-condition_D; i.e. a relevant set of close possible worlds from which to apply sub-condition_F. In order for modal theorists to offer an account that differs meaningfully from probabilistic accounts of luck, they need to explain how we can judge whether an event is modally (in)frequent without knowing the probability of that event. How can we judge whether an event obtains in any portion of relevantly close possible worlds if not by some kind of (implicit) reference to its probability? In other words, how can modal theorists apply sub-condition_F without reference to probability?

For example, compare the following two lottery events. In the first lottery, six numbers are randomly drawn from a pool of 49 numbers. Any and all tickets that match all six numbers will win the jackpot. In the second lottery, only one number is randomly drawn from a pool of 49 numbers. Any and all tickets that match the selected number will win the same amount of money as the winner of the first lottery.¹⁵ Holding the value of the lottery prizes fixed, surely winning the first lottery is more of a matter of luck than winning the second lottery. Probability theorists about luck can explain this in terms of the probability of winning each lottery. It is more of a matter of luck to win the first lottery (where the probability of winning is 1 in 14 million) because, compared to the second lottery, the chances of winning are much lower, probabilistically speaking. But how might a modal theorist explain this without reference to, or knowledge of, the probability of winning each lottery? With respect to the second lottery, there is, presumably, a large portion of close possible worlds where the actual winners do not win the lottery. But, presumably, there is a substantially larger portion of close possible worlds where the actual winners of the first lottery do not win the lottery. Why should

¹⁵ We can assume that the winnings are fixed so the pot does not need to be split between multiple winners.

we not think that our judgements about modal frequency is inferred from and ultimately reduces to judgments about probabilistic likelihood?¹⁶

The problem with modal accounts of luck is this. A descriptively accurate modal account must include sub-condition_F, otherwise the account would be over-inclusive. But it appears that our judgements about the modal frequency of events are guided by our judgments about probability.¹⁷ This is the sense in which the modal account of luck is parasitic on the probabilistic account of luck. What is wrong with modal accounts of luck is that we would not be able to apply sub-condition_F unless we already have some judgment about the probability of an event. But why should we not think that our judgements about probability are enough to tell us whether the event is a matter of luck?

This challenge leads us back to sub-condition_D—i.e. the need to specify a relevant set of possible worlds, which modal theorists could argue is crucial for accurate ascriptions of luck. Assuming that this is true, this line of argument still fails to differentiate a modal account from a probabilistic account in any meaningful way. Ironically, on some interpretations of probability, this step can make a modal account look more like a probabilistic account of luck. For instance, a popular interpretation of objective probability is frequentism, where probability expresses the frequency of an event relative to the size of a specified reference class. Modal theorists' concern with relevant sets of possible worlds resembles frequentists' attempt to situate events within appropriate reference classes. And when modal theorists

¹⁶ A reviewer has pointed out that this argument might be turned around. How can we determine the probability of E if not by appealing to modal considerations? For instance, on the classical interpretation of probability, the probability of E can be determined *a priori* by considering the range of possible outcomes, and what probability expresses is the proportion of equipossible outcomes that contain E. Why should we not think that probabilistic accounts of luck are parasitic on modality? One reason is that our judgments about luck are not limited to events where the probabilities are knowable *a priori*. We consider it a matter of good luck when an unreliable train reaches its destination on time, when a novice archer shoots a bullseye, or when a student aces her exams without studying. Such events are matters of good luck because they are infrequent, but the way we come to know that such events are infrequent is not through *a priori* reflections on the range of possible outcomes.

¹⁷ This is merely a claim about how we seem to make judgements regarding the modal frequency of an event. It does not point to any metaphysical thesis about the relationship between modal frequency and objective probability. The emphasis on judgments is also in keeping with the nature of luck as an inherently subjective phenomenon, which is wholly dependent on how competent speakers tend to use the concept.

disagree about which sets of possible worlds are relevant, their disagreements resemble disagreements between frequentists over the appropriateness of various reference classes.

For example, suppose we want to determine the extent to which finding a buried treasure is a matter of luck. Different modal calculations using different sets of possible worlds will yield different verdicts, just as different probabilistic calculations using different reference classes will yield different verdicts. If modal theorists only consider nearby possible worlds, where the treasure is buried in the same spot and where the subject digs for treasure in *that* spot, then the event would not be a matter of luck because it occurs in most of the specified worlds. But if we consider all possible worlds, including ones where the treasure has not been buried, then the event would be a matter of luck since it only occurs in a small portion of all possible worlds.¹⁸ Probabilistically speaking, if frequentists situate the event within the reference class of individuals who look for treasure in the particular spot where the treasure is buried, then we might think that the event is not a matter of luck since it has a probability of 1. But if we situate the event within the reference class of all individuals (including those who have never looked for treasure), then we would get a significantly different verdict. And if we situate the event within the reference class of individuals who look for treasure on that island, we may get a different answer still.

So, on some interpretations of probability, the modal account of luck looks like a probabilistic account of luck in disguise, where the disguise comes from sub-condition_D. But the question of how modal theorists can apply sub-condition_F without reference to probability is meant to reveal a fundamental problem for the modal account of luck. If modal theorists must rely on probability judgments in order to apply half of their account (while the other half of their account can arguably be accommodated by some versions of a probabilistic account), then it is unclear why we need a modal account of luck at all, especially when probabilistic accounts are more intuitive and easier to apply.

¹⁸ See Carter and Peterson (2017) who argue that the modal account should consider events across all possible worlds.

2 Lottery Propositions and the Modal Account of Luck

In response to my line of criticism, Pritchard points to one instance of luck in particular which he claims can only be accounted for by a modal conception of luck. Suppose a subject has purchased a ticket to a fair lottery with long odds—fourteen million to one. The draw has been made but the subject has not heard the winning numbers and so has no idea whether her ticket is a winner or a loser. In one variation of this scenario, the subject knows the odds of winning this lottery and on this basis forms the (accidentally) true belief that her ticket is a loser. In a second variation, the subject does not know the odds but reads the winning numbers in a reliable newspaper and thereby forms the true belief that her ticket is a loser.¹⁹ Here, Pritchard's argument for the modal account of luck goes as follows:

- (1) The first subject does not know that her ticket is a loser.
- (2) The first subject does not have knowledge because her belief is true as a matter of luck.²⁰
- (3) A modal account of luck counts the first subject's true belief as a matter of luck.
- (4) A probabilistic account of luck does not count the first subject's true belief as a matter of luck.
- (5) Therefore, luck is a modal (rather than probabilistic) notion.²¹

The soundness of this argument depends on premise (2). It depends on whether subject one's true belief is a matter of luck, and I will show that it is not.

To see why it is not an instance of luck, we need to begin by separating questions about luck from questions about knowledge. The intuition we are supposed to have in this scenario is that the first subject does not possess knowledge whereas the second subject does. What makes our judgments about these two subjects puzzling, allegedly, is that

¹⁹ Pritchard (2014, p. 596).

²⁰ *Ibid.*

²¹ *Ibid.*, p. 596-8.

probabilistically speaking, it is extremely likely for the first subject to form the true belief that she has lost the lottery and yet we do not think she possesses knowledge. On the contrary, assuming that even reliable newspapers are subject to misprints, it is less likely, probabilistically speaking, for the second subject to form the true belief that she has lost the lottery, and yet we think she does possess knowledge. If our intuitions are correct, it would seem that knowledge is not merely “a straightforward function of the strength of one’s evidence, probabilistically conceived”.²² So far, so good.

Now, we must note that our judgments about knowledge should not affect our judgments about luck. If a subject does not possess knowledge, it does not necessarily mean that the truth of her belief is a genuine matter of luck. Surely, our judgments about luck should be made independently of our judgments about knowledge. This point concerns a broader issue that is relevant to those who discuss luck in relation to other philosophical questions. Pritchard, for instance, undertook an analysis of luck with the aim of developing a theory of knowledge known as anti-luck epistemology. According to him, the first step in developing such a theory is to offer an account of luck; that is, a descriptively accurate, metaphysical account of the phenomenon of luck. Equipped with such an account, the next steps are to specify the sense in which luck is incompatible with knowledge, and to develop an anti-luck condition on knowledge that captures this specific incompatibility.²³

²² Ibid., p. 596.

²³ Ibid., p. 594-5. Note that, at best, this approach can only form a small part of a much bigger philosophical inquiry. Before anyone offered an analysis of the concept of “luck”, philosophers were using “luck” rather loosely in debates about knowledge, justice, free will, and responsibility. In these debates, philosophers were not necessarily talking about genuine matters of luck, but rather using “luck” to denote a wide range of conditions such as accidents, chance occurrences, fortunate circumstances, lack of control, etc. Some or all of these conditions are thought to be in some sense incompatible with knowledge, justice, free will, or responsibility. But insofar as these conditions can be (and have been) differentiated from luck, the focus on finding a descriptive account of luck has ruled out these conditions from the philosophical discussion. Consequently, those who attempt to understand the nature of luck before using it to investigate why luck is incompatible with knowledge, justice, free will, or responsibility, are at best only addressing a small part of the question. Their theory cannot amount to an adequate explanation because they do not explain why knowledge, justice, free will, or responsibility is sometimes also incompatible with accidents, chance, fortune, lack of control, etc. If this criticism is valid, then we should likewise be

If, like Pritchard, one's aim is to offer a descriptively accurate account of luck before using it to better understand knowledge, then the first step of doing conceptual analysis on luck must be an isolated inquiry, independent of our preconceived understanding that some conditions might be incompatible with knowledge. This is to ensure that said conditions do not end up in our account of luck unless they are genuine defining features of luck. For example, suppose knowledge is incompatible with the truth of one's belief being an accident. Whether the quality of [being an accident] is a feature of luck is a question that must be addressed independently of considerations relating to its incompatibility with knowledge. Otherwise, we risk assuming a specific notion of the phenomenon of which we are trying to give a descriptively accurate account.

With this in mind, we can turn to the issue of luck. Unlike the paradigm case of luck I discussed earlier—that of [winning a fair lottery with long odds]—here, the matter in question is [the truth of a subject's belief] with respect to her ticket.²⁴ Intuitively, we think [winning a fair lottery] is a genuine matter of luck, but do we really think that about [the truth of subject one's belief]? Pritchard seems to think so, and some epistemologists might readily agree. But why should we think this? Obviously, the answer cannot be because her belief could have easily been false, or that her belief is false in a sufficient portion of relevantly nearby possible worlds. This response merely presupposes the modal account of luck that is under scrutiny. In order for Pritchard's argument to be sound, what is meant by "luck" in premise (2) must reflect an ordinary or colloquial understanding of the term luck. This is because (our understanding of) the nature of luck is tied to how we would ordinarily use the concept.

Here is the fallacy in Pritchard's argument. On the one hand, the meaning of "luck" that is familiar to epistemologists is that some characteristics of the subject's belief formation process is incompatible with

careful about the emergence of anti-risk theories of knowledge, which employ the same approach as its anti-luck predecessors.

²⁴ Pritchard (2007a, p. 30): "the core worry about luck as regards knowledge possession relates to luck in the truth of the relevant belief".

knowledge. On the other hand, the ordinary or colloquial meaning of “luck” is not necessarily equivalent to whatever characteristics that make a true belief incompatible with knowledge. In arguing for his modal account of luck, Pritchard equivocates on these two different meanings. The conclusion in (5) concerns an ordinary, colloquial notion of luck because Pritchard’s account is meant to be a descriptively accurate account of the phenomenon. But the notion of luck that is in play in premise (2) is not an ordinary, colloquial notion. There, “luck” is a stipulated term, which merely conforms to epistemologists’ rough understanding of the concept prior to any conceptual analysis. Subject one’s true belief is a matter of luck merely in the sense that her belief is true in some way that is incompatible with knowledge. Outside of some discussions in contemporary epistemology, we would not ordinarily describe subject one’s true belief as a matter of luck.²⁵ For instance, consider this parallel example.

Suppose I possess a big book of Edinburgh phone numbers. The book lists phone numbers from Edinburgh only, without any indication as to whom each number belongs. Suppose I randomly select a number from this book and try to determine whether it belongs to Pritchard. Given that this book contains hundreds of thousands of phone numbers, the likelihood that I just happen to have selected Pritchard’s number is probabilistically unlikely, though modally close. Based on the probability alone, I form the belief that I have not selected Pritchard’s number. If my belief turns out to be true, we would not say that my true belief is a matter of luck. For why would we say such a thing? What else should we expect? Given the odds, we should

²⁵ To make the point about equivocation clearer, consider how “luck” has been treated in discussions about legal responsibility. For legal philosophers, “luck” is a stipulated term, used to designate an agent’s lack of (some sense of) control over an outcome. It is specifically lack of control—as opposed to any other characteristics of luck—that is thought to be incompatible with responsibility. When an incompetent shooter tries but fails to cause harm, legal theorists might describe his failure as a matter of luck because he lacks adequate control over his shot. But just because legal theorists would describe the failed attempt as a matter of luck does not necessarily mean that it really was a matter of luck. For instance, if the probabilistic account of luck is right, then the failed attempt is not a matter of luck because the likelihood of failure is quite high for an incompetent shooter. We cannot simply take legal theorists’ use of the term “luck” as evidence for luck in the ordinary, colloquial sense of the term. See, e.g., Moore (1994; 2009) who thinks the problem of moral luck would be better cast in terms of causation.

expect my belief to be true.²⁶ On the other hand, if my belief turns out to be false, that I have in fact selected Pritchard's number, *then* we would say that my belief is a matter of (bad) luck. And it would be appropriate to say such a thing, given that it was extremely unlikely for my belief to be false.

The same thing can be said about subject one's belief in the lottery ticket scenario. Why should we say that her true belief is a matter of luck when that is exactly what we should expect?²⁷ Pritchard's methodology, one of taking ordinary ascriptions of luck to be good evidence for the phenomenon, can get us a descriptively accurate account of luck only when the ascription is in fact an ordinary or colloquial ascription. The lottery ticket scenario is a paradigmatic instance where the acquisition of a true belief is incompatible with knowledge, and it is appropriate for epistemologists to say that "subject one's true belief is a matter of luck" *if* "luck" is used as a stipulated term. But we cannot take epistemologists' use of "luck" as evidence for the nature of luck. At best, what we can say about subject one's true belief is that reasonable people can disagree about whether it is an instance of luck. However, if subject one's true belief is *not* an instance of luck, then what the lottery ticket scenario illustrates, contrary to Pritchard's claim, is not that luck is a distinctively modal notion. What it illustrates is that the probabilistic account of luck gets us the right result here precisely because it does *not* count subject one's true belief as a matter of luck.²⁸

3 What's Wrong with Modal Conceptions of Risk

Recently, modal theorists about luck have turned their attention to risk, and anti-luck theories of epistemology have turned into anti-risk epistemology. This transition is partly motivated by the suggestion that luck and risk are

²⁶ See, e.g., Rescher (2001, p. 24-8) who characterizes matters of luck as matters that we could not have rationally expected to occur.

²⁷ It might be the case that in this type of situation, we would actually say that the subject's true belief is a matter of luck. But even if we do say such a thing, I am skeptical that the ascription of luck here is really attributable to the closeness of the counterfactual event. Instead, I think we consider subject one's true belief to be lucky because we attach (whether explicitly or implicitly) some kind of significance to true beliefs.

²⁸ Of course, according to earlier formulations of the modal account, the truth of subject one's belief would also not be considered a matter of luck since her belief would be true in a wide range of nearby possible worlds.

essentially the same concept, except judgments about luck are made *ex post* while risk assessments are made *ex ante*.²⁹ If this suggestion is sensible, it is sensible to the extent that the matter of luck in question involves some degree of risk. However, since matters of luck do not always involve risk—e.g. finding money on the side of the road—at best, we can only get a partial understanding of luck from an account of risk.³⁰ In the rest of this paper, I will show that a modal approach to risk assessment and reduction is either unnecessary or trivial. Before I explain why this is the case, I want to highlight a difficulty for modal theorists who think we can understand luck in terms of risk and vice versa.

Recall that the modal condition on luck has two sub-conditions: sub-condition_D and sub-condition_F. The modal account of risk, however, appears to include only sub-condition_D. According to Pritchard, there is a high risk that a potential negative event E will occur if E is modally close, and there is a low risk if E is modally far off.³¹ The notion of modal frequency is missing from this formulation of risk, where modal distance alone is meant to tell us whether or not E is risky. Earlier, I argued that the frequency condition is indispensable to a modal account of luck because the account would be over-inclusive without it. But if the frequency condition is not essential to a modal account of risk, then it is unclear how modal conceptions of luck and risk are related to each other.

Modal theorists can either amend the modal condition on risk to include a sub-condition_F, or forfeit their claim that luck and risk are basically the same concept. The first option would make their account of risk vulnerable to the same criticism I levied against the modal account of luck—i.e. their account would be parasitic on a probabilistic account of risk. The second option allows the modal account of risk to remain as it is, where degrees of risk are measured in terms of modal distances. I will explore the second option here since it deals with the modal account of risk as it has

²⁹ Coffman (2007); Pritchard (2014; 2015; 2016); Broncano-Berrocal (2015).

³⁰ Of course, there is a *chance* that the subject will walk on without spotting the money on the sidewalk, but it is a stretch to say there is a risk that she will not spot the money.

³¹ Pritchard (2016, p. 563).

been formulated by its proponents. As we will see, depending on the level of probabilistic risk involved, the modal approach to risk assessment and reduction is either unnecessary or trivial.

To begin, we can understand the modal account of risk by contrasting it with the familiar probabilistic way of thinking about risk. Typically, we understand and express degrees of risk in terms of our evidential probability. A potential unwanted event *E* carries a high risk if its probability is close to 1, and *E* carries a low risk if its probability is close to 0. For instance, the risk of being selected to meet some dreaded fate in a typical national lottery style lottery is very low, at roughly one in fourteen million. In contrast, the risk of losing a typical national lottery is very high, at roughly close to one. For modal theorists, the risk of *E* does not depend (solely) on its probability but also on its modal distance to the actual world.³² *E* is a high risk event if it is modally close, and it is a low risk event if it is modally far off. Accordingly, all lottery events are high risk events for modal theorists since not much has to change about the actual world in order for a winner to become a loser and vice versa.

In many cases, probabilistic likelihood will match modal distance, but probability and modality can come apart. Here are the possibilities:

- (1) *E* is probabilistically unlikely but modally close (e.g. being selected in a national lottery)
- (2) *E* is probabilistically unlikely and modally far off³³
- (3) *E* is probabilistically likely and modally close (e.g. not being selected in a national lottery)
- (4) *E* is probabilistically likely but modally far off

I will discuss scenarios (1) and (2) before turning to (3) and (4).

If we are following probability, we will think that the risk of *E* is extremely low in both (1) and (2). Modal theorists, on the other hand, claim that the risk in (1) is much higher because the event is modally close and, in

³² See at fn. 38.

³³ See Pritchard (2015) for his examples of scenarios (1) and (2).

that sense, could easily occur. The probability of E can be identical across (1) and (2), but modal theorists would still believe we should be more worried about scenario (1).³⁴ Note that modal theorists are advancing two theses here. First, there is a descriptive thesis regarding our judgements about risk. Here, modal theorists believe we judge (1) to be more risky than (2), and that generally speaking, our judgements about risk tend to track counterfactual closeness rather than probability.³⁵ Second, there is a normative thesis about what we should do. Modal theorists tell us we should be significantly more concerned about (1) than about (2), and that we should add safeguards to (1) in order to make E modally far off.

For the purpose of this paper, I will assume that the descriptive thesis is true. The problem I want to raise for modal theorists concerns their normative thesis. In scenario (1), where the probabilistic risk is low, it is unnecessary for us to make E modally far off. There is no need for us to make (1) look more like (2) because it is unclear why we *should* be concerned with the modal distance of a risk event. Probability expresses the frequency with which we can expect E to occur. And if E is sufficiently infrequent, why should we care that E is modally close? For modal theorists, a modally close event is one that could easily occur. But for those who are inclined to follow probability, the counterfactual language in the modal account of risk is unpersuasive. After all, in a perfect world, if we conduct the lottery fourteen million times, only one of those fourteen million draws will result in an explosion.

In response to this point, modal theorists will often refer back to their descriptive claim and reiterate the point that people do in fact care about counterfactual closeness in their judgements about risk. But given that the descriptive claim is simply about people's perception of risk, it cannot lend any support to the normative thesis. The mere fact that we tend to follow modality in our judgements about risk does not show that we *should* follow

³⁴ Pritchard (2015, p. 441-2).

³⁵ Pritchard (2015, p. 444-8). Note that Pritchard makes an additional and much stronger claim, which is that (1) is *in fact* more risky than (2).

modality.³⁶ In fact, what is interesting about the descriptive thesis is precisely the suggestion that people tend to ignore probability, a framework which we think is efficacious and *should* be followed. Until modal theorists offer a better explanation for why we should care about counterfactual closeness in our risk assessments, we have no reason to think that scenario (1) is more risky than scenario (2). And until modal theorists explain why (1) is more risky than (2), it is unnecessary to try to make (1) look more like (2), especially when the probability of E is already very low.

We can now consider scenarios (3) and (4), where the initial risk is high probabilistically speaking.

(3) E is probabilistically likely and modally close (e.g. not being selected in a national lottery)

(4) E is probabilistically likely and modally far off.³⁷

Here, probability theorists and modal theorists would agree that the level of risk is high in both (3) and (4). Modal theorists may argue that we should be concerned with the modal distance of a risk event, but they do not argue we should therefore disregard probability.³⁸ If the probability of E is identical across (3) and (4), a modal theorist would argue that (4) is less risky

³⁶ A reviewer has pointed out that given the discussion in section 2, we might want to note the relevance of ordinary judgements to our understanding of luck and risk. Earlier I rejected Pritchard's argument for the modal account of luck on the basis that we would not consider subject one's true belief about her lottery ticket to be a matter of luck in the ordinary, colloquial sense of the term. But in this section, I am rejecting the relevance of ordinary people's judgments for how we should understand risk. There is no tension between using ordinary judgments to understand luck while rejecting such judgments for risk. The difference between luck and risk is that luck does not typically have implications for how we should act. Additionally, it is not obvious there is anything more to luck than how we ordinarily understand and use the concept. On the other hand, we are primarily interested in risk because we want to understand how we should act under uncertainty. We would not want to rely on lay people's judgments about risk because they are often imprecise, mistaken, and thus not very useful in guiding our actions.

³⁷ An example of (4) might look as follows. Suppose a bomb will go off in a city if the majority of its residents take a different route to work tomorrow. Unbeknownst to everyone, the information regarding the bomb has been miscommunicated and the residents are under the impression that the bomb will go off *unless* most people take a different route to work. There is no way to correct this misunderstanding. The probability that most people will take a different route is high since they do not want the bomb to explode. But this scenario is modally far off since, presumably, a great deal would have to change in order for most people to take a different route to work.

³⁸ See Pritchard (2015, p. 449); Whittington (2017, p. 13).

than (3), but given that E is probabilistically likely, we would still want to reduce the probabilistic risk involved in both (3) and (4).

To reduce the risk in (3) and (4), probability theorists would implement safeguards to reduce the probability of E. Presumably, modal theorists would agree with this, except they would also implement safeguards to make E modally far off. As I mentioned earlier, modal closeness will often match probabilistic likelihood so that when we decrease the modal risk of E, we also decrease its probabilistic likelihood. However, modal risk and probabilistic risk may go in different directions. To see why the modal approach to risk reduction is trivial, consider the following options:

- (a) We can increase modal risk and decrease probabilistic risk.
- (b) We can decrease modal risk and increase probabilistic risk.
- (c) We can decrease modal risk and decrease probabilistic risk.

It might look as if we ought to favour option (c) since it decreases both modal and probabilistic risks. But which option we implement should surely depend on the extent to which probabilistic risk can be reduced.

Suppose the initial probability of E is 0.7, and suppose that options (a), (b), and (c) will bring the probability of E to 0.5, 0.9, and 0.1 respectively. Under option (b), we would implement a number of safeguards, all of which would have to fail in order for E to occur, but the probability of failure is quite high. Surely, this would offer us no comfort. Between options (a) and (c), it seems quite clear that we should favour (c). Of course, option (a) would involve implementing something like a lottery-style safeguard, where E would only occur if a certain set of numbers came up in the national lottery. This makes E a very close possibility. But surely this has nothing to do with why we would favour option (c). We would go with (c) because it is more effective at reducing probabilistic risk than (a), regardless of its effect on modal risk. If we modified option (a) so that it now reduced the probability of E to 0.0001, then we ought to favour (a), even if it means implementing a lottery style system. In situations where the initial (probabilistic) risk is high, the modal

approach to risk reduction is trivial because our decisions are ultimately guided by probability rather than the modal distance of a risk event.³⁹

This raises a question for modal theorists. They believe that the modal approach to risk assessment and reduction should supplement the probabilistic conception of risk, but how is this to be done in practice?⁴⁰ Counterfactual reasoning certainly plays a key role in our judgments and deliberations, but it rarely plays a decisive role. Modal theorists have argued that if the probability of a risk event is identical across two scenarios, then the notion of modal closeness is able to determine which scenario poses a greater risk. Their approach seems to get the deliberative process backwards.⁴¹ In practice, we often use counterfactual reasoning to identify several options—i.e. we think about what could happen if we did such and such. We think about different preventive measures that can be adopted in order to create a situation where the risk event would not easily occur. However, aside from policy considerations, whether we choose to adopt any of these measures ultimately depends (and should depend) on whether the measure is effective at reducing the probability of the risk event. Counterfactual thinking might help us identify possible measures to adopt, but modal considerations should not be the final arbiter when it comes to adopting any such measures.

³⁹ This is not a knockdown argument against the modal account of risk. Modal theorists can concede that since the decrease in probabilistic risk is so significant, we can favour modified option (a) without having to abandon the modal account of risk. After all, the modal account does not rule out the importance of probability. What this means though is that modal theorists need to give an account of when modality matters, and the extent to which modal risk matters more than probabilistic risk. At what point is a decrease in probabilistic risk sufficient for us to ignore modal risk, and at what point is a decrease in modal risk sufficient for us to ignore probabilistic risk?

⁴⁰ See at fn. 38.

⁴¹ Thanks to Rosa Hardt for this point.

Acknowledgements I am grateful to Andrew Cornford, Antony Duff, Duncan Pritchard, Martin Smith, Ben Sworn, and Lee Whittington for their helpful comments on previous drafts of this paper, and to Adam Carter, Ana-Maria Cretu, Giada Fratantonio, and Rosa Hardt for their discussions. This paper also benefited greatly from comments and suggestions made by two anonymous referees for *Erkenntnis*. I am very thankful for their input.

Chapter 2

Luck and Culpability

In the previous chapter, I offered a detailed discussion of what luck is not. It is not a modal concept, which we can understand with reference to close possible worlds. Although I have hinted that luck may be a probabilistic notion, I will not try to offer a positive view either in this thesis or elsewhere. As I suggested earlier, I am skeptical about the extent to which a positive account of luck can inform various problems in philosophy.¹ And given the amount of discussion that has already been had on this topic, it is not clear to me that much more need to be said about it. In this chapter (and the rest of this thesis), I will proceed on the basis that “luck” is a stipulated term in legal philosophy.² And that, when it comes to the problem of outcome luck, we are interested in luck to the extent that it connotes lack of control since the legal philosophical problem concerns the alleged incompatibility between liability and lack of control.

In this chapter, I explore the problem of outcome luck against a point which arose in the previous chapter. From the discussion on lottery cases, we see that luck is not a binary concept. Not only can outcomes be matters of luck, but outcomes can be matters of luck to varying degrees. Traditionally, the problem of outcome luck concerns the relevance of outcomes for liability.³ If outcomes are subject to luck, then one way to insulate the criminal law from luck is to ground liability solely in terms of an agent’s culpability or fault. But once we recognize degrees of luck, we will see that luck can also come from the level of culpability. I do not mean to point out that there is luck with respect to the choices that an agent makes (although our choices, much

¹ See Chapter 1.2, fn. 23. See also Ballantyne (2014).

² Legal theorists are certainly aware that they have adopted a non-colloquial understanding of luck in their discussions about responsibility (see Enoch & Marmor (2007)). Their response, however, is not to look for a descriptively accurate conception of luck since the problem of outcome luck can be framed in terms of causation and control (e.g. Moore (1994; 2009, 24). See Zimmerman (1987) for a discussion on luck and control.

³ See Ashworth (1993); Enoch (2010).

like the consequences of our choices, are to some extent also a matter of luck).⁴ What I mean by luck at the level of culpability is that the degree to which an outcome is a matter of luck can be determined by factors that are relevant to an agent's culpability. And the topic of this chapter is to explore the extent to which culpability can be sensitive to degrees of risk.

I begin in section 1 by introducing the problem of outcome luck in criminal law and by explaining the sense in which liability is thought to be incompatible with lack of control. Then in section 2, I offer some examples to show how outcomes can be matters of luck to varying degrees and identify the factors that affect degrees of luck. As we will see, degrees of outcome luck are dependent on the degree of risk we impose on others, which is something that speaks to our culpability. So in addition to making outcomes irrelevant for liability, we can further insulate criminal law from luck by making culpability sensitive to degrees of risk. Alexander and Ferzan have argued for a theory of culpability which does just this. On their view, culpability is in part a function of the level of risk an agent (knowingly) imposes on others. In section 3, I will show that this view of culpability is attractive for being sensitive to luck, especially in the context of reckless actions. However, this view is also problematic because, as I will argue, when it comes to purposeful actions, degrees of risk are irrelevant for culpability. The implications of what I have to say here is that, in the context of purposeful actions, it is, in one sense, acceptable for the criminal law to admit some degrees of luck. I will say more about this in section 4.

1 Outcome Luck and Control

The problem of outcome luck arises out of a conflict between two independently plausible principles.⁵ On the one hand, we find great intuitive appeal in the Control Principle, which roughly is the idea that agents should be responsible only for what is appropriately within their control. On the other hand, we also appear to accept the long-standing legal tradition where

⁴ See Ohana (2007).

⁵ Greco (1995); Nagel (1979).

criminal liability and punishment are graded in proportion to the actual amount of harm caused, which is never wholly within anyone's control.⁶

Suppose D shoots at V intending to kill her. D does everything he believes necessary to kill V—e.g. positioning himself for a clear shot at V; taking aim at her; pulling the trigger, etc. Even if D manages to do everything he believes is necessary to kill V, his success in killing V still depends on various events that are beyond his control. A stray bird may fly in front of V and block the bullet; V may be wearing a bulletproof vest; V may move suddenly; D's shot may miss. Should V survive D's murderous attempt, this fact is likely, all else being equal, to result in a lesser sentence for D than if V had not survived. However, why should V's survival, which is not to D's credit, be a ground for imposing a lesser punishment on D? He intended to kill V and tried to kill her by shooting at her; that he failed to accomplish what he intended to do is external to him and does not detract from his culpability.

Those who deny that outcomes should matter for criminal liability argue that liability should be determined by an agent's culpability only. And culpability, in turn, should only be determined by what is (and can be) within an agent's control.⁷ These theorists deny that we are able to exercise control over the outcomes of our actions, and argue that to allow outcomes to influence our liability is to subject criminal liability and punishment to "luck".⁸ Instead, culpability and liability should be based on matters that are within an agent's control, such as his choice to act in a certain way.

There is plenty of room for disagreement among legal theorists about whether we can exercise control over the outcomes of our actions (or anything at all).⁹ Some would argue, for instance, that if I pick up my glass of water, intending to drink from it and I do drink from it, then surely I have control over the outcome of my action, which is to drink water from my glass. We may not have control over all of the outcomes of our actions, but that does not mean we cannot have control over some outcomes of our actions.

⁶ Duff (1996, 327-29).

⁷ See Alexander & Ferzan (2009, chapters 2 & 5).

⁸ See for examples Alexander & Ferzan (2009); Ashworth (1988); Feinberg (1995).

⁹ See Duff (1996, 327-347); Moore (2009, 20-33).

One of the things we need to clarify within this debate is the kind of control that is required for legal responsibility, and what different theorists mean when they talk about lack of control. For starters, we can differentiate between restricted control and unrestricted control. Restricted control is satisfied if the agent has the ability to bring about an event and to prevent its occurrence.¹⁰ For instance, at this moment, I have the ability to either get up or remain seated. If asked, I can demonstrate my control over my bodily movements by getting up (if I am asked to get up) or remaining seated (if I am asked to remain seated). I enjoy control over this matter, but the kind of control I enjoy is, in a sense, restricted. My control over getting up or remaining seated is restricted by numerous other events which I cannot control. I cannot control, for instance, whether I will be asked to demonstrate my ability to get up or remain seated, which action(s) I will be asked to demonstrate, or whether I would be inclined to cooperate with the demonstration. Furthermore, suppose the postman knocked on my door at this moment. I would, out of habit, decide to get up and answer the door. And if a fire was breaking out, I would, quite sensibly, get up to try to put out the fire. These events, which would make me get up from my seat, are all beyond my control.

Here, we can begin to see the difference between restricted and unrestricted control. An agent has unrestricted control over an event E when he is able to exercise restricted control over both E as well as all other events on which E 's occurrence or non-occurrence is contingent.¹¹ In the example I just described, I may have restricted control over the act of getting up or remaining seated, but I do not have unrestricted control over either event since I do not have restricted control over all other events on which these two events are dependent—e.g. which action I will be asked to demonstrate, the absence of the postman, fire not breaking out in my office, etc.

Among those who deny that we have control over the outcomes of our actions, some deny that we have unrestricted control.¹² They argue that not

¹⁰ See Zimmerman (1987).

¹¹ Moore (2009, 20-33); Zimmerman (1987).

¹² See Levy (2011).

only do we lack unrestricted control over the outcomes of our actions, but we also lack unrestricted control in everything we do. We do not have restricted control over all the events that determine the way our actions turn out, nor do we have restricted control over all the events that determine the choices that we make. For anyone who argues that outcomes are irrelevant for liability because we lack unrestricted control over outcomes, they must explain why legal responsibility should require this kind of control. If criminal responsibility requires unrestricted control, then not only can we not be responsible for the outcomes of our actions, but we can never be responsible for our choices either. Furthermore, we can never be responsible for anything since our choices, characters, and dispositions are always beyond our unrestricted control.¹³

Legal theorists seem to agree that this notion of control is too stringent for legal responsibility. For instance, suppose D exercises restricted control over his decision to shoot V, as well as numerous other events on which the outcome of his shot depends. That is, D has positioned himself to get a clear view of V; he has practiced the shot many times in order to maximize his chances of success; he has cleared the area in which he plans to shoot V of any birds that might interfere with his shot. Despite all this, D still lacks restricted control over whether or not V is wearing a bullet proof vest. If it turns out that V is not wearing a bullet proof vest and D successfully kills V, it is rather implausible to hold that D lacks control over V's death.¹⁴

This leads us to a further differentiation among those who take legal responsibility to require restricted control. What I have just said reflects the views of theorists like Duff and Moore, who would argue that D had control over V's death; that we can have some (e.g. restricted) control over the outcomes of our actions. Others like Alexander and Ferzan take a position that I mentioned earlier, which is that D had (restricted) control over his choice to kill V, but that D does not have control over V's death. So for

¹³ This is where the problem of outcome luck turns into the problem of free will. See Nagel's classic piece on moral luck; Nagel (1979).

¹⁴ See Duff (1996, chapter 12).

Alexander and Ferzan, we have restricted control over our choices but not the outcomes of those choices.¹⁵

This disagreement stems from the kinds of things over which we can exercise restricted control. For Moore, when we exercise control over our choices, we also control the outcomes of our chosen actions (assuming that we chose to bring about those outcomes).¹⁶ For instance, I exercise restricted control over where to work this afternoon. I have the ability to work from my office, and I have the ability to work elsewhere. I exercise restricted control over my choice to work from my office and, in acting on that choice, I also exercise restricted control over *what* I choose to do—i.e. that I am working from my office this afternoon.¹⁷ In contrast, Alexander and Ferzan endorse a reason responsiveness view, which holds that we can only exercise (restricted) control over things that are capable of responding to our reasons. It is our choices, and only our choices, which in virtue of being reason responsive, that are within our control. The consequences of our choices cannot be within our control since they are not the sort of things that can be reason responsive.¹⁸

When it comes to outcome luck, whether an outcome is a matter of luck depends on the kind of control that we are talking about. If we are talking about unrestricted control, then outcomes are always a matter of luck since we never have unrestricted control over anything.¹⁹ It is a matter of luck whether D kills V, just as it is a matter of luck whether D chooses to kill V, just as it is a matter of luck what any of us chooses or does.²⁰ But if we are talking about restricted control, then outcomes may or may not be a matter of luck depending on the kind of things that we think can be within our restricted control. For theorists like Duff and Moore, some outcomes may be matters of luck while others may not. But for those like Alexander and Ferzan, outcomes

¹⁵ See Alexander & Ferzan (2009, chapter 5).

¹⁶ Moore (2009, 20-33)

¹⁷ Moore (2009, 20-33); Duff (1996, 331-333).

¹⁸ Alexander & Ferzan (2009, 171-96).

¹⁹ Moore (2009, 20-33). Unless we think that libertarian free will is possible. See Stanford Encyclopaedia of Philosophy entry on Free Will.

²⁰ As Nagel (1979, 35) puts it, agency “shrinks to an extensionless point”.

are always a matter of luck since they are not the sort of things that can be within our control.²¹

Without taking a stance on whether or not we have restricted control over the outcomes of our actions, we can say with respect to both camps that the way we have been talking about outcome luck is incomplete.²² In addition to talking about whether an outcome is a matter of luck, we can often talk about the degree to which an outcome is a matter of luck. For those like Alexander and Ferzan, who argue that outcomes are always a matter of luck, we can always ask about the degree of luck that is inherent in a particular outcome. In the next section, I will elaborate on this point and show how degrees of luck can come from the level of culpability. As we will see, the degree to which an outcome is a matter of luck is the degree to which it is beyond the agent's control. From this, we will see that degrees of luck are tied to the degree of risk an agent imposes through his action.

2 Luck as Lack of (Degrees of) Control

I mentioned in the introduction that, from this chapter onwards, when we talk about outcome luck in the law, I will simply take "luck" to mean lack of control. One of the supposed disadvantages of talking about luck only in terms of lack of control is that we would not be able to capture degrees of luck.²³ Once again, consider the example of winning a fair lottery with long odds. Typically, we are able choose whether we want to participate in any given lottery, but participants of a lottery lack control over the outcome of the draw; they lack control over whether they win or lose the lottery.²⁴ If luck is merely the absence of control, then both winning and not winning the lottery would, counterintuitively, be a matter of luck since both outcomes are beyond the agent's control. But even if we accept that both winning and losing are matters of luck, we must still admit that there is a higher degree of luck

²¹ See Alexander & Ferzan (2009, 171-96).

²² But see Duff (1996,) chapter 12.

²³ See Pritchard (2014).

²⁴ Granted, one could buy a large number of lottery tickets, thus increasing one's chances of winning the lottery. However, unless the participant purchases all possible tickets (in which case, it would not be a lottery anymore since the outcome is guaranteed to obtain) he lacks complete control over the outcome of the lottery.

involved in winning than there is in losing, something that the notion of control is unable to capture.

Those who defend chanciness conceptions of luck can easily capture degrees of luck, for instance, by considering the ex-ante probability of the matter of luck in question. So the degree to which an event is a matter of luck would depend on the ex-ante likelihood of that event. It is more of a matter of luck to win the lottery than it is to lose the lottery given that the chances of winning are much lower than that of losing. We can certainly account for degrees of luck through probability, but surely we can account for the probability of some events through the level of control we have over those events. In the lottery case, suppose we have control over the number of tickets that we can purchase (including the means to purchase them). The more tickets we purchase the higher chances we have of winning. There is less luck involved if I win the lottery after having purchased more than half the tickets, compared to if I had only purchased one of thousands of tickets.

Agents can exercise varying degrees of control when they engage in the kind of conducts that concern the criminal law. This is true even if we suppose, in agreement with Alexander and Ferzan, that whatever an agent sets out to do—whether he intends to cause harm or whether he risks harm through recklessness—the outcome of his action is *a/ways* beyond his control. Agents may not have complete control over the outcomes of their actions, but they do exercise some control over what they choose to do, which affects the degree to which the outcome is a matter of luck. I will illustrate this point in relation to reckless actions first, before moving on to purposeful actions.

Consider two reckless agents—S and R—both of whom decide to drive home after having enjoyed a few too many drinks. Each agent is aware of the fact that his blood alcohol level is above the prescribed legal limit and that he should not operate a motor vehicle. S and R each have two routes available to them. They can take a well-populated route through the city centre or take a much quieter detour. Driver R, who is prone to taking risks, decides to take the populated route and Driver S, who is typically a safe

driver, decides to take the quiet detour. It is beyond each agent's control whether he encounters other drivers or pedestrians on his chosen route, but it is within each agent's (restricted) control which route he chooses to take and when he chooses to take it.

In choosing to take the quiet route, Driver S runs a lower risk of encountering anyone on his drive home, whereas Driver R runs a higher risk by going through the city centre.²⁵ We could suppose, in agreement with Alexander and Ferzan, that the drivers lack control over the presence of other drivers and pedestrians on the road, and thus lack control over whether or not they will get into an accident. But these drivers surely exercise some control over the likelihood that they will encounter someone on the drive home in virtue of the route they decide to take. This means that even if the drivers lack control over the outcomes of their actions, they do not lack control over the likelihood that a particular outcome will occur.

Now suppose that Drivers S and R meet the same fate; either they both make it home safely or they both get into an accident and cause serious injury or death to other drivers or pedestrians. We could accept that, for each driver, either outcome would be a matter of luck, but surely it would be *more* of a matter of bad luck for Driver S to get into an accident than it is for Driver R, and surely it would be *more* of a matter of good luck for Driver R to make it home safely than it is for Driver S.²⁶ Again, the probability of getting into an accident is higher for Driver R than it is for Driver S, so it is more of a matter of luck when S gets into an accident compared to R. Conversely, the probability of making it home safely is higher for Driver S than it is for Driver R, so it is more of a matter of luck for R to get home safely compared to S.

²⁵ There is disagreement among legal theorists regarding the kind of risk—i.e. objective or subjective risk—that should ground culpability. Subjective risk refers to the agent's assessment of the risk that is associated with his conduct, whereas objective risk is typically construed in terms of a reasonable person's assessment of that risk. See Alexander & Ferzan (2009, 28-31). I will not address any disagreements between subjectivists and objectivists here, and will simply talk of *the* level of risk imposed. I will assume that subjective risk tracks objective risk, which I construe in terms of a reasonable person's ex-ante assessment of the probability of the risk event. Furthermore, we can also bracket any alleged difficulties that come with adopting a reasonable person's perspective for the purpose of this chapter.

²⁶ See Rescher (2014).

What accounts for differences in degrees of luck has to do with differences in the probability of a particular outcome for a particular driver—i.e. the level of risk a driver imposes on others.²⁷ And agents can exercise some degree of control over the level of risk they impose through their actions. So even if we think that outcomes are wholly beyond our control, the level of risk we unleash (and thus the likelihood of bringing about a particular outcome) is not wholly beyond our control. This is true in the context of reckless actions when agents have some control over the means with which they carry out a potentially harmful conduct, such as choosing which route to take when driving intoxicated. The same is also true in the context of purposeful actions.

Consider a pair of shooters—Competent Shooter and Incompetent Shooter—both of whom decide to commit murder by shooting at their respective victims from atop abandoned buildings. Competent Shooter is an Olympic shooting champion and Incompetent Shooter is someone who has never fired a gun before. Upon seeing their victims, both shooters take aim and pull the trigger. Nothing gets in their way; their victims stand still; the agents execute their shots. Alexander and Ferzan would argue that, whatever the outcome may be for each agent, that outcome is beyond the agent's control and thus a matter of luck. But like the case of the reckless drivers, surely the same outcome is not equally a matter of luck for both shooters.

Presumably, compared to Incompetent Shooter, Competent Shooter has greater control over his shot in virtue of his abilities as a skilled marksman. Competent Shooter is able to impose (purposely) a higher risk on his victim than what Incompetent Shooter is able to do.²⁸ If both shooters succeed, surely success is *more* of a matter of luck for Incompetent Shooter

²⁷ In this example, the level of risk depends on the route that a driver decides to take, but it could also depend on other factors such as degree of intoxication, ability to drive while intoxicated, the kind of car the driver has, etc. All of these factors can affect the level of risk an agent imposes, but not all of them are relevant for culpability.

²⁸ Note that Incompetent Shooter may actually impose a higher risk than Competent Shooter by accident. And Incompetent Shooter may also impose higher risks on those other than the intended victim (e.g. by-standers). But only Competent Shooter is able to impose a higher risk purposely because he has greater control over his shot whereas Incompetent Shooter does not.

than it is for Competent Shooter. The latter has greater control over his shot, and thus greater control over the outcome of his attempt. There is a higher chance of success for Competent Shooter compared to Incompetent Shooter. Conversely, if both shooters fail, surely failure is *more* of a matter of luck for Competent Shooter. Given his chances of success, it was more likely than not that he would succeed.

Again, even if both Shooters lack control over the outcomes of their attempts, they do have some control over the level of risk they impose on their respective victims, and thus some control over the probability that the intended outcome would occur. The degree to which an outcome is a matter of luck for a particular shooter, just as it is for a reckless driver, depends on the level of risk his action imposes on others, which is tied to the level of control that he has over his actions. I mentioned earlier that the level of risk an agent imposes is something that speaks to his culpability. So factors that affect culpability can affect the degree to which an outcome is a matter of luck. We might then think that in order to insulate criminal liability from luck, culpability needs to be sensitive to varying degrees of risk.

This is a departure from the way that the problem of outcome luck has typically been framed in the legal philosophical debate. For one thing, we do not typically think that there is a connection between luck and culpability. For another, in our discussion about outcome luck, we explicitly stipulate that the culpability of two agents is identical in order to examine the relevance of outcomes for liability. In the next section, I discuss the issue of culpability in light of what I have just said about luck, risk, and control. The main question has to do with whether and when degrees of risk are relevant to culpability. I will begin by discussing culpability for reckless actions, and why degree of risk is thought to matter for recklessness, before moving on to discuss culpability for purposeful actions.

3 Risk and Culpability

For offences that are committed recklessly, it is easy to appreciate why culpability is sensitive to the level of risk that the defendant imposed through

his action. Consider again the example of the two reckless drivers from earlier. All else being equal, we think badly of Driver S, who imposes an unjustifiable risk on others by driving drunk, but we think even worse of Driver R, who imposes an even greater risk on others by driving drunk on a populated route. Similarly, we think badly of a reckless driver who drives twenty kilometres over the speed limit, but all else being equal, we would think even worse of another reckless driver who drives forty kilometres over the speed limit.

While both drivers are at fault for driving over the prescribed speed limit, the one who drives forty kilometres over the speed limit is *more* at fault than the driver who drives twenty kilometres over the speed limit. And while Driver R and Driver S are both at fault for driving while intoxicated, Driver R is *more* at fault than Driver S because R's conduct exposes more people to the dangers of drunk driving. The actions of both S and R display insufficient concern toward others in virtue of the unjustifiable risks that their actions create.²⁹ All else being equal, a reckless action that unleashes a greater risk on others displays a greater lack of concern for their interests and wellbeing than a reckless action that unleashes a lesser risk.

Matters are different in cases of purposeful actions, or so I argue, where an agent acts with an intention to cause harm. Consider the pair of shooters from earlier: Competent Shooter (who is an Olympic shooting champion) and Incompetent Shooter (who has never used a gun prior to his attempt). Competent Shooter skilfully aims at his victim and, in doing so, imposes a very high risk on his victim's life. Incompetent Shooter clumsily aims at his victim and, in doing so, imposes a low risk on his victim's life. Suppose that each shooter knows, before he fires the shot, the level of risk he is about to impose on his victim. We think badly of Competent Shooter for acting on his intention to commit murder, but do we not think *just as badly* of Incompetent Shooter for acting on the same intention? Both shooters are at fault for acting on their intention to commit murder, but is Competent Shooter

²⁹ See Alexander & Ferzan (2009, chapter 2).

really more at fault than Incompetent Shooter for imposing a higher level of risk on his victim?

Alexander and Ferzan would argue that Competent Shooter is more culpable than Incompetent Shooter in virtue of the risks they impose on their victims. Before I get to their argument, it would be useful to clarify the issue here. For reckless actions, it is intuitively plausible that an agent's culpability is sensitive to the level of risk he imposes on others. But for purposeful actions, it is not so intuitive that the level of risk an agent imposes matters for judgments about the agent's level of fault.³⁰ The issue here is that either the intuitive view is mistaken (and that culpability for both reckless and purposeful actions are sensitive to risk), or there is something significant about purposeful actions that makes risk irrelevant to culpability. In the rest of this section, I argue that risk is irrelevant to the culpability of purposeful actions because culpability here has to do solely with the agent's reasons for action, rather than degrees of risk imposed. To see why this might be the case, we can begin by considering arguments for the opposing view—i.e. that culpability for any offence is sensitive to degrees of risk.

On Alexander and Ferzan's theory of culpability, they take the four culpable mental states in the Model Penal Code—i.e. purpose, knowledge, recklessness, and negligence—and collapse them into one mental state, namely recklessness.³¹ Recklessness (and culpability) in turn is a straightforward function of two factors: the level of risk an agent believes he imposes on others and his reasons for imposing that risk. A reckless agent who imposes a high risk might be less culpable than another reckless agent who imposes a low risk if the former agent's reasons for imposing that risk are better (e.g. more justifiable) than the latter's. The same goes for agents who commit crimes purposely since, on Alexander and Ferzan's view, purpose is merely a form of recklessness. Specifically, they argue that purpose is a special form of recklessness where the agent's reasons for imposing risks on others are presumptively unjustifiable.³² An agent who

³⁰ See Alexander & Ferzan (2009), chapter 2.

³¹ Ibid.

³² Ibid., 39.

purposely imposes a certain level of risk on others is presumptively more culpable than a reckless agent who imposes the same level of risk because the former acts on reasons that are presumptively worse than the latter's. However, if two agents act on identical reasons then culpability is graded according to the level of risk they each impose on others.

An important aspect of Alexander and Ferzan's theory, which helps to explain why they believe culpability is sensitive to degrees of risk imposed for any culpable action, is their notion that culpability is determined by the risk(s) an agent believes he imposes and his reasons for imposing *that particular level of risk*. Returning to our pair of shooters from earlier, if Competent Shooter and Incompetent Shooter attempt to commit murder for identical reasons, then according to Alexander and Ferzan, Competent Shooter is more culpable than Incompetent Shooter since the former imposes a greater risk than the latter.³³ Assuming that both shooters know the level of risk they impose on their victims, Alexander and Ferzan would argue that Competent Shooter is more culpable because he *chose* to impose a higher degree of risk, and Incompetent Shooter is less culpable because he *chose* to impose a lower degree of risk. They argue that Incompetent Shooter could have imposed a higher risk, for example, by shooting his victim up-close, or by taking shooting lessons before attempting to commit murder. For Alexander and Ferzan, the fact that Incompetent Shooter did not choose (but presumably could have chosen) a means that could increase his chance of success is a factor that affects his culpability, even if only minutely.³⁴

Here is the problem with their view of culpability for purposeful actions. The fact that a purposeful actor could have imposed (but did not actually impose) a greater/lesser risk than the one he actually imposed is irrelevant to his culpability.³⁵ For example, suppose I intend to damage your delicate vase. I can throw the vase to you from across the room, believing there is a

³³ Ibid.

³⁴ Ibid., 38-9.

³⁵ Note that if the purposeful actor chose to shoot his victim from afar (thus giving her a higher chance of survival) rather than from up-close because he wants her to survive, then there are questions about half-hearted attempts, and whether they are less culpable than whole-hearted attempts. This point is not relevant here though, where I stipulate that the purposeful actors are wholeheartedly committed to their intended course of action.

good chance you will not be able to catch it in time, or I can simply throw the vase on the floor, believing it will shatter upon hitting the floor. I impose a higher risk on your vase by throwing it on the floor since there is a small but significant chance that you will catch it if I were to throw it to you. However, if I try to damage the vase by throwing it to you, still believing and hoping that you will not catch it in time, then the fact that I could have thrown the vase on the floor, thus almost guaranteeing that I would damage your vase, should be irrelevant to my culpability. In order to see why this might be the case, it will be useful to discuss culpability for reckless actions first, and to explain why culpability is sensitive to risk in such cases.

Reckless actions create substantial and unjustifiable risks to protected interests. Those who act recklessly are culpable because they fail to respond to the right kind of reasons; that is, they have reasons against imposing substantial risks on others and they are at fault for disregarding those reasons.³⁶ In the case of Driver S (who takes the quiet detour) and Driver R (who takes a populated route), neither driver intends to cause harm to other people. Their purpose is to drive home, though they realize that by choosing to drive in their intoxicated states, they create a substantial risk to other drivers and pedestrians on the road. They have reasons not to drive, which include the risks they would impose on others by driving drunk. They undoubtedly have other reasons for driving home, such as not wanting to take a taxi or not wanting to leave their cars behind. However, should they choose to drive, these reasons do not justify or excuse their actions, and they would be at fault for not responding to the reasons they have against driving drunk.

Drivers S and R, like any other agent who acts recklessly, *can* minimize the risk they create despite the fact that their actions impose some substantial and unjustifiable risk on others. There is room, within reckless actions, for agents to act with greater caution even if such agents choose to act in ways that would impose unjustifiable risks on others.³⁷ And when reckless agents fail to minimize the risk their actions create, their culpability

³⁶ See Duff (2008) for a discussion on attacks vs. endangerments.

³⁷ See Duff (2008).

increases in accordance with the level of risk they impose on others. For instance, it is possible for Drivers S and R to drive home drunk and, at the same time, do it cautiously. Each driver can try to minimize the risk his action creates by avoiding well-populated routes and by opting for quiet routes instead. Should a driver choose to take the least populated route, his choice would be compatible with the successful pursuit of his aim, which is to drive home safely. Given that his aim does not include harming others, there would be no inconsistencies within his action should he choose to take the least populated route. In fact, there are good reasons for the drivers to opt for quiet routes even if they disregard any reasons against drunk driving; that is, taking a quiet route would decrease the number of pedestrians and other drivers they would encounter on the road, thereby decreasing their chances of being involved in an accident with other drivers and pedestrians.

Driver R, who decides to take the well-populated route, fails to respond to the right reasons more often than Driver S, who decides to take the quiet route. Driver S's action reflects a disregard for reasons against drunk driving, but it appears to respond to other reasons that are relevant to him. Driver R's action likewise reflects a disregard for reasons against drunk driving, but it further reflects a disregard for reasons he has for taking care even though he decides to drive drunk. The differences in the risks these drivers impose are relevant to their culpability because each driver could have imposed a different risk, and whether or not they do impose a different risk speaks to the set(s) of reasons their actions are responding to. Driver R is more culpable than Driver S because the former's action reflects a greater lack of concern for others and a greater disregard for the range of reasons he should respond to.³⁸

In reckless actions, Alexander and Ferzan's formula for culpability seems to capture the right idea; that an agent's culpability depends on his imposition of *this* level of risk for *these* reasons. Matters are different,

³⁸ Even if there are no quiet routes available to driver R such that he must take a well-populated route, the lack of alternatives does not affect his culpability in any way. He can still choose not to drive and he has reasons against driving drunk. Should he choose to disregard those reasons, he is culpable for drunk driving and for the risks that come with his action. .

however, when we turn to purposeful actions where an agent's aim is to bring about a prohibited harm. Like reckless actions, purposeful criminal actions (typically)³⁹ impose substantial and unjustifiable risks on protected interests. However, those who try to cause harm are not at fault for failing to respond to the right reasons, but are rather at fault for acting on the wrong reasons.⁴⁰ Those who commit crimes purposely are at fault because they are acting on reasons which they should not follow.⁴¹ Consider again the example of Competent Shooter and Incompetent Shooter. Their purpose is to kill their respective victims, and each agent attempts to commit murder by shooting at his victim from atop an abandoned building. Undoubtedly, they have reasons for their actions—e.g. they hate their victims and want them dead—but those are not the kind of reasons that the agents should follow and act on.

Unlike those who commit crimes recklessly, the pair of shooters (as well as any other agent who tries to cause harm purposely) *cannot* minimize the risk their actions create.⁴² This is not to say that the shooters lack the capacity to minimize the risk they unleash, since they can purposefully take a bad shot or refrain from shooting altogether. It is rather that, if the shooters are trying to kill their victims, then it is not logically possible for them to try to minimize the risk they unleash on their victims. There is no room within purposeful actions for agents to act with greater caution if they intend to cause some particular harm.⁴³ When an attacker tries to minimize the risk his action imposes on his intended victim(s), it is questionable whether he is wholeheartedly committed to his attack. His culpability may decrease if he is not wholeheartedly committed. But if he is committed to his attack, then taking steps to minimize the risk his action imposes on an intended victim is incompatible with his action as an attack.

³⁹ Objectively speaking, attempting to commit murder by voodoo or to commit arson by incantation do not impose any risk on others, but there is dispute whether such actions are count as attacks of purposeful actions. See Duff (1996), chapter 8.

⁴⁰ See Duff (2007), chapter. 7.

⁴¹ See Duff (2007), chapter 7.

⁴² See Duff (2008).

⁴³ An agent who attempts to murder his victim can take greater care not to injure other people during his attempt, but he cannot take greater care not to injure his intended victim.

The pair of shooters cannot pursue their intended aim, which is to kill their respective victims, by trying to fire shots that would kill their victim on the one hand, and by trying to fire in such a way that would minimize the risk they create for their victims on the other hand. An attacker, insofar as he aims to inflict harm, cannot carry out an attack that deliberately tries to minimize the level of risk involved.⁴⁴ An agent who attacks property cannot try to minimize the risk of damage to property by deliberately starting a fire with insufficient fuel; and an agent who attacks the physical wellbeing of another person cannot try to minimize the risk of wounding by deliberately throwing a light punch.

Alexander and Ferzan would disagree with what I have said about purposeful actions thus far. Presumably, agents can choose to execute their attacks in different ways, all of which would create different levels of risks, and an agent can choose to execute his attack in a way that creates a very low risk. For Alexander and Ferzan, should an agent choose a low risk means of carrying out an attack, this choice should be relevant to his culpability. In the case of the shooters, Competent Shooter's attempt is likely to succeed since he is a very skilled and his shot will most likely hit and kill his victim. Incompetent Shooter's chances of success, on the other hand, are much lower. Instead of trying to shoot his victim from afar, Incompetent Shooter can approach his victim and try to shoot him from up-close, which would significantly increase his chances of success. Incompetent Shooter might have a variety of reasons against choosing this alternative. He might be a coward and would rather make his attempt from far away; he fears being caught by the victim's bodyguards; he fears he might have a change of heart from up close.⁴⁵ Although these reasons may figure in Incompetent Shooter's decision to try to shoot his victim from afar—i.e. they are relevant to his decision to take a particular course of action—they are irrelevant to his

⁴⁴ It could be argued that matters are different in cases of self-defence, where an agent tries to harm his attacker up to a point that is necessary to defend himself, but he can take care not to inflict more harm than is necessary. However, this does not refute my point. Suppose the agent thinks he needs to wound his attacker in order to defend himself. He can take care not to kill his attacker by wounding him in the leg or arm, but avoiding wounding him in the chest, but he cannot take care not to wound his attacker at all.

⁴⁵ Alexander & Ferzan (2009), chapter. 2

culpability. These reasons help to explain why Incompetent Shooter does not choose a means that would increase his chances of success, but they do not detract from other reasons that are guiding his action. The other reasons are that Incompetent Shooter hates his victim and wishes to see him dead, reasons that he should not act on. What makes the shooter culpable is that he does act on these reasons. His reasons against shooting his victim from up close are different from his reasons to commit murder; they do not detract from his purpose nor from the unjustifiability of his intended aim, which is what speaks to his culpability.

So far, I have argued that culpability is sensitive to degrees of risk in reckless actions but not in purposeful actions. The action of the reckless driver who creates a higher risk reflects a greater lack of concern for the wellbeing of others and is thus more culpable than a reckless action that creates a lower degree of risk. But in purposeful actions, an agent is culpable because his action stems from the wrong reasons, which the agent should not follow.⁴⁶ Incompetent Shooter is just as culpable as Competent Shooter despite the fact that the former's action imposes a lower risk since both actions, and the risk they create, are equally unjustifiably or guided by equally bad reasons. The degree of risk their actions create, as well as the reasons they have against imposing a higher or lower risk, are irrelevant as long as they are committed to bringing about the intended harm. Differences in the level of risk they impose do not detract from or add to the wrongfulness of their actions.

4 Outcome Luck, Risk, and Fair Impositions of Liability

In the preceding sections, I have argued that we can account for degrees of luck in terms of the level of risk an agent imposes (which in turn can be accounted for by the level of control an agent exercises over his action). I have also argued that degrees of risk are irrelevant for the culpability of purposeful actions. So while Alexander and Ferzan's theory of culpability is attractive for insulating criminal liability from degrees of luck, we should

⁴⁶ See Duff (2008).

nonetheless reject their view on the grounds that it gives us a problematic view of culpability for purposeful actions. What I have said thus far has implications for fair assignments of liability. If degrees of luck are attributable to degrees of risk, but degrees of risk are irrelevant for culpability, do we not face the problem of subjecting criminal liability for purposeful actions to the influence of luck?

In order to address this point, we must clarify the language of luck and its role throughout different points of the outcome luck debate. That is, we must identify the various issues that “luck” is meant to be tracking. First, there is luck in outcomes; that is, the outcome of a culpable action or omission can be a matter of luck.⁴⁷ It was a matter of luck that Driver R made it home safely, and it was a matter of luck that Competent Shooter failed to kill his victim. Here, the language of luck is meant to be tracking the issue of control. As in, if Competent Shooter tried but failed to kill V, his failed attempt is a matter of luck to the extent that it is beyond his control. In addition to luck in outcomes, there is also luck in holding an agent responsible for the outcomes of his actions. If the consequences of an agent’s actions are to some extent a matter of luck, then it is also to that extent a matter of luck to be held responsible for those consequences. Here, the language of luck is again tracking the issue of control, but what agents lack is complete control over the extent of their liability for the outcomes of their actions. This leads us to another way in which the language of luck is used. It is used to track issues of justice and fairness in the way that we assign criminal liability.⁴⁸ If Competent Shooter tries but fails to harm his victim, is it fair for him to receive a more lenient sentence for his failed attempt than if he had succeed?⁴⁹ The failed attempt does not properly reflect the shooter’s intentions and desires, so is it acceptable for the shooter to be punished less—if he is punished less—for an outcome that is to some extent beyond his control?

⁴⁷ See for examples Alexander & Ferzan (2009, 191); Duff (1996, 327); Kadish (1994, 685); Sverdlik (1988, 79); Zimmerman (2002, 559).

⁴⁸ See Lewis (1989).

⁴⁹ See Nagel (1979).

The issue of luck in outcomes and in holding agents responsible for those outcomes—i.e. the issue of lack of control—opens up questions about justice and fairness in the way that we assign criminal responsibility. But while the extent to which an outcome is a matter of luck necessarily affects the extent to which being held responsible for that outcome is a matter of luck, neither of these issues necessarily affect the question about fairness. For instance, if it is unfair or unjust for an agent to be liable and punished for a particular outcome, it does not mean that the outcome was necessarily a matter of luck. Conversely, if an outcome (and being held responsible for that outcome) is to some extent beyond an agent's control and thus to some extent a matter of luck, it does not necessarily mean that it would be unfair (i.e. unlucky) for that agent to be liable and punished for the outcome. It would be fair, for instance, to be held responsible for outcomes if the exercise of some (as opposed to complete) control is sufficient for criminal liability, *and* if we are capable of having some control over outcomes.

The problem of outcome luck is a question about the extent of the connection between control and fair assignments of responsibility. To say that an outcome is to some degree a matter of luck is not yet to admit that it would be unfair to be held responsible for that outcome. Going back to the discussion on purposeful actions, there may be degrees of luck that come from the level of culpability. But given that degrees of risk is irrelevant for culpability here, it would not be unfair for the agent's culpability to be insensitive to risk (and therefore to the influence of luck). For subjectivists like Alexander and Ferzan, who argue that liability should be grounded in an agent's culpability only, there is nothing unfair about imposing the same level of punishment on Competent Shooter and Incompetent Shooter, even though the same outcome is more of a matter of luck for one of them.

There is still the question of whether it would be fair to hold agents responsible for the outcomes of their actions, something that I have not commented on throughout this chapter. Attempts to resolve this debate requires discussions on the justification of punishment, which I will offer in Chapter 4 and 5, and it requires discussions on the debate between

subjectivists and objectivists about criminal liability, which I do not offer in this thesis. I will say something about the relevance of outcomes for liability in the conclusion of this thesis. For now, I want to turn our attention to a problem in the philosophy of criminal law, that of impossible attempts. The next chapter will illuminate some foundational issues in the *mens rea* requirement for criminal attempts. These issues are particularly relevant when it comes to the problem of outcome luck in the context of failed attempts.

Chapter 3

Recklessness and Circumstances at the Attack– Endangerment Divide

When we consider the influence of luck on the criminal law, the problem of impossible attempts inevitably surfaces. These are cases in which, given the actual state of affairs, it is impossible for an agent to commit the intended crime.¹ For instance, it is impossible for D to commit murder given that, unbeknownst to D, the would-be victim is already dead. Similarly, it is impossible for D to damage V's property, given that, unbeknownst to D, the property in question belongs to D and not to V. And it is impossible for D to commit the crime of adultery given that, unbeknownst to D, adultery is not a criminal offence. In these examples, what makes the attempt impossible—e.g. the status of the victim, the ownership of the property, and the legality of the intended act—are to some extent a matter of luck. Had D found V a few minutes earlier, V would still have been alive and D might have committed murder. And had D decided to damage another piece of property, he might have actually damaged something that belonged to V. And had D been in a different jurisdiction, one that criminalized adultery, D would have committed an offence.

In Canada² and the United Kingdom³, impossibility alone does not preclude criminal liability. However, there is wide agreement that cases of “legal impossibility” should be exempt from liability. And, in practice, cases of “inherent impossibility” are typically exempt as well. Legal impossibility occurs where the intended act is not prohibited by the criminal law. Given that what is criminalised in criminal attempts are attempts to commit other offences, cases of legal impossibility are exempt from liability simply because it is not an offence to attempt to commit acts which have not been

¹ See Alexander (1993); Duff (1996), chapter 3; Fletcher (1986); Hart (1981); Westen (2008).

² See *United States of America v. Dynar*, [1997] 2 SCR 462.

³ See Simester & Sullivan (2013, 359-66).

criminalised. In other words, there is no offence with which to charge an agent who has merely committed what he believes is a criminal offence. Inherent impossibility describes cases where the manner with which an agent tries to commit an offence is utterly misguided and inherently doomed (e.g. murder by voodoo or arson by incantations).⁴ Normally, these cases are either not prosecuted because it not in the public's interest to hold such hopelessly misguided criminals to account, or they would be left unreported given the absurdity of what the agent tried to do.

Other types of impossibility generate much more controversy. And even when legal theorists agree on how the criminal law should treat a particular instance or type of impossibility, there is nonetheless disagreement over the rationale for treating the case in that way. This chapter will look at a variation of "attempts" which are thought to be impossible of success given that a relevant circumstance element of the completed offence is missing. For instance, suppose that D tries to damage V's property or that D tries to rape V. The completed offence of criminal damage requires D to damage property that belongs to another person.⁵ And the completed offence of rape requires that D sexually penetrates V who is non-consenting.⁶ Given these requirements, it is impossible for D to commit criminal damage by damaging property that does not belong to another person, and it is impossible for D to commit rape by sexually penetrating a person who is actually consenting. But the fact that, unbeknownst to D, the property does not belong to another person and that V is actually consenting does not detract from his culpability nor (as some would argue) should it exempt him from criminal liability.

Before I say anything about liability, we need to distinguish two variations of the examples that I just described. In one variation, D acts with an intention [to damage V's property] or [to have non-consensual intercourse with V].⁷ Here, D intends *both* the consequence element of the offence (e.g.

⁴ See Alexander (1993) on why we could think of all failed attempts as cases of inherent impossibility. See also Yaffe (2010), chapter 9.

⁵ Simester & Sullivan (2013, 593-603)

⁶ *Ibid.*, 465-66.

⁷ The brackets are meant to mark out the *content* of an agent's intention and to differentiate it from the *context* of that intention. For instance, there is a difference between D who

to damage property or to have intercourse⁸) and the circumstance element of the offence (e.g. that the property belongs to V, or that V does not consent to intercourse). That is, D does not merely intend [to damage property] or [to have intercourse with V], but he also intends that the property in question belongs to V and that V does not consent to intercourse.

This chapter will not be concerned with the variation of impossibility that I have just described. Instead, I will focus on another variation where D acts with an intention to the consequence element of the offence, but only with recklessness as to the circumstance element of that offence.⁹ Here, D intends [to damage property] or [to have intercourse], but he does not intend that the property belongs to another person, nor does he intend that the person with whom he is having intercourse is non-consenting. Nonetheless, he foresees that there is a risk that the property belongs to another person, or that the other person may not consent. And he acts with recklessness as to these circumstances when he acts with an awareness of the risks involved.

Under English law, the *mens rea* requirement for criminal attempts has become somewhat confusing in recent years. Many substantive or completed offences can be committed either intentionally or recklessly, but the *mens rea* for attempt requires an *intention* to commit the completed offence.¹⁰ The precise meaning of acting with an intention to commit a completed offence is open to interpretation. There is consensus that D acts with an intention to commit an offence when he intends both the consequence and circumstance element of an offence. But it is contentious whether D acts with intention

intends [to cut down this tree] and D who intends [to cut down his neighbour's tree]. To the former D, ownership of the tree is no part of his intention, while for the latter D, it is part of his intention that the tree belongs to his neighbour. For the latter D, who intends to [cut down his neighbour's tree], his action would be an offence if the context of his action is such that his neighbour does not consent to the tree being cut down, and that it is an offence to cut down another's tree. See Duff (1996), chapter 1 on the distinction between the content and context of an agent's intention.

⁸ In the case of rape, to intend [to have intercourse] is to intend the conduct element rather than the consequence element of the offence. This distinction is immaterial for the purpose of this chapter.

⁹ See also Duff (2012).

¹⁰ See Simester & Sullivan (2016, 346-55).

when he intends the consequence element but acts with recklessness as to the circumstance element of an offence.

In 1990, the Court of Appeal ruled in *Khan*—a case in which D tried but failed [to have intercourse with V] while aware of the risk that V may not consent—that recklessness as to consent would suffice for intention.¹¹ The Court noted that, given that V was actually non-consenting, D would have committed the completed offence of rape had he succeeded in having intercourse with V. It would not be in the interest of public protection to absolve D of liability, so the Court ruled that the *mens rea* requirement (at least) for attempted rape should not depart too much from the requirement for the completed offence.¹²

Then, in 2014, the Court of Appeal considered *Pace*, a case where D purchased goods which he suspected were stolen, but were, unbeknownst to him, not stolen (since they belonged to the police who had set up the sting operation). Here, the Court ruled that the *mens rea* for attempts required intention as to every element of the offence.¹³ In other words, in order for D to be liable for an attempt, he had to intend the goods to be stolen. But the court also argued that knowledge or belief that the goods are stolen would suffice for an intention that the goods be stolen.¹⁴ In making this ruling, it appears as though the Court has contradicted its earlier decision in *Khan*, but the decision in *Khan* was not overruled in *Pace*.¹⁵

In light of the decisions in *Khan* and *Pace*, one way to interpret the *mens rea* requirement for attempt under English law is that, for certain offences such as rape, D must act with intention as to the consequence (or conduct) element of the offence and recklessness as to the circumstance element of that offence (assuming that recklessness as to circumstance suffices for the completed offence). But, also, that for certain other offences such as converting stolen property, D must act with intention as to the

¹¹ *Ibid.*, 347-50.

¹² *Ibid.*, 349.

¹³ *Ibid.*, 347-50.

¹⁴ This is in keeping with the definition of intention in English law, where the consequences that D foresees as being virtually certain are part of D's intention. See Simester & Sullivan (2013, 132-7).

¹⁵ Simester & Sullivan (2016, 350).

consequence element of the offence and with knowledge or belief as to the circumstance element of that offence. Under this doctrine, D is liable for an attempt to convert stolen property if he tries but fails [to convert property] with knowledge or belief that the property is stolen. But D would be guilty of attempted rape when he tries but fails [to have sexual intercourse] and is merely reckless about V's lack of consent.

This chapter argues that, with respect to some cases, such as criminal damage or crimes involving bodily harm, the ruling in *Pace* is correct; that knowledge or belief as to circumstances is required for an attempt. However, the offence of rape may or may not be an exception to this rule. It is possible that recklessness as to consent would suffice for attempted rape. The aim of this chapter is to explain why knowledge or belief as to circumstance is required for an attempt to commit some offences, and to explore why attempted rape may have a lower culpability requirement. The discussion on rape will be exploratory, but if we have principled reasons for the ruling in *Pace*, then at least with respect to some offence, we will be able to deal with some types of "impossible attempts" in a more straightforward manner.

Before I outline the structure of this chapter, it would be useful to list the range of cases that will be discussed throughout.

- (1) D₁ intends [to damage property]; D₁ is reckless as to ownership of the property; the property belongs to V₁; D₁ damages the property.
- (2) D₂ intends [to damage property]; D₂ is reckless as to ownership of the property; the property belongs to V₂; D₂ tries but fails to damage the property.
- (3) D₃ intends [to damage property]; D₃ is reckless as to ownership of the property; the property belongs to D₃; D₃ damages the property.
- (4) D₄ intends [to damage property]; D₄ is reckless as to ownership of the property; the property belongs to D₄; D₄ tries but fails to damage the property.

For all of these cases, we can substitute property damage and ownership of property with sexual intercourse and consent to sexual intercourse. And,

again, we should distinguish these cases from ones in which D intends [to damage V's property] or [to have non-consensual intercourse].¹⁶

Scenario (1) is a case of criminal damage, which, like many completed offences, can be committed either intentionally or recklessly. There is no controversy that D₁ is guilty of criminal damage. Some theorists, most notably Duff, have argued, contrary to the decision in *Pace*, that scenario (2) amounts to a case of attempted criminal damage.¹⁷ For Duff, recklessness as to circumstances suffices for an attempt, whether the offence attempted is rape or criminal damage (again assuming that recklessness as to circumstances suffices for the completed offence). Scenarios (3) and (4) are sometimes considered to be "impossible attempts" since it is impossible for D to commit criminal damage by damaging his own property. Note that, if we think (2) is an attempt, we have reason to think the same about scenarios (3) and (4). The only difference between (2) and (4) has to do with ownership of the property in question. And the only difference between (4) and (3) has to do with whether D succeeds in damaging the property. Both ownership of the property and D's success are, to some extent, matters of luck.

In Duff's writing on this, he argues that (2) amounts to an attempt, but denies that (3) and (4) should be seen as attempts.¹⁸ On the face of it, this certainly appears to be a rather odd view since there do not appear to be any principled differences between (2), (3), and (4). And once we admit that these differences can be (and often are) a matter of luck, there is a strong *prima facie* case for judging D₃ and D₄ to be guilty of an attempt if we hold D₂ liable for an attempt. I share Duff's view that D₃ and D₄ should be acquitted of an attempt, but unlike Duff, I am not committed to holding D₂ liable for an attempt. Assuming that D₃ and D₄ should be acquitted of an attempt, those who defend Duff's view face the difficult task of explaining why they should be acquitted, especially when D₃ and D₄ are just as culpable as D₂, and it is

¹⁶ See Duff (2012).

¹⁷ See Duff (1991; 1996, 374).

¹⁸ See Duff (1996), chapter. 3. Whatever reasons we might have for acquitting the defendant in (3)—a defendant who has damaged property—of an attempt, those reasons should acquit the defendant in (4)—a similar defendant who has *not* damaged property—of an attempt as well. See Duff (2012).

not immediately obvious what principled reasons we have for exempting D_3 and D_4 from liability if D_2 is liable for an attempt.

In this chapter, I argue that part of the difficulty lies in Duff's analysis of (2). If (2) amounts to an attempt, then the task of drawing a coherent and principled line between scenario (2) on the one hand and scenarios (3) and (4) on the other hand will be difficult. But, as I will argue, we should not see (2) as an attempt, but rather as a case of non-consummated endangerment. I begin in sections 1 and 2 by restating Duff's distinction between attacks and endangerment, and by reiterating his view that we should see the law of attempts as a law of attacks. Then in section 3, I explain Duff's rationale for seeing (2) as a failed attack before offering my argument, in section 4, for seeing (2) as a case of non-consummated endangerment. Once we clarify why (2) amounts to an instance of endangerment, we can easily explain why the defendants in (3) and (4) should be exempt from criminal liability. Lastly, in section 5, I explore a possible exception to my argument. In the case of rape, it is not clear to me whether a defendant attacks or endangers another's sexual integrity if he intends [to have sexual intercourse] and is reckless as to the other person's consent. Much of this discussion will be exploratory as I try to identify relevant questions for further research.

1 Attacks and Endangerments

Attacks and endangerments mark out different kinds of moral wrongs. The distinction between these kinds of wrongs turns on the intentional structure of an agent's action (or omission) and the attitude he takes towards the prospect of harm.¹⁹ An attack on a legally protected interest requires an intention on the part of the attacker [to cause harm], which is prohibited by the criminal law. D attacks V's property if he purposely damages or appropriates her property, and D attacks V's life if he purposely kills her. The attacker's attitude towards the interests or people he attacks is one of practical hostility. His action (or omission) is guided by a direct intention to cause harm and it is directed against the object or subject that the agent

¹⁹ Duff (2008).

intends to harm.²⁰ An attacker would only count his endeavour a success if the intended harm occurs, and he would count it a failure if the harm does not occur.²¹

While attacks require a direct intention to cause harm coupled with an action or omission that executes the harmful intention, endangerments merely require the creation of a significant risk of harm to a person or interest.²² For instance, D endangers V's car if he takes it to a race track thereby creating a risk that V's car will be damaged. And a hunter endangers another's physical wellbeing if he absent-mindedly fires his rifle in the woods thereby creating a risk that other hunters will be injured by a stray bullet. The paradigm fault element for endangerment is recklessness, where the agent consciously disregards a significant and unjustifiable risk that his action unleashes on others. Unlike attacks, the attitude of someone who recklessly endangers the rights or interests of others is not one of practical hostility, but one of practical indifference. His action or omissions displays an unjustifiable willingness to risk harm to others, but the action is not aimed at harming a protected interest.²³

Many consummated offences can be committed either as an attack or as an endangerment.²⁴ Under English law, a defendant is guilty of assault occasioning bodily harm whether he intentionally or recklessly causes bodily injury to his victim.²⁵ The defendant attacks his victim if he acts with an intention to cause injury, but the defendant endangers his victim if he is reckless as to the risks he creates against the victim's bodily integrity.²⁶ Similarly, a defendant is guilty of criminal damage whether he intentionally or recklessly damages another's property.²⁷ The defendant attacks another's property if he acts with an intention to damage her property, but the defendant endangers another's property if he is reckless as to the risks he

²⁰ Ibid., 943-4.

²¹ Ibid., 945.

²² Ibid., 943-4.

²³ Ibid., 944.

²⁴ Ibid., 949.

²⁵ Simester & Sullivan (2013, 436; 439-42).

²⁶ See Duff (2008).

²⁷ Ibid., 593-603.

creates against her property. An attacker, in deliberately (and unjustifiably) trying to harm a protected interest, is acting for the wrong reasons.²⁸ As Duff puts it, his action is *essentially* harmful since the causation of harm is intrinsic to his action, and he ought not to be guided by whatever reasons that motivate him to harm others.²⁹ In the context of endangerment, an agent, in carrying out a risky action, may be pursuing legitimate ends. The risk of harm his action creates is not the object but rather a side-effect of his action; he would not count his endeavour a failure if the potential harm does not occur.³⁰ However, by disregarding the unjustifiable risks his action creates against others, he is failing to respond to reasons against imposing unjustifiable risks on others.³¹ Unlike attacks, endangerments are *potentially* harmful since the threat of harm is not intrinsic, but incidental, to the action.

These differences between attacks and endangerments do not imply that attacks are necessarily more culpable than endangerments. An agent who endangers others may be (and often is) less culpable than an attacker, but an agent who is extremely reckless as to the wellbeing of others may be as culpable as another who attacks others' wellbeing. All else being equal, an attack may or may not be more culpable than an act of endangerment. I will not dwell on this issue here,³² but I mention it in order to clarify a point that is relevant to my main argument later in this paper. Specifically, when I argue that the defendants in (1)—(4) endanger (and harm) the property of others when they act with an intention [to damage property], I do not suggest that these cases are necessarily less culpable than clear cases of attacks, where the defendants act with an intention [to damage another person's property]. The attack and endangerment distinction is significant not because it marks differences in degrees of culpability, but because it marks different *kinds* of moral wrongs. When I argue that an agent endangers another's property interests when he acts with an intention [to damage property], I am arguing

²⁸ Duff (2008).

²⁹ Duff (1996, 364-5).

³⁰ Duff (2008), 14.

³¹ *Ibid.*

³² See Simester (1996).

that this agent perpetrates a different kind of moral wrong than another who acts with an intention [to damage another person's property].³³

Cases that resemble (1) and (2) are, according to Duff, attacks and failed attacks. Duff defends an objective account of attempts liability which counts as attacks actions that are *directed against* what is *actually* a protected interest.³⁴ So the defendants' actions in (1) and (2) count as attacks because they are directed at an outcome which is (or would be) harmful given the actual circumstances.³⁵ The defendants may not intend [to attack V's property], but their actions are directed at an outcome (i.e. property damage) which *in fact* harms (or would harm) V's property given that the property belongs to her. The context in which the defendants execute their intention is one where the property belongs to another, such that what actually happened is that the defendant in (1) did [damage property] belonging to another and the defendant in (2) did try [to damage property] belonging to another.³⁶ From the victims' perspectives, the defendants did not just damage or attempt to damage property, but rather attacked or tried to attack what is in fact their property.³⁷ This is a very brief summary of Duff's objective account of attempts liability, which I mention here in order to hint at why (1) and (2) might be thought of as (failed) attacks. I will discuss Duff's view in detail in section 3. For now, I will turn to the law of attempts and its relation to attacks.

2 Law of Attempts as a Law of Attacks

Many jurisdictions have a general law of attempts, where it is an offence to attempt to commit most indictable offences (excluding inchoate offences such as conspiracy, aiding or abetting, etc.).³⁸ As we have already seen, in

³³ Note that we need not accept Duff's distinction between attacks and endangerments to think that it would nonetheless be worthwhile to differentiate intention and recklessness as significantly different *kinds* of culpability. See Ferzan (2009) for a critique of the attack-endangerment distinction. But see also Alexander & Ferzan (2009) chapter 2, who argue that purpose and knowledge are merely forms of recklessness.

³⁴ Duff (1996, 372).

³⁵ See Duff (1996, 371-4; 1995, 319-24).

³⁶ See Duff (1996), chapter 1 on the content and the context of an intention.

³⁷ Duff (1996, 372-4).

³⁸ Simester & Sullivan (2013), 343.

England (and many Anglo-American jurisdictions), attempts require an intention to commit the relevant offence. This is in contrast to a broader law of inchoate offences, which may require the same fault element for an inchoate offence as it does for the corresponding completed offence.³⁹ This means, for example, that not only is it an offence to endanger and harm another's property, but it would also be an offence to merely endanger another's property without causing harm. Two questions are relevant here. First, why should a law of attempts require intention to commit an offence, especially when intention is not required for many completed offences that might be attempted? And second, given that many offences can be committed either intentionally or recklessly, why should we only have a law of attempts (which requires intent) rather than a broader law of inchoate offences (e.g. a law of endangerment)?⁴⁰

One explanation for why attempts should require intention is that it fits with our ordinary understanding of "attempt". To attempt to ϕ is to act with an intention as to the occurrence of ϕ . We would normally say that D attempts to damage V's property if D acts with an intention to damage V's property, but not if D merely foresees damage to V's property as a side-effect of his action.⁴¹ If we take the principle of fair labelling seriously—"that the label applied to an offence ought fairly to represent the offender's wrongdoing"⁴²—then the law of attempts should require intention to cause harm, even if the corresponding completed offence does not. If "attempt" was taken in criminal law to include both intentional and reckless harming, both offenders and victims could justifiably complain that the offence label distorts what they have done or what they have suffered. And when offence labels, rather than descriptions or particulars of offences, receive the most attention (as they often do) from the general public (including potential employers), it is in the offender, the victim, and the public's interest to have elements of an offence

³⁹ See Duff (1996, 362).

⁴⁰ See Duff (1996), chapter 13.

⁴¹ Duff (1996, 20).

⁴² Ashworth, A. (1981, 53). See also Chalmers & Leverick (2008).

align as closely as possible to our ordinary understanding of the offence label.⁴³

For Duff, to have a general law of attempts as opposed to a broader law of inchoate offences is to have a general law of attacks rather than a general law of endangerment.⁴⁴ Given an ordinary language understanding of “attempt”, the agent’s action (or omission) in an attempt, much like the action in an attack, is directed towards a particular wrong. What differentiates an attack from an attempt is merely causation of the intended harm. On this view, attempts differ from non-consummated acts of endangerment in the same way that attacks differ from consummated acts of endangerment. In attempts and attacks, the action is intrinsically harmful since it aims at causing a legally relevant harm, while in acts of (non-consummated) endangerment the action is potentially harmful as it does not aim at, but rather risks as a side-effect, the occurrence of harm.⁴⁵

If we take attempts to require intention to cause harm, then our law of criminal attempts is a law of attacks, and it differentiates attempts from (non-consummated) acts of endangerment as a different kind of moral wrong.⁴⁶ If less stringent fault elements such as knowledge or recklessness were sufficient for attempt, then the law of attempts would become a law of endangerments. We would then either wrongfully label as criminal “attempters” those who have not tried to commit a criminal offence, or we would have to replace the law of attempts with a broader—and descriptively accurate—law of inchoate offences which includes both failed attacks and non-consummated endangerments. This brings us to the second question. Why should we only have a general law of attacks without a general law of endangerments?

On Duff’s distinction between attacks and endangerments, attacks are, paradigmatically, the kind of action that should be of concern to the criminal law. Not only are completed attacks wrongs that cause harm, but

⁴³ See Chalmers & Leverick (2008).

⁴⁴ Duff (1996, 366).

⁴⁵ See Duff (2008).

⁴⁶ *Ibid.*

what makes them paradigmatically wrongful is that they unjustifiably *aim* at legally relevant harms. An attacker, in carrying out his intention to harm another, relates himself, in Duff's words, "as an agent as closely as he can to that harm".⁴⁷ His action, when it causes harm, is both wrongful and harmful. But it is still paradigmatically wrongful when it does not cause harm because the action was nonetheless intended to be harmful.

In a consummated act of endangerment, the action is still both harmful and wrongful. However, the agent, in acting not with an intention but the knowledge or foresight that he may cause harm, does not relate himself as closely as he can to that harm. In acting with the knowledge or foresight of harm, the agent's action is not directed towards the causation of harm and is, again in Duff's words, potentially harmful rather than intrinsically harmful. He can hope that his reckless act of endangerment will not cause harm whereas an attacker cannot, and he can be relieved when his action does not cause harm whereas an attacker cannot.⁴⁸

These differences between attacks and endangerments should matter, again, because they convey significant differences in what an agent has done. Such differences matter not only for fair labelling concerns, but also for the purpose of punishment.⁴⁹ An agent who has endangered another's property has to answer for an action that is distinctly different from that of an agent who has attacked another's property. The process of reform, for the agent who has endangered another, requires that he give sufficient attention and concern to the rights and interests of others. But for an attacker, the process of reform requires that he appreciate why we should not purposely violate the rights and interests of others.

To have a general law of attempts—i.e. a law that criminalizes failed attacks—is to criminalize intrinsically wrongful actions. Insofar as we have reasons to criminalize wrongful actions that cause harm, we have the same reasons to criminalize wrongful actions that try but fail to cause harm. Arguably, the same can be said for a general law of non-consummated

⁴⁷ Duff (1996, 366)

⁴⁸ See Duff (2016).

⁴⁹ See Duff (2008, 957).

endangerments—that we also have the same reasons to criminalize potentially harmful actions, ones in which the agent recklessly risks harm. And that in addition to attacks, we would recognize endangerment as a different paradigm of criminal wrongs, one that is characterized by the agent’s lack of sufficient concern for the rights and interests of others.

Many (if not most) jurisdictions have specific endangerment offences but few have enacted general endangerment offences and for good reason. The rationale for enacting specific endangerment offences such as dangerous driving or possession of firearms is that these acts pose serious risks and, without criminalization and policing, would be prevalent enough to cause substantial harm to others.⁵⁰ Specific endangerment offences can provide citizens with guidance and fair warning about the kind of conduct that they may or may not engage in since what is criminalized is a specific act which is thought to be unacceptably risky. Or, they can provide citizens with guidance about the kind of harm that is to be avoided, which is usually quite serious, such as death or serious bodily injury.⁵¹ In contrast, general endangerment offences would do neither and would thus have the potential to be overly broad and vague.

For instance, many of our routine activities such as driving or handling kitchen knives can risk serious harm to others, but the chances that harm would occur from driving or handling kitchen knives are very low compared to dangerous driving and possessing firearms. These routine activities could come under the umbrella of a general endangerment offence, where it would be an offence to create a substantial and unjustifiable risk to others through the pursuit of these activities. Quite often though, it is impossible to judge whether the risk is unjustifiable, or whether the activity actually poses any risks until harm actually occurs. This would not only make it impossible to offer citizens guidance with regard to their conduct, but it also has the potential of being abused through selective and over-zealous enforcement.

⁵⁰ See Clarkson (2005).

⁵¹ Specific acts of endangerment can also target specific agent such as those with special responsibility, or especially vulnerable victims. See Duff (2008).

Given that many specific endangerment offences already cover serious cases of risk creation, we have no need for a general law of endangerment, especially in light of the costs of criminalisation. Later in this chapter, when I argue that the defendant in (2) is not guilty of an *attempt* to commit an offence (e.g. criminal damage), the question of whether he should be exonerated will depend on whether he should be liable for a specific act of non-consummated endangerment, and whether we have good reasons to criminalize that act.

3 Recklessness, Circumstances, and Attempts

Having described the differences between attacks and endangerments, we can return to our scenarios from the beginning of this paper. The focus here will be on cases (1) and (2).

- (1) D_1 intends [to damage property]; D_1 is reckless as to ownership of the property; the property belongs to V_1 ; D_1 damages the property.
- (2) D_2 intends [to damage property]; D_2 is reckless as to ownership of the property; the property belongs to V_2 ; D_2 tries but fails to damage the property.

In this section, I explain why some theorists, most notably Duff, see (1) as an attack and (2) as a failed attempt before offering my argument, in the next section, for why we should see these cases as acts of endangerment.

Let us start with (1), which is a straightforward case of criminal damage, which can be committed either intentionally or recklessly. The *mens rea* for this completed offence requires that the defendant acts with intention or *recklessness* to damaging property, and that he is at least reckless as to whether the property belongs to another.⁵² Under current English law, the defendant in (1) is guilty of criminal damage since he intentionally damaged property and was reckless as to ownership of the property. This does not yet tell us whether (1) is an attack or an endangerment, but it is clear that, under English law, D_1 is guilty of the completed offence of criminal damage.

⁵² Simester & Sullivan (2013, 593).

Now consider scenario (2). Criminal attempts require an intention to commit the completed offence, but an intention to commit the completed offence does not require that the agent act with intention to every aspect of the offence. As we have seen in the introduction, the Court in *Pace* ruled that an agent acts with an intention to commit the completed offence when he acts with intention to the consequence element of the offence, but only with knowledge or belief as to the circumstance elements of that offence. And that in *Kahn*, the Court ruled that attempt (in the case of rape) requires intention as to the consequence element of an offence, but only reckless as to the circumstance element of that offence. Following the ruling in *Pace*, the defendant in (2) would not be guilty of an attempt to commit criminal damage. But if scenario (2) concerned an agent who tried but failed to have sexual intercourse while being reckless about the other person's consent, then according to *Kahn*, the defendant would be guilty of attempted rape.

Duff's view of the *mens rea* for attempts, for any kind of offence, is aligned with the decision in *Kahn*.⁵³ And both Duff's view and the ruling in *Kahn* are attractive for several reasons.⁵⁴ For one thing, we can contrast it with the doctrine that criminal attempts require intention as to every element of the completed offence. This more stringent doctrine would acquit the defendant in (2) of an attempt, whether he tried but failed [to damage property] or [to have sexual intercourse], since it is no part of his intention that the property belongs to another or that the would-be victim be non-consenting. Proponents of Duff's view and of the ruling in *Kahn* would argue that surely we should reject the more stringent doctrine (and perhaps the ruling in *Pace*) since we would not want the defendant to be acquitted.⁵⁵ (Indeed, this reasoning appeared to guide the decision in *Khan*.) The defendant's conduct (especially in the case of rape) is a serious wrong and it is not in the interest of justice to acquit the defendant.⁵⁶

⁵³ See Duff (1991; 1996).

⁵⁴ Duff (1991, 103).

⁵⁵ *Ibid*.

⁵⁶ Duff's view is also attractive when compared with a broader law of inchoate offences, which would require the same *mens rea* for an attempt as it would for the completed offence (see section 2 above).

The intuitive appeal of this argument disappears quickly once we differentiate being acquitted of a *criminal attempt* from being acquitted of a *criminal offence* altogether. The fact that most jurisdictions have a general law of attempts means that if an agent's action is intended to cause a prohibited harm, he is guilty of the completed offence if he succeeds, and of an attempt to commit the completed offence if he fails. If an agent is reckless as to whether his conduct may harm others, then he is guilty of a completed offence if he causes harm. But the absence of a general law of endangerment means that, depending on the conduct in question, he may or may not be guilty of an offence of (non-consummated) endangerment if he does not cause harm. So the defendant in (2) may not be guilty of any offence in jurisdictions that do not have either a general law of endangerment or a specific endangerment offence relating to criminal damage. But the fact that D₂ may not be convicted of any existing criminal offence should not be seen as a reason to convict the defendant of attempted criminal damage if his action is really one of endangerment. Instead, we should assess whether endangering property should be dealt with by a specific endangerment offence.

Returning to Duff's view and the ruling in *Kahn*, where the defendant in (2) would be convicted of attempted criminal damage or attempted rape, Duff seems to think that his view is attractive for another reason. He argues that to convict the defendant in (2) of attempted criminal damage or attempted rape—i.e. to label him as someone who has *attempted* to damage another's property or *attempted* to rape someone—is consistent with our ordinary understanding of "attempt".⁵⁷ This argument gets to the heart of Duff's view on the circumstance elements of attempts. It requires detailed explanation and will occupy the rest of this section. Given that I think the offence of rape is not very clear cut, the rest of this section will focus on the example of criminal damage.

Recall that for Duff, attempts are failed attacks. To label a defendant as an attempter is to indicate that he attacked, but failed to harm, a protected

⁵⁷ Duff (1991, 103).

interest. This is Duff's view with respect to our pair of cases; that (1) amounts to a successful attack on another's property interest and that (2) amounts to a failed attack on that interest. These two cases do not represent what we typically think of as paradigm cases of (failed) attacks. The defendants in each case do not intend [to damage another's property]. They intend [to damage property], which is not wrongful in and of itself, but what makes their actions wrongful is the ownership of the property and their recklessness as to that fact.

It is not immediately clear—at least not to me—that the actions of D_1 and D_2 constitute (failed) attacks. Our ordinary understanding of “attempt” might offer some evidence in support of Duff's view, for it is quite natural to say that D_1 attacked V_1 's property since he intentionally damaged V_1 's property. However, ordinary language is unclear on this matter since it is also quite natural, and perhaps more accurate, to say that D_1 did not attack V_1 's property because he intentionally damaged some property which unfortunately turned out to belong to V_1 . In the next section, I offer my argument for why cases like (1) and (2) typically constitute acts of endangerment. Before then, we need to look at Duff's principled argument for why (1) and (2) constitute an attack and a failed attempt.

We can begin by contrasting (1) and (2) with paradigm cases of attacks. In (1) and (2), the defendants act with a direct intention to the consequence element of an offence, but only with recklessness as to the circumstance element of that offence. In a paradigm case of an attack, the attacker acts with a direct intention to both the consequence and circumstance elements of an offence. That is, the attacker does not merely intend [to damage property], but [to damage another's property]. The paradigmatic attacker's action is guided by a direct intention to harm another person's interests, and it is this intention that marks his action as a clear case of an attack. The success or failure of his action is determined not only by the act of damaging property, but also by the fact that what he damages is *another's* property. If the attacker manages to damage property, he would count his action a success if the property belongs to another and a failure if

the property does not belong to another.⁵⁸ In a paradigm case of an attack, not only does the defendant harm another's rights or interests, but he also directs his action at [harming another's rights or interests].

The defendants in (1) and (2), on the other hand, do not direct their actions at [harming others]. They merely intend [to damage property], though the circumstances are such that D₁ actually damages another's property. His action is made criminal by the fact that the property belongs to another and by his recklessness as to ownership of the property. But the success of his action is only marked by the completed act of damaging and not by the fact that the damaged property belongs to another. D₁ and D₂ take unjustifiable risks that their actions will harm the property rights of others, and they take this risk by being reckless with respect to ownership of the property. However, the crucial different between their actions and the action of a paradigmatic attacker is that the defendants in (1) and (2) do not intend to attack V₁ or V₂ through their properties. That is, they do not intend [to harm another's rights or interests].

Duff argues that even though the defendants in (1) and (2) do not intend to attack V₁ or V₂ through their properties, their actions can nonetheless amount to attacks.⁵⁹ This is because, for Duff, the nature of an action 'depends not just on the intentions or beliefs with which it is done, but on its actual connection to the world'.⁶⁰ The actual facts of our world make up the objective dimension of an action, which along with our subjective intentions determine what our actions really amount to. On Duff's view, an action which is not intended to attack any protected rights or interests may amount to an attack if certain objective facts obtain.⁶¹ In the case of (1), the fact that the damaged property belongs to another is enough for D₁'s action to count as an attack. For Duff, the objective dimension of an action is sufficient to make that action an attack, given that D acts with intention to the

⁵⁸ See section 4, where I discuss cases in which D intends to [to damage this property], knowing that it belongs to V, but whether or not the property belongs to V makes no difference to his action.

⁵⁹ Duff (1996, 371).

⁶⁰ Ibid.

⁶¹ Duff (1996, 371-2).

consequence element of an offence and recklessness as to the circumstance element of that offence.

While Duff argues that an objective connection to the world is necessary for an attempt, his argument need not concern me here. In the next section, I argue that attacks require more than recklessness as to the circumstance element of an offence.⁶² This point applies whether we are objectivists or subjectivists about liability since what defines an attack is that the action is guided by an intention to cause harm. If recklessness as to circumstance does not suffice for the kind of intention that is required for an attack, then an objective dimension is not enough to make the action an attack. It might be worthwhile to note that my argument in the next section is meant to show that the defendants' actions in (1) and (2) do not amount to attacks, and that if we think the law of attempts should be a law of attacks, we cannot hold D₂ liable for attempted criminal damage if his action does not amount to an attack. It is a separate question, which I will not address in this chapter, whether D₂ should be liable for endangering another's property. As I mentioned earlier, we should be careful that our intuition about whether to hold D₂ liable should not cloud our judgment about whether to hold him liable for an *attempt*.

4 From Attempts to Endangerments

The structures of cases (1) and (2) are of course not limited to property offences, and apply to a wide range of criminal conduct. In this section I argue that for many offences that conform to the structures of (1) and (2), the acts constitute endangerments rather than attacks. I will make this argument by discussing criminal damage and crimes involving bodily harm before turning, in the next section, to explore the case of rape. In making my argument, I do not take issue with Duff's characterization of either attacks or endangerment.

Consider criminal damage first. Suppose an agent intends [to cut down his neighbour's tree] and carries out his intention by cutting down a tree

⁶² This is only to say that the defendants' actions in (1) and (2) do not amount to attacks.

which he knows or believes belongs to his neighbour. By acting with an intention [to cut down his neighbour's tree], he directs his action at [damaging his neighbour's property] (assuming his neighbour does not consent to the tree being cut down). The agent commits criminal damage if he cuts down the tree and the tree in fact belongs to his neighbour, and he commits attempted criminal damage if he does not succeed in cutting down the tree or if the tree is actually his own. The agent acts with the paradigmatic intention for an attack and his action is a paradigm case of an attack.⁶³

Now suppose the agent acts with an intention [to cut down this tree] while knowing or believing that the tree belongs to his neighbour. According to Duff's view (and the ruling in *Pace*), the agent attacks his neighbour's property. He would, quite rightly, be guilty of attempted criminal damage if he fails to cause damage. His action does not represent a paradigmatic attack since he does not direct his action at [cutting down or damaging another's tree]. It is removed from paradigmatic attacks because he does not intend [to cut down another person's tree] and the actual ownership of the tree makes no difference to the success or failure of his action.

However, the agent's action is an attack precisely because the issue of ownership makes no difference to his course of action *given* his knowledge or belief. He is committed to cutting down the tree *even though* he knows or believes that it belongs to his neighbour. If attacks are actions directed against (what are actually) protected interests, then it is the agent's willingness to cut down the tree, despite knowing or believing that it belongs to his neighbour, which commits him to damaging his neighbour's property. It is his willingness to purposely damage what he knows or believes is another's property that makes his action an attack. Given his knowledge or belief, he is committed to harming another's property, and he is thus committed even if his belief turns out to be false.⁶⁴

⁶³ By acting with an intention [to cut down his neighbour's tree], the agent directs his action at [damaging his neighbour's property] and he commits an attack even if he is merely reckless as to ownership of the property.

⁶⁴ See Yaffe (2010), chapter 5, where he argues that in these cases, the agent normally has, but does not necessarily have, the *mens rea* for an attempt.

Compare the defendant who intends [to cut down this tree] and knows or believes that it belongs to his neighbour with the defendant who intends [to cut down this tree] and realises that it might belong to his neighbour. Both defendants can be relieved if it turns out that the tree belongs to them rather than their neighbours since it is no part of their intention that the tree belongs to their neighbours, and the success or failure of their actions do not depend on it. But, in comparing these two defendants, we can say that the reckless defendant might have refrained from cutting down the tree had he known that it belonged to his neighbour, whereas we cannot say the same about the other defendant. The reckless defendant acts with an intention which is situated in an uncertain context—e.g. an awareness that the tree may belong to his neighbour—so that as he carried out his action, he can still hope that the tree does not belong to his neighbour. But for the defendant who knows or believes that the tree belongs to his neighbour, his intention is situated in a specific and defined context—e.g. that the tree belongs to his neighbour.⁶⁵ It is not possible for him to hope that the tree does not belong to his neighbour, and it is not possible to separate his intention [to cut down this tree] from his belief that the tree belongs to his neighbour. Unlike his reckless counterpart, the defendant who acts with knowledge or belief as to ownership acts with an intention [to cut down this tree] which belongs to his neighbour, and it is true that “he tried to cut down his neighbour’s tree” whomever the tree belongs to.

An agent who is reckless as to certain relevant circumstances, who does not have knowledge or firm beliefs regarding ownership of the tree, is not committed to damaging his neighbour’s property. Duff would agree that the action of an agent who is reckless as to ownership of the property is even further removed from paradigmatic attacks than the action of an agent who knows or believes that the property belongs to another. However, we disagree about whether recklessness as to circumstances still counts as an attack. For Duff, an agent who cuts down a tree and is reckless as to ownership of the tree is attacking another’s property since the action is directed against the tree, which is *in fact* another’s property.

⁶⁵ See Duff (1996), chapter 1 on the content and context of an intention.

The fact that the tree is another person's property is the objective fact that suffices to make the agent's action an attack, even though he is merely reckless as to that fact. Duff argues that should his neighbour intervene right before he takes an axe to the tree, the neighbour could legitimately accuse the agent of trying to cut down his tree—"you tried to damage my tree!". The moral force of the neighbour's accusation, according to Duff, is "qualified, but not negated" if the agent responds honestly "I thought the tree belongs to me".⁶⁶ In contrast, an agent who is truly reckless—e.g. who intends to carry out an extensive renovation project in his garden and realizes that his project may damage his neighbour's tree—can respond to such an accusation by saying "I was not trying to damage your tree". For Duff, a significant difference between attacks and endangerment lies in the charges we can legitimately make against the defendant.⁶⁷ The charge "you tried to damage my property" is appropriately levied only against those who attack others' property whereas the only charge we can appropriately levy against those whose actions endanger others' property is "you might have damaged my property".

I think Duff is in large part correct; that in clear cases of attacks and endangerments, the difference between the accusations "you tried to damage my property" and "you might have damaged my property" is the difference between attacks and endangerments. However, things are not so clear in borderline cases.⁶⁸ First of all, it is not clear that we would accuse an agent of trying to damage another's property if he was merely reckless as to ownership of the property. Secondly, even if it is appropriate for us to accuse this agent of trying to damage another's property, it is just as appropriate for us to make demands and accusations that are typically associated with endangerments—e.g. "you should have taken more care because you might have damaged *my* property". Suppose that Duff is right, that we can accuse

⁶⁶ Duff (1995, 324).

⁶⁷ See Duff (1996) chapter 13.

⁶⁸ There is also the question of how sharply we can distinguish attacks and endangerments. These two categories may not be exclusive or exhaustive, so there can be cases that do not fall neatly into either category. Irrespective of this point, I take it that cases that conform to the structures of (1) and (2) typically constitute acts of endangerments, with the possible exception of rape. See Duff (2007), chapter 7.

an agent who is reckless as to relevant circumstances with “you tried to damage my property”. If my claim is also correct, that we can also appropriately accuse the agent with “you might have damaged my property”, then it is not clear how decisive—if at all—appeals to ordinary language can be in borderline cases.

Suppose the defendant intends [to cut down this tree]; is reckless as to ownership of the tree; and the tree in fact belongs to another person. The tree belongs to the defendant’s neighbour, who intervenes just as the defendant swings his axe towards the tree. Of course the neighbour could justify his intervention and accuse the defendant of committing an attack by saying “I intervened because you were trying to damage my tree!”. The defendant could explain himself by saying, honestly, any or all of the following: “I did not mean to cut down *your* tree”; “I thought there was a chance that the tree belonged to me”; “I would not have attempted to cut the tree had I known it belonged to you”. The neighbour can accept any or all of these claims; he can acknowledge that the defendant did not intend [to damage his (the neighbour’s) tree] but merely intended [to damage this tree]. The neighbour can nonetheless say to the defendant “well, the tree *is* mine so you *were* trying to damage *my* tree”. I think this is what Duff has in mind when he argues that cases similar to this one amount to attacks, because it is certainly possible (and perhaps even appropriate) for the neighbour to accuse the defendant of trying to damage *his* (the neighbour’s) property irrespective of the defendant’s explanations.

There are a couple of issues with Duff’s reasoning here. First, although Duff is right in thinking that there is nothing inappropriate about the neighbour’s accusation of an attack, such an accusation, by itself, fails to give us a complete account of what the defendant was trying to do. Let us deviate from the scenario in (2) for a moment and suppose, for instance, that the defendant was trying [to cut down his own tree]. We can even suppose that this defendant was not reckless about ownership of the tree; that he justifiably believed that the tree belongs to him when it in fact belongs to his neighbour. This defendant has done nothing wrong, but it is nonetheless

appropriate for his neighbour to accuse him of trying to damage his (the neighbour's) property. The difference between this case and a clear case of an attack is that the accusation—which, on the face of it, suggests that the defendant has attacked the neighbour's tree—does not remain wholly appropriate or correct in light of the defendant's motivations for his action. This defendant can explain his justified, though mistaken, belief to his neighbour, who can then interpret the defendant's action as an honest mistake rather than an attempt to cut down his (the neighbour's) tree. Likewise, the defendant in (2), who was reckless about ownership of the tree, can explain his recklessness and lack of concern as to ownership of tree, and it would be appropriate for the neighbour to interpret the defendant's action as one that lacks practical concern for others' property rather than an attempt to damage other's property. What is different about clear cases of attacks is that there is nothing that an attacker can say that would allow the neighbour to interpret the attacker's action as anything other than an attempt to cut down someone else's tree.

The second issue with Duff's reasoning has to do with why we should think, as Duff seems to think, that an accusation of an attack is the only appropriate response in such situations. The defendant in (2) did not intend [to damage another's property] though he did in fact attempt [to damage] another's property. As things turned out, he attempted [to damage] another's property partly because he was reckless and did not take steps to ascertain ownership of the tree. It would not be inappropriate for the neighbour to demand of the defendant that he (the defendant) should take more care in the future; that the defendant should inquire about ownership of a tree before he makes any attempts to cut it down. The rationale behind such a demand targets the inherent recklessness in the defendant's action, and the demand expresses the neighbour's disapproval of the defendant's indifference and carelessness towards other people's property. This demand and its corresponding rationale treat the agent's action as an instance of endangerment. The demand is not "you should take more care next time because you tried to damage my property", but rather "you should take more

care next time because you might have damaged someone else's (or my) property", or "you should inquire about ownership before you try to cut down a tree because you might damage someone else's property".

The same kind of response is appropriate in cases involving physical harm. Suppose a hunter intends to kill a deer and sees what he thinks might be a deer in the woods. He knows that there are other hunters around and realizes that what he sees might be a person. He has no intention of killing a person, though he takes aim at the living thing without determining whether it really is a human or a deer. As it turns out, the hunter had aimed at and shot a human being, though luckily his shot had missed. Duff argues that the victim could accuse the hunter of trying to kill her since the living thing which the hunter tried to kill is really her. Furthermore, Duff argues that the victim's accusation may be qualified though it could not be negated by the hunter's response "I thought you were a deer".

Again, it is not obvious why the victim's accusation would be a response that signals an attack (e.g. "you tried to kill me") rather than one that signals an endangerment (e.g. "you might have killed me"). We can certainly imagine the victim rejecting the hunter's explanation—that he thought she was a deer—and say to the hunter "well, I'm not a deer and you just tried to kill *me!*". However, we can just as easily imagine a different response from the victim, one that addresses the hunter's explanation with the legitimate demand, "why did you not check before you fired the shot?", or "you should have checked before you fired at me". The second kind of response—which treats the hunter's action as a case of endangerment—is just as appropriate, if not more so, as the first kind of response (which treat the hunter's action as an attack). The second kind of response emphasizes the agent's recklessness by demanding not that "you should check next time because you tried to kill someone", but rather that "you should check next time because you might have killed someone".⁶⁹

At best, our ordinary responses to borderline cases cannot help us determine whether such cases are attacks or endangerments. Luckily, our

⁶⁹ See Alexander (2010) for a similar argument, and see Duff's (2010) reply to Alexander .

response to such cases—i.e. what we can appropriately accuse the defendants of in such cases—is only one of several considerations that differentiates attacks from endangerments. In Chapter 2, I argued that, all else being equal, culpability for purposeful actions does not vary according to the level of risk an agent imposes while culpability for reckless actions does. Purposeful actions and reckless actions constitute the paradigm of attacks and endangerments respectively, and what I have argued in Chapter 2 is applicable here. In paradigm cases of attacks—e.g. D attempts [to damage his neighbour’s tree]—the defendant is guided by the wrong reasons.⁷⁰ Given the intention with which it is done, D’s action is wrongful independent of the risks he imposes. If D intends [to damage his neighbour’s tree] and tries to damage a tree which he thinks might belong to his neighbour, the risk he imposes on his neighbour’s property can vary depending on the means he employs to damage the tree or the likelihood that the tree belongs to his neighbour. However, his intention [to damage his neighbour’s tree] supersedes any of the other considerations. Should he adopt an inefficient means of damaging the tree, his choice of means may speak to his competence or lack thereof, but it does not speak to, nor detract from, the blameworthiness of his action.

Matters are different in paradigm cases of endangerment. If D is renovating his garden and is reckless towards his neighbour’s tree, then he is failing to respond to the right kind of reasons, such as taking appropriate care to avoid damage to others’ properties. Suppose, for instance, that in renovating his garden, D is choosing between two different methods of accomplishing a particular task. Both methods will accomplish the task, though they carry different levels of risk to the tree. If both methods pose substantial risks to the tree, then D has reasons to refrain from trying to accomplish that task. However, should he choose one of the two options, he is more blameworthy for choosing the riskier option. If he chooses the less risky option, his actions speaks badly of him because is not responding to reasons against undertaking the renovation task altogether—i.e. reasons

⁷⁰ See Duff (2008).

against exposing other people's property to substantial risk of damage. However, if he chooses the more risky option, his actions speaks *more* badly of him since (compared to the less risky option) he imposes a higher risk on another's property.⁷¹

Culpability in attacks is not sensitive to the level of risk imposed because the agent aims to accomplish something which is inherently harmful; his action is wrong because it specifically aims at an outcome that is harmful. The level of risk an attacker imposes on another is irrelevant to the wrongfulness of his action. Culpability in endangerments is sensitive to the level of risk imposed because what the agent is blameworthy for is his lack of proper concern for a protected interest; he is more or less blameworthy depending on the level of risk he imposes, as greater risk imposition reveal greater lack of proper concern for others. The action of a defendant who intends [to damage this tree] and is reckless as to ownership of the tree is not inherently harmful; it would be harmful if the tree belongs to another but it would not be harmful if the tree belongs to the defendant. What makes this defendant's action wrongful is not that he tried [to damage a tree] which actually belongs to another, for this suggests that D's action would not be wrongful if the tree did not belong to another. Instead, what makes D's action wrongful is that he is reckless as to ownership of the tree, which is to suggest that D's action is wrongful whether or not the tree belongs to another.⁷²

My view regarding (1) and (2) fits with what I have said about risk and culpability in attacks and endangerments. Suppose in one variation of case (2), the defendant thinks there is a twenty percent chance that the property

⁷¹ There is a possible exception to what I have said in this (and the previous) paragraph, which is that of half-hearted attempts. Suppose the defendant intends [to damage his neighbour's tree] and has two options for doing so. One option is less efficient and exposes the tree to a lower risk. If the defendant chooses the less efficient option because he is not wholeheartedly committed to the attack, then he may be less culpable than another agent who is wholeheartedly committed. Whether the former agent is less culpable is a question about whether half-hearted attempts are less culpable than wholehearted attempts, and the difference in culpability should not be taken as evidence that the agents' actions amount to endangerments rather than attacks.

⁷² This means that the actions of the defendants in (1) and (3) are equally wrongful, and that the actions of the defendants in (2) and (4) are also equally wrongful. It does not, however, mean that cases (1)—(4) should all be criminalized. But when we ask whether (3) and (4) should be criminalized, I think the right response involves treating (3) and (4) as acts of non-consummated endangerments.

he intends to damage belongs to his neighbour, and in another variation the defendant thinks there is an eighty percent chance that the property belongs to his neighbour. In each variation, the defendant should make inquiries about ownership and determine whether the property belongs to his neighbour. Given that he does not make the appropriate inquiries, we think badly of the defendant in the first variation (where the risk is low) but we think worse of the defendant in the second variation (where the risk is high). The defendant should refrain from damaging the property regardless of whether he thinks there is a twenty percent chance or an eighty percent chance that the property belongs to his neighbour, but the defendant's attempt to damage property is less objectionable in the first variation than in the second variation. We can say the same about the hunter. He should refrain from shooting the living thing if he thinks there is a chance that it is a human being. However, if he shoots at the living thing anyway, his action is more objectionable if he thinks there is an eighty percent chance the being is a human than if he think there is a twenty percent chance that the being is a human.

So far in this section I have argued that scenarios (1) and (2) amount to acts of (non-consummated) endangerment. Contrary to Duff's view, my position suggests that the *mens rea* for attempts should at least require that the agent act with a direct intention to the consequence element of an offence and knowledge or belief as to circumstance elements of that offence. I mentioned earlier that an agent who acts with knowledge or belief as to circumstance elements of an offence does not direct his action towards harm. But his willingness to carry out his action knowing or believing that it will cause harm commits him to an attack on what he knows or believes to be another's rights or interests. Thus far, this discussion has been limited to cases of criminal damage and assault involving bodily harm. But unlike criminal damage or assault, my argument that attempts require more than recklessness as to the circumstances of an offence may or may not apply to cases of rape. I admit that it is not entirely clear to me whether a defendant attacks another's sexual integrity if he intends [to have intercourse] and is

reckless as to the other person's consent. In the rest of this chapter, I turn to the offence of rape and explore why rape may be different from criminal damage or assault in the context of this discussion.

5 Recklessness and Rape

With respect to criminal damage and assault, I have argued that when an agent acts with an intention [to damage property] or an intention [to cause injury] and is reckless as to ownership of the property or to the fact that the living thing is a human being, what makes such actions wrongful is not that the intended aim is inherently wrongful, but that the agent's recklessness as to relevant circumstances surrounding his action reveals his indifference to the interests or well-being of others. His action is blameworthy because he ought to have taken more care and he ought to have made inquiries about ownership of the property or the status of the living thing before acting as he did.

We can say similar things in the context of rape. An agent who acts with an intention [to have intercourse] and is reckless as to the other person's consent ought to take more care and be sensitive about consent before attempting to have intercourse with that person. We can accuse the agent of taking an unjustifiable risk with regards to another person's consent, and we can demand that he take more care in his sexual encounters. Perhaps (though I am not sure) we think badly of the defendant who is aware of a 20 percent risk that his "partner" may be non-consenting, and think worse of the defendant who is aware of a 80 percent risk that the other person may be non-consenting. Everything I have said in the context of criminal damage and of causing bodily harm can be applied to the case of rape, so we at least have some reasons to think that recklessness as to consent constitutes an act of endangerment.

The problem with rape though is that, as Duff argues, it would be odd and morally reprehensible if the accusations that I mentioned above are all that we had to say to the agent. Surely we can accuse the agent of much more than simply taking an unjustifiable risk as to consent, and surely we can

demand much more from him than simply that he take more care in his sexual encounters. When an agent is reckless as to another person's consent, he consciously disregards a risk that his sexual "partner" may be non-consenting.⁷³ He may hope that the other person does consent, but when he proceeds with intercourse while disregarding the risk of non-consent, he is no longer treating the other person as a partner in what should be a mutually gratifying activity between consenting agents. To merely hope that he is having [consensual intercourse] is to take an unacceptable normative view of sex.⁷⁴ Against this view, we *should* see it as a mutually gratifying activity between consenting agents. Neither agent is simply a means to the other's end. Instead, each person's interest in the activity is seen as important and each person's autonomy is respected within the activity.

Duff argues that to be reckless about consent is to see sexual intercourse as an activity that serves to gratify him or herself, and that the agent sees his or her sexual "partner" as a mere means to that gratification. Consent is not merely a contingent circumstance of consensual sex. Without it, the non-consenting agent is in fact a *mere* means which the defendant uses for his own end. One way in which recklessness about consent could amount to an instance of an attack then is if using others as a mere means constitute a species of an attack, where the object of the attack is the person's sexual integrity. Recklessness as to consent suggests an indifference to the other person's integrity and that defendant is willing to use the other person as a passive object. Such actions and attitudes are inherently wrongful. Even if the person does consent, she may feel violated if she found out that her "partner" was indifferent to her interests. And even if the intercourse is consensual, we may still object that the defendant has committed a wrong, though luckily no harm has occurred.

What is needed here is an account of the wrongness of rape, and whether Duff's normative view of sex should be adopted by the criminal law. I

⁷³ For various definitions of recklessness, see Duff (1990), chapter 7.

⁷⁴ For Duff's discussion on recklessness and rape, see Duff (1990, 167-73); also see Duff (1981).

will not address these issues here and will simply conclude by noting that there is a third option if we remain undecided on this issue. I mentioned earlier that in making the ruling in *Khan*, the court was partly guided by the recognition that if the defendant was not convicted of attempted rape, the alternative would be exonerating an agent who has done something that is clearly wrong. Intuitively, this would be an unappealing verdict, and given the seriousness of rape (which would have been the offence in *Khan* had the defendant succeeded in having intercourse), we may have pragmatic reasons for convicting the defendant of attempted rape even if we think the act is one of non-consummated endangerment. The defendant is just as culpable as someone who has committed the completed offence; and that when it is unclear whether the defendant has committed an endangerment or an attack, it would be permissible, on pragmatic grounds, to convict him of an attempt and to label him as an attempted rapist.

6 Conclusion

In this chapter I have argued that cases like (1) and (2) constitute acts of endangerment rather than attacks. In both cases, the defendant acts with an intention [to damage property], is reckless as to ownership of the property, and the property belongs to another person. The defendant in (1) successfully damages property while the defendant in (2) fails to do so. I have argued that with the possible exception of rape, crimes that conform to the structure of (1) constitute acts of consummated endangerment while those that conform to the structure of (2) constitute acts of non-consummated endangerment. I have not commented on cases (3) and (4), where the property belongs to the defendant himself. As I mentioned earlier, (3) generates some controversy since it is sometimes considered to be an impossible attempt. Given my arguments in this paper, (3) does not constitute an impossible attempt because, much like the defendants' actions in (1) and (2), the action in (3) does not amount to an attack. If anything, (3) and (4) constitute acts of non-consummated endangerment.

Subjectivists and objectivists about criminal liability can disagree about whether (3) and (4) amount to acts of endangerment. Subjectivists may argue that we have grounds to hold (3) and (4) liable for specific acts of endangerment since the defendants in these cases realise that there may be a risk to others' property (which, by luck, turned out not to be the case). Objectivists may disagree and argue that because the property in fact belongs to the defendants themselves, we should not hold the defendants liable for anything since they have not actually created any risks to anyone's property. Irrespective of deep disagreements between subjectivists and objectivists, we may, on pragmatic grounds, favour not criminalizing or prosecuting acts of endangerment that conform to (3) and (4).⁷⁵

⁷⁵ Note, though, that in the context of rape, if we think the defendant's action in (3) constitutes an attack, then it is still an instance of impossibility, and we are again faced with the question of whether we should hold the defendant liable for an attempt to commit rape.

Chapter 4

Recognition and the Significance of Formal Apologies

Having looked at some philosophical questions surrounding the notion of luck, culpability, and the *mens rea* requirement for criminal attempts, Chapters 4 and 5 will turn to a question that must be addressed if we are to make progress on the problem of outcome luck. Part of the problem has to do with the extent to which an agent should be punished for the consequences of his actions, and we need a normative theory of punishment in order to address this issue.¹

In these last two chapters, I will defend Duff's communicative account of punishment from an important criticism from instrumentalists; namely, that it fails to take crime prevention seriously.² I argue that communicative theorists can overcome this objection by showing that the communication of censure entails not only special prevention (as Duff has shown) but general prevention as well.³ This requires that I begin, in Chapter 4, with an overview of communicative punishment before moving on, in Chapter 5, to the connection between censure and general prevention.

In addition to giving an overview of Duff's account, this chapter defends several key points in his view, including the idea that we should see the offender's punishment as a formal apology.⁴ I do not argue that formal apologies are sufficiently valuable to justify our penal institutions. On the contrary, throughout this chapter and the next, I concede instrumentalists' point that general deterrence must figure in the justification of punishment.⁵ As we will see in this chapter, the idea of punishment as apology affects

¹ See Feinberg (1995).

² See Tadros (2011), chapter 5.

³ See Duff (2001, 88-89; 107-9).

⁴ See Duff (2001).

⁵ NB As we will see in the next chapter, there is a distinction to be drawn between general deterrence and general prevention as distinct modes of crime prevention. On the communicative view that I defend (and Duff defends), we are interested in the latter rather than the former.

neither the justification nor the distribution of communicative punishment. However, the idea should nonetheless be defended as an ideal account of how we should view punishment and how the community should regard offenders.

I will begin in section 1 with a brief overview of Duff's account, and of the role that apologies play within it. For Duff, punishment has three interrelated aims: recognition, self-reform, and reconciliation.⁶ Talk of recognition often focuses on *offenders'* recognition of their wrongdoing. In sections 2 and 3, I highlight the importance of the *community's* recognition of the victim, and the sense in which punishment is necessary for this kind of recognition.⁷ I argue that the community's recognition can be an important aim in itself, but, in the next chapter, we will also see that it is the community's recognition of wrongdoing that entails general prevention. This emphasis on recognition does not support the idea that punishment should be seen as an apology; recognition (especially by the community) can be achieved without an apology or reconciliation. In section 4, I offer some reasons for why we should nonetheless see punishment as an apology and explain why punishment as apology does not raise any problems about unjustified state coercion.

1 Communicative Punishment and Formal Apologies

Central to Duff's theory is the idea that criminal punishment should be understood and justified in terms of communicating deserved censure.⁸ Punishment should communicate to the offender that what he did was wrong, to the wider political community that we take such wrongdoing seriously, and to the victim that the political community takes her legally protected interests

⁶ Duff (2001, 107-12).

⁷ It might be thought that symbolic and non-punitive measures are sometimes sufficient for recognition. So the question is why or when is punishment needed for adequate recognition? See Duff (2001, 82-85); Matravers (2011); Narayan, (1993); von Hirsch (1993), chapter 2.

⁸ This is slightly different from what Duff says, which is that we should "understand and justify criminal punishment...as a species of secular penance". Punishment is a form of secular penance where "a burden [is] imposed on an offender for his crime, through which, it is hoped, he will come to repent his crime, to begin to reform himself, and thus reconcile himself with those he has wronged." Duff (2001, 106).

seriously.⁹ Censure, most distinctly, includes the backward-looking judgment that what the offender did was wrong, and on Duff's account, it is communicated with the intention that the offender will recognize, understand, and accept that he has committed a wrong. Furthermore, censure also entails forward-looking aims such as reform and reconciliation; the offender should, through a genuine recognition and acceptance of his wrongdoing, be motivated to make a commitment not to re-offend and to apologise for his action.¹⁰

On Duff's view, censure or moral criticism is the intrinsically appropriate response to criminal wrongdoing, and punishment is the *necessary* means of communicating censure adequately.¹¹ When we censure an offender, we should not just *express* to him that what he did was wrong. The expression of censure does not require the offender to be responsive to our moral criticism.¹² But if we are to treat the offender as a responsible moral agent, we must try to persuade him to recognize that his conduct was wrong, to understand why it was wrong, and to accept his wrong as something which he should not have done and will not do in the future.¹³ And if we are to treat him as a fellow citizen, with whom we must continue to share our political community after the offence, he should be expected to

⁹ Duff (2001, 80-1; 114).

¹⁰ *Ibid.*, 107-10.

¹¹ *Ibid.*, 88-89; 82. Duff now puts more emphasis on the idea of punishment as an appropriate kind of apologetic reparation, which the offender owes to those whom he has wronged (private correspondence with Antony Duff, July 23rd, 2018).

¹² *Ibid.*, 79-80.

¹³ For a criticism of the communicative account, see Joshua Glasgow (2015). One of Glasgow's criticisms is that of the uncommunicative offender, who is so incapacitated in his abilities to communicate that he cannot engage in communications with society. He is not capable of receiving society's message and, for him, the aims of communication are not achievable. We might have the intuition that many of these offenders—especially those who commit serious wrongs—should be punished, but communicative punishment is not justified for these offenders. Glasgow's argument assumes that an initially uncommunicative offender cannot become communicative throughout the penal process. This is to assume that the offender cannot learn to communicate through exposure to paradigmatic communicative practices such as victim-offender mediation; that the offender is incapable of learning from their social context. I think this is a contentious assumption which requires empirical support, and it would be premature to dismiss the communicative account based on this assumption. Additionally, if the offender is so incapacitated that he cannot communicate, then there might be questions about whether he is fit to be tried (i.e. to answer for his actions) and fit to be punished. See Duff (2007).

apologise for his wrongdoing and the community should be ready to reconcile with him after an appropriate apology has been made.¹⁴

Communicative punishment is thus justified by its pursuit of recognition, reform, and reconciliation. Recognition is internal to censure, and reform and reconciliation follow naturally from recognition of wrongdoing. In order to pursue these aims, punishment is meant to perform the following functions.¹⁵ First, it should provide a structure within which the offender can focus on his wrongdoing and thus come to understand why it was wrong. This is not to suggest that offenders, prior to being punished, do not know that their conducts are wrongful.¹⁶ The idea is that punishment should appeal to the offender's sense of right and wrong, and it should focus the offender's attention on a deeper appreciation for the wrongfulness of his conduct.¹⁷

As we will see shortly, punishment ought to pursue the aim of recognition even if we think the offender already recognizes his wrong. This is because genuine recognition of one's wrong—especially for serious wrongs that warrant punishment—requires the performance of reformatory or reparative acts. If the offender truly appreciates the wrongfulness of his conduct, he should naturally make a commitment to refrain from similar wrongdoing in the future. If he were to re-offend, we would typically have grounds to question whether he did recognise the wrongfulness of his conduct. Note, though, that the process of self-reform is not just about making a commitment to act differently in the future; it involves figuring out what one must do to be able to act differently. This is the second function of punishment. It should stimulate and help the offender through the process of self-reform; it should help him understand *how* he is to reform himself and what it takes to refrain from future wrongdoing.¹⁸

Finally, the offender's punishment should also be *seen* as an apology through which the offender can seek reconciliation with whom he has

¹⁴ Duff (2001, 109; 113-4).

¹⁵ See Duff (2001, 107-9).

¹⁶ See Narayan (1993, 173).

¹⁷ Duff (2001, 91; 97-8). See also von Hirsch (1993, 11).

¹⁸ Duff (2001, 108-9).

wronged.¹⁹ This is especially pertinent if we take seriously the idea that the offender and the victim are to continue living with each other as fellow citizens after the offence. In many cases, adequate apologies are less about what one says and more about what one does, and we should not think that mere verbal apologies would be enough to repair serious wrongs. If I recklessly damaged your record player, it would not be enough for me to say that I am sorry. I ought to make it up to you, for example, by buying you a new record player. In order to repair the wrong that I have committed against you, I ought to make adequate reparations for that wrong, and a verbal apology alone is often not enough to constitute the kind of *apologetic* reparation that is required in response to wrongdoing.²⁰ Here, punishment is meant to give weight to the apology, as an adequate response that an offender must make in order for reconciliation to be possible.²¹

It might be tempting to think about the aims of censure in terms of successive stages, where we use punishment first to achieve recognition, followed by self-reform, and ending in the process of reconciliation. If this was the case, then the same offence may warrant different levels of punishment when committed by equally culpable offenders.²² An offender who already appreciates the wrongfulness of his conduct would, in principle, receive a lesser sentence than another relevantly similar offender who has no such appreciation. It is easy enough to see why we might think of punishment in this way. The connection between recognition and self-reform is an obvious example. Recognition has to do with an offender's response to the backward-looking message of censure, where the political community is attempting to convey to the offender the reprobation and disapproval that his wrongdoing warrants. If we think that genuine recognition of wrongdoing

¹⁹ The emphasis here is that we should see the offender's punishment as an apology. Communicative theorists need not argue that punishment actually *is* an apology. But if all the requisite acts of an apology have been performed, we can act as if the offender has apologised for his wrong. See section 4 below.

²⁰ Note that for Duff, it may not be enough for me to simply repair the harm that has been caused by buying you a new record player. The act of repairing the harm should resemble an *apologetic* reparation, which may require more or less than the mere reparation of the harm.

²¹ Duff (2001, 109).

²² See Duff (2001, 118-21).

should motivate an offender to avoid similar wrongdoing in the future, and to work out what he must do to avoid such wrongdoing, then it is surely tempting to think of recognition as the precursor to reform. But this is not how we should interpret Duff's view.

On Duff's account, the entire sentence should be seen as an apology for the wrong that has been committed, which simultaneously provides both the structure for recognition and the conditions for undertaking self-reform.²³ This way of interpreting the aims of punishment might have led some theorists to think that we are headed for communicative confusion.²⁴ Their reasoning is that if punishment is meant to stimulate recognition of wrongdoing, then it must be imposed before the offender has recognized his wrong. But if punishment is meant to act as an apology, then it must be imposed after the offender has recognized his wrong. How is punishment meant to play both of these roles? Again, a tempting response would be to say that punishment is a process, where recognition is followed by an apology. But this is not the case.

According to Duff, punishment does not just *do* different things to different offenders, such as stimulating recognition in unrepentant offenders but serving only as an apology for reformed offenders. Instead, punishment also *means* different things to different offenders.²⁵ For an offender who already appreciates the nature of his wrongdoing, he will see his punishment as something which he should undertake in order to reform himself and to apologise for his wrong. *That* is what it means to truly recognise one's wrong. For an offender who does not appreciate the wrongfulness of his conduct and feels no need to reform himself, punishment can still be seen as an apology which involves recognition and a commitment to self-reform.²⁶ The offender may not see his punishment in this way while he is being punished, but (we hope) he will come to see it as an apology after his sentence.

²³ See Duff (2001, 118-21).

²⁴ Tadros (2011, 101).

²⁵ *Ibid.*, 121.

²⁶ Again, this is not to say that punishment actually *is* an apology. See fn. 19.

The problem with interpreting communicative punishment in terms of successive stages is that it fails to capture the intimate connection between the aims of recognition, reform, and reconciliation. A commitment to self-reform and the need to apologise do not just follow from recognition of wrongdoing; they are necessary for genuine recognition.²⁷ And the actions which constitute an apology also constitute self-reform and recognition of wrongdoing. Suppose, for instance, that you committed a wrong against your friend. You promised and then forgot to pick him up from the hospital after his operation. The appropriate response is for you to issue an apology. A fitting apology should include an adequate acknowledgement of the wrong you have committed, such as saying that you are sorry for what you have done. But a mere verbal apology is not enough.

Adequate recognition often requires appropriate reparations. You failed to keep an important promise and you ought to *do* something in order to make it up to him. But this reparative dimension is not something in addition to an acknowledgement of wrongdoing. Crucially, it is the vehicle through which you can adequately acknowledge your wrongdoing. A failure to do something in order to make it up to your friend would suggest, at least evidentially, that you do not genuinely recognize the seriousness of your wrong. But we have reason to think that you do not genuinely recognize your wrongdoing *because* a failure to repair serious wrongs (in the absence of extenuating circumstances) actually constitutes a lack of adequate recognition.

Similar things can be said about the relationship between recognition and reform. Suppose you make no effort to apologise to your friend, and due to your lack of initiative, your partner intervenes and makes you apologise, for instance, by making you visit him and help him around the house while he recovers. You may not recognize the wrongfulness of your actions, but your

²⁷ This is not to say that a *verbal* apology is necessary for genuine recognition. We can easily imagine a situation in which there are no victims to apologise to—perhaps all of the victims have passed away. Of course the offender can verbally apologise to the community, but (1) the verbal apology is not required for the offender to be truly repentant, and (2) genuine repentance is about what the offender does after the offence rather than about what he says.

partner certainly does. The forced apology constitutes her recognition of your wrong, and it constitutes what she thinks you ought to do if you did recognize your wrongdoing. By forcing you to visit your friend and help him around the house, your partner hopes that you will come to appreciate *why* your action was wrong. In this case, the wrong stems from your failure to keep an important promise—i.e. to treat other people’s interests with proper concern. Someone who does treat another’s interest with proper concern would not have acted so carelessly and, had they broken a similar promise, they would have voluntarily done the things that your partner is now making you do.

Here, the forced apology also constitutes a process of reform. You ought to repair the wrong you have committed because that is what is required of someone who shows proper concern for other people’s interest. Again, assuming that you are able to visit your friend and help him around the house, your failure to undertake such tasks would suggest, evidentially, that you do not truly understand *why* your action is wrong. And we have evidence to think that you do not truly understand why your action is wrong *because* genuine recognition—barring extenuating circumstance that would make reform impossible or unnecessary—actually requires reformative acts. The process of reform, which is materially identical to an apology, is necessary for genuine recognition.

Of course, *you* may not realize the meaning of your apology while you are making it, but it nonetheless suggests to your friend that you recognize your wrongdoing. And even though *you* may not realize it, your apology also constitutes what you must do if you are to reform yourself. What is going on here is not that you recognize or acknowledge the wrong you have committed against your friend at t_1 and take steps to address the wrong and reform yourself at t_2 . The apology, which must involve (measures akin to) repairing the wrong you have committed, constitutes simultaneously both your recognition of wrongdoing and a necessary step in reforming yourself as

someone who had acted wrongly.²⁸ Again, the reparative act is not just a means of self-reform, but a necessary element of genuine recognition.

Duff's account of punishment captures what an offender should do in the aftermath of an offence. He should make an apology that would be sufficient to redress the wrong he has committed. And in performing the requisite acts for an apology, the offender both recognizes the wrong he has committed and takes the necessary steps to reform himself. For an offender who already recognizes and regrets his wrongdoing, he will see his punishment as embodying both his recognition of the wrong and the necessary steps he needs to take for self-reform and reconciliation. For an offender who does not recognize or regret his wrongdoing, his punishment embodies the community's recognition of his wrongdoing and the necessary steps for reform. But such things are imposed on him, with the hope that he will come to understand his punishment as an appropriate response to his wrongdoing.

Many critics find it difficult to reconcile the value of an apology with the nature of punishment (as something that is imposed on offenders). If we think that apologies are valuable because they convey the offender's remorseful recognition of their wrongdoing, then the forced nature of punishment would seem to undermine the sincerity of that recognition. I will say more about this in section 4. For now, it would be worthwhile to note that although apologies are central to Duff's account, the idea of punishment as apology and the aim of reconciliation have no impact on either the justification or distribution of communicative punishment.²⁹

²⁸ In this sense, punishment is less like criminal trials where the sentencing stage follows the fact-finding stage, and more like political demonstrations where forceful expressions of disapproval for a certain state of affairs is, at the same time, an attempt to alter that state of affairs. The expression of disapproval itself includes a desire for change, and the act of protesting embodies both aspects simultaneously.

²⁹ On Duff's account, reconciliation is an aim of censure and a part of the justification for communicative punishment. The point I want to press is that even without the aim of reconciliation, we could still justify communicative punishment in terms of recognition and reform, without altering how we distribute punishment to individual offenders. On the view that I am defending, communicative punishment is justified at least in part on general preventive grounds, which is entailed by the aim of recognition. Whether or not punishment aims at reconciliation has no effect on its preventive aims.

The notion of an apology has no impact on the distribution of punishment because the same kind of burden is involved in the apology (i.e. the aim of reconciliation), the process of self-reform, and the recognition of wrongdoing.³⁰ The entire sentence is meant, simultaneously, to act as an apology and aid in the process of self-reform, which is necessary for genuine recognition. In the previous example, the act of visiting your friend may be enough to serve as an apology, but it is also what you must do in the process of self-reform. Even if you have already apologised (verbally) to your friend and reconciled with him, you still need to repair the wrong you have committed because that is what is required of you as someone who shows proper concern for another's interests. It is not the case that, by getting rid of the apology, you only need to repair half of the wrong you have committed. Our decision to pursue the aim of reconciliation—or not, as the case may be—does not affect how we ought to pursue the aims of recognition and reform. Genuine recognition still requires a commitment to self-reform. And while adequate means of self-reform can constitute an apology, they do not require anything more or less than what an adequate apology would require.

In the next chapter, we will also see that the aim of reconciliation does not affect the justification of communicative punishment. There, I argue that the aims of communicative punishment entail general prevention, but it will be clear that recognition—without reconciliation—is sufficient to entail general prevention. Nonetheless, in section 4 I will try to convince critics that we should see punishment as a formal apology. For those who remain unconvinced, they can simply see communicative punishment as an appropriate means of recognition and self-reform.

Before I defend punishment as an apology, I want to comment on the aim of recognition. Both critics and proponents of the communicative account often discuss the aim of recognition in terms of the offender's recognition of his wrongdoing. This is especially the case when the discussion focuses on the idea of punishment as apology, which is meant to convey the offender's remorseful recognition for his offence. But as we have seen in the previous

³⁰ Duff (2001, 111) seems to make this point as well.

example, a forced apology does not just convey the offender's recognition of his wrongdoing. It also conveys the recognition of the party who is imposing the punishment. That is, punishment can convey the community's moral criticism of the offender's conduct. In the next two sections, I will highlight why such recognition is valuable and explain why punishment is necessary for recognition.

2 The Content of Recognition

Whether recognition is an important aim of punishment depends on exactly *what* we think should be recognized and by *whom*. For instance, instrumentalists may think that an offender's recognition of wrongdoing is important because it leads to self-reform or special deterrence. In this case, the value of recognition would be contingent on its deterrent effects. In the next chapter, I argue that recognition entails general prevention, but it is specifically recognition by the political community that entails general prevention. In this section here, I will discuss *what* the political community can and should recognise through punishment.

One of the things that a community should recognise is the victim. She has been wronged through the offence and it is important for the state or community to recognise the wrong she has suffered. Many theorists have acknowledged the importance of recognizing the victim.³¹ It is, for instance, present in Duff's account of crime and punishment. According to him, crimes are public wrongs. They are public in the sense that the wrongdoing *concerns* the public; it is not merely a private matter between the wrongdoer and the victim. The political community shares the wrong that the victim has suffered, and in this sense, the community owes it to the victim to recognize her suffering by censuring the offender.³² Several German legal theorists have also noted this point, and they seem to share Duff's view that punishment is able to recognise the wrong that the victim has suffered *by* censuring the offender for his offence.³³

³¹ See Narayan (1993, 169-71); Roxin (2014); Tadros (2011, 91-2).

³² Duff (2001, 114).

³³ See Roxin (2014, 39-41).

The interest that victims have in being recognised as having been wronged is something that a theory of punishment should take seriously. This is not to say that the aim of recognition is enough to justify our penal institutions. Not all victims require recognition, especially if the wrong she has suffered is relatively minor. And we should not think that recognition of the victim's suffering is always worth pursuing in light of the costs and harms that are associated with punishment. But it would be wrong to think that this kind of recognition could not contribute to the justification of individual tokens of punishment. In some cases, especially cases involving egregious wrongs, recognition is paramount for the victim (and those who identify with the victim). It is more significant than any preventive or retributive aims that punishment could pursue because victims have an interest in telling the truth about what has happened to them and having that truth recognize by their community. It is, as some theorists put it, an existential need for victims to receive public confirmation about the wrong they have suffered and the responsibility of those who perpetrated the wrong.³⁴

We should not think of victims' need for recognition as some kind of irrational preoccupation, which we might lump together with our retributive urges as something that should be overcome.³⁵ Recognition is important to our identity. It is important to the narrative we tell ourselves about our lives and what has happened to us. By addressing the offender and his conduct, recognition offers a satisfactory closure to the victim's ordeal and allows her to move on with her life. It publicly confirms what the victim considers to be an important truth about her, and it shows victims that the political community is in solidarity with her. Victims can pursue their case through various means other than (and in addition to) the criminal justice system, such as tort law and public inquiries, which can give some recognition of the wrong she has suffered. I will discuss this point in the next section. But if the victim is continuously ignored through all available means, we should not be surprised

³⁴ *Ibid.*, 40.

³⁵ Narayan, for instance, identifies one of the roles of censure as satisfying the victim's animus against wrongdoers (1993, 170).

if she comes to feel that she is insignificant in the eyes of the state or the political community.

For some theorists, including Duff and Roxin, it is quite clear that recognition of the victim's suffering is tied to recognition of the offender's wrongdoing.³⁶ But this point has not always been made clearly, and other theorists seem to suggest that it is possible to recognize (1) the victim's suffering without recognizing (2) the offender's wrongdoing.³⁷ Victor Tadros, for instance, argues that we can affirm the victim without censuring the offender.³⁸ He starts from the premise that the state has a duty to protect its citizens, and that when it fails to do so it ought to affirm the rights of the citizen as a victim of criminal wrongdoing. Otherwise she might come to believe that the state does not take her wellbeing seriously. One way for the state to affirm her rights is, as Duff and other theorists have pointed out, to censure the offender for what he has done. Through punishment, we communicate to the offender that his conduct was wrong and that he should have treated the victim differently. But Tadros rejects this option. Instead, he argues that we can affirm the victim by showing her that adequate measures were taken to protect her, and that appropriate responses will be taken to protect her in the future, such as incapacitating the offender.³⁹

The problem with this suggestion is that it fails to appreciate *why* victims may come to feel insignificant, and why censure is the appropriate means of affirming their significance. What makes many criminal acts wrong is precisely the fact that the offender does not treat the victim with a degree

³⁶ von Hirsch (1993, 10) also seems to take this view.

³⁷ See, e.g., von Hirsch (1993), chapter 2; Narayan (1993); Tadros (2011).

³⁸ See Tadros (2011, 91-2). It's not clear to me how Joshua Glasgow understands the relationship between (1) and (2). He recognizes that (1) is the basis for (2), but also claims that these are two separate judgments; see Glasgow (2015, 610).

³⁹ Narayan seems to argue that recognition of the victim can be separated from censuring wrongdoers. She differentiates denunciation and censure: denunciation has to do with the *wrongfulness of the conduct* whereas censure addresses the wrongdoer. Denunciation identifies (and "takes a negative moral stance to") the offender as being responsible for the victim's wrong, but does not censure him. Censure is not required because denunciation is enough to address the victim's grievance. One context in which Narayan's position is clearly unsatisfactory is that of historical wrongs, both in the domestic and international context. It is not enough for the community to merely denounce the perpetrator's wrongs. To victims of such wrongs, denunciation without censure (i.e. recognition, reform, and reparations) would only amount to a hollow gesture. See Narayan (1993, 169-170; 171).

of moral significance that they should grant to other people.⁴⁰ Part of what it means to recognize an offender's wrongdoing is to recognize that he treated the victim as if she was less significant. This point is quite obvious when we consider recognition by the offender. If he truly appreciates the wrongfulness of his conduct, he should understand that he treated the victim as if she was less significant. And if he appreciates that he failed to treat the victim with a sufficient degree of moral significance, then he should understand that his conduct was wrong.

If the state had failed to take very basic steps to protect its citizens, then it would be partly to blame for the wrongs that its citizens have suffered. Here, the state has committed a wrong for failing to carry out its duty towards its own citizens. If the state is to blame for the victim's suffering, then the state should take steps to protect her in the future, and it would help to show her that such steps have been taken. This would be an appropriate response to its prior failure to protect its citizens because it offers adequate redress for its wrongdoing. The state recognizes its prior failing by making a commitment not to do so in the future. But in cases where the state has done its duty, the offender alone is at fault for his wrong against the victim. The state can of course show citizens that appropriate measures were taken to protect her, but this would not do anything to address the offender's wrong, which is what manifests a denial of the victim's significance.

Censure is the appropriate response because it conveys, to the victim, that the political community recognizes the offender's conduct as wrongful. *That* is the appropriate way to show victims that they are morally significant. Why should victims be satisfied with knowing that the state took adequate measures to protect her prior to the offence? When we are wronged, our indignation is against the wrongdoer for what he has done. It is not against those who have an obligation to protect us, especially if such an obligation has been met. But we may come to feel indignant against the political community when they fail to censure the offender because it would suggest that the polity does not take the wrong we have suffered with the seriousness

⁴⁰ See Hampton (1992).

that it deserves. Barring extenuating circumstance that might militate against the use of punishment, lack of censure suggests that the polity is willing to condone such wrongdoing⁴¹, which is what *would* make victims feel as if they lacked moral significance.⁴²

At this point, critics like Tadros may accept that recognition of the victim requires censuring the offender, but they will nonetheless deny that punishment is necessary for censure. They argue that censure can be communicated or expressed through non-punitive measures; that we can censure the offender (and thus recognize the victim) through formal convictions, symbolic punishments, civil judgments and remedies, public inquiries, restorative justice, etc. The next section examines this line of argument. I will show why such arguments are mistaken and explain the sense in which “non-punitive” measures are insufficient for adequate recognition.

3 Punishment and Recognition

We need to start by noting the difference between punishment and what we tend to think of as non-punitive measures. Punishment necessarily involves the communication or expression of censure and the imposition of hard treatment.⁴³ Both of these elements require justification because each one involves the purposeful infliction of something unpleasant by the state on its citizens.⁴⁴ Non-punitive measures will typically include either the expression of censure or the imposition of hard treatment, but not both. Formal convictions, symbolic punishments, and public inquiries are non-punitive in the sense that they aim to express censure or assign blame without imposing

⁴¹ See Chapter 5.2 for a discussion on the kind of things that impunity may imply.

⁴² Perhaps it is possible to recognize the victim without censuring the offender *when recognition is given by those who do not have the standing to censure the offender*. We often see something like this in the international context. For example, suppose Country A has committed a serious wrong against Country B. A neighbouring country, Country C, does not have the standing to censure Country A, but it can recognize Country B's suffering through other measures, such as the creation of memorials. This, of course, falls short of the kind of recognition that victims are interested in, which is that of the offender and the community to which they belong (including, in this case, countries other than Country C).

⁴³ See Chapter 5 for a detailed explanation of these two elements.

⁴⁴ See Chapter 5.1 for a brief discussion of the need to justify both the censuring and hard treatment elements of punishment.

hard treatment.⁴⁵ And civil remedies may impose hard treatment in the form of damages without censuring those who have committed a wrong.

Legal theorists and lay people often think of restorative approaches to justice as an alternative to punishment, along with probations. But insofar as these measures include hard treatment and are imposed in response to criminal wrongdoing (e.g. following a criminal conviction), they constitute a form of punishment rather than an alternative to punishment.⁴⁶ Restorative approaches to justice are often optional and undertaken in addition to an offender's court imposed sentences, but there is no reason to think that victim-offender mediations cannot be imposed as a mode of punishment. The mediation process, much like Community Payback/Service Orders, extract time from offenders and force them to confront their wrongdoing by at least listening to (and hopefully engaging with) victims and their community. Irrespective of any reparative actions that may result from such mediations, the process is itself a form of hard treatment, especially if it is imposed on offenders who have no options but to participate in (or be confronted by) the process.

So what does it mean to say that punishment is necessary for censure? For starters, it means that hard treatment punishment is the best or only way to communicate or express censure adequately.⁴⁷ Accordingly, punishment is the best or only way to recognize the wrong that the offender has committed against the victim. This is not to say that punishment is the best or only way to communicate censure for *all* criminal wrongdoing. Formal convictions and symbolic punishment may be enough to communicate censure for very minor offences, although I suspect reasonable minds can disagree about which offences are included in this category. But for serious offences, non-punitive measures would not be enough to communicate the appropriate level of censure that is warranted for such wrongdoing.

⁴⁵ Censure itself may be painful and constitute a form of hard treatment for the person who is censured. The hard treatment that comes with punishment is imposed in addition to any pain that may arise from being censured.

⁴⁶ See Duff (2003c; 2001, 92-106).

⁴⁷ Duff (2001, 29).

Critics who deny the necessary connection between punishment and censure are thus committed to the thesis that it is possible for us to communicate censure for serious offences through non-punitive measures. Such measures include formal convictions, symbolic punishment, civil remedies, and public inquiries; and they exclude fines, probations, community service, and victim-offender mediations. Each of these non-punitive measures fails, in different ways, to fully communicate the level of censure that is warranted for an offence.⁴⁸

Civil judgments may impose damages, which could equal the kind of hard treatment that would be involved in criminal punishment, but such judgments do not include a censuring element. Tort remedies can offer compensation to victims for the wrongful harm they have suffered, but it does not blame the offender for wrongfully harming the victim. The judgements issue a finding of liability rather than fault, i.e. an agent or party is found responsible but not deserving of moral criticism for their conduct.⁴⁹ If we think that full recognition of the victim requires the imposition of censure, then tort remedies fail to achieve full recognition because it lacks the censuring element of criminal punishment. There will be a significant portion of cases where adequate recognition is not possible because censure, rather than compensation, is the appropriate form of redress for the kind of wrong that has been committed. But we might also think that for some types of wrongs, tort remedies may be able to give some (and even adequate) recognition to victims because compensation, rather than censure, is enough to remedy the wrong they have suffered.

Public inquiries differ from civil judgements in that the former can issue both a finding of wrongdoing and a judgment of fault. They can also recommend reform measures, which are meant to help prevent similar harm or wrongdoing in the future. When such recommendations are implemented,

⁴⁸ This is not to say that non-punitive measures can never give *adequate* recognition to victims. I suspect this is another issue where reasonable persons can differ. For some, adequate recognition of the victim must involve a level of censure that is proportionate to the offence. For others, adequate recognition can involve a level of censure that is slightly less than what is proportionate to the offence.

⁴⁹ See Robinson (1996).

the measures may be onerous for those who have been found guilty of wrongdoing. But unlike punishment, the measures are not intended to be onerous, nor will they necessarily be onerous. Nonetheless, in these cases, public inquiries can offer full recognition of the victim in the same way that criminal punishment recognizes the victim, by censuring the offender (through hard treatment) for their wrongdoing. Quite often though, when the recommendations are not implemented, public inquiries are more akin to formal convictions than punishment, and they fail to censure the offender in the same way that formal convictions fail to censure the offender.

Formal convictions and symbolic punishment are intended to convey censure, but they do not include the use of (additional) hard treatment. Moral criticism would be expressed through words alone (i.e. formal convictions) or through the creation of new conventions which would express our criticism (i.e. symbolic punishment). We can tell offenders just how wrong their conduct was at the sentencing hearing, or we can start a new convention, such as giving offenders weeds or scarlet letters, which would symbolise our criticism of their wrongdoing.⁵⁰

Critics often suggest that this is possible; that hard treatment is unnecessary for censure, which can be expressed adequately through these non-punitive measures. These suggestions betray a lack of appreciation for the complexities of our moral communication, and how such communications evolve over time. Differences in the level of moral criticism we want to communicate for murder and minor assault are not something we can capture through mere words alone. It is not enough to say to those who commit murder that their conduct was most seriously wrongful, however vehemently we say it and however clear we are in explaining why the conduct was wrong. The adequate communication of certain judgments—including moral criticism—requires actions. For instance, think about expressions of gratitude or apologies.⁵¹

⁵⁰ See Tadros (2011, 102-104); see also Nathaniel Hawthorne's 1850 novel, *The Scarlet Letter*.

⁵¹ See Glasgow (2015).

Suppose you needed a kidney transplant and your colleague donates one of his kidneys to you. It is not enough for you to say a quick “thank you” as you might do to a stranger who holds the door for you, nor would it be enough for you to passionately say that you are eternally grateful for his generosity. Perhaps your verbal communication may be enough for your colleague, who is not expecting anything for his generosity. But would *you* really think that mere verbal expressions are enough to convey your deep gratitude?⁵² Or suppose I have deliberately and unjustifiably caused severe bodily harm to you, which requires you to undergo extensive treatment. You would surely think that the wrong you have suffered is not being taken seriously if the court merely disciplined me verbally and was willing to close the matter after I issued you a verbal apology. In the case of serious wrongdoing, formal convictions without punishment suggest to us that the courts do not consider such wrongdoing to be so serious after all. It conveys to us a false recognition of the wrong that has been committed.

That is how we communicate the range of our moral judgments or sentiments—often through actions rather than words alone. Critics agree that certain actions such as the giving of gifts, the undertaking of burdens, or the imposition of punishment are central to how we in fact communicate gratitude, remorse, and moral criticism. And they rightly point out that our way of communicating such sentiments is shaped by social conventions, which can change overtime. What they argue is that, when it comes to censure and punishment, we should hasten the evolution of our social conventions so that the giving of weeds or scarlet letters may be enough for us to communicate the full range of moral judgments. So rather than letting social conventions evolve organically, they argue that we ought to be active

⁵² It could be objected that, in this case, many non-verbal expressions of gratitude such as giving your colleague gifts or money would trivialise what your colleague has done for you. This may be true, but the objection then would not be that non-verbal expressions of gratitude are unnecessary because mere verbal expressions would be enough. The objection is about something else, i.e. the appropriateness of responding to certain acts of kindness with non-verbal expressions of gratitude.

in establishing new conventions of communicating censure, one that would be less harmful and costly than punishment.⁵³

I think critics are right to suggest that we have reasons to establish new conventions for communicating censure⁵⁴, but it is nonetheless unrealistic to think that we can adopt *drastically* different conventions such as the giving of weeds or scarlet letters as adequate means of communication. Before I defend this point, we also need to note that, by suggesting we find new conventions for communicating censure, critics have already conceded that punishment is necessary for *us* to communicate censure adequately.⁵⁵ Those who argue that punishment is necessary for censure do not argue that punishment will *always* be necessary for expressing censure. Many of us would argue that capital punishment is unnecessary for expressing *our* moral criticism for the most egregious wrongs,⁵⁶ but it does not mean that it was always unnecessary. Similarly, hard treatment punishment may be necessary for *us* to express censure, but it may not be necessary for more expressively advanced societies. It is quite possible for us to one day use non-punitive measures to communicate censure adequately. But this possibility does not show that punishment is unnecessary within our current legal system.

For critics then, a revised claim would be that punishment is necessary within our current legal system as a matter of convention, but we ought to establish new conventions so that punishment would no longer be necessary for censure. The plausibility of this suggestion depends significantly on the kinds of new conventions that critics argue should replace our current forms of punishment. The giving of weeds and scarlet letters are examples of drastically different conventions, and to think that we can

⁵³ Tadros (2011, 102).

⁵⁴ See Narayan (1993, 179).

⁵⁵ See Glasgow (2015).

⁵⁶ Even if capital punishment is necessary for expressing moral criticism, we may have reasons not to do use it, e.g. if it is not something that the state should do to its own citizens. Similar things can be said about many other forms of punishment, some of which are still in use today. What we should do is to look for alternative, more humane forms of punishment. These alternatives may not be enough to fully convey the level of censure that is warranted for an offence, but they constitute the best available means of expressing the highest level of censure that we can express, all things considered.

replace punishment with such conventions is to fail to appreciate how society actually moves forward.

Moral progress does not take place within a theoretical vacuum and it certainly does not occur when philosophers will it to occur. To paraphrase Paul Robinson, even if good instrumentalist academics are themselves ready to adopt new social conventions, the world is not made up of good instrumentalist academics.⁵⁷ The critics' suggestion that we can learn to communicate our moral sentiments through drastic new social conventions ignores the fact that moral progress is tied to socioeconomic, political, cultural, and technological advances. Philosophers can help speed up moral progress through intellectual discourse, public engagement, or social, political activism. We can advocate for more humane and meaningful ways of engaging with offenders. But genuine moral progress requires rational agents to actually be part of such progress; it requires time and often intellectual, social, and political struggles. The fact that, at this moment, many lay people compare restorative approaches to justice—which is arguably a form of punishment—to a slap on the wrist for offenders shows how difficult it will be to move away from our penal conventions. The sudden imposition of different conventions, which may be out of sync with the rest of society, will likely be met with confusion and backlash rather than moral progress.⁵⁸

Some theorists who defend the necessary connection between punishment and censure assign critics the task of coming up with new social conventions for the expression or communication of censure. This task should not be shouldered by critics alone; those who work on the philosophy

⁵⁷ Robinson was discussing whether we should keep the criminal-civil law distinction. Both the criminal law and the civil law can be seen as a system of prudential disincentives against a range of conducts, but the criminal law has an additional blaming function. Some utilitarians would be happy to merge the two systems into one, where disincentives are used to regulate behaviour without the assignment of blame. Against this view, Robinson wrote: "even if good utilitarian academics have themselves overcome the irrationality of blaming, the world is not made up of good utilitarian academics. Criminal policy must take account of the fact that the world is made up of people who see blame, condemnation, and punishment as natural and necessary aspects of human interactions." See Robinson (1996, 209)

⁵⁸ We don't even need to consider what would happen if we were to use absurd symbolic punishments, such as giving offenders weeds or scarlet letters, to censure offenders. For offences like murder, assault, and even property damage, a formal conviction or symbolic punishment (i.e. a negligible fine) would be enough to elicit outrage from the community and generate disrespect for our legal system.

of punishment must seriously consider both alternatives to punishment and alternative forms of punishment. But we need to appreciate *how* we should attempt to change social conventions, which are often deeply entrenched and widely shared among the general public.

4 The Significance of Formal Apologies

So far I have argued that the political community's recognition of the victim's suffering is an important aim, and that this is to be pursued by censuring the offender. I have also explained the sense in which punishment is necessary for censure. Formal convictions and symbolic punishments cannot convey an adequate range of moral judgements, while tort remedies and public inquiries (often) do not impose censure, which is what would be required for recognition in cases of serious wrongdoing. In this section, I discuss why we should see punishment as a formal apology. If what I have said in section 1 is correct—that the aim of reconciliation does not affect either the distribution or justification of communicative punishment⁵⁹—then we need to consider why it is important to see punishment as an apology, especially given all the criticisms that have been raised against it. What reasons do we have for pursuing the aim of reconciliation if it appears to be negligible within our account of punishment?

The reason I want to explore here has to do with the kind of penal institutions that we should pursue. As Duff points out, too many instances of punishment are, in fact, exclusionary. It divides our society into *them* (the offenders), whom the state needs to punish or segregate in order to protect

⁵⁹ I suspect this is not entirely correct. The aims of recognition and reform may be enough to entail general prevention, so the aim of reconciliation does not make a difference to whether communicative punishment can be justified. But reconciliation may bolster the preventive effect of punishment. The idea is that reconciliation—which repairs the civic relationship between offenders, their victims, and their community—can help the offender reintegrate into their community, and to perhaps receive support from the community after serving their sentence. Rather than seeing the offender as an outsider, who must be kept away from the rest of the law abiding citizens, the community should see him as a fellow citizen with whom they must continue to live. This can have an effect on recidivism rates. For instance, many states in the US have legislations that make it impossible for convicted sex offenders to find jobs or housing after serving their sentences. Society's reluctance to allow the offender to reintegrate into society can often serve as a stressor, which contributes to such individual's choice to reoffend. If this is true, then we would have (instrumental) reasons to pursue the aim of reconciliation in at least some types of offences.

us (the law abiding citizens) from harm.⁶⁰ Even after serving their sentences, ex-offenders are portrayed as an outside group, who often do not enjoy the same privileges as the rest of the polity. Against this exclusionary practice, we should strive for an inclusionary institution where offenders, regardless of their offence, are still part of the political community. To this end, the offender should take steps to reconcile with those whom he has wronged, and this requires that he make an appropriate apology, one that is enough to repair the damage that his wrongdoing has caused to their relationship as fellow citizens. In return, the victim and the political community should also be ready and willing to accept his apology and reconcile with him if they are to treat him as a fellow citizen, with whom they must continue to live after the offence.⁶¹

I suspect that Duff's idea about reconciliation and inclusionary punishment resonates with many of his critics, but these critics nonetheless find the idea of punishment as apology odd and difficult to accept. One of the criticisms that I mentioned earlier revolves around the nature of an apology (as a voluntary acknowledgement of one's wrongdoing) and its apparent conflict with the nature of punishment (as an unwanted burden that is imposed on an offender). An appropriate apology—one that is adequate for reconciliation—should address the *wrong* that has been committed rather than merely compensating any *harm* that has been caused.⁶² And in order to address the wrongdoing, the apology should convey recognition and regret for that wrong and a commitment to avoid such wrongdoing in the future.⁶³ Critics argue that we typically consider a *sincere* apology to be valuable because it conveys such sentiments and commitments, but it is doubtful that such judgments can be expressed through punishment. The fact that punishment is imposed on offenders, who have no choice but to serve their sentences, gives us reason to doubt the sincerity of the apology. Rather than

⁶⁰ See Duff (2001, 75-9).

⁶¹ *Ibid.*, 98.

⁶² Duff (2001, 94).

⁶³ *Ibid.*

constituting an adequate apology, punishment undermines the sincerity of the apology which it is meant to convey.⁶⁴

Duff's response to this line of criticism relies on the kind of relationship which the apology is meant to repair.⁶⁵ The offender and the victim are fellow citizens who owe each other certain levels of concern and respect. Whatever other relationships they may have with each other, criminal law and punishment are only interested in their civic relationship, which is formal and impersonal. This kind of relationship can of course be repaired by a sincere apology, but sincerity would not be necessary. It would be sufficient for the offender to issue a formal apology, to play his part, as a repentant offender would, and to go through the motions of an adequate apology (i.e. to carry out the kind of penal burdens that would be required for a fitting apology).⁶⁶

Punishment serves as an apologetic ritual, in which the offender is required to participate. The victim and the community are likewise required to play their role in this ritual, which is to accept the offender's formal apology. The apology may not be sincere and the offender may not genuinely recognize or regret his offence. But out of a concern for his privacy, we should neither insist on nor inquire about his sincerity.⁶⁷ The victim and the political community should treat the offender as *if* he regrets his wrong and to treat the damage to their civic relationship as having been repaired. As Duff puts it, we should act as *if* the apology is sincere.⁶⁸

But in order for the apology to remedy the wrong that has been committed—i.e. in order for us to act as *if* the apology is sincere—it must be appropriately weighty. If I had damaged your bike and merely said to you that I am sorry for the damages, my apology would be insufficient because it fails

⁶⁴ Tadros (2011, 101-2).

⁶⁵ See Duff (2001, 109-10).

⁶⁶ In talking about punishment and formal apologies, both Duff and his critics tend to use examples which involve personal relationships and the apologies that occur within that context. I think a more apt analogy may be the relationship between nation states, where the relationships are always formal, and where we have witnessed that mere verbal apologies for serious wrongdoing are not enough.

⁶⁷ See Tadros (2011), chapter 5, on the limits of our respect for the offender's privacy over his judgments.

⁶⁸ Duff (2001, 95). According to Duff, the point here is not that insincere apologies are valuable, but that we should assume that formal apologies are sincere out of a respect for the offender's privacy over his attitudes and judgments.

to address the seriousness of the wrong. The penal burdens that offenders bear are meant to give adequate weight to their apology, as an appropriate normative response that would suffice to repair the damages that result from the wrong.⁶⁹ Just as punishment is necessary for censure, penal hard treatment is in the same sense necessary for adequate apologies.

Critics object that it is not clear how suffering penal burdens can add weight to an apology, or how it is meant to deepen a verbal apology.⁷⁰ And why should we think that the (apparent) sincerity of an apology would depend on the suffering of penal burdens rather than a commitment not to reoffend?⁷¹ The second point is right. A sincere apology requires a genuine recognition of one's wrongdoing, and genuine recognition requires a commitment to self-reform. As for the first point, rather than trying to convince critics that punishment can deepen an apology, I will simply point out that the offender would receive the same form and level of punishment whether or not he makes an apology. Punishment is meant to help the offender appreciate why his conduct was wrong, and to help him work out what he needs to do in order to reform himself. His punishment may constitute an apology, but the state does not impose any additional punishment for the purpose of an apology. The burdens that are necessary for recognition and reform just happen to be the same burdens that would figure in an adequate apology.⁷²

Some critics nonetheless worry about the sort of attitude that offenders are required to take if punishment is meant to act as an apology. These critics argue that it would be inappropriate for the state or the community to force offenders to apologise, which involves taking a particular attitude towards his wrongdoing which he may not agree with. Duff's response to this objection is that offenders are not asked to *mean* what they

⁶⁹ Ibid., 109.

⁷⁰ Tadros (2011, 101-2)

⁷¹ Matravers (2011, 79).

⁷² Duff and I hold differing views on this matter. On Duff's view, we should first work out the kind of apologetic reparations that would be required for a particular offence, and in undertaking the apology, the offender is able to achieve adequate recognition and undergo a process of self-reform. I am inclined to suggest a reversed picture, where we first work out the level/kind of punishment that would be required for adequate recognition, and in serving his punishment, the offender can also be *seen* as undertaking an apology for his wrongdoing.

say since that would indeed be unjustifiably intrusive. And proponents of the apology ritual certainly do not require (or even ask) offenders to issue any kind of verbal apology, which would force them to say things that they may not want to say.⁷³

I am not sure if this kind of response would be enough to satisfy staunch critics. Even if offenders are not asked to do anything for an apology—e.g. to issue a verbal apology, or take a particular attitude towards his offence, or bear additional burdens for the sake of an apology—the fact that punishment is meant to convey an apology implicitly suggests that we would assume certain things about the offender which may not be true. We would, for instance, assume that the offender regrets his wrongdoing, or that he disowns his conduct as something that he should not have done and will not do in the future. All of these things may be false. Should we not refrain from making such assumptions?

I think these kinds of assumptions are exactly the ones that a community should make; it should assume that an offender regrets his wrongdoing unless he announces otherwise.⁷⁴ This is not an issue about what the offender must do or say. The idea that punishment constitutes an apology depends on how the political community chooses to interpret the offender's punishment. And we should interpret his punishment as an apology because that is what is required of us by an inclusive penal institution which aims to reconcile the offender with his community.⁷⁵ If we have any confidence in the effectiveness of our penal practices, then we should think that an offender has recognized his wrongdoing after serving his

⁷³ See Bennett (2008), chapter 7.

⁷⁴ Duff suggests that we should also give offenders an opportunity to say explicitly that they do not undertake their punishment as an apology. Irrespective of questions about what kind of forum we should give to offenders to allow them to make such declarations, the option not to apologise must be made available. Many conducts are in fact (and will be) unjustifiably criminalized in many jurisdictions. A subset of these crimes includes conducts that are not wrongful, and those who engage in such conducts should be allowed to say to their community that they have not committed any wrongs. Another subset includes actions that are morally right. Ag-gag laws in some US jurisdictions, which prohibit whistleblowing within the agriculture industry, are an example. We need to allow those who commit acts of civil disobedience to disown the apology. See Duff (2001, 110-1).

⁷⁵ See Duff (2001, 123-4).

sentence.⁷⁶ If we have any confidence in our penal practice, we should treat him as if he would not offend in the future, and as if he has done all that is necessary to apologise for his wrong.⁷⁷

People's behaviours are to a large extent shaped by their communities. The way we treat other people affects both their character and our own. It also affects those who share their lives with us, whose worldviews we can influence through our actions. Our attitudes thus speak to the kind of punishment and community that we wish to maintain. It can tell ex-offenders that they are still part of our community and that they should abide by our norms.⁷⁸ Furthermore, it tells members of our community that we should treat him with the same level of concern that is due to any citizens. By treating ex-offenders as if they have apologised, we treat them with an attitude of acceptance, forgiveness, and respect—as fellow citizens within our community.

⁷⁶ Of course, we could have very little confidence in our corrections system, but if that is the case, surely the appropriate response is to look for more effective practices rather than subjecting offenders to ineffective ones.

⁷⁷ There is a separate question here about dangerous offenders—for whom we have good evidence that they will reoffend—and how we should treat such offenders. I do not address this issue in this thesis.

⁷⁸ *Ibid.*, 113.

Chapter 5

Censure and Prevention in Communicative Punishment

In this chapter, I defend the communicative account of punishment against instrumentalists' objection that communication alone is not enough to justify punishment. Instead, crime prevention—specifically general deterrence—must be part of the normative justification, especially in light of the costs that are associated with our penal institutions. Against this objection, I will show that communicative punishment *does* have an essential preventive dimension. Unlike instrumentalist accounts, where punishment is intended (in part) as a set of prudential disincentives that are meant to *deter* potential offenders from committing crimes, communicative punishment reinforces the moral reasons against wrongdoing. And by reinforcing the moral reasons against wrongdoing, communicative punishment aims at general prevention through moral dissuasion rather than prudential calculation.

I begin in section 1 by giving a brief overview of Duff's communicative account of punishment, especially of the justifications for censure and hard treatment. Unlike instrumentalist accounts—or accounts that feature instrumentalist aims—the hard treatment element of punishment, on Duff's view, is justified by its communicative rather than deterrent function. Then, in section 2, I explain exactly what is communicated through punishment, and show how general prevention is part of the communicative aims of punishment. What I have to say in section 2 will not be enough to satisfy instrumentalist critics, since they can still argue that moral dissuasion can only amount to an ineffective means of crime prevention. This is an empirical point, which I will leave largely unaddressed in this chapter and the thesis, but I argue in section 3 that, irrespective of the efficacy of different modes of prevention (i.e. moral dissuasion vs. prudential disincentives), we may have good reasons, which go beyond the justification of punishment, to favour a system of moral dissuasion. The discussion in section 3 will be tentative and

incomplete, but it will highlight a point for further discussion in the philosophy of punishment.

1 Censure and Hard Treatment

The consensus among legal theorists is that punishment, as a conceptual matter, necessarily involves the expression of censure and the imposition of hard treatment.¹ When a supposed offender is punished, he is both censured and subject to treatment that is materially burdensome. Without reprobation, the impositions of hard treatment or material burdens would not be punishment but more akin to taxation, and without hard treatment, reprobation is not sufficient to amount to punishment.² Both of these elements require justification if we are to justify the use of punishment.³

Quite often though, normative theories of punishment—i.e. theories about what can justify the use of punishment—tend to focus on hard treatment without commenting much on censure.⁴ Two explanations come to mind for why this might be the case. First, hard treatment, rather than censure, is seen as being more in need of justification, presumably because we find it more objectionable for the state to impose hard treatment on its citizens. Second, many normative theories put forth instrumentalist justifications, where punishment is justified (at least in part) as a means to crime prevention, and the preventive function of punishment is typically associated with the use of hard treatment rather than the expression of reprobation. But given that censure can be painful for the person who is being censured—even if the pain of censure may be much less than that of hard treatment—some kind of justification is nonetheless required when the

¹ See Feinberg (1965).

² On Duff's view, where hard treatment is necessary for censure, we cannot have adequate reprobation without hard treatment.

³ See Narayan (1993).

⁴ Theories that tend to focus on the hard treatment element include self-defence or duty views of punishment and public health or quarantine models of punishment. See Caruso (Forthcoming). For self-defence views, see Tadros (2011); Farrell (1989).

state censures its citizens.⁵ Otherwise, the normative justification would be incomplete.⁶

In contrast to instrumentalist theories, expressivist theories typically justify punishment as the best or only means of expressing or communicating the censure, condemnation, disapproval, or reprobation an offender deserves for his wrongdoing.⁷ For expressivists, hard treatment has an expressive or communicative function; it is the necessary vehicle for the adequate expression of different degrees of censure. Critics and proponents of expressivist theories alike have argued (or at least concede) that the value of expressing deserved censure is not enough, by itself, to offer a positive justification of punishment. Punishment inflicts a great deal of pain on offenders, who must suffer the hard treatment that is inherent in the punishment. But it also comes at a cost to offenders' families, who may depend on the offender for support; to those who work within our corrections system, who are tasked with administering punishment on a daily basis; and to citizens who pay for the maintenance of our corrections system either directly through taxation or indirectly, because part of the state's resources are being allocated to corrections rather than to some other worthwhile social services. Critics argue that given the associated costs, punishment must produce some socially beneficial effects (such as crime prevention) in order to be justifiable.

A noteworthy feature of punishment is that however it is designed and whatever its aims may be, it will most likely have some deterrent effect.

⁵ Narayan (1993, 167).

⁶ If a supposed justification of punishment does not provide a justification for either censure or hard treatment, then it will be incomplete. But we can remedy the incompleteness by offering a justification for the missing element. However, if the normative theory denies that one of these elements can be justified, then the theory cannot amount to a justification of *punishment*—it is rather denying that punishment *can* be justified. This, I think, is the case with public health or quarantine models of “punishment”. See Caruso (Forthcoming), who offers a justification of hard treatment as a means of crime prevention but denies that agents can be deserving of praise and blame. On his account, offenders may be segregated from the community until they have reformed, but they do not deserve reprobation for their wrongdoing. This undoubtedly is a proposal for how we might respond to criminal offending, but the proposal is not a punitive response. Punishment presupposes responsibility and blameworthiness. A justification for using hard treatment to respond to crime does not amount to a justification of *punishment* if it denies that agents can be responsible and deserving of blame for their conduct.

⁷ See Glasgow (2015).

Insofar as punishment involves the use of hard treatment, even purely retributive institutions will deter a significant number of otherwise would-be offenders. The fact that rational agents are motivated to avoid pain, and that the use or threat of hard treatment can discourage some rational agents from offending, gives punishment a natural preventive function.⁸ For those who defend non-instrumentalist accounts of punishment, the question then is not whether their proposed penal institution will lead to prevention, but whether the institution is justified *if* it did not produce any preventive effects.

In response to this thought, some theorists have developed a hybrid view where punishment should be used within a censoring framework, but its use is justified to the extent that it acts as an effective general deterrent.⁹ That is, the level of punishment we inflict on offenders must be proportionate to the level of censure he deserves for his offence, but we should abolish our penal institutions if they cannot or need not be used for preventive purposes.¹⁰ For these expressivists, the hard treatment element of punishment is not only the necessary vehicle for expressing various degrees of censure, but it also acts as a prudential disincentive against criminal offending. If the criminal law communicates to citizens through moral persuasion, then punishment is meant to act as a prudential supplement to the moral voice of the law.¹¹

Duff's communicative account is a type of expressivist account, where the communication of censure is justified as an intrinsically appropriate moral response to the offender for the wrong he has committed.¹² But it is not a hybrid view since communication is meant to provide a *complete* justification for punishment. Censure or moral criticism is what we ought to communicate to the offender if we are to treat him as a responsible agent, who is capable of understanding the moral reasons against his wrongdoing. We ought to

⁸ See von Hirsch (1993), chapter 2; Narayan (1993).

⁹ See Hart (1959); Rawls (1955); von Hirsch (1993, 14); Narayan (1993).

¹⁰ von Hirsch and Narayan maintain that, on their accounts, crime prevention is meant to be pursued within a censoring framework, such that the state is not permitted to inflict disproportionate punishment for the sake of prevention. But they also argue that punishment cannot be justified on the basis of censure alone, and that we ought not to use punishment if it cannot or need not be used to prevent crimes.

¹¹ See von Hirsch (1993), chapter 2, and von Hirsch (2003).

¹² Duff (2001).

hold him to account for his wrongdoing if we are to treat him as someone who is still a fellow member of our political community after the offence.¹³ Duff's view agrees with other expressivist views about the relationship between censure and hard treatment; the latter is needed to give weight to the former so as to convey the degree of moral criticism that is warranted in response to the offender's wrong. But on Duff's view, hard treatment is justified in terms of its communicative function only; it is justified as the necessary means of communicating appropriate degrees of censure, and as constituting an appropriate reparative apology that the offender owes to whom she has wronged, irrespective of its preventive function.¹⁴

Two popular lines of criticism have been raised against Duff's account, and both converge on the objection that communication alone is not enough to justify punishment. Something else, specifically crime prevention, is needed to justify an institution that is as costly and harmful as criminal punishment. The first line of criticism, which is typically raised by instrumentalists (and is also applicable to other expressivist theories), denies that hard treatment is necessary for censure. Critics who take this line agree that hard treatment punishment is only one *possible* means of communicating censure, which we can do just as well through symbolic punishment or formal convictions.¹⁵ These critics may still see censure as an appropriate moral response to wrongdoing, which may play a role in adjudicating between different forms of punishment. However, if censure can be communicated through less harmful means, then communication is not enough to justify the hard treatment element of punishment and it alone cannot provide a complete justification for punishment. The use of hard treatment must be justified by some other function, which instrumentalists take to be that of crime prevention or general deterrence.¹⁶

¹³ See the discussion in Chapter 4.4.

¹⁴ See Duff (2001).

¹⁵ See Tadros (2011), chapter 5.

¹⁶ The first line of criticism can also lead to an abolitionist argument—e.g. censure is the *only* appropriate response to wrongdoing, and we can abolish punishment since it is not necessary for censure.

In the previous chapter, I defended the view that hard treatment is necessary for the communication of censure. It might be possible for communicatively advanced societies to convey censure without the use of hard treatment, but this does not mean that hard treatment is not necessary for *us* to communicate censure adequately. Given the discussion in the previous chapter, I will not dwell on this issue, except to point out that instrumentalists can accept the necessity of hard treatment for censure and still argue that the communication of censure is not enough, by itself, to justify the use of hard treatment and consequently the use of punishment.

This takes us to the second line of criticism, which is usually raised by expressivists who endorse some kind of a hybrid account, where punishment is distributed within a censuring framework, but is only justified to the extent that it acts as an effective general deterrent.¹⁷ These critics typically agree with Duff that the censuring element of punishment is justified as an appropriate response to wrongdoing, which must be communicated through hard treatment (though they may disagree about the message of censure and how that message should be expressed or communicated). But given the costs that are associated with hard treatment punishment, effective crime prevention—even if it is not an essential justifying *aim* of punishment—must be a condition or side-constraint on the use of hard treatment and thus of punishment. So rather than justifying hard treatment wholly in terms of its censuring function, these critics agree with instrumentalists that hard treatment (and consequently punishment) should also be justified by their preventive function.

Both criticisms point to Duff's reluctance to assign a normative role to general deterrence (or *deterrence* of any kind) in his justification of punishment. The notion of deterrence that is in play throughout much of the literature on punishment revolves around the idea that the use or threat of hard treatment can discourage some rational agents from offending. It involves offering rational agents prudential reasons to refrain from criminal wrongdoing. The aim of general deterrence involves punishing offenders in

¹⁷ See von Hirsch (1993); Narayan (1993).

order to deter *other* rational agents from committing an offence, while special deterrence involves punishing offenders in order to deter *them* from re-offending. As a mechanism of crime prevention, deterrence is in contrast to moral dissuasion, where (potential) offenders are made aware of the moral reasons for refraining from various conducts. The kind of crime prevention that Duff is aiming for is moral dissuasion rather than deterrence.

Note that on Duff's account, the communication of censure entails special prevention in the sense that punishment communicates to individual offenders the moral reasons they have for refraining from future wrongdoing.¹⁸ Censure or moral criticism is communicated to an offender with the hope that he will come to a genuine recognition of his wrongdoing and make a commitment not to reoffend.¹⁹ The criticisms that are raised against Duff is not that crime prevention of any kind is missing from his communicative account, but that a compelling justification for hard treatment punishment should include *deterrence* (i.e. an appeal to prudential disincentives) since a system of pure moral dissuasion would be unacceptably ineffective. And in order for our penal institution to be justified, especially given its associated costs, it must act as an effect *general* deterrent.

In section 3, I will say more about the alleged ineffectiveness of a penal system that is modelled purely on moral dissuasion. For now, we can simply note the intuitive appeal of the above line of criticism against Duff. That is, we generally share the intuition that punishment should (at least in part) perform a preventive function. As we have seen, the goal of crime prevention is so important that many expressivists (and self-professed retributivists) see it as an essential part of their justification of punishment.²⁰ For those who do not share this intuition, instrumentalists would ask them to consider the following hypothetical.

¹⁸ On Duff's view, this kind of special prevention follows from censure and is part of the communicative justification for punishment.

¹⁹ See Duff (2001, 88-89; 107-9).

²⁰ See Husak, (1992); von Hirsch, (1993), chapter 2. For Duff's view on this, see his debate with von Hirsch in Duff (2003a; 2003b) and von Hirsch (2003), specifically von Hirsch (2003, 413-414) and Duff (2003b, 426-27).

Suppose that punishment has both a preventive function and a non-preventive function (e.g. a communicative, expressive, or retributive function), and that when we abolish our penal institutions we inevitably lose both a preventive tool and a communicative, expressive, or retributive tool.²¹ Instrumentalists argue that if we were to abolish punishment, surely we would be much more concerned about losing its preventive effects than about its censuring or retributive function. And if we are more concerned with prevention, it would seem to suggest that punishment should primarily be justified on preventive rather than non-preventive grounds.²²

I believe instrumentalists are right that punishment must have a preventive function if it is to be justified. But their point does not imply (i) that preventive aims must take precedence over communicative aims, nor does it imply (ii) that punishment must aim at prevention since a purely communicative account may be justified as long as it can give us prevention as a side-effect. Furthermore, it does not imply (iii) that we should be satisfied with crime prevention of any kind. Instrumentalists and expressivists appear to agree that the use of hard treatment can lead to general prevention by offering rational agents prudential disincentives against criminal wrongdoing. But some also appear to believe that the communication or expression of censure bears little or no relation to general prevention. This second point is false. In the next section, I will illustrate the intimate connection between censure and general prevention, and show how the expression of censure can entail general prevention as a matter of moral dissuasion. If I am right, then communicative theorists can respond to instrumentalist critics by pointing out that general prevention is part of the communicative justification for punishment, and that it is possible to pursue general prevention without conceptualising punishment as a set of prudential disincentives. In section 3, I will comment on the criticism that a system of

²¹ Although, as I have noted above, the kind of communication that Duff aims for does include a preventive purpose—i.e. punishment aims to communicate to the offender that his past conduct was wrong and to persuade him that he ought not to engage in similar conducts in the future. (And, as I will argue in the next section, the communicative aim of punishment also includes a *general* preventive purpose.)

²² Tadros (2011), chapter 5.

pure moral dissuasion would be unacceptably ineffective and explain why we should nonetheless aim for a system of moral dissuasion even if it proves to be ineffective within our society.

2 Censure and Prevention in the Justification of Punishment

In arguing for the expressive function of punishment, Feinberg pointed out that punishment can perform various acts, some of which are only possible if punishment did in fact express reprobation. He highlighted four such acts, most of which can help reveal the connection between censure and general prevention. These include *authoritative disavowal*, where, by punishing and condemning an offender for an action, the state publicly disavows that action and recognizes it as being wrongful; *symbolic non-acquiescence*, where punishment conveys a shared attitude among citizens that they do not endorse or condone the action in question; *vindication of the law*, where punishment “emphatically reaffirms” the moral judgements that are laid down in the law; and *absolution of others*, where the conviction and punishment of one supposed offender absolves other suspects (and sometimes the victim) of any blame.²³

The first three acts share a common theme, which is that punishment is a well-understood and established means of reaffirming norms and disowning violations of such norms. When the state does not punish an offender, it could signal that the state or his fellow citizens condone his conduct, or that his conduct is not really judged to be wrongful. Punishment can play a role in either reaffirming or repudiating provisions of the criminal law, and it can do so through moral dissuasion (as well as prudential dissuasion). The connection between censure and general prevention rests on the moral message that is communicated to citizens through punishment. And in order to illustrate this connection, we need to clarify exactly *what* is being communicated and to *whom*.

Within the interests of this chapter, the question of *what* is being communicated and to *whom* is not a normative question about what

²³ See Feinberg (1965, 404-8).

punishment *should* communicate, and to whom it should be communicated. Nor is it a question about the specifics of what various legal theorists have had to say about the message of censure, and what that entails. The question is about a broad, common sense notion of censure, one that ordinary, reasonable people would accept. This ordinary notion of censure is common to expressivist accounts of punishment, though different theorists will have different things to add to it. In showing that general prevention is entailed by an ordinary common sense notion of censure, I will also show that the connection between censure and prevention is not specific to any particular expressivist account of punishment.

Within the Anglo-American context, punishment *typically* expresses judgements of wrongdoing and blame. Punishment can, of course, express other judgements and sentiments, including ones which many of us would consider to be inappropriate, such as sentiments that are associated with pure retribution and revenge. And depending on how punishment is carried out, we may (mistakenly) interpret it as expressing certain sentiments even when it is not intended to express such sentiments. Regardless of what else might be expressed through punishment, and how we might interpret the content of penal expression, I think we can agree that punishment, at the very least, expresses judgements of wrongdoing and blame.

Generally speaking, when the state punishes an offender for a supposed offence, it expresses the condemnatory judgment that the conduct in question is wrong, and that the offender deserves moral criticism or blame for having engaged in that conduct. Most citizens, I think, see punishment as expressing judgments of wrongdoing and blame regardless of why an offender is being punished. An offender may be punished for retributive or instrumentalist reasons, but in punishing the offender, the state publicly recognizes that he has acted wrongly and criticizes him for having done so. If punishment did not express judgements of wrongdoing and blame, the state would not be able to use it to disavow the offender's conduct, nor would it be able to use punishment to give official recognition to victims for the wrongs they have suffered.

Assuming a legal system in which the state, on behalf of its citizens, dictates a set of shared norms, whose violations are punished by the state, which is again acting on behalf of its citizens, the judgements that are expressed through state punishment are often expressed implicitly by its citizens who make up the political community.²⁴ This is true at least in non-controversial cases where there is shared agreement among citizens regarding the wrongfulness of a conduct. If punishment did not typically express judgements of wrongdoing and blame, the political community would not be able to use punishment to show solidarity with victims. As Feinberg argues, punishment is a conventional way for the community to show that they do not endorse or condone the offender's conduct. It is possible for the community to show, through punishment, that they do not condone the offender's conduct precisely because punishment includes judgements of wrongdoing and blame.

So when the state punishes an offender, it expresses judgments of wrongdoing on behalf of its citizens; it censures the offender on behalf of citizens, who, in sharing the victim's wrong, have a legitimate interest in holding the offender to account.²⁵ But censure is also communicated to

²⁴ There are ways for citizens to show disagreement with what the state expresses (e.g. protests or speaking publicly against what punishment is meant to express). When citizens show no sign of dissent, especially when various means of dissent are available and tolerated within the political community, citizens *implicitly* express the judgments that are expressed through state punishment. This point is more salient when citizens have an obligation to voice their disagreement with the state's decisions and actions.

²⁵ I take this to be a fairly uncontroversial description of what is going on in serious and paradigmatic criminal wrongs, e.g. murder, rape, assault, fraud, etc. There is agreement between the state and an overwhelming majority of citizens that such conducts ought to be criminalized, and we have a legitimate shared interest in such cases when they occur. Some theorists have sketched a slightly different, idealized view of our legal system. See e.g. Duff & Marshall (2016). On their idealized view, offenders should see their punishment as an active fulfilment of their civic duty rather than an inconvenient burden or an imposition which they must endure passively. Offenders ought to *make* an apology and commit to a process of self-reform rather than seeing punishment as an apologetic ritual which is *imposed* on them. This requires a certain conception of the law, as a common law, one that "free and equal citizens can make their own" (Duff & Marshall (2016, 39)). Such a conception of the criminal law requires a set of shared norms, not as something imposed on citizens by the state, but as something that citizens can accept as theirs upon critical reflection—as something that reflects their shared values. This conception obviously does not apply to non-paradigmatic wrongs, or controversial criminal prohibitions, where the state acts more like a separate sovereign which imposes certain norms on its citizens, rather than speaking on behalf of its citizens. I would not expect the connection between censure and general deterrence to hold when it comes to punishing conducts which might be considered

citizens, who have a legitimate interest in seeing that the offender is held to account. There is nothing odd about a political community who communicates censure to itself. We often do this in our daily activities. I can censure myself for being tempted to do something that I believe I ought not to do, thereby preventing myself from doing something wrong. And I can censure myself after the fact for having done something that I believe I should not have done. Likewise, when I see another person doing something that I believe he ought not to do, I can tell myself (and perhaps him) that his action is wrong and that he should not have acted thus. By telling myself that his action was wrong, I also confirm to myself that I should not act as he did.

So far we see that censure can be communicated or expressed to a specific individual or the general public. It is expressed to a specific individual when it addresses and criticizes the offender as the agent who is responsible for the wrong that has been committed. And censure is expressed to the general public when it addresses the community, who have a legitimate interest in the offender's wrongdoing.²⁶ Not only is censure often expressed both to individual offenders and the general public, but the message of censure also contains both general and specific judgments. In blaming the offender for his offence, censure applies a general norm (i.e. that ϕ is wrong) to a specific individual as the agent who is blameworthy for an instance of ϕ . And in communicating censure, we are simultaneously saying that the conduct in question is wrong and that the offender is responsible and blameworthy for that particular wrongdoing.

The general claim includes the criminal law's declaration that certain conducts are wrong. It also includes the judgement that those who engage in such conducts have done something wrong and deserve moral criticism for what they have done. Judgements of wrongdoing and blame are reinforced by the criminal law's declaration about wrongful conducts. And in punishing and censuring an offender for his offence, the message of censure reiterates

controversial, but the disconnect between censure and general deterrence has more to do with the nature of the conduct than the conceptual relationship between censure and prevention.

²⁶ At the very least, the community is capable of hearing the communication, which is made in public.

to citizens the moral reasons not to engage in those conducts (i.e. that such conducts are wrong), and conveys the message that the state really does consider such conducts to be wrongful. When the state punishes offenders, this message is conveyed publicly so that it can be heard by citizens. But it is also communicated to citizens, who have a legitimate interest in how we respond to the offender. By reiterating the message of the law to citizens, communicative punishment aims at general prevention in a way that is similar to that of the criminal law.²⁷

Before we get to general prevention, it will be useful to start with the connection between censure and special prevention. Censure entails special prevention when the state *communicates* the message of censure to the offender. When censure is communicated to offenders, the state or political community engages in a moral dialogue with the offender, who is expected to participate in the communicative process. This is in contrast to mere expressions of censure, which do not require offenders to play any role in the penal process. When censure is merely expressed against the offender, he need not engage with or even be aware of the moral criticism that is directed against him. For expressivists, censure need not be expressed with the hope that the offender will acknowledge or atone for his wrong; it can be expressed purely for the sake of the expresser, who wishes to convey his disapproval or condemnation of the offender's conduct.

On Duff's account, the communication of censure has implications for what the offender should do. Through the communicative process, he should gain an appreciation for why his conduct was wrong. He should accept his wrongdoing as something that he should not have done and make a commitment not to reoffend in the future. For Duff, the communication of censure aims to dissuade the offender from reoffending by bringing him to a deep moral appreciation for the wrongfulness of his conduct. Here, the distinction between communication and expression is significant. The communication of censure entails special prevention while expression does

²⁷ von Hirsch seems to recognise this point, but does not go on to argue that a purely communicative account like that of Duff's can give us general prevention. See von Hirsch (1993), chapter 2.

not because the former requires active participation by the offender, as someone who (at some point during the communicative process) will be receptive to the message of censure.

The connection between censure and general prevention rests on the communication of censure to the wider community. The message of censure is essentially the same whether it is communicated to the offender, the victim, or the community. It concerns the wrongfulness of the conduct in question and the offender who is responsible and blameworthy for that conduct. In communicating this message to the community, censure entails general prevention in the same way that it entails special prevention, as a matter of moral rather than prudential dissuasion. Except in the case of general prevention, moral dissuasion is directed at the political community rather than any individual offender. When it comes to special prevention, an offender could not come to a deep appreciation for the wrongfulness of his conduct (which is necessary for motivating a process of self-reform) if he is not receptive to the message of censure. We might therefore think that the difference between communicative and expressivist accounts of punishment is significant for the connection between censure and general prevention. If an entire political community is deaf to the message of censure, then censure could not entail general prevention either. But expressivist accounts of punishment need not exclude the kind of communication that is required for general prevention.

Expressivist accounts of punishment presuppose a community that is at the very least aware of and interested in how and whether we respond to offenders. For expressivists, we punish offenders in part to express our strong disapproval for what they have done. If the community had no interest in or awareness of how the state responds to offenders' wrongdoing, then the expression of censure could not entail general prevention since the expression of censure would not be heard by the community. But to the extent that expressions of censure are heard by citizens, it is possible for censure to bear some relationship to general prevention on expressivist accounts of punishment.

Note that earlier in section 1, I distinguished between two conceptions of crime prevention: that is, prevention through moral dissuasion and deterrence through prudential disincentives. The idea that censure is necessarily communicated through hard treatment coupled with the fact that hard treatment has a natural preventive effect makes it difficult to pinpoint whether prevention comes from moral or prudential appeals. There are empirical questions here about the extent to which censure can have a preventive effect, and about the socio-economic and political conditions that would be necessary for censure to have any preventive effect. Without getting into the empirical research, it is possible to illustrate the connection between censure and prevention by considering the implications of not punishing offenders. Specifically, we need to consider what the state fails to communicate to citizens when it fails to punish offenders.

When the state does not punish an offender, one of the things that it fails to do is to recognize the offender as being blameworthy for his wrongdoing.²⁸ And depending on whether the state chooses to adopt any other remedies, by not punishing an offender, it fails to affirm—after the fact—that the offender’s conduct was wrongful and that the offence should not occur in the future.²⁹ Impunity can imply many things, and depending on the reasons for not punishing offenders, impunity can imply that the state does not take the wrongdoing seriously or that it is not sufficiently wrongful to warrant criminalisation and punishment.³⁰

In light of the natural preventive effect of hard treatment, when the state repeatedly withholds punishment, the prudential disincentives against criminal wrongdoing start to disappear. For anyone who sees punishment in part as a set of prudential disincentives, their prudential reasons against offending disappears when the state fails to punish offenders on a consistent

²⁸ See Chapter 4 for why civil remedies and formal convictions fall short of this kind of recognition.

²⁹ This is similar to what Feinberg was talking about when we discussed authoritative disavowal. By not punishing the offender, the state does not disown the offender’s conduct, and acts as if it condones the conduct in question. See Feinberg (1965).

³⁰ This point is similar to Feinberg’s point about symbolic non-acquiescence, where a failure to punish can suggest that the community endorses the conduct in question as being morally acceptable. See Feinberg (1965).

basis. But given that punishment can also reiterate the moral reasons against criminal wrongdoing, impunity also erodes the clarity with which the criminal law communicates the moral reasons against offending.

The state can choose not to punish offenders on many levels. There are (at least) two types of cases that can illustrate the link between censure and prevention. These involve situations where the state chooses not to punish anyone who commits a particular offence, and where it does not punish certain groups of offenders for certain offences. An example of the first case would be if a jurisdiction decided not to enforce a particular criminal prohibition—e.g. no one who commits drug possession offences are convicted or punished. If the state never punishes those who commit ϕ , then it never communicates to its citizens, *through punishment*, that ϕ is a criminal wrong, or that it is sufficiently wrongful to concern the criminal law. The state, by not punishing any instances of ϕ , does not communicate clearly and fully to citizens that they ought not to ϕ , or that those who ϕ ought not to do so in the future.

Insofar as there is still a criminal prohibition on ϕ , there is still an official declaration that ϕ is wrong and is of concern to the criminal law. But when that prohibition is not being enforced through the courts or corrections, the state implicitly sends the message that it is not really concerned with ϕ or those who engage in ϕ .³¹ Without punishment, the state is no longer communicating to citizens a clear message that ϕ is a criminal wrong and that they should not engage in it. For citizens who are capable of responding to the moral voice of the law, the moral case against ϕ is no longer being made clearly by the state; it is no longer trying to persuade citizens not to commit ϕ . If everything else about punishment remains the same—i.e. if punishment is still imposed on those who commit other offences, but just not

³¹ This is similar to what Feinberg refers to as vindication of the law, where impunity for wrongful conducts suggests that the law is merely paying lip service to morality. Note though that my point about the state sending a message of condonation through impunity only holds if (1) prohibitions are normally enforced, and (2) punishment is normally seen as sending a condemnatory message. In our current system, the percentage of crime that results in guilty verdicts is extremely low, but it is reasonable to think that we at least have the perception that (1) and (2) holds at least in serious offences. See Feinberg (1965).

on those who commit ϕ —we would start to think that the state no longer considers ϕ to be a prohibited conduct, and that it is now permissible to ϕ .

Depending on citizens' attitudes toward ϕ , whether the state chooses to punish those who commit ϕ will have implications for crime prevention that go beyond the general deterrent effect of censure. If citizens believe that ϕ is not a wrong which should concern the criminal law, then the state's choice not to punish those who commit ϕ will be supported by its citizens. However, if ϕ is a rather serious moral wrong—e.g. murder, rape, assault, fraud, etc.—or if citizens believe that the state should be concerned with ϕ , then widespread impunity for ϕ can erode citizens' respect for the law, which can in turn affect citizen's responsiveness to the state's attempt to morally dissuade them from other criminal wrongs.³²

We can now turn to the second type of cases where the consequences of impunity can help highlight the connection between censure and general prevention. Here, the state does not punish certain groups of offenders for a subset of criminal wrongs, but it nonetheless punishes other groups of offenders for those wrongs. There are situations where the state may have legitimate reasons to withhold punishment for deserving offenders. These reasons may stem from social policy considerations such as serious negative social or political consequences that can result from punishing a particular offender. They can also include considerations having to do with old age or assisting with other criminal investigations.³³ When the state has legitimate reasons to withhold punishment, it is not treating the offender's conduct as if it is not wrongful and as if it is not in need of censure. Instead, the conduct is still treated as being wrongful, but not one that warrants the imposition of censure all things considered. In contrast to these scenarios, the type of cases I have in mind involves illegitimate reasons for not punishing offenders.

Suppose that a particularly serious offence, say rape, is normally punished with a prison sentence upon conviction. However, for a particular

³² See Robinson & Darley (1996).

³³ Husak (1992, 448-9); see also 1987 Federal Sentencing Guidelines Manual, chapter 5: <https://www.ussc.gov/guidelines/archive/1987-federal-sentencing-guidelines-manual>

group of convicted offenders, say white, middle-class, varsity football players, the sentence is, more often than not, community service. Community service can constitute a form of punishment, but it is hardly a fitting sentence for the offence in question, and compared to the prison sentence that other rapists would normally receive, we cannot, in all seriousness, say that the privileged group has been punished for their offence. When we object to the state's failure to punish offenders in such cases, our objection has to do with the state's apparent willingness to condone serious criminal wrongs when they are perpetrated by certain groups. We object because the state sends the message that the wrong, which has been perpetrated against and suffered by the victim, is not so serious when it is committed by a member of the privileged group.

In a sense, instrumentalists are right that without punishment, we would be concerned with the absence of a preventive mechanism. In the cases I have just described, we would be concerned that impunity would lead other members of the offender's group to commit similar crimes since the prudential disincentives against doing so are no longer in place. But without punishment, we should also be concerned with the absence of *another* preventive mechanism, one that is intimately connected to punishment's censuring function. Impunity also erodes this second kind of preventive function because the message of censure is no longer communicated to those who identify with the offender.

One of the reasons we have for objecting to the state's failure to punish offenders from the privileged group stems from our sense of justice. Not only is it unfair to apply the law in a discriminatory manner, but the injustice comes from our judgement that the wrongs which have been perpetrated are in fact wrongful and that the state should treat them as such. By not punishing the offenders, the state fails to show sufficient concern for the victims by not recognising the wrongs they have suffered. And in doing so, it fails to uphold a norm which we believe should be upheld and applied to everyone within our political community. By not punishing the offenders, the state fails to communicate to them that what they did was wrong, that they

ought not to repeat such wrongdoing in the future. This is why impunity suggests an implicit condonation of the wrong that has been perpetrated. For the privileged group, impunity weakens the message of the criminal law; it suggests to them that certain prohibitions are only wrong when committed by other people, but it is not (seriously) wrongful when they themselves commit those wrongs.³⁴

What I have said so far does not imply that impunity of any kind will weaken the message of the criminal law, or that impunity always implies condonation or endorsement of the wrong that is in question. Suppose, for instance, that I believe it is wrong to eat meat. When I choose not to censure strangers in a restaurant for eating meat, it does not imply that I condone their dietary choices. This is true even if I have good reasons to think that my fellow patrons will appreciate the moral case against eating meat (and consequently be motivated to make different choices in the future). On the contrary, it would be a different matter if I *never* censured my partner, my close friends, or my fellow ethicists for eating meat.³⁵

Whether impunity implies condonation of a wrong depends on the relationship between the party who censures and the party who is being censured and the obligations they have toward each other. I have an obligation toward my partner, my close friends, and perhaps my fellow ethicists to hold them to a certain ethical standard, but I do not have an obligation to hold strangers to the same standard. I also have an obligation to my partner and my close friends to be truthful about how I choose to live my life and an obligation to my fellow ethicists to speak my mind on what I think are important ethical issues. A consistent failure to censure those whom we have an obligation to censure would imply condonation of the wrong that is in question. If I censured my friends for eating meat but never my partner, it would suggest (to my friends and to outside observers) that I condone meat eating when it is done by my partner but not my friends.

³⁴ Again, in this case, not only does impunity weak the message of the law for members of the privileged group, it can also delegitimize the law for the rest of the citizens by eroding citizens' respect for the law and their responsiveness to the moral voice of the law.

³⁵ For anyone who does not like the example of eating meat, simply substitute other moral wrongs.

Political communities presuppose a set of shared standards and an obligation on members to hold each other accountable for actions that violate those shared standards. The state, on behalf of its citizens, has an obligation to hold citizens to account for certain wrongs. When it fails to fulfil this obligation on a consistent basis and in the absence of extenuating circumstances (e.g. during wartime), it could imply condonation or endorsement of criminal wrongdoing. In light of this obligation, if the state decided not to punish any offenders, it *could* suggest that the state no longer sees itself as having an obligation to hold citizens to account, or that it no longer sees punishment as an appropriate means of holding citizens to account.³⁶ And if the state decided not to hold citizens who commit certain types of wrongs to account, impunity *could* suggest state condonation of that wrong. Furthermore, if the state decided not to hold certain groups of citizens to account when they commit criminal wrongdoing, impunity again *could* suggest state condonation of wrongdoing when it is committed by particular groups. Whether impunity actually does imply these things depends in large part on how citizens perceive the state's decision not to punish.

One instance where impunity need not suggest condonation of wrongdoing relates to the opportunity principle of prosecution in common law jurisdictions. On an opportunity model of prosecution, the police and prosecutors are allowed discretion in deciding which cases to prosecute and which cases to divert away from criminal court hearings. Decisions are often made by considering factors relating to cost or availability of resources, public interest, seriousness of the offence, and proportionality. This is in contrast to many civil law jurisdictions which, in theory, operate on the mandatory principle of prosecution, which is intended to be fairer and to prevent matters of justice from political motivations and discriminatory practices.³⁷ There are two questions here. One has to do with the issue of

³⁶ We can also couch the state's obligation to censure offenders in terms of an obligation to protect citizens from harm. If we think that censure entails general prevention, then one way in which the state can protect citizens is by censuring offenders. The question then is why the state should attempt to protect citizens through the imposition of censure (as opposed to other modes of deterrence, e.g. prudential appeals). See section 3 for this discussion.

³⁷ See Ashworth & Redmayne (2010), chapter 6.

justice and whether an opportunity model of prosecution is in any sense unjust for allowing unelected officials to decide which offenders are held to account for which offences.³⁸ Another question has to do with the impact of selective prosecution on prevention, and whether decisions not to prosecute some cases lead to a weakening of the moral voice of the law.

In the two scenarios I discussed earlier, I argued that impunity suggests condonation of wrongdoing, which can both weaken the moral message of the law and erode citizens' respect for the law. The first type of scenario is defined by the state's decision not to punish a particular offence on a consistent basis. And the second type of scenario is defined by the state's decision to privilege certain groups of citizens. Neither of these features is entailed by the opportunity principle of prosecution. It is certainly possible for jurisdictions that operate on the opportunity principle to end up in one of these scenarios, but the scenario will not occur by necessity. On the contrary, allowing prosecutors the discretion to determine which cases to pursue could give the perception of a fair and just legal system, where wrongs that garner more of the community's attention are prosecuted and censured.³⁹

Given what I have argued thus far, we see that censure can be a means of crime prevention. Punishment reiterates the message of the criminal law, and in communicating to citizens the moral reasons against offending, censure aims at general prevention through moral persuasion. If I am right about the connection between censure and prevention, we need to rethink current disagreements between instrumentalists and expressivists. For one thing, it no longer makes sense for instrumentalists to ask, with respect to abolition, whether we would be more concerned about losing the preventive or censoring function of punishment since prevention comes in part from the expression of censure. In addition, when instrumentalists

³⁸ Although, in practice, jurisdictions that operate on the mandatory principle of prosecution also allow for significant discretion, and many cases are likewise diverted from criminal courts.

³⁹ This is not to say that this kind of practice does in fact lead to a fair and just legal system, or that it is always an appropriate measure of a fair and just legal system. In some cases, a just legal system would require prosecutors to lead public opinion by prosecuting serious crimes that tend to be ignored, especially on discriminatory grounds.

challenge expressivists on whether they would still use punishment if it was ineffective or unnecessary for crime prevention, we need to consider the possibility that punishment may be ineffective or unnecessary because censure has not been communicated adequately or because there are non-punitive means of expressing censure.⁴⁰

3 Censure and Prevention in the Distribution of Punishment

The conception of crime prevention that I have been discussing throughout this chapter is one of pure moral dissuasion. And as I mentioned in section 1, this is to be distinguished from prevention through prudential disincentives (i.e. general deterrence). Instrumentalists and hybrid expressivists like von Hirsch and Narayan—all of whom see punishment as offering prudential disincentives against criminal wrongdoing—will argue that what I have offered in this chapter amounts to an unacceptably ineffective means of crime prevention. They would argue that, if we take prevention seriously, we must use prudential disincentives to deter potential offenders since too many of us would be unmoved by moral dissuasion alone.

The disagreement over how we should conceptualize punishment—i.e. as a system of moral dissuasion only or a set of prudential disincentives—runs much deeper than the issue of prevention. It rests on disagreements about human nature and about how the political community should address potential offenders. These disagreements cannot be resolved here, but I can raise a couple of points in defence of a system of moral dissuasion. The first point goes back to the idea that punishment necessarily involves the imposition of hard treatment. As we have seen in chapter 4, many paradigmatic modes of communicative punishment (such as victim-offender mediations, fines, probations, community payback orders) involve treatments that are materially burdensome. This aspect of punishment is

⁴⁰ If I am right about the connection between censure and general deterrence, then symbolic reprobation should lead to general deterrence in expressively advanced societies which do not require hard treatment to express censure. And for critics who reject that hard treatment is necessary for censure, they need to explain the sense in which hard treatment is not necessary for censure if symbolic reprobation is not enough to produce a deterrent effect in our society.

necessary for the adequate communication of censure (i.e. moral dissuasion), but there is nothing to prevent those who wish to see punishment as a set of prudential disincentives from seeing such material burdens as a means of prudential rather than moral dissuasion.

Critics may argue that it is not enough to simply see the material burdens that are involved in communicative punishment as a set of prudential disincentives since they may not be severe enough to act as an effective deterrent. The aims of censure may dictate different levels of punishment than the aim of general deterrence, and when the latter calls for a higher level of punishment than what is required for the former, how we conceptualize the hard treatment element of punishment would no longer just be a matter of perception. Censure dictates that we must punish an offender to the extent that is necessary to communicate the moral criticism that is warranted in response to his offence. If the punishment is too lenient, then he will not receive the full message. And if the punishment is too harsh, he will receive more criticism than is appropriate, which may drown out any legitimate moral criticism that could have been communicated to him. But in order for punishment to act as an effective deterrent, it might have to be set at a level that would drown out the moral message of censure.⁴¹

The empirical research on prevention seems to suggest that the aims of censure and prevention need not pull in different directions. What matters for deterrence has less to do with the sentencing levels and more to do with offenders' perception of the risk or certainty of being detected and punished. Punishment can have a deterrent effect when offenders believe there is a substantial risk of being detected and punished, but there is limited evidence in support of any correlation between crime rates and the severity of punishment.⁴² At best, there are limited cases in which sentencing levels affect crime rates if offenders are aware of the sentence involved *and* believe that there is a substantial risk of being detected. But even in such limited cases, increases in penalty will eventually be met with diminishing returns.

⁴¹ Duff (2001, 87-88).

⁴² See Bottoms & von Hirsch (2010).

So it would seem that, in practice, there is a fairly weak case for raising sentencing levels for the sake of deterrence.⁴³

In theory though, we can still ask what we should do *if* there is ever a conflict between the aims of censure and deterrence. Do we give up on censure in favour of prevention, or vice versa? This issue partly concerns the two conceptions of prevention and which one we should adopt. But an adequate treatment of this issue also requires us to look beyond legal theory. On the issue of moral vs. prudential dissuasion, Duff maintains that the criminal law should communicate to citizens in a moral voice; it ought to offer citizens moral reasons to refrain from criminal wrongdoing. To use punishment as a prudential supplement to the moral voice of the law—especially if the prudential disincentives are severe enough to replace the moral voice of the law—is to fail to treat citizens with the respect that is due to them as rational moral agents.⁴⁴ There is something to be said for Duff's reluctance to conceptualize punishment as a set of prudential disincentives, but it does not depend on his point about treating citizens with due respect.

Many have argued that the state does not automatically fail to treat citizens with due respect when it offers them prudential disincentives against criminal wrongdoing.⁴⁵ If the state is justified in punishing D for ϕ ing, it does not treat D with disrespect if it warns D in advance that he will be punished if he ϕ s. To warn D in advance of the punishment he would receive for ϕ ing *is* to treat him as a rational agent, who is capable of exercising his autonomy in light of the consequences of ϕ ing. This line of argument is meant to show that the state can use punishment as a prudential disincentive without treating its citizens with any lack of respect. But even if this is true, there is still a significant difference between prudential appeals and moral persuasion which pushes us in favour of Duff's view.

To see punishment as an attempt at pure moral persuasion is to supply citizens with moral reasons against offending in the hope that moral

⁴³ See also Robinson & Darley (1996). They offer an instrumentalist defence of retributive punishment, i.e. if we care about prevention, an effective way to prevent criminal offending is to punish an offender according to a community's perception of retributive desert.

⁴⁴ Duff (2001, 82-85). But see Edwards & Simester (2014).

⁴⁵ See *ibid.*; Hoskins (2011); Matravers,(2011).

reasons alone would be enough to dissuade them from committing crimes. In other words, citizens refrain from wrongdoing not out of a concern for their self-interest, but because it is the right thing to do. On the other hand, to see punishment as a prudential supplement to the moral voice of the law is to admit that some members of the society may not be persuaded by moral reasons alone, and that for these citizens, prudential disincentives are necessary to deter them from wrongdoing. Prudential disincentive presupposes that rational agents are motivated to avoid pain, but moral persuasion presupposes (and requires) much more. For instance, it presupposes that citizens will be responsive to moral reasons, and it requires that the state has legitimate standing to ask citizens to follow the moral voice of the law. We may not be treating these citizens with any disrespect when we offer them prudential reasons against wrongdoing, but there are deeper issues about why these citizens would not be dissuaded by moral appeals alone. Here, questions about the justification of punishment connect to questions in social and political philosophy.

Before we say anything about what we should do when moral persuasion is not enough to dissuade citizens from offending, we should at least understand why it is not enough. If citizens are committing theft because the state has failed to meet its basic obligations to provide for citizens, then the appropriate response is not to increase prudential disincentives against theft, but for the state to do something about its failure to care for its citizens. In the real world, where society is far from being perfectly just, crimes are often a symptom of wider socio-economic and political inequalities. A society in which pure moral persuasion is enough to convince citizens to refrain from wrongdoing is more likely to be a just society, where the state treats every citizen with equal concern and provides every citizen with an opportunity to live a good life. Indeed, effective moral persuasion may require such a society.

This is not to say that within just societies, all citizens would be responsive to the moral reasons against criminal wrongdoing or that the prudential incentives to commit crimes will completely disappear. But the

extent to which criminal wrongdoing is a product of deeper systemic injustices and problematic worldviews should factor into our approach to punishment. To try to build a system of punishment in which we communicated to offenders and the general public in a moral voice is not to be any less concerned with crime prevention. It is to care about the kind of prevention that we may achieve. Even when we have good evidence that moral dissuasion alone may not be as effective as prudential disincentives, we should not be ready to give up on moral communication so easily. Before we give up on it, we should examine and address the systemic injustices and problematic worldviews that make people ignore the moral reasons against criminal wrongdoing.⁴⁶

The point here is that, if we recognize criminal wrongdoing as a symptom of deeper, structural problems, we should address the structural problems rather than simply treating the symptoms. This is the same kind of criticism that is often raised against those who try to deal with global poverty through effective altruism rather than dismantling inherently exploitative practices that lead to poverty; or against those who advocate for safeguards to protect women from harassment and exploitation without dismantling the patriarchal structure that tend to promote harassment and exploitation; or against those who try to solve our environmental problems by advocating for electric cars without addressing deeper problems about overconsumption. Of course, our attempts to deal with the symptoms of structural problems are compatible with addressing those structural problems. We can donate to charities while addressing exploitation, and we can protect women from harassment while dismantling the patriarchy. Likewise, we can use punishment to deter potential offenders while addressing the structural inequalities that tend to promote criminal wrongdoing. But until we make a concerted effort to address those structural inequalities, we cannot so easily

⁴⁶ Whether or not we should give up on punishment as moral dissuasion and when we should give up on it will depend on the costs of pursuing such a system. For instance, if the costs of addressing structural problems is too costly, or if there is an unacceptable and sustained increase in crime rates, then we need to consider alternatives (e.g. switch back to a deterrent system of punishment). These issues require further reflection, and I will not try to address them here.

dismiss moral dissuasion as an ineffective means of crime prevention. And while we work to eliminate the kind of criminogenic inequalities that exist within our society, we can strive toward a system of moral dissuasion both as a reminder of the work we should be doing to address various structural problems and as a step towards the cohesive society that we want to build, where we communicate to our fellow citizens primarily in a voice moral.

For instrumentalists who advocate what I think is an admirable idea—i.e. that punishment is only justified as an effective preventive tool or not at all—it is important for them to tell us why we are concerned with general deterrence. That is, *to what end* are we concerned with general deterrence? Instrumentalists would surely agree that it matters not just that there *is* less crime in society, but *why* there is less crime, and the means by which we build a society in which there is less crime. For instance, we can have a perfectly safe or secure society in which potential offenders are deterred by draconian laws and intrusive policing practices, but that is not the kind of society that instrumentalists would want to build. Their view of society is somewhere in between a secure community maintained by draconian means and a peaceful and cohesive community where the moral reasons against offending are enough to dissuade citizens. Their view of society fits with what our society actually looks like. But, as I'm sure they would agree, we should work towards the kind of peaceful and cohesive society where moral dissuasion would be an effective means of prevention. Again, contrary to instrumentalists' criticisms, a reluctance to prioritize general deterrence in the justification of punishment is not to care any less about crime prevention. It is to care about the kind of prevention and the kind of society that we would want. We ought to strive toward Duff's view of punishment, where moral communication is the only kind of prevention we need, because it is more in line with the kind of society that we ought to build, where citizens are treated with equal concern and respect beyond the realm of the criminal law.

Conclusion

In the Introduction, I stated that this thesis should not be treated as offering any solutions to the problem of outcome luck, but rather as an attempt to address some of the foundational questions that are implicated in that problem. Having examined some of the foundational questions in metaphysics and legal theory, I can now offer a brief commentary on how we might treat the problem of outcome luck in light of some of the positions I defend in this thesis.

Most of what I have to say stem from my views on punishment—i.e. our rationale for using punishment as a response to criminal wrongdoing. In Chapters 4 and 5, I defended an account of punishment where we punish an offender in order to communicate to him the moral criticism he deserves for his offence. The severity of his punishment should be graded in proportion to the level of criticism or censure that is warranted by his offence. Furthermore, the aims of communication is meant to provide a complete justification of punishment such that other considerations like retribution or deterrence are not relevant for either the justification or distribution of punishment.

So far, this is not enough to tell us how we should punish an offender. Should he be censured for what he has actually done (such that the consequences of his actions are relevant to sentencing) or simply for what he has tried to do regardless of the actual consequences? But if we take seriously the idea of punishment as apologetic reparations, through which the offender can adequately recognise his wrongdoing and take steps to reform himself, then outcomes constitute an essential component of that for which the offender needs to apologise. A reckless driver who luckily caused no injury to anyone has less to apologise for than a similarly culpable driver who did cause bodily injury to another. And the second driver has less to apologise for than a third similarly culpable driver who caused death to another.

As I pointed out in Chapter 2, the mere fact that some degree of luck was involved in the actual outcome does not necessarily mean that it would be unfair to allow luck to influence our judgments of culpability or liability. When we criticise an agent and ask him to apologise for the harm that he has wrongfully caused, our criticism is often legitimate—unless the outcome was a *sheer* matter of luck, in which case the agent would not have committed a wrong.¹ For instance, suppose I purposely throw your delicate vase on the ground. If the vase shatters, you can legitimately ask me to apologise for what I have done, in part, by getting you a new vase. I purposely broke your vase, something that I should not have done, and it would be appropriate for me to buy you a new vase.

But if, luckily, you caught the vase before it hit the floor and, despite my efforts, your vase is still intact, it would be illegitimate for you to demand that I get you a new vase. I should still apologise for what I tried to do, and I can still be asked to apologise for it, but the apology need not include getting you a new vase. It would still require other actions, such as verbally apologising for what I have done, or making a promise not to attack your property again. The verbal apology and the promise would also be appropriate in addition to buying you a new vase if I did in fact damage your vase. But to demand that I get you a new vase when I failed to damage it would be to demand an apology that is disproportionate to the wrong that has been committed.

I explained in Chapter 4 that, at least in the context of criminal law, the apology that an offender is seen to undertake through his punishment is also meant to constitute an adequate recognition of the wrong he has committed and serve as a process through which he can reform himself. For the victim who has been wronged *and harmed*—as well as the community who takes an interest in the wrong that she has suffered—what needs to be recognised is not just the wrong that the offender tried to perpetrate against her, but also the harm that was caused to her. We would not consider it sufficient or appropriate for a rapist to merely acknowledge that he *tried* to rape the victim

¹ See Duff, (1996), chapter 12.

when he in fact *raped* the victim. And for the would-be rapist who, by chance, failed to penetrate his victim, it would simply be wrong for him to acknowledge that he had raped her. The false recognition would be in excess of what he has actually done and it would distort both his action and what the victim has suffered.

The occurrence or non-occurrence of harm is also relevant for the process of self-reform. An offender who has caused harm experiences something significantly different from another similarly culpable offender who has not caused any harm. The actual outcomes of our actions, even when they are down to luck, are an essential part of our experience in the world. It is part of who we are and the narrative we tell ourselves about our lives. The harmful offender knows what it is like to have wronged *and harmed* someone, an experience that the non-harmful offender cannot fully appreciate. We might think that, having harmed someone, the harmful offender can reform himself in a shorter process than the offender who has not harmed anyone. After all, the harmful offender knows how terrible it is to have harmed another agent, and that, in itself, should be enough to make him appreciate the wrongfulness of his action and make him refrain from similar actions in the future.

Even if this is true, the harmful offender is nonetheless in greater need of the self-reformative process. He has in fact harmed someone and, in doing so, has done something worse than the offender who has not caused harm. The latter understands his action as one of trying and failing to cause harm. He can reflect on why it was wrong to try to harm others, but there is less need for him to reflect on any actual harm that has been caused (since none were caused in his case). But for the harmful offender, even if he fully appreciates (and is sensitive to) the terribleness of what he has done, he should still reflect on those things. As Duff reminds us, we may appreciate the wrongfulness of our actions immediately after the fact, but, as fallible agents, it is easy to forget such things and to carry on as if nothing has happened.² If we truly appreciate the wrongfulness of what we had done, we

² See Duff (2001, 118-21).

should voluntarily commit ourselves to an adequate period of reform, and to give ourselves enough time and space to make sure that we do not repeat our past wrongs. For the harmful offender, his appreciation of the terribleness of what he done—if he appreciates it—only confirms (to himself) the need for a greater process of self-reform.

What I have said here suggests that, insofar as we subscribe to a communicative theory of punishment, outcomes should be relevant to the severity of an offender's sentence. But that is not yet to say anything about *how* severe the sentence should be, or how much of a difference there should be between agents who have caused harm and those who have not. As I indicated in Chapter 2, some theorists have argued that we should not think that the answer to this question will be the same across purposeful actions and reckless actions. What is needed here is an examination of the kind of wrong that is perpetrated in purposeful and reckless actions, and how they relate to the issue of punishment. These questions are beyond the scope of this thesis, but, once again, they highlight further areas of research for the problem of outcome luck and for legal theory more generally.

References

- Alexander, L. & Ferzan, K. (2009). *Crime and Culpability: A Theory of Criminal Law*. Cambridge: Cambridge University Press.
- Alexander, L. (1993). "Inculpatory and Exculpatory Mistakes and the Fact/Law Distinction: An Essay in Memory of Myke Balyes." *Law and Philosophy* 12(1): 33-70.
- Alexander, L. (2011). "Duff on Attempts." In *Crime, Punishment, and Responsibility: The Jurisprudence of Antony Duff*. Eds. R. Cruft, M. H. Kramer & M. R. Reif. Oxford: Oxford University Press, 215-238.
- Ashworth, A. & Redmayne, M. (2010). *The Criminal Process*. New York: Oxford University Press.
- Ashworth, A. (1981). "The Elasticity of Mens Rea." In *Crime, Proof and Punishment: Essays in Memory of Sir Rupert Cross*. Eds. C. Tapper. London: Butterworths, 45-70.
- Ashworth, A. (1993). "Taking the Consequences." In *Action and Value in Criminal Law*. Eds. S. Shute, J. Gardner & J. Horder. Oxford: Oxford University Press, 107-124.
- Ashworth, A. (1998). "Criminal Attempts and the Role of Resulting Harm Under the Code, and in the Canon Law." *Rutgers Law Journal* 19: 725-772.
- Ballantyne, N. (2012). "Luck and Interests." *Synthese* 185(3): 319-334.
- Ballantyne, N. (2014). "Does Luck Have a Place in Epistemology?" *Synthese* 191: 1391-1407.
- Bennett, C. (2008). *The Apology Ritual*. Cambridge: Cambridge University Press.
- Bottoms, A. & von Hirsch, A. (2010). "The Crime Preventive Impact of Penal Sanctions." In *The Oxford Handbook of Empirical Legal Research*. Eds. P. Cane & H. M. Kritzer. Oxford: Oxford University Press, 96-124.
- Broncano-Berrocal, F. (2015). "Luck as Risk and the Lack of Control Account of Luck." *Metaphilosophy* 46(1): 1-25.
- Carter, J. A., & Peterson, M. (2017). "The Modal Account of Luck Revisited." *Synthese* 194(6): 2175-2184.
- Caruso, G. D. (Forthcoming). "The Public Health-Quarantine Model." In *The Oxford Handbook of Moral Responsibility*. Eds. D. Nelkin & D. Pereboom. New York: Oxford University Press.

- Chalmers, J. & Leverick, F. (2008). "Fair Labelling in Criminal Law." *Modern Law Review* 71(2): 217-246.
- Clarkson, C. M. V. (2005). "General Endangerment Offences: The Way Forward?" *University of Western Australia Law Review* 32: 131-144.
- Coffman, E. J. (2007). "Thinking About Luck." *Synthese* 158(3): 385-398.
- Cornford, A. (2011). "Resultant Luck and Criminal Liability." In *The Structures of the Criminal Law*. Eds. R. A. Duff, L. Farmer, S. E. Marshall, M. Renzo & V. Tadros. Oxford: Oxford University Press, 36-58.
- Duff, R. A. & Marshall, S. E. (2016). "Civic Punishment." In *Democratic Theory and Mass Incarceration*. Eds. A. Dzur, I. Loader & R. Sparks. New York: Oxford University Press, 33-58.
- Duff, R. A. (1981). "Recklessness and Rape." *Liverpool Law Review* 3(2): 49-64.
- Duff, R. A. (1990). *Intention, Agency and Criminal Liability: Philosophy of Action and the Criminal Law*. Oxford: Basil Blackwell.
- Duff, R. A. (1991). "The Circumstances of an Attempt." *Cambridge Law Journal* 50(1): 100-119.
- Duff, R. A. (1995). "Recklessness in Attempts (Again)." *Oxford Journal of Legal Studies* 15(2): 309-325.
- Duff, R. A. (1996). *Criminal Attempts*. Oxford: Oxford University Press.
- Duff, R. A. (2001). *Punishment, Communication, and Community*. Oxford: Oxford University Press.
- Duff, R. A. (2003a). "Punishment, Communication and Community." In *Debates in Contemporary Political Philosophy: An Anthology*. Eds. D. Matravers & J. Pike. London: Routledge, 387-407.
- Duff, R. A. (2003b). "Response to von Hirsch." In *Debates in Contemporary Political Philosophy: An Anthology*. Eds. D. Matravers & J. Pike. London: Routledge, 423-427.
- Duff, R. A. (2003c). "Restoration and Retribution." In *Restorative Justice and Criminal Justice: Competing or Reconcilable Paradigms?* Eds. A. E. Bottoms, K. Roach, J. V. Roberts, M. Schiff & A. von Hirsch. Oxford: Hart Publishing, 43-60.
- Duff, R. A. (2007). *Answering for Crime: Responsibility and Liability in the Criminal Law*. Oxford: Hart Publishing.

- Duff, R. A. (2008). "Criminalizing Endangerment." *Louisiana Law Review* 65: 941-965.
- Duff, R. A. (2010). "In Response." In *Crime, Punishment, and Responsibility: The Jurisprudence of Antony Duff*. Eds. R. Cruft, M. H. Kramer & M. R. Reif. Oxford: Oxford University Press, 351-379.
- Duff, R. A. (2012). "Guiding Commitments and Criminal Liability for Attempts." *Criminal Law and Philosophy* 6(3): 411-427.
- Duff, R. A. (2016). "Whose Luck is it Anyway?" In *Criminal Liability for Non-Aggressive Death*. Eds. C. M. V. Clarkson & S. Cunningham. London: Routledge, 73-90.
- Edwards, J. R. & Simester, A. P. (2014). "Prevention with a Moral Voice." In *Liberal Criminal Theory: Essays for Andreas von Hirsch*. Eds. A. du Bois-Pedain, U. Neumann & A. P. Simester. Oxford: Hart Publishing, 43-65.
- Enoch, D. & Marmor, A. (2007). "The Case Against Moral Luck." *Law and Philosophy* 26(4): 405-436.
- Enoch, D. (2010). "Moral Luck and the Law." *Philosophy Compass* 5(1): 42-54.
- Farrell, D. (1989). "On Threats and Punishments." *Social Theory and Practice* 15(2): 125-154.
- Feinberg, J. (1965). "The Expressive Function of Punishment." *The Monist* 49(3): 397-423.
- Feinberg, J. (1995). "Equal Punishments for Failed Attempts: Some Bad but Instructive Arguments Against It." *Arizona Law Review* 37: 117-133.
- Ferzan, K. (2009). "The Structure of Criminal Law." *Criminal Justice Ethics* 28(2): 223-237.
- Fletcher, G. P. (1986). "Constructing a Theory of Impossible Attempts." *Criminal Justice Ethics* 5(1): 53-69.
- Glasgow, J. (2015). "The Expressivist Theory of Punishment Defended." *Law and Philosophy* 34(6): 601-631.
- Greco, J. (1995). "A Second Paradox Concerning Responsibility and Luck." *Metaphilosophy* 26(1-2): 81-96.

- Hacking, I. (2001). *An Introduction to Probability and Inductive Logic*. Cambridge: Cambridge University Press.
- Hampton, J. (1992). "Correcting Harms Versus Righting Wrongs: The Goal of Retribution." *UCLA Law Review* 39: 1659-1702.
- Hart, H. L. A. (1959). "The Presidential Address: Prolegomenon to the Principles of Punishment." *Proceedings of the Aristotelian Society* 60: 1-26.
- Hart, H. L. A. (1981). "The House of Lords on Attempting the Impossible." *Oxford Journal of Legal Studies* 1(2): 149-166.
- Hart, H. L. A. (2008). *Punishment and Responsibility: Essays in the Philosophy of Law*. 2nd ed. Oxford: Oxford University Press.
- Hoskins, Z. (2011). "Deterrent Punishment and Respect for Persons." *Ohio State Journal of Criminal Law* 8: 369-384.
- Husak, D. (1992). "Why Punish the Deserving?" *Noûs* 26(4): 447-464.
- Kadish, S. H. (1994). "The Criminal Law and the Luck of the Draw." *The Journal of Criminal Law and Criminology* 84: 679-702.
- Latus, A. (2000). "Moral and Epistemic Luck." *Journal of Philosophical Research* 25: 149-172.
- Levy, N. (2011). *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: Oxford University Press.
- Lewis, D. K. (1979). "Counterfactual Dependence and Time's Arrow." *Noûs* 13(4): 455-476.
- Lewis, D. K. (1989). "Punishment that Leaves Something to Chance." *Philosophy and Public Affairs* 18(1): 53-67.
- Matravers, M. (2011). "Duff on Hard Treatment." In *Crime, Punishment, and Responsibility: The Jurisprudence of Antony Duff*. Eds. R. Cruft, M. H. Kramer & M. R. Reif. Oxford: Oxford University Press, 68-83.
- McKinnon, R. (2013). "Getting Luck Properly Under Control." *Metaphilosophy* 44(4): 496-511.
- Mellor, D. H. (2004). *Probability: A Philosophical Introduction*. New York: Routledge.
- Moore, M. S. (1994). "The Independent Moral Significance of Wrongdoing." *Journal of Contemporary Legal Issues* 5: 237-281.

- Moore, M. S. (2009). *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics*. Oxford: Oxford University Press.
- Nagel, T. & Williams, B. A. O. (1976). "Moral Luck." *Proceedings of the Aristotelian Society Supplementary Volumes* 50: 115-151.
- Nagel, T. (1979). "Moral Luck." In *Mortal Questions*. Cambridge: Cambridge University Press, 24-38.
- Narayan, U. (1993). "Appropriate Responses and Preventive Benefits: Justifying Censure and Hard Treatment in Legal Punishment." *Oxford Journal of Legal Studies* 13: 166-182.
- Ohana, D. (2007). "Desert and Punishment for Acts Preparatory to the Commission of a Crime." *Canadian Journal of Law and Jurisprudence* 20(1): 113-142.
- Peels, R. (2015). "A Modal Solution to the Problem of Moral Luck." *American Philosophical Quarterly* 52(1): 73-87.
- Pritchard, D. & Smith, M. (2004). "The Psychology and Philosophy of Luck." *New Ideas in Psychology* 22(1): 1-28.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, D. (2007a). "Knowledge, Luck, and lotteries." In *New Waves in Epistemology*. Eds. V. Hendricks & D. Pritchard. London: Palgrave Macmillan, 28-51.
- Pritchard, D. (2007b). "Anti-Luck Epistemology." *Synthese* 158(3): 277-297.
- Pritchard, D. (2014). "The Modal Account of Luck." *Metaphilosophy* 45(4-5): 594-619.
- Pritchard, D. (2015). "Risk." *Metaphilosophy* 46(3): 436-461.
- Pritchard, D. (2016). "Epistemic Risk." *The Journal of Philosophy* 113(11): 550-571.
- Rawls, J. (1955). "Two Concepts of Rules." *The Philosophical Review* 64(1): 3-32.
- Rescher, N. (2001). *Luck: The Brilliant Randomness of Everyday Life*. Pittsburgh: University of Pittsburgh Press.
- Rescher, N. (2014). "The Machinations of Luck." *Metaphilosophy* 45(4-5): 620-626.

- Riggs, W. (2009). "Knowledge, Luck, and Control." In *Epistemic Value*. Eds. A. Haddock, A. Millar & D. Pritchard. Oxford: Oxford University Press, 204-221.
- Robinson, P. H. (1996). "The Criminal-Civil Distinction and the Utility of Desert." *Boston University Law Review* 76: 201-214.
- Robinson, P. H. & Darley, J. M. (1996). "Utility of Desert." *Northwestern University Law Review* 91: 453-499.
- Roxin, C. (2014). "Prevention, Censure and Responsibility: The Recent Debate on the Purposes of Punishment." In *Liberal Criminal Theory: Essays for Andreas von Hirsch*. Eds. A. du Bois-Pedain, U. Neumann & A. P. Simester. Oxford: Hart Publishing, 23-42.
- Simester, A. P., Spencer, J. R., Sullivan, G. R., & Virgo, G. J. (2013). *Simester and Sullivan's Criminal Law: Theory and Doctrine*. 5th ed. Oxford: Hart Publishing.
- Simester, A. P., Spencer, J. R., Sullivan, G. R., & Virgo, G. J. (2016). *Simester and Sullivan's Criminal Law: Theory and Doctrine*. 6th ed. Oxford: Hart Publishing.
- Simester, A. P. (1996). "Why Distinguish Intention from Foresight?" In *Harm and Culpability*. Eds. A. P. Simester & A. T. H. Smith. Oxford: Clarendon Press, 71-102.
- Steglich-Peterson, A. (2010). "Luck as an Epistemic Notion." *Synthese* 176(3): 361-377.
- Sverdlik, S. (1988). "Crime and Moral Luck." *American Philosophical Quarterly* 25(1): 79-86.
- Tadros, V. (2011). *The Ends of Harm*. Oxford: Oxford University Press.
- von Hirsch, A. (1993). *Censure and Sanctions*. Oxford: Oxford University Press.
- von Hirsch, A. (2003). "Punishment, Penance and the State: A Reply to Duff." In *Debates in Contemporary Political Philosophy: An Anthology*. Eds. D. Matravers & J. Pike. London: Routledge, 408-422.
- Westen, P. (2008). "Impossibility Attempts." *Ohio State Journal of Criminal Law* 5: 523-565.
- Whittington, E. J. (2017). "Counterfactual Chanciness and Risk: Giving Possibility a Chance." Unpublished Manuscript.

- Williams, B. A. O. (1981). "Moral Luck." In *Moral Luck*. Cambridge: Cambridge University Press, 20-39.
- Yaffe, G. (2010). *Attempts: In the Philosophy of Action and Criminal Law*. Oxford: Oxford University Press.
- Zimmerman, M. J. (1987). "Luck and Moral Responsibility." *Ethics* 97(2): 374-386.
- Zimmerman, M. J. (2002). "Taking Luck Seriously." *The Journal of Philosophy* 99(11): 553-576.

